Contents

Third-party copyrights

Exchange admin center keyboard shortcuts

Privacy statement

Exchange content updates

8/3/2020 • 41 minutes to read • Edit Online

This topic lists Exchange Server and Exchange Online topics that have been changed over the last several weeks.

Week of July 06, 2020

PUBLISHED ON	TOPIC TITLE	CHANGE
6/30/2020	Exchange Server hybrid deployments	modified
6/30/2020	Configure Microsoft 365 Groups with on-premises Exchange hybrid	modified
6/30/2020	Enable or disable SMTP AUTH	modified
6/30/2020	Fix email delivery issues for error code 451 4.7.500-699 (ASxxx) in Exchange Online	modified
6/30/2020	Use Directory Based Edge Blocking to reject messages sent to invalid recipients	modified
6/30/2020	Change how long permanently deleted items are kept for an Exchange Online mailbox	modified
6/30/2020	Clean up or delete items from the Recoverable Items folder	modified
6/30/2020	About Exchange documentation	modified
6/30/2020	Accessibility for people with disabilities	modified
6/30/2020	Third-party copyright notices	modified
6/30/2020	The Best Practices Analyzer for Exchange Server is no longer available	modified
6/30/2020	Exchange ActiveSync device settings with Exchange hybrid deployments	modified
6/30/2020	Certificate requirements for hybrid deployments	modified
6/30/2020	How and when to decommission your on-premises Exchange servers in a hybrid deployment	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
6/30/2020	Edge Transport servers with hybrid deployments	modified
6/30/2020	Edge Transport servers in Exchange 2013/Exchange 2007 hybrid deployments	modified
6/30/2020	Hybrid deployments with Exchange 2013 and Exchange 2007	modified
6/30/2020	Hybrid management in Exchange 2013/Exchange 2007 hybrid deployments	modified
6/30/2020	Server roles in Exchange 2013/Exchange 2007 hybrid deployments	modified
6/30/2020	Transport options in Exchange 2013/Exchange 2007 hybrid deployments	modified
6/30/2020	Transport routing in Exchange 2013/Exchange 2007 hybrid deployments	modified
6/30/2020	Edge Transport servers in Exchange 2013/Exchange 2010 hybrid deployments	modified
6/30/2020	Hybrid deployments with Exchange 2013 and Exchange 2010	modified
6/30/2020	Hybrid management in Exchange 2013/Exchange 2010 hybrid deployments	modified
6/30/2020	Server roles in Exchange 2013/Exchange 2010 hybrid deployments	modified
6/30/2020	Transport options in Exchange 2013/Exchange 2010 hybrid deployments	modified
6/30/2020	Transport routing in Exchange 2013/Exchange 2010 hybrid deployments	modified
6/30/2020	Hybrid Configuration wizard FAQs	modified
6/30/2020	Hybrid Configuration wizard options	modified
6/30/2020	Hybrid Configuration wizard	modified
6/30/2020	Hybrid deployment prerequisites	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
6/30/2020	Create a cloud-based archive for an on-premises primary mailbox in an Exchange hybrid deployment	modified
6/30/2020	Create a hybrid deployment with the Hybrid Configuration wizard	modified
6/30/2020	Microsoft Hybrid Agent	modified
6/30/2020	Hybrid deployment procedures	modified
6/30/2020	Hybrid deployments with multiple forests	modified
6/30/2020	Move mailboxes between on-premises and Exchange Online organizations in hybrid deployments	modified
6/30/2020	Configure Exchange to support delegated mailbox permissions in a hybrid deployment	modified
6/30/2020	Configure document collaboration with OneDrive for Business and Exchange 2016 on-premises	modified
6/30/2020	Configure Exchange Online public folders for a hybrid deployment	modified
6/30/2020	Configure legacy on-premises public folders for a hybrid deployment	modified
6/30/2020	Configure Exchange Server public folders for a hybrid deployment	modified
6/30/2020	Simplify the Outlook Web App URL for Microsoft 365 or Office 365 Hybrid	modified
6/30/2020	Troubleshoot a hybrid deployment	modified
6/30/2020	Hybrid management in Exchange hybrid deployments	modified
6/30/2020	IRM in Exchange hybrid deployments	modified
6/30/2020	Use the Microsoft 365 and Office 365 mail migration advisor	modified
6/30/2020	ActiveSync device access rule attributes	modified
6/30/2020	Active Sync Mailbox Policy attributes	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
6/30/2020	ActiveSync organization settings attributes	modified
6/30/2020	Address list attributes	modified
6/30/2020	Dlp Policy attributes	modified
6/30/2020	Malware Filter policy attributes	modified
6/30/2020	Mobile Device Mailbox Policy attributes	modified
6/30/2020	Organization configuration transfer attributes	modified
6/30/2020	Organization Config	modified
6/30/2020	OWA Mailbox Policy attributes	modified
6/30/2020	Policy tip config attributes	modified
6/30/2020	Remote Domains attributes	modified
6/30/2020	Retention Policy Tags attributes	modified
6/30/2020	Retention Policy attributes	modified
6/30/2020	Sharing Policy attributes	modified
6/30/2020	Smime Config attributes	modified
6/30/2020	Transport Config attributes	modified
6/30/2020	Organization Configuration Transfer Attributes	modified
6/30/2020	Performance factors and best practices for hybrid migrations	modified
6/30/2020	Permissions in Exchange hybrid deployments	modified
6/30/2020	Server roles in Exchange hybrid deployments	modified
6/30/2020	Shared free/busy in Exchange hybrid deployments	modified
6/30/2020	Single sign-on with hybrid deployments	modified
6/30/2020	Transport options in Exchange hybrid deployments	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
6/30/2020	Transport routing in Exchange hybrid deployments	modified
6/30/2020	About Exchange documentation	modified
6/30/2020	Accessibility in the Exchange admin center in Exchange Online	modified
6/30/2020	Accessibility in Exchange Online	modified
6/30/2020	Get started using a screen reader in the Exchange admin center in Exchange Online	modified
6/30/2020	Keyboard shortcuts for the Exchange admin center in Exchange Online	modified
6/30/2020	Use a screen reader to manage anti-spam protection in Exchange Online	modified
6/30/2020	Use a screen reader to add a new equipment mailbox in the Exchange admin center	modified
6/30/2020	Use a screen reader to add a new mail contact in the Exchange admin center in Exchange Online	modified
6/30/2020	Use a screen reader to add members to a distribution group in the Exchange admin center in Exchange Online	modified
6/30/2020	Use a screen reader to add a new room mailbox in the Exchange admin center in Exchange Online	modified
6/30/2020	Use a screen reader to add a new shared mailbox in the Exchange admin center in Exchange Online	modified
6/30/2020	Use a screen reader to archive mailbox items in the Exchange admin center in Exchange Online	modified
6/30/2020	Use a screen reader to configure collaboration in the Exchange admin center in Exchange Online	modified
6/30/2020	Use a screen reader to configure mail flow rules in the Exchange admin center in Exchange Online	modified
6/30/2020	Use a screen reader to configure voice mail in the Exchange admin center in Exchange Online	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
6/30/2020	Use a screen reader to create a new distribution group in the Exchange admin center in Exchange Online	modified
6/30/2020	Use a screen reader to define rules that encrypt or decrypt email messages in the Exchange admin center in Exchange Online	modified
6/30/2020	Use a screen reader to edit the mailbox display name in the Exchange admin center in Exchange Online	modified
6/30/2020	Use a screen reader to export and review audit logs in the Exchange admin center	modified
6/30/2020	Use a screen reader to identify your admin role in the Exchange admin center	modified
6/30/2020	Use a screen reader to manage anti-malware protection in the Exchange admin center in Exchange Online	modified
6/30/2020	Use a screen reader to open the Exchange admin center in Exchange Online	modified
6/30/2020	Use a screen reader to run an audit report in the Exchange admin center in Exchange Online	modified
6/30/2020	Use a screen reader to trace an email message in the Exchange admin center in Exchange Online	modified
6/30/2020	Use a screen reader to work with mobile clients in the Exchange admin center in Exchange Online	modified
6/30/2020	Address book policies in Exchange Online	modified
6/30/2020	Address book policy procedures in Exchange Online	modified
6/30/2020	Assign an address book policy to users in Exchange Online	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
6/30/2020	Change the settings of an address book policy in Exchange Online	modified
6/30/2020	Create an address book policy in Exchange Online	modified
6/30/2020	Remove an address book policy in Exchange Online	modified
6/30/2020	Turn on address book policy routing in Exchange Online	modified
6/30/2020	Address books in Exchange Online	modified
6/30/2020	Address list procedures in Exchange Online	modified
6/30/2020	Address lists in Exchange Online	modified
6/30/2020	Configure global address list properties in Exchange Online	modified
6/30/2020	Create a global address list in Exchange Online	modified
6/30/2020	Manage address lists in Exchange Online	modified
6/30/2020	Remove a global address list in Exchange Online	modified
6/30/2020	Recipient filters for address lists in Exchange Online PowerShell	modified
6/30/2020	Enable or disable hierarchical address books in Exchange Online	modified
6/30/2020	Hierarchical address books in Exchange Online	modified
6/30/2020	Add an address list to or remove an address list from an offline address book in Exchange Online	modified
6/30/2020	Change the default offline address book in Exchange Online	modified
6/30/2020	Configure offline address book distribution properties	modified
6/30/2020	Create an offline address book	modified
6/30/2020	Offline address book procedures	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
6/30/2020	Offline address books in Exchange Online	modified
6/30/2020	Provision recipients for offline address book downloads in Exchange Online	modified
6/30/2020	Remove an offline address book from Exchange	modified
6/30/2020	Back up email in Exchange Online	modified
6/30/2020	Add-ins for Outlook in Exchange Online	modified
6/30/2020	Install or remove add-ins for Outlook for your Exchange Online organization	modified
6/30/2020	Manage user access to add-ins for Outlook in Exchange Online	modified
6/30/2020	Using third-party add-ins for online meetings in Outlook for iOS and Android	modified
6/30/2020	Specify the administrators and users who can install and manage add-ins for Outlook	modified
6/30/2020	Client Access Rules in Exchange Online	modified
6/30/2020	Procedures for Client Access Rules in Exchange Online	modified
6/30/2020	Clients and mobile in Exchange Online	modified
6/30/2020	Disable Basic authentication in Exchange Online	modified
6/30/2020	Enable or disable modern authentication for Outlook in Exchange Online	modified
6/30/2020	Exchange ActiveSync in Exchange Online	modified
6/30/2020	Mobile device mailbox policies in Exchange Online	modified
6/30/2020	Configure custom MailTips for recipients	modified
6/30/2020	Configure the large audience size for your organization	modified
6/30/2020	MailTips over organization relationships	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
6/30/2020	MailTips	modified
6/30/2020	Manage MailTips for organization relationships	modified
6/30/2020	Configure mobile phones to access email	modified
6/30/2020	Mobile access in Exchange Online	modified
6/30/2020	Perform a remote wipe on a mobile phone	modified
6/30/2020	Managing Outlook for iOS and Android in Exchange Online	modified
6/30/2020	Deploying Outlook for iOS and Android app configuration settings	modified
6/30/2020	Outlook for iOS and Android in Exchange Online: FAQ	modified
6/30/2020	Using Outlook for iOS and Android in the Government Community Cloud	modified
6/30/2020	Outlook for iOS and Android in Exchange Online	modified
6/30/2020	Securing Outlook for iOS and Android in Exchange Online	modified
6/30/2020	Sensitivity labeling and protection in Outlook for iOS and Android	modified
6/30/2020	Account setup with modern authentication in Exchange Online	modified
6/30/2020	Apply or remove an Outlook on the web mailbox policy on a mailbox in Exchange Online	modified
6/30/2020	View or configure Outlook on the web mailbox policy properties in Exchange Online	modified
6/30/2020	Create an Outlook on the web mailbox policy in Exchange Online	modified
6/30/2020	Modify the space used by Inbox rules in Exchange Online	modified
6/30/2020	Outlook on the web in Exchange Online	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
6/30/2020	Outlook on the web mailbox policies in Exchange Online	modified
6/30/2020	Outlook on the web mailbox policy procedures in Exchange Online	modified
6/30/2020	Public attachment handling in Exchange Online	modified
6/30/2020	Remove an Outlook on the web mailbox policy from Exchange Online	modified
6/30/2020	Enable or Disable POP3 or IMAP4 access for a user	modified
6/30/2020	POP3 and IMAP4	modified
6/30/2020	Set POP3 or IMAP4 settings for a user	modified
6/30/2020	Remote Connectivity Analyzer tests for Exchange Online	modified
6/30/2020	Collaboration in Exchange Online	modified
6/30/2020	Accessing public folders with Outlook 2016 for Mac	modified
6/30/2020	Assign "Send As" or "Send on Behalf" permissions for mail-enabled public folders	modified
6/30/2020	Use batch migration to migrate Exchange Online public folders to Microsoft 365 Groups	modified
6/30/2020	Use batch migration to migrate legacy public folders to Microsoft 365 or Office 365 and Exchange Online	modified
6/30/2020	Create a public folder mailbox	modified
6/30/2020	Create a public folder	modified
6/30/2020	Mail-enable or mail-disable a public folder	modified
6/30/2020	Migrate your public folders to Microsoft 365 Groups	modified
6/30/2020	Public folder procedures in Microsoft 365, Office 365, and Exchange Online	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
6/30/2020	Public folders in Microsoft 365, Office 365, and Exchange Online	modified
6/30/2020	Recover a deleted public folder mailbox	modified
6/30/2020	Remove a public folder	modified
6/30/2020	Restore a deleted public folder	modified
6/30/2020	Configure Exchange Online public folders for a hybrid deployment	modified
6/30/2020	Configure legacy on-premises public folders for a hybrid deployment	modified
6/30/2020	Configure Exchange Server public folders for a hybrid deployment	modified
6/30/2020	Set up public folders in a new organization	modified
6/30/2020	Update the public folder hierarchy	modified
6/30/2020	Use favorite public folders in Outlook on the web	modified
6/30/2020	View statistics for public folders and public folder items	modified
6/30/2020	Shared mailboxes in Exchange Online	modified
6/30/2020	Exchange admin center in Exchange Online	modified
6/30/2020	Exchange Online	modified
6/30/2020	Configure the external postmaster address in Exchange Online	modified
6/30/2020	Fix issues with printers, scanners, and LOB applications that send email using Microsoft 365 or Office 365	modified
6/30/2020	How to set up a multifunction device or application to send email using Microsoft 365 or Office 365	modified
6/30/2020	Mail flow best practices for Exchange Online, Microsoft 365, and Office 365 (overview)	modified
6/30/2020	Enable mail flow for subdomains in Exchange Online	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
6/30/2020	Manage accepted domains in Exchange Online	modified
6/30/2020	Manage mail flow with mailboxes in multiple locations (Exchange Online and on-premises)	modified
6/30/2020	Manage mail flow using a third-party cloud service with Exchange Online and on-premises mailboxes	modified
6/30/2020	Manage mail flow using a third-party cloud service with Exchange Online	modified
6/30/2020	Manage all mailboxes and mail flow using Microsoft 365 or Office 365	modified
6/30/2020	Message format and transmission in Exchange Online	modified
6/30/2020	Fix email delivery issues for error code 5.1.8 in Exchange Online	modified
6/30/2020	Fix email delivery issues for error code 5.4.6 or 5.4.14 in Exchange Online	modified
6/30/2020	Fix email delivery issues for error code 5.7.12 in Exchange Online	modified
6/30/2020	Fix email delivery issues for error code 5.7.124 in Exchange Online	modified
6/30/2020	Fix email delivery issues for error code 5.7.13 or 5.7.135 in Exchange Online	modified
6/30/2020	Fix email delivery issues for error code 5.7.133 in Exchange Online	modified
6/30/2020	Fix email delivery issues for error code 5.7.134 in Exchange Online	modified
6/30/2020	Fix email delivery issues for error code 5.7.136 in Exchange Online	modified
6/30/2020	Fix email delivery issues for error code 5.7.23 in Exchange Online	modified
6/30/2020	Fix email delivery issues for error code 5.7.57 in Exchange Online	modified
6/30/2020	Fix email delivery issues for error code 5.7.64 in Exchange Online	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
6/30/2020	Fix email delivery issues for error codes 5.7.700 through 5.7.750 in Exchange Online	modified
6/30/2020	Fix email delivery issues for error code 550 4.4.7 in Exchange Online	modified
6/30/2020	Fix email delivery issues for error code 550 5.0.350 in Exchange Online	modified
6/30/2020	Fix email delivery issues for error code 550 5.1.0 in Exchange Online	modified
6/30/2020	Fix email delivery issues for error code 550 5.1.1 through 5.1.20 in Exchange Online	modified
6/30/2020	Fix email delivery issues for error code 550 5.1.10 in Exchange Online	modified
6/30/2020	Fix email delivery issues for error code 550 5.4.1 in Exchange Online	modified
6/30/2020	Fix email delivery issues for error code 550 5.6.11 in Exchange Online	modified
6/30/2020	Fix email delivery issues for error code 550 5.7.1 in Exchange Online	modified
6/30/2020	Email non-delivery reports in Exchange Online	modified
6/30/2020	Manage remote domains in Exchange Online	modified
6/30/2020	Remote domains in Exchange Online	modified
6/30/2020	Supported character sets for remote domains in Exchange Online	modified
6/30/2020	Test mail flow by validating your connectors	modified
6/30/2020	Troubleshoot mail flow	modified
6/30/2020	Scenario Conditional mail routing in Exchange Online	modified
6/30/2020	Do I need to create a connector in Exchange Online?	modified
6/30/2020	Enhanced filtering for connectors	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
6/30/2020	Scenario Integrate Microsoft 365 or Office 365 with an email add-on service	modified
6/30/2020	Set up connectors for secure mail flow with a partner organization	modified
6/30/2020	Set up connectors to route mail between Microsoft 365 or Office 365 and your own email servers	modified
6/30/2020	Configure mail flow using connectors	modified
6/30/2020	Validate connectors	modified
6/30/2020	Add an SSL certificate to Exchange 2007	modified
6/30/2020	Add an SSL certificate to Exchange 2010	modified
6/30/2020	Add an SSL certificate to Exchange 2013	modified
6/30/2020	Assign Exchange permissions to migrate mailboxes to Microsoft 365 or Office 365	modified
6/30/2020	CSV files for mailbox migration	modified
6/30/2020	Migrate email using the Exchange cutover method	modified
6/30/2020	Decide on a migration path	modified
6/30/2020	Enable your Gmail account for IMAP	modified
6/30/2020	Ways to migrate multiple email accounts to Microsoft 365 or Office 365	modified
6/30/2020	Manage migration batches	modified
6/30/2020	Migrate from Lotus Notes to Microsoft 365 or Office 365	modified
6/30/2020	How to migrate mailboxes from one Microsoft 365 or Office 365 organization to another	modified
6/30/2020	CSV files for IMAP migration batches	modified
6/30/2020	Enable 2-step verification for your Google apps users	modified
6/30/2020	IMAP migration in the Microsoft 365 admin center	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
6/30/2020	Migrate consumer G Suite mailboxes to Microsoft 365 or Office 365	modified
6/30/2020	Migrate other types of IMAP mailboxes to Microsoft 365 or Office 365	modified
6/30/2020	What you need to know about migrating your IMAP mailboxes to Microsoft 365 or Office 365	modified
6/30/2020	Migrating your Outlook.com account to Microsoft 365 or Office 365	modified
6/30/2020	Tips for optimizing IMAP migrations	modified
6/30/2020	Prepare your Gmail or G Suite account for connecting to Outlook and Microsoft 365 or Office 365	modified
6/30/2020	Learn more about setting up your IMAP server connection	modified
6/30/2020	Migration users status report	modified
6/30/2020	Microsoft 365 and Office 365 migration performance and best practices	modified
6/30/2020	Perform a G Suite migration	modified
6/30/2020	Convert Exchange 2003 mailboxes to mail-enabled users	modified
6/30/2020	Convert Exchange 2007 mailboxes to mail-enabled users	modified
6/30/2020	Perform a staged migration of email	modified
6/30/2020	Plan to coexist with a third-party messaging system using Active Directory Domain Services	modified
6/30/2020	Track and Prevent Migration Data Loss	modified
6/30/2020	Use Minimal Hybrid to quickly migrate Exchange mailboxes to Microsoft 365 or Office 365	modified
6/30/2020	What you need to know about a cutover email migration	modified
6/30/2020	What you need to know about a staged email migration	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
6/30/2020	Customize and schedule mail protection reports to be automatically sent to your inbox	modified
6/30/2020	Monitoring, reporting, and message tracing in Exchange Online	modified
6/30/2020	Message Trace FAQ	modified
6/30/2020	Run a message trace and view the results in the Exchange admin center	modified
6/30/2020	Trace an email message	modified
6/30/2020	Use mail protection reports to view data about malware, spam, and rule detections	modified
6/30/2020	What happened to delivery reports?	modified
6/30/2020	Feature permissions in Exchange Online	modified
6/30/2020	Permissions in Exchange Online	modified
6/30/2020	Role assignment policies in Exchange Online	modified
6/30/2020	Manage role groups in Exchange Online	modified
6/30/2020	Configure a moderated recipient in Exchange Online	modified
6/30/2020	Create user mailboxes in Exchange Online	modified
6/30/2020	Delete or restore user mailboxes in Exchange Online	modified
6/30/2020	Create a distribution group naming policy	modified
6/30/2020	Create and manage distribution groups	modified
6/30/2020	Override the distribution group naming policy	modified
6/30/2020	Manage dynamic distribution groups	modified
6/30/2020	View members of a dynamic distribution group	modified
6/30/2020	Manage equipment mailboxes	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
6/30/2020	Manage Facebook contact sync in your organization	modified
6/30/2020	Manage guest access to Microsoft 365 groups	modified
6/30/2020	Manage LinkedIn contact sync in your organization	modified
6/30/2020	Manage mail contacts	modified
6/30/2020	Manage mail-enabled security groups	modified
6/30/2020	Manage mail users	modified
6/30/2020	Manage permissions for recipients in Exchange Online	modified
6/30/2020	Create and manage room mailboxes	modified
6/30/2020	Add or remove email addresses for a mailbox	modified
6/30/2020	Automatically save sent items in delegator's mailbox	modified
6/30/2020	Change the branding of Clutter notifications	modified
6/30/2020	Clutter notifications in Outlook	modified
6/30/2020	Configure email forwarding for a mailbox	modified
6/30/2020	Configure message delivery restrictions for a mailbox	modified
6/30/2020	Convert a mailbox	modified
6/30/2020	Enable or disable Exchange ActiveSync for a mailbox	modified
6/30/2020	Enable or disable MAPI for a mailbox	modified
6/30/2020	Enable or disable Outlook on the web for a mailbox	modified
6/30/2020	Enable or disable single item recovery for a mailbox	modified
6/30/2020	Mailbox plans in Exchange Online	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
6/30/2020	Manage user mailboxes	modified
6/30/2020	Recover deleted messages in a user's mailbox	modified
6/30/2020	Use Exchange Online PowerShell to display mailbox information	modified
6/30/2020	Message and recipient limits in Exchange Online	modified
6/30/2020	Recipients in Exchange Online	modified
6/30/2020	Create or remove an In-Place Hold	modified
6/30/2020	Create a custom DLP policy	modified
6/30/2020	Create a DLP policy from a template	modified
6/30/2020	Data loss prevention	modified
6/30/2020	DLP policy templates supplied in Exchange	modified
6/30/2020	How DLP rules are applied to evaluate messages	modified
6/30/2020	Integrating sensitive information rules with mail flow rules in Exchange Online	modified
6/30/2020	Manage policy tips	modified
6/30/2020	Policy Tips in Exchange Online	modified
6/30/2020	Exchange Online auditing reports	modified
6/30/2020	Export mailbox audit logs	modified
6/30/2020	Run a non-owner mailbox access report	modified
6/30/2020	Run a per-mailbox litigation hold report	modified
6/30/2020	Search the role group changes or administrator audit logs in Exchange Online	modified
6/30/2020	View the administrator audit log	modified
6/30/2020	View and export the external admin audit log	modified
6/30/2020	In-Place Hold and Litigation Hold	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
6/30/2020	Assign eDiscovery permissions in Exchange	modified
6/30/2020	Create a discovery mailbox	modified
6/30/2020	Create a custom management scope for In-Place eDiscovery searches	modified
6/30/2020	Create an In-Place eDiscovery search	modified
6/30/2020	Delete and re-create the default discovery mailbox in Exchange	modified
6/30/2020	Export eDiscovery search results to a PST file	modified
6/30/2020	In-Place eDiscovery	modified
6/30/2020	Message properties and search operators for In-Place eDiscovery	modified
6/30/2020	Reduce the size of a discovery mailbox in Exchange	modified
6/30/2020	Search limits for In-Place eDiscovery in Exchange Online	modified
6/30/2020	Configure Journaling in Exchange Online	modified
6/30/2020	Journaling in Exchange Online	modified
6/30/2020	Manage journaling in Exchange Online	modified
6/30/2020	Common attachment blocking scenarios for mail flow rules in Exchange Online	modified
6/30/2020	Common message approval scenarios in Exchange Online	modified
6/30/2020	Mail flow rule conditions and exceptions (predicates) in Exchange Online	modified
6/30/2020	Best practices for configuring mail flow rules in Exchange Online	modified
6/30/2020	Organization-wide message disclaimers, signatures, footers, or headers in Exchange Online	modified
6/30/2020	Enable message encryption and decryption	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
6/30/2020	Use mail flow rules to inspect message attachments in Exchange Online	modified
6/30/2020	Mail flow rule actions in Exchange Online	modified
6/30/2020	Mail flow rule procedures in Exchange Online	modified
6/30/2020	Mail flow rules (transport rules) in Exchange Online	modified
6/30/2020	Manage mail flow rules in Exchange Online	modified
6/30/2020	Manage message approval in Exchange Online	modified
6/30/2020	Test a mail flow rule in Exchange Online	modified
6/30/2020	Manage and troubleshoot message approval in Exchange Online	modified
6/30/2020	Use mail flow rules to automatically add meetings to calendars in Exchange Online	modified
6/30/2020	Use mail flow rules so messages can bypass Clutter in Exchange Online	modified
6/30/2020	Use mail flow rules to route email based on a list of words, phrases, or patterns in Exchange Online	modified
6/30/2020	Add retention tags to or remove retention tags from a retention policy	modified
6/30/2020	Apply a retention policy to mailboxes	modified
6/30/2020	Create a Retention Policy	modified
7/1/2020	Sensitivity labeling and protection in Outlook for iOS and Android	modified
7/2/2020	Clean up or delete items from the Recoverable Items folder	modified

Week of July 06, 2020

PUBLISHED ON	TOPIC TITLE	CHANGE
7/7/2020	CSV files for mailbox migration	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
7/7/2020	Perform a G Suite migration	modified
7/7/2020	Manage on-premises mailbox moves in Exchange Server	modified
7/7/2020	Use batch migration to migrate Exchange 2010 public folders to Exchange 2016	modified
7/7/2020	Use batch migration to migrate Exchange Server public folders to Microsoft 365 Groups	modified
7/7/2020	Exchange Deployment Assistant	modified
7/7/2020	CSV files for mailbox migration: Exchange 2013 Help	modified
7/7/2020	Deploy multiple forest topologies for Exchange 2013: Exchange 2013 Help	modified
7/7/2020	Manage on-premises moves: Exchange 2013 Help	modified
7/7/2020	About Exchange documentation	modified
7/7/2020	Create a hybrid deployment with the Hybrid Configuration wizard	modified
7/7/2020	Move mailboxes between on-premises and Exchange Online organizations in hybrid deployments	modified
7/7/2020	Configure document collaboration with OneDrive for Business and Exchange 2016 on-premises	modified
7/7/2020	Troubleshoot a hybrid deployment	modified
7/7/2020	About Exchange documentation	modified
7/7/2020	Assign an address book policy to users in Exchange Online	modified
7/7/2020	Change the settings of an address book policy in Exchange Online	modified
7/7/2020	Create an address book policy in Exchange Online	modified
7/7/2020	Remove an address book policy in Exchange Online	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
7/7/2020	Turn on address book policy routing in Exchange Online	modified
7/7/2020	Configure global address list properties in Exchange Online	modified
7/7/2020	Create a global address list in Exchange Online	modified
7/7/2020	Manage address lists in Exchange Online	modified
7/7/2020	Remove a global address list in Exchange Online	modified
7/7/2020	Enable or disable hierarchical address books in Exchange Online	modified
7/7/2020	Add an address list to or remove an address list from an offline address book in Exchange Online	modified
7/7/2020	Change the default offline address book in Exchange Online	modified
7/7/2020	Configure offline address book distribution properties	modified
7/7/2020	Create an offline address book	modified
7/7/2020	Remove an offline address book from Exchange	modified
7/7/2020	Specify the administrators and users who can install and manage add-ins for Outlook	modified
7/7/2020	Procedures for Client Access Rules in Exchange Online	modified
7/7/2020	Configure custom MailTips for recipients	modified
7/7/2020	Configure the large audience size for your organization	modified
7/7/2020	Manage MailTips for organization relationships	modified
7/7/2020	Perform a remote wipe on a mobile phone	modified
7/7/2020	Apply or remove an Outlook on the web mailbox policy on a mailbox in Exchange Online	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
7/7/2020	View or configure Outlook on the web mailbox policy properties in Exchange Online	modified
7/7/2020	Create an Outlook on the web mailbox policy in Exchange Online	modified
7/7/2020	Modify the space used by Inbox rules in Exchange Online	modified
7/7/2020	Public attachment handling in Exchange Online	modified
7/7/2020	Remove an Outlook on the web mailbox policy from Exchange Online	modified
7/7/2020	Enable or Disable POP3 or IMAP4 access for a user	modified
7/7/2020	Set POP3 or IMAP4 settings for a user	modified
7/7/2020	Assign "Send As" or "Send on Behalf" permissions for mail-enabled public folders	modified
7/7/2020	Use batch migration to migrate legacy public folders to Microsoft 365 or Office 365 and Exchange Online	modified
7/7/2020	Create a public folder mailbox	modified
7/7/2020	Create a public folder	modified
7/7/2020	Mail-enable or mail-disable a public folder	modified
7/7/2020	Remove a public folder	modified
7/7/2020	Set up public folders in a new organization	modified
7/7/2020	Update the public folder hierarchy	modified
7/7/2020	View statistics for public folders and public folder items	modified
7/7/2020	Shared mailboxes in Exchange Online	modified
7/7/2020	Configure the external postmaster address in Exchange Online	modified
7/7/2020	Enable mail flow for subdomains in Exchange Online	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
7/7/2020	Manage accepted domains in Exchange Online	modified
7/7/2020	Fix email delivery issues for error code 451 4.7.500-699 (ASxxx) in Exchange Online	modified
7/7/2020	Fix email delivery issues for error code 5.1.8 in Exchange Online	modified
7/7/2020	Fix email delivery issues for error code 5.4.6 or 5.4.14 in Exchange Online	modified
7/7/2020	Fix email delivery issues for error code 5.7.12 in Exchange Online	modified
7/7/2020	Fix email delivery issues for error code 5.7.124 in Exchange Online	modified
7/7/2020	Fix email delivery issues for error code 5.7.13 or 5.7.135 in Exchange Online	modified
7/7/2020	Fix email delivery issues for error code 5.7.133 in Exchange Online	modified
7/7/2020	Fix email delivery issues for error code 5.7.134 in Exchange Online	modified
7/7/2020	Fix email delivery issues for error code 5.7.136 in Exchange Online	modified
7/7/2020	Fix email delivery issues for error code 5.7.23 in Exchange Online	modified
7/7/2020	Fix email delivery issues for error code 5.7.57 in Exchange Online	modified
7/7/2020	Fix email delivery issues for error code 5.7.64 in Exchange Online	modified
7/7/2020	Fix email delivery issues for error codes 5.7.700 through 5.7.750 in Exchange Online	modified
7/7/2020	Fix email delivery issues for error code 550 4.4.7 in Exchange Online	modified
7/7/2020	Fix email delivery issues for error code 550 5.0.350 in Exchange Online	modified
7/7/2020	Fix email delivery issues for error code 550 5.1.0 in Exchange Online	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
7/7/2020	Fix email delivery issues for error code 550 5.1.1 through 5.1.20 in Exchange Online	modified
7/7/2020	Fix email delivery issues for error code 550 5.1.10 in Exchange Online	modified
7/7/2020	Fix email delivery issues for error code 550 5.4.1 in Exchange Online	modified
7/7/2020	Fix email delivery issues for error code 550 5.6.11 in Exchange Online	modified
7/7/2020	Fix email delivery issues for error code 550 5.7.1 in Exchange Online	modified
7/7/2020	Email non-delivery reports in Exchange Online	modified
7/7/2020	Manage remote domains in Exchange Online	modified
7/7/2020	Scenario Integrate Microsoft 365 or Office 365 with an email add-on service	modified
7/7/2020	Use Directory Based Edge Blocking to reject messages sent to invalid recipients	modified
7/7/2020	Run a message trace and view the results in the Exchange admin center	modified
7/7/2020	Role assignment policies in Exchange Online	modified
7/7/2020	Manage role groups in Exchange Online	modified
7/7/2020	Configure a moderated recipient in Exchange Online	modified
7/7/2020	Create user mailboxes in Exchange Online	modified
7/7/2020	Delete or restore user mailboxes in Exchange Online	modified
7/7/2020	Create a distribution group naming policy	modified
7/7/2020	Override the distribution group naming policy	modified
7/7/2020	Manage dynamic distribution groups	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
7/7/2020	View members of a dynamic distribution group	modified
7/7/2020	Manage equipment mailboxes	modified
7/7/2020	Manage mail contacts	modified
7/7/2020	Manage mail-enabled security groups	modified
7/7/2020	Manage mail users	modified
7/7/2020	Manage permissions for recipients in Exchange Online	modified
7/7/2020	Create and manage room mailboxes	modified
7/7/2020	Add or remove email addresses for a mailbox	modified
7/7/2020	Change how long permanently deleted items are kept for an Exchange Online mailbox	modified
7/7/2020	Configure email forwarding for a mailbox	modified
7/7/2020	Configure message delivery restrictions for a mailbox	modified
7/7/2020	Convert a mailbox	modified
7/7/2020	Enable or disable Exchange ActiveSync for a mailbox	modified
7/7/2020	Enable or disable MAPI for a mailbox	modified
7/7/2020	Enable or disable Outlook on the web for a mailbox	modified
7/7/2020	Enable or disable single item recovery for a mailbox	modified
7/7/2020	Mailbox plans in Exchange Online	modified
7/7/2020	Manage user mailboxes	modified
7/7/2020	Create a custom DLP policy	modified
7/7/2020	Create a DLP policy from a template	modified
7/7/2020	Manage policy tips	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
7/7/2020	Export mailbox audit logs	modified
7/7/2020	Run a non-owner mailbox access report	modified
7/7/2020	Run a per-mailbox litigation hold report	modified
7/7/2020	Search the role group changes or administrator audit logs in Exchange Online	modified
7/7/2020	View the administrator audit log	modified
7/7/2020	View and export the external admin audit log	modified
7/7/2020	Assign eDiscovery permissions in Exchange	modified
7/7/2020	Create a discovery mailbox	modified
7/7/2020	Create an In-Place eDiscovery search	modified
7/7/2020	Export eDiscovery search results to a PST file	modified
7/7/2020	Journaling in Exchange Online	modified
7/7/2020	Manage journaling in Exchange Online	modified
7/7/2020	Organization-wide message disclaimers, signatures, footers, or headers in Exchange Online	modified
7/7/2020	Manage mail flow rules in Exchange Online	modified
7/7/2020	Manage and troubleshoot message approval in Exchange Online	modified
7/7/2020	Use mail flow rules to automatically add meetings to calendars in Exchange Online	modified
7/7/2020	Add retention tags to or remove retention tags from a retention policy	modified
7/7/2020	Apply a retention policy to mailboxes	modified
7/7/2020	Create a Retention Policy	modified
7/7/2020	Place a mailbox on retention hold	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
7/7/2020	Retention tags and retention policies	modified
7/7/2020	Modify archive policies	modified
7/7/2020	Clean up or delete items from the Recoverable Items folder in Exchange Online	modified
7/7/2020	Create an organization relationship in Exchange Online	modified
7/7/2020	Modify an organization relationship in Exchange Online	modified
7/7/2020	Remove an organization relationship in Exchange Online	modified
7/7/2020	Apply a sharing policy to mailboxes in Exchange Online	modified
7/7/2020	Create a sharing policy in Exchange Online	modified
7/7/2020	Add an auto attendant extension number in Exchange Online	modified
7/7/2020	Configure an auto attendant for users who have similar names in Exchange Online	modified
7/7/2020	Configure business hours in Exchange Online	modified
7/7/2020	Configure a DTMF fallback auto attendant in Exchange Online	modified
7/7/2020	Configure the time zone in Exchange Online	modified
7/7/2020	Configure the group of users that can be contacted in Exchange Online	modified
7/7/2020	Create a holiday schedule in Exchange Online	modified
7/7/2020	Create a UM auto attendant in Exchange Online	modified
7/7/2020	Create business hours navigation menus in Exchange Online	modified
7/7/2020	Create menu navigation in Exchange Online	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
7/7/2020	Create non-business hours navigation menus in Exchange Online	modified
7/7/2020	Delete a UM auto attendant in Exchange Online	modified
7/7/2020	Disable a UM auto attendant in Exchange Online	modified
7/7/2020	Enable a customized business hours greeting in Exchange Online	modified
7/7/2020	Enable a customized business hours menu prompt in Exchange Online	modified
7/7/2020	Enable a customized non-business hours greeting in Exchange Online	modified
7/7/2020	Enable a customized non-business hours menu prompt in Exchange Online	modified
7/7/2020	Enable an informational announcement in Exchange Online	modified
7/7/2020	Enable or disable directory lookups in Exchange Online	modified
7/7/2020	Enable or disable sending voice messages to users in Exchange Online	modified
7/7/2020	Enable or disable automatic speech recognition in Exchange Online	modified
7/7/2020	Enable or prevent transferring calls from an auto attendant in Exchange Online	modified
7/7/2020	Enable a UM auto attendant in Exchange Online	modified
7/7/2020	Enter a business name in Exchange Online	modified
7/7/2020	Manage a UM auto attendant in Exchange Online	modified
7/7/2020	Set a business location in Exchange Online	modified
7/7/2020	Change the audio codec in Exchange Online	modified
7/7/2020	Configure a dial plan for users who have similar names in Exchange Online	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
7/7/2020	Configure a fully qualified domain name in Exchange Online	modified
7/7/2020	Configure the IP address in Exchange Online	modified
7/7/2020	Configure the listening port in Exchange Online	modified
7/7/2020	Configure the maximum call duration in Exchange Online	modified
7/7/2020	Configure the maximum recording duration in Exchange Online	modified
7/7/2020	Configure the recording idle time-out value in Exchange Online	modified
7/7/2020	Configure the VoIP security setting in Exchange Online	modified
7/7/2020	Create a UM dial plan in Exchange Online	modified
7/7/2020	Create a UM hunt group in Exchange Online	modified
7/7/2020	Create a UM IP gateway in Exchange Online	modified
7/7/2020	Delete a UM dial plan in Exchange Online	modified
7/7/2020	Delete a UM hunt group in Exchange Online	modified
7/7/2020	Delete a UM IP gateway in Exchange Online	modified
7/7/2020	Disable a UM IP gateway in Exchange Online	modified
7/7/2020	Enable a UM IP gateway in Exchange Online	modified
7/7/2020	Manage a UM dial plan in Exchange Online	modified
7/7/2020	Manage a UM IP gateway in Exchange Online	modified
7/7/2020	View a UM hunt group in Exchange Online	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
7/7/2020	Enable custom prompt recording using the telephone user interface in Exchange Online	modified
7/7/2020	Investigate the audio quality of voice calls for a user in Exchange Online	modified
7/7/2020	Investigate the audio quality of voice calls in your organization in Exchange Online	modified
7/7/2020	Disable common PIN patterns for voice mail in Exchange Online	modified
7/7/2020	Enable common PIN patterns for voice mail in Exchange Online	modified
7/7/2020	Include text with the email message sent when a PIN Is reset in Exchange Online	modified
7/7/2020	Reset a voice mail PIN in Exchange Online	modified
7/7/2020	Retrieve voice mail PIN information in Exchange Online	modified
7/7/2020	Set the minimum PIN length for voice mail in Exchange Online	modified
7/7/2020	Set the number of previous voice mail PINs to recycle in Exchange Online	modified
7/7/2020	Set the number of sign-in failures before a voice mail user Is locked out in Exchange Online	modified
7/7/2020	Set the number of sign-in failures before a voice mail PIN is reset in Exchange Online	modified
7/7/2020	Set the PIN lifetime for voice mail in Exchange Online	modified
7/7/2020	Set Outlook Voice Access PIN policies in Exchange Online	modified
7/7/2020	Allow Message Waiting Indicator (MWI) on a UM IP gateway in Exchange Online	modified
7/7/2020	Allow users in the same dial plan to receive faxes in Exchange Online	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
7/7/2020	Authorize calls for a group of users in Exchange Online	modified
7/7/2020	Authorize calls for auto attendant callers in Exchange Online	modified
7/7/2020	Authorize calls for users in a dial plan in Exchange Online	modified
7/7/2020	Call answering rules in the same mailbox policy in Exchange Online	modified
7/7/2020	Call answering rules in Exchange Online	modified
7/7/2020	Configure dial codes in Exchange Online	modified
7/7/2020	Configure the number of input failures before Outlook Voice Access users are disconnected in Exchange Online	modified
7/7/2020	Configure the number of sign-in failures before Outlook Voice Access users are disconnected in Exchange Online	modified
7/7/2020	Configure an Outlook Voice Access number in Exchange Online	modified
7/7/2020	Configure the limit on personal greetings for Outlook Voice Access users in Exchange Online	modified
7/7/2020	Configure the primary way for Outlook Voice Access users to search in Exchange Online	modified
7/7/2020	Configure Protected Voice Mail from authenticated callers in Exchange Online	modified
7/7/2020	Configure Protected Voice Mail from unauthenticated callers in Exchange Online	modified
7/7/2020	Configure the secondary way for Outlook Voice Access users to search in Exchange Online	modified
7/7/2020	Configure the group of users that Outlook Voice Access users can contact in Exchange Online	modified
7/7/2020	Configure Voice Mail Preview partner services for users in Exchange Online	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
7/7/2020	Create a call answering rule in Exchange Online	modified
7/7/2020	Create dialing rules for users in Exchange Online	modified
7/7/2020	Disable faxing for a group of users in Exchange Online	modified
7/7/2020	Disable missed call notifications for a user in Exchange Online	modified
7/7/2020	Disable Message Waiting Indicator (MWI) for users in Exchange Online	modified
7/7/2020	Disable outgoing calls on UM IP gateways in Exchange Online	modified
7/7/2020	Disable selected features for Outlook Voice Access users in Exchange Online	modified
7/7/2020	Disable Voice Mail Preview for users in Exchange Online	modified
7/7/2020	Enable a customized greeting for Outlook Voice Access users in Exchange Online	modified
7/7/2020	Enable a user to receive faxes in Exchange Online	modified
7/7/2020	Enable an informational announcement for Outlook Voice Access users in Exchange Online	modified
7/7/2020	Enable faxing for a group of users in Exchange Online	modified
7/7/2020	Enable missed call notifications for a user in Exchange Online	modified
7/7/2020	Enable Message Waiting Indicator (MWI) for users in Exchange Online	modified
7/7/2020	Enable or disable a call answering rule for a user in Exchange Online	modified
7/7/2020	Enable or disable automatic speech recognition for an Outlook Voice Access user in Exchange Online	modified
7/7/2020	Enable or disable multimedia playback of protected voice messages in Exchange Online	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
7/7/2020	Enable or disable Outlook Voice Access for users in Exchange Online	modified
7/7/2020	Enable or disable Play on Phone for Outlook Voice Access users in Exchange Online	modified
7/7/2020	Enable or disable sending voice messages from Outlook Voice Access in Exchange Online	modified
7/7/2020	Enable or prevent transferring calls from Outlook Voice Access in Exchange Online	modified
7/7/2020	Enable outgoing calls on UM IP gateways in Exchange Online	modified
7/7/2020	Enable Voice Mail Preview for users in Exchange Online	modified
7/7/2020	Include text with the email message sent when a fax message is received in Exchange Online	modified
7/7/2020	Prevent a user from receiving faxes in Exchange Online	modified
7/7/2020	Prevent Message Waiting Indicator (MWI) on a UM IP gateway in Exchange Online	modified
7/7/2020	Prevent users in the same dial plan from receiving faxes in Exchange Online	modified
7/7/2020	Remove a call answering rule for a user in Exchange Online	modified
7/7/2020	Set mailbox features for an Outlook Voice Access user in Exchange Online	modified
7/7/2020	Set mailbox features for Outlook Voice Access users in Exchange Online	modified
7/7/2020	Set the maximum delivery delay for a Voice Mail Preview partner in Exchange Online	modified
7/7/2020	Set the maximum message duration for a Voice Mail Preview partner in Exchange Online	modified
7/7/2020	Set the partner fax server URI to allow faxing in Exchange Online	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
7/7/2020	Set the Voice Mail Preview partner address in Exchange Online	modified
7/7/2020	Set the Voice Mail Preview partner ID in Exchange Online	modified
7/7/2020	Specify the text to display for email clients that don't support Windows Rights Management in Exchange Online	modified
7/7/2020	View and manage a call answering rule in Exchange Online	modified
7/7/2020	Add an E.164 number in Exchange Online	modified
7/7/2020	Add an extension number in Exchange Online	modified
7/7/2020	Add a SIP address in Exchange Online	modified
7/7/2020	Allow callers without a caller ID to leave a voice message in Exchange Online	modified
7/7/2020	Assign a UM mailbox policy in Exchange Online	modified
7/7/2020	Change an E.164 number in Exchange Online	modified
7/7/2020	Change an extension number in Exchange Online	modified
7/7/2020	Change a SIP address in Exchange Online	modified
7/7/2020	Change the UM dial plan in Exchange Online	modified
7/7/2020	Create a UM mailbox policy in Exchange Online	modified
7/7/2020	Delete a UM mailbox policy in Exchange Online	modified
7/7/2020	Disable calls from users who aren't UM-enabled in Exchange Online	modified
7/7/2020	Disable voice mail for a user in Exchange Online	modified
7/7/2020	Enable a user for voice mail in Exchange Online	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
7/7/2020	Enable calls from users who aren't UM-enabled in Exchange Online	modified
7/7/2020	Include text with the email message sent when a voice message Is received in Exchange Online	modified
7/7/2020	Include text with the email message sent when a user Is enabled for voice mail in Exchange Online	modified
7/7/2020	Manage a UM mailbox policy in Exchange Online	modified
7/7/2020	Manage voice mail settings for a user in Exchange Online	modified
7/7/2020	Prevent callers without a caller ID from leaving a voice message in Exchange Online	modified
7/7/2020	Remove an E.164 number in Exchange Online	modified
7/7/2020	Remove an extension number in Exchange Online	modified
7/7/2020	Remove a SIP address in Exchange Online	modified
7/7/2020	About Exchange documentation	modified
7/7/2020	Procedures for antimalware protection in Exchange Server	modified
7/7/2020	Download antimalware engine and definition updates	modified
7/7/2020	Enable antispam functionality on Mailbox servers	modified
7/7/2020	Attachment filtering procedures on Edge Transport servers	modified
7/7/2020	Configure Exchange antispam settings on mailboxes	modified
7/7/2020	Configure a spam quarantine mailbox	modified
7/7/2020	Connection filtering procedures on Edge Transport servers	modified
7/7/2020	Content filtering procedures	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
7/7/2020	Recipient filtering procedures on Edge Transport servers	modified
7/7/2020	Release quarantined messages from the spam quarantine mailbox	modified
7/7/2020	Safelist aggregation procedures	modified
7/7/2020	Sender filtering procedures	modified
7/7/2020	Sender ID procedures	modified
7/7/2020	Sender reputation procedures	modified
7/7/2020	Configure Outlook to show the original sender in the spam quarantine mailbox	modified
7/7/2020	View antispam stamps in Outlook	modified
7/7/2020	Assign certificates to Exchange Server services	modified
7/7/2020	Configure the Availability service for cross-forest topologies	modified
7/7/2020	Configure client-specific message size limits	modified
7/7/2020	Complete a pending Exchange Server certificate request	modified
7/7/2020	Create an Exchange Server certificate request for a certification authority	modified
7/7/2020	Create a new Exchange Server self-signed certificate	modified
7/7/2020	Turn off access to the Exchange admin center	modified
7/7/2020	Export a certificate from an Exchange server	modified
7/7/2020	Import or install a certificate on an Exchange server	modified
7/7/2020	Renew an Exchange Server certificate	modified
7/7/2020	Address rewriting procedures on Edge Transport servers	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
7/7/2020	Configure internet mail flow through Edge Transport servers without using EdgeSync	modified
7/7/2020	Procedures for Edge Subscriptions	modified
7/7/2020	Import address rewrite entries on Edge Transport servers	modified
7/7/2020	Manage mailbox databases in Exchange Server	modified
7/7/2020	Enable the MRS Proxy endpoint for remote moves	modified
7/7/2020	Prepare mailboxes for cross-forest moves using the Exchange Management Shell	modified
7/7/2020	Recreate missing arbitration mailboxes	modified
7/7/2020	Procedures for Client Access Rules in Exchange 2019	modified
7/7/2020	Enable or disable Exchange ActiveSync access to mailboxes in Exchange Server	modified
7/7/2020	Perform a remote wipe on a mobile phone	modified
7/7/2020	Enable or disable MAPI access to mailboxes in Exchange Server	modified
7/7/2020	Using hybrid Modern Authentication with Outlook for iOS and Android	modified
7/7/2020	Use AD FS claims-based authentication with Outlook on the web	modified
7/7/2020	Customize the Outlook on the web sign-in, language selection, and error pages in Exchange Server	modified
7/7/2020	Configure http to https redirection for Outlook on the web in Exchange Server	modified
7/7/2020	Enable or disable Outlook on the web access to mailboxes in Exchange Server	modified
7/7/2020	View or configure Outlook on the web mailbox policy properties	modified
7/7/2020	Create a theme for Outlook on the web in Exchange Server	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
7/7/2020	View or configure Outlook on the web virtual directories in Exchange Server	modified
7/7/2020	Configure authenticated SMTP settings for POP3 and IMAP4 clients in Exchange Server	modified
7/7/2020	Enable and configure IMAP4 on an Exchange server	modified
7/7/2020	Enable or disable POP3 or IMAP4 access to mailboxes in Exchange Server	modified
7/7/2020	Enable and configure POP3 on an Exchange server	modified
7/7/2020	Create a public folder mailbox in Exchange Server	modified
7/7/2020	Create a public folder	modified
7/7/2020	Mail-enable or mail-disable a public folder	modified
7/7/2020	Migrate public folders from Exchange 2013 to Exchange 2016 or Exchange 2019	modified
7/7/2020	Batch migrate Exchange Server public folders to Microsoft 365 or Office 365	modified
7/7/2020	Set up public folders in a new organization	modified
7/7/2020	View statistics for public folders and public folder items	modified
7/7/2020	Create shared mailboxes in the Exchange admin center	modified
7/7/2020	Procedures for address book policies in Exchange Server	modified
7/7/2020	Procedures for address lists in Exchange Server	modified
7/7/2020	Procedures for email address policies in Exchange Server	modified
7/7/2020	Procedures for offline address books in Exchange Server	modified
7/7/2020	Exchange Server 2010 load balancer deployment	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
7/7/2020	Exchange Server	modified
7/7/2020	Use Windows Server Backup to back up Exchange	modified
7/7/2020	Create a recovery database	modified
7/7/2020	Move a mailbox database using database portability	modified
7/7/2020	Perform a dial tone recovery	modified
7/7/2020	Recover a database availability group member server, recover Exchange DAG member, Exchange DAG server recovery, DAG server recovery, Exchange DAG failover	modified
7/7/2020	Recover Exchange server, recover lost Exchange Server, Lost Exchange Server recovery	modified
7/7/2020	Restore data using a recovery database	modified
7/7/2020	Use Windows Server Backup to restore a backup of Exchange	modified
7/7/2020	Activate a mailbox database copy	modified
7/7/2020	Activate a lagged mailbox database copy	modified
7/7/2020	Add a mailbox database copy	modified
7/7/2020	Configure activation policy for a mailbox database copy	modified
7/7/2020	Configure AutoReseed for a database availability group	modified
7/7/2020	Configure database availability group network properties	modified
7/7/2020	Configure database availability group properties	modified
7/7/2020	Configure mailbox database copy properties	modified
7/7/2020	Create a database availability group network	modified
7/7/2020	Create a database availability group	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
7/7/2020	Manage database availability group membership	modified
7/7/2020	Move the mailbox database path for a mailbox database copy	modified
7/7/2020	Pre-stage the cluster name object for a database availability group	modified
7/7/2020	Remove a database availability group	modified
7/7/2020	Remove a mailbox database copy	modified
7/7/2020	Perform a server switchover	modified
7/7/2020	Suspend or resume a mailbox database copy	modified
7/8/2020	What's new in Exchange admin center	modified
7/8/2020	Managing Outlook for iOS and Android in Exchange Online	modified
7/8/2020	Exchange Online	modified
7/8/2020	Exchange Deployment Assistant release notes	modified
7/9/2020	Perform a G Suite migration	modified
7/10/2020	Fix issues with printers, scanners, and LOB applications that send email using Microsoft 365 or Office 365	modified
7/10/2020	Using hybrid Modern Authentication with Outlook for iOS and Android	modified

Week of July 13, 2020

PUBLISHED ON	TOPIC TITLE	CHANGE
7/14/2020	How and when to decommission your on-premises Exchange servers in a hybrid deployment	modified
7/14/2020	Manage mail contacts	modified
7/14/2020	Mail flow rule conditions and exceptions (predicates) in Exchange Online	modified
7/16/2020	Exchange Online	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
7/16/2020	Exchange Server	removed
7/16/2020	Exchange dev/test environment in Azure	modified
7/17/2020	Configure global address list properties in Exchange Online	modified
7/17/2020	Exchange admin center in Exchange Online	modified
7/17/2020	How to set up a multifunction device or application to send email using Microsoft 365 or Office 365	modified

Week of July 20, 2020

PUBLISHED ON	TOPIC TITLE	CHANGE
7/21/2020	MailTips	modified
7/21/2020	POP3 and IMAP4	modified
7/21/2020	Perform a G Suite migration	modified
7/21/2020	Mail flow rule conditions and exceptions (predicates) in Exchange Online	modified
7/21/2020	Exchange Server build numbers and release dates	modified
7/21/2020	Mail flow rule conditions and exceptions (predicates) in Exchange Server	modified
7/21/2020	Transport rule conditions and exceptions (predicates) in Exchange 2013	modified
7/23/2020	Disable Basic authentication in Exchange Online	modified
7/23/2020	Create a hybrid deployment with the Hybrid Configuration wizard	modified
7/23/2020	Move mailboxes between on-premises and Exchange Online organizations in hybrid deployments	modified
7/23/2020	Configure document collaboration with OneDrive for Business and Exchange 2016 on-premises	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
7/23/2020	Troubleshoot a hybrid deployment	modified
7/23/2020	Deploying Outlook for iOS and Android app configuration settings	modified
7/23/2020	Sensitivity labeling and protection in Outlook for iOS and Android	modified
7/24/2020	Fix email delivery issues for error code 5.1.8 in Exchange Online	modified
7/24/2020	Email non-delivery reports in Exchange Online	modified

Week of July 27, 2020

PUBLISHED ON	TOPIC TITLE	CHANGE
7/28/2020	Microsoft Hybrid Agent	modified
7/28/2020	Recreate missing arbitration mailboxes	modified
7/28/2020	Create shared mailboxes in the Exchange admin center	modified
7/28/2020	Create a database availability group	modified
7/28/2020	Procedures for mail flow rules in Exchange Server	modified
7/28/2020	Exchange 2013 Performance Counters: Exchange 2013 Help	modified
7/29/2020	Client Access Rules in Exchange Online	modified
7/29/2020	Managing Outlook for iOS and Android in Exchange Online	modified
7/29/2020	Sensitivity labeling and protection in Outlook for iOS and Android	modified
7/29/2020	Fix email delivery issues for error codes 5.7.700 through 5.7.750 in Exchange Online	modified
7/29/2020	Managing devices for Outlook for iOS and Android for Exchange Server	modified
7/29/2020	Using Basic authentication with Outlook for iOS and Android	modified

PUBLISHED ON	TOPIC TITLE	CHANGE
7/29/2020	Using hybrid Modern Authentication with Outlook for iOS and Android	modified
7/29/2020	Microsoft Hybrid Agent	modified
7/29/2020	Permissions in Exchange hybrid deployments	modified
7/29/2020	Manage mail flow rules in Exchange Online	modified
7/29/2020	Recreate missing arbitration mailboxes	modified
7/29/2020	Connection filtering on Edge Transport servers	modified
7/30/2020	Public attachment handling in Exchange Online	modified
7/30/2020	Delete or restore user mailboxes in Exchange Online	modified
7/30/2020	Create and manage room mailboxes	modified
7/31/2020	Microsoft Hybrid Agent	modified
7/31/2020	Create an organization relationship in Exchange Online	modified
7/31/2020	Public folder procedures	modified

What's new in Exchange Server

8/3/2020 • 54 minutes to read • Edit Online

Exchange Server 2019 brings a new set of technologies, features, and services to Exchange Server, the messaging platform that provides email, scheduling, and tools for custom collaboration and messaging service applications. Its goal is to support people and organizations as their work habits evolve from a communication focus to a collaboration focus. At the same time, Exchange 2019 helps lower the total cost of ownership whether you deploy Exchange 2019 on-premises or provision your mailboxes in the cloud.

Choose the section below that matches the version of Exchange that you're upgrading from. If you want to know about features that have been removed or replaced in Exchange 2019, see What's discontinued in Exchange Server.

For more information about deploying Exchange 2019, see Planning and deployment for Exchange Server.

What's new when upgrading from Exchange 2016 to Exchange 2019?

Security

- **Windows Server Core support**: Running Exchange on a Windows deployment with less surface area means less attack surface area and fewer components to service.

- **Block external access to Exchange admin center (EAC) and the Exchange Management Shell**: You can use Client Access Rules to only allow administration of Exchange from the internal network instead of using complex network and firewall rules.

- **TLS 1.2 is the only version that's enabled by default**: Exchange Server 2019 includes important changes to improve the security of client and server connections. The default configuration for encryption will enable TLS 1.2 only and disable support for older algorithms (namely, DES, 3DES, RC2, RC4 and MD5). It will also configure elliptic curve key exchange algorithms with priority over non-elliptic curve algorithms. In Exchange Server 2016 and later, all cryptography settings are inherited from the configuration specified in the operating system. For additional information, see Exchange Server TLS Guidance.

Performance

- **Improved search infrastructure**: The completely rebuilt search infrastructure for cloud scale and reliability in Exchange Online is now available in Exchange 2019. This new search infrastructure allows for indexing of bigger files, simpler management, and better search performance.

- **Faster, more reliable failovers**: The changes to the search architecture result in significantly faster and more reliable failover over between servers.

- **Metacache database**: Improvements at the core of Exchange's database engine enable better overall performance and take advantage of the latest storage hardware, including larger disks and SSDs.

- **Modern hardware support**: Exchange now supports up to 256 GB of memory and 48 CPU cores.

- **Dynamic database cache**: The information store process employs dynamic memory cache allocation optimizing memory usage to active database usage.

Clients

- **Calendar - Do Not Forward**: This is similar to Information Rights Management (IRM) for calendar items without the IRM deployment requirements. Attendees can't forward the invitation to other people, and only the organizer can invite additional attendees.

- **Calendar - Better Out of Office**: Additional options when you won't be in the office. Key options include:

add an event to your calendar that shows you as Away/Out of Office, and a quick option to cancel/decline meetings that will happen while you're away.

- **Calendar - Remove-CalendarEvents cmdlet**: Enables administrators to cancel meetings that were organized by a user that has left the company. Previously, conference rooms or meeting attendees would have these defunct meetings permanently on their calendars.
- **Email address internationalization (EAI)**: Email addresses that contain non-English characters can now be routed and delivered natively.

What's new when upgrading from Exchange 2013 to Exchange 2019?

Exchange 2019 architecture

Today, CPU horsepower is significantly less expensive and is no longer a constraining factor. With that constraint lifted, the primary design goal for Exchange 2019 is for simplicity of scale, hardware utilization, and failure isolation. With Exchange 2019, we reduced the number of server roles to two: the Mailbox and Edge Transport server roles.

Unified Messaging (UM) has been removed from Exchange 2019. Other than that, the Mailbox server in Exchange 2019 includes all of the server components from the Exchange 2013 Mailbox and Client Access server roles:

- Client Access services provide authentication, limited redirection, and proxy services. Client Access services don't do any data rendering and offer all the usual client access protocols: HTTP, POP and IMAP, and SMTP.
- Mailbox services include all the traditional server components found in the Exchange 2013 Mailbox server role except Unified Messaging: the backend client access protocols, Transport service, and Mailbox databases. The Mailbox server handles all activity for the active mailboxes on that server.

The Edge Transport role is typically deployed in your perimeter network, outside your internal Active Directory forest, and is designed to minimize the attack surface of your Exchange deployment. By handling all Internet-facing mail flow, it also adds additional layers of message protection and security against viruses and spam, and can apply mail flow rules (also known as transport rules) to control message flow.

For more information about the Exchange 2019 architecture, see Exchange architecture.

Along with the new Mailbox role, Exchange 2019 now allows you to proxy traffic from Exchange 2013 Client Access servers to Exchange 2019 mailboxes. This new flexibility gives you more control in how you move to Exchange 2019 without having to worry about deploying enough front-end capacity to service new Exchange 2019 servers.

Clients

Outlook on the web (formerly known as Outlook Web App)

Outlook Web App is now known as Outlook on the web, which continues to let users access their Exchange mailbox from almost any web browser.

> **NOTE**
>
> Supported Web browsers for Outlook on the web in Exchange 2019 are Microsoft Edge, Internet Explorer 11, and the most recent versions of Mozilla Firefox, Google Chrome, and Apple Safari.

The former Outlook Web App user interface has been updated and optimized for tablets and smart phones, in addition to desktop and laptop computers. New Exchange 2019 features include:

- **Platform-specific experiences for phones** for both iOS and Android.
- **Premium Android experience** using Chrome on devices running Android version 4.2 or later.
- **Email improvements**, including a new single-line view of the Inbox with an optimized reading pane,

archiving, emojis, and the ability to undo mailbox actions like deleting a message or moving a message.

- **Contact linking** the ability for users to add contacts from their LinkedIn accounts.
- **Calendar** has an updated look and new features, including email reminders for Calendar events, ability to propose a new time in meeting invitations, improved search, and birthday calendars.
- **Search suggestions and refiners** for an improved search experience that helps users find the information they want, faster. Search suggestions try to anticipate what the user's looking for and returns results that might be what the user is looking for. Search refiners will help a user more easily find the information they're looking for by providing contextually-aware filters. Filters might include date ranges, related senders, and so on.
- **New themes** Thirteen new themes with graphic designs.
- **Options** for individual mailboxes have been overhauled.
- **Link preview** which enables users to paste a link into messages, and Outlook on the web automatically generates a rich preview to give recipients a peek into the contents of the destination. This works with video links as well.
- **Inline video** player saves the user time by keeping them in the context of their conversations. An inline preview of a video automatically appears after inserting a video URL.
- **Pins and Flags** which allow users to keep essential emails at the top of their inbox (Pins) and mark others for follow-up (Flags). Pins are now folder specific, great for anyone who uses folders to organize their email. Quickly find and manage flagged items with inbox filters or the new Task module, accessible from the app launcher.
- **Performance improvements** in a number of areas across Outlook on the web, including creating calendar events, composing, loading messages in the reading pane, popouts, search, startup, and switching folders.
- **New Outlook on the web action pane** that allows you to quickly click those actions you most commonly use such as New, Reply all, and Delete. A few new actions have been added as well including Archive, Sweep, and Undo.

MAPI over HTTP

MAPI over HTTP is now the default protocol that Outlook uses to communicate with Exchange. MAPI over HTTP improves the reliability and stability of the Outlook and Exchange connections by moving the transport layer to the industry-standard HTTP model. This allows a higher level of visibility of transport errors and enhanced recoverability. Additional functionality includes support for an explicit pause-and-resume function, which enables supported clients to change networks or resume from hibernation while maintaining the same server context.

Note: MAPI over HTTP isn't enabled in organizations where the following conditions are both true:

- You're installing Exchange 2019 in an organization that already has Exchange 2013 servers installed.
- MAPI over HTTP wasn't enabled in Exchange 2013.

While MAPI over HTTP is now the default communication protocol between Outlook and Exchange, clients that don't support it will fall back to Outlook Anywhere (RPC over HTTP).

For more information, see MAPI over HTTP in Exchange Server.

Document collaboration

Exchange 2019, along with SharePoint Server 2019, enables Outlook on the web users to link to and share documents that are stored in OneDrive for Business in an on-premises SharePoint server instead of attaching files to messages. Users in an on-premises environment can collaborate on files in the same manner.

For more information about SharePoint Server 2019, see New and improved features in SharePoint Server 2019.

When an Exchange 2019 user receives a Word, Excel, or PowerPoint file in an email attachment, and the file is stored in OneDrive for Business or on-premises SharePoint, the user will now have the option of viewing and editing that file in Outlook on the web alongside the message. To do this, you'll need a separate computer in your on-premises organization that's running Office Online Server. For more information, see Install Office Online Server in an Exchange organization.

Exchange 2019 also brings the following improvements to document collaboration:

- Saving files to OneDrive for Business.
- Uploading a file to OneDrive for Business.
- Most Recently Used lists populated with both local and online files.

Microsoft 365 or Office 365 hybrid

The Hybrid Configuration Wizard (HCW) that was included with Exchange 2013 is moving to become a cloud-based application. When you choose to configure a hybrid deployment in Exchange 2019, you'll be prompted to download and install the wizard as a small app. The wizard will function the same in previous versions of Exchange, with a few new benefits:

- The wizard can be updated quickly to support changes in the Microsoft 365 or Office 365 service.
- The wizard can be updated to account for issues detected when customers try to configure a hybrid deployment.
- Improved troubleshooting and diagnostics to help you resolve issues that you run into when running the wizard.
- The same wizard will be used by everyone configuring a hybrid deployment who's running Exchange 2013 or later.

In addition to Hybrid Configuration Wizard improvements, multi-forest hybrid deployments are being simplified with Azure Active Directory Connect (AADConnect). AADConnect introduces management agents that will make it significantly easier to synchronize multiple on-premises Active Directory forests with a single Microsoft 365 or Office 365 organization. For more information about AADConnect, see What is hybrid identity with Azure Active Directory?.

Exchange ActiveSync clients will be seamlessly redirected to Microsoft 365 or Office 365 when a user's mailbox is moved to Exchange Online. To support this, ActiveSync clients need to support HTTP 451 redirect. When a client is redirected, the profile on the device is updated with the URL of the Exchange Online service. This means the client will no longer attempt to contact the on-premises Exchange server when trying to find the mailbox.

Messaging policy and compliance

There are several new and updated message policy and compliance features in Exchange 2019.

Data loss prevention

To comply with business standards and industry regulations, organizations need to protect sensitive information and prevent its inadvertent disclosure. Examples of sensitive information that you might want to prevent from leaking outside your organization include credit card numbers, social security numbers, health records, or other personally identifiable information (PII). With a DLP policy and mail flow rules (also known as transport rules) in Exchange 2019, you can now identify, monitor, and protect 80 different types of sensitive information with new conditions and actions:

- With the new condition **Any attachment has these properties, including any of these words**, a mail flow rule can match messages where the specified property of the attached Office document contains specified words. This condition makes it easy to integrate your Exchange mail flow rules and DLP policies with SharePoint, Windows Server 2012 R2 File Classification Infrastructure (FCI), or a third-party classification system.

- With the new action **Notify the recipient with a message**, a mail flow rule can send a notification to the recipient with the text you specify. For example, you can inform the recipient that the message was rejected by a mail flow rule, or that it was marked as spam and will be delivered to their Junk Email folder.
- The action **Generate incident report and send it to** has been updated to enable the notification of multiple recipients by allowing a group address to be configured as the recipient.

To learn more about DLP, see Data loss prevention in Exchange Server.

In-place Archiving, retention, and eDiscovery

Exchange 2019 includes the following improvements to In-Place Archiving, retention, and eDiscovery to help your organization meet its compliance needs:

- **Public folder support for In-Place eDiscovery and In-Place Hold**: Exchange 2019 integrates public folders into the In-Place eDiscovery and Hold workflow. You can use In-Place eDiscovery to search public folders in your organization, and you can put an In-Place Hold on public folders. And similar to placing a mailbox on hold, you can place a query-based and a time-based hold on public folders. Currently, you can only search and place a hold on all public folders. In later releases, you'll be able to choose specific public folders to search and place on hold. For more information, see Search and place a hold on public folders using In-Place eDiscovery.
- **Compliance Search**: Compliance Search is a new eDiscovery search tool in Exchange 2019 with new and improved scaling and performance capabilities. You can use it to search very large numbers of mailboxes in a single search. In fact, there's no limit on the number of mailboxes that can be included in a single search, so you can search all mailboxes in your organization at once. There's also no limit on the number of searches that can run at the same time. For In-Place eDiscovery in Exchange 2019, the limits are the same as in Exchange 2013: you can search up to 10,000 mailboxes in a single search and your organization can run a maximum of two In-Place eDiscovery searches at the same time.

  In Exchange 2019, Compliance Search is only available by using the Exchange Management Shell. For information about using the Compliance Search cmdlets, see the following topics:

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

  > **NOTE**
  >
  > To have access to the Compliance Search cmdlets, an administrator or eDiscovery manager must be assigned the Mailbox Search management role or be a member of the Discovery Management role group.

For more information, see Messaging policy and compliance in Exchange Server.

Improved performance and scalability

In Exchange 2019, the search architecture has been redesigned. Previously, search was a synchronous operation that was not very fault-tolerant. The new architecture is asynchronous and decentralized. It distributes the work across multiple servers and keeps retrying if any servers are too busy. This means that we can return results more reliability, and faster.

Another advantage of the new architecture is that search scalability is improved. The number of mailboxes you can search at once using the console has increased from 5k to 10k for both mailboxes and archive mailboxes, allowing you to search a total of 20k mailboxes at the same time.

Microsoft Exchange Server 2016 brings a new set of technologies, features, and services to Exchange Server, the messaging platform that provides email, scheduling, and tools for custom collaboration and messaging service applications. Its goal is to support people and organizations as their work habits evolve from a communication focus to a collaboration focus. At the same time, Exchange 2016 helps lower the total cost of ownership whether you deploy Exchange 2016 on-premises or provision your mailboxes in the cloud.

Choose the section below that matches the version of Exchange that you're upgrading from. If you want to know about features that have been removed or replaced in Exchange 2016, see What's discontinued in Exchange Server.

For more information about deploying Exchange 2016, see Planning and deployment.

What's new when updating from Exchange 2016 RTM to Exchange 2016 CU1?

When you update to Exchange 2016 Cumulative Update 1 (CU1) from Exchange 2016 RTM, you'll get the following new features:

- **SHA-2 compliant S/MIME in Outlook on the web**: We've updated the certificate that's used by the S/MIME control in Outlook on the web. The certificate is now SHA-2 compliant. Users who downloaded the control from Exchange 2016 RTM will need to download the control again after you install CU1.

- **Additional languages for Outlook on the web**: With CU1, we're adding 17 new languages to Outlook on the web.

- **Improved download package**: Exchange 2016 releases, starting with CU1, are now packaged as ISO files instead of a self-extracting EXE file. The ISO file can be mounted directly in Windows Server 2012 or later. If you need to install Exchange across a network, you can create a network share from the mounted ISO drive.

What's new when upgrading from Exchange 2013 to Exchange 2016 RTM?

Exchange 2016 architecture

Clients

Outlook on the web (formerly Outlook Web App)

MAPI over HTTP

Document collaboration

Microsoft 365 or Office 365 hybrid

Messaging policy and compliance

Data loss prevention

In-place Archiving, retention, and eDiscovery

Exchange 2016 architecture

Today, CPU horsepower is significantly less expensive and is no longer a constraining factor. With that constraint lifted, the primary design goal for Exchange 2016 is for simplicity of scale, hardware utilization, and failure isolation. With Exchange 2016, we reduced the number of server roles to two: the Mailbox and Edge Transport server roles.

The Mailbox server in Exchange 2016 includes all of the server components from the Exchange 2013 Mailbox and Client Access server roles:

- Client Access services provide authentication, limited redirection, and proxy services. Client Access services don't do any data rendering and offer all the usual client access protocols: HTTP, POP and IMAP, and SMTP.
- Mailbox services include all the traditional server components found in the Exchange 2013 Mailbox server role: the backend client access protocols, Transport service, Mailbox databases, and Unified Messaging. The Mailbox server handles all activity for the active mailboxes on that server.

The Edge Transport role is typically deployed in your perimeter network, outside your internal Active Directory forest, and is designed to minimize the attack surface of your Exchange deployment. By handling all Internet-facing mail flow, it also adds additional layers of message protection and security against viruses and spam, and can apply mail flow rules (also known as transport rules) to control message flow.

For more information about the Exchange 2016 architecture, see Exchange 2016 architecture.

Along with the new Mailbox role, Exchange 2016 now allows you to proxy traffic from Exchange 2013 Client Access servers to Exchange 2016 mailboxes. This new flexibility gives you more control in how you move to Exchange 2016 without having to worry about deploying enough front-end capacity to service new Exchange 2016 servers.

Clients

Outlook on the web (formerly Outlook Web App)

Outlook Web App is now known as Outlook on the web, which continues to let users access their Exchange mailbox from almost any web browser.

> **NOTE**
>
> Supported Web browsers for Outlook on the web in Exchange 2016 are Microsoft Edge, Internet Explorer 11, and the most recent versions of Mozilla Firefox, Google Chrome, and Apple Safari.

The former Outlook Web App user interface has been updated and optimized for tablets and smart phones, in addition to desktop and laptop computers. New Exchange 2016 features include:

- **Platform-specific experiences for phones** for both iOS and Android.
- **Premium Android experience** using Chrome on devices running Android version 4.2 or later.
- **Email improvements**, including a new single-line view of the Inbox with an optimized reading pane, archiving, emojis, and the ability to undo mailbox actions like deleting a message or moving a message.
- **Contact linking** the ability for users to add contacts from their LinkedIn accounts.
- **Calendar** has an updated look and new features, including email reminders for Calendar events, ability to propose a new time in meeting invitations, improved search, and birthday calendars.
- **Search suggestions and refiners** for an improved search experience that helps users find the information they want, faster. Search suggestions try to anticipate what the user's looking for and returns results that might be what the user is looking for. Search refiners will help a user more easily find the information they're looking for by providing contextually-aware filters. Filters might include date ranges, related senders, and so on.
- **New themes** Thirteen new themes with graphic designs.
- **Options** for individual mailboxes have been overhauled.
- **Link preview** which enables users to paste a link into messages, and Outlook on the web automatically generates a rich preview to give recipients a peek into the contents of the destination. This works with video

links as well.

- **Inline video** player saves the user time by keeping them in the context of their conversations. An inline preview of a video automatically appears after inserting a video URL.
- **Pins and Flags** which allow users to keep essential emails at the top of their inbox (Pins) and mark others for follow-up (Flags). Pins are now folder specific, great for anyone who uses folders to organize their email. Quickly find and manage flagged items with inbox filters or the new Task module, accessible from the app launcher.
- **Performance improvements** in a number of areas across Outlook on the web, including creating calendar events, composing, loading messages in the reading pane, popouts, search, startup, and switching folders.
- **New Outlook on the web action pane** that allows you to quickly click those actions you most commonly use such as New, Reply all, and Delete. A few new actions have been added as well including Archive, Sweep, and Undo.

MAPI over HTTP

MAPI over HTTP is now the default protocol that Outlook uses to communicate with Exchange. MAPI over HTTP improves the reliability and stability of the Outlook and Exchange connections by moving the transport layer to the industry-standard HTTP model. This allows a higher level of visibility of transport errors and enhanced recoverability. Additional functionality includes support for an explicit pause-and-resume function, which enables supported clients to change networks or resume from hibernation while maintaining the same server context.

Note: MAPI over HTTP isn't enabled in organizations where the following conditions are both true:

- You're installing Exchange 2016 in an organization that already has Exchange 2013 servers installed.
- MAPI over HTTP wasn't enabled in Exchange 2013.

While MAPI over HTTP is now the default communication protocol between Outlook and Exchange, clients that don't support it will fall back to Outlook Anywhere (RPC over HTTP).

For more information, see MAPI over HTTP in Exchange 2016.

Document collaboration

Exchange 2016, along with SharePoint Server 2016, enables Outlook on the web users to link to and share documents that are stored in OneDrive for Business in an on-premises SharePoint server instead of attaching files to messages. Users in an on-premises environment can collaborate on files in the same manner that's used in Microsoft 365 and Office 365.

For more information about SharePoint Server 2016, see New and improved features in SharePoint Server 2016.

When an Exchange 2016 user receives a Word, Excel, or PowerPoint file in an email attachment, and the file is stored in OneDrive for Business or on-premises SharePoint, the user will now have the option of viewing and editing that file in Outlook on the web alongside the message. To do this, you'll need a separate computer in your on-premises organization that's running Office Online Server. For more information, see Install Office Online Server in an Exchange 2016 organization.

Exchange 2016 also brings the following improvements to document collaboration:

- Saving files to OneDrive for Business.
- Uploading a file to OneDrive for Business.
- Most Recently Used lists populated with both local and online files.

Microsoft 365 or Office 365 hybrid

The Hybrid Configuration Wizard (HCW) that was included with Exchange 2013 is moving to become a cloud-based application. When you choose to configure a hybrid deployment in Exchange 2016, you'll be prompted to download and install the wizard as a small app. The wizard will function the same in previous versions of Exchange, with a few new benefits:

- The wizard can be updated quickly to support changes in the Microsoft 365 or Office 365 service.
- The wizard can be updated to account for issues detected when customers try to configure a hybrid deployment.
- Improved troubleshooting and diagnostics to help you resolve issues that you run into when running the wizard.
- The same wizard will be used by everyone configuring a hybrid deployment who's running Exchange 2013 or Exchange 2016.

In addition to Hybrid Configuration Wizard improvements, multi-forest hybrid deployments are being simplified with Azure Active Directory Connect (AADConnect). AADConnect introduces management agents that will make it significantly easier to synchronize multiple on-premises Active Directory forests with a single Microsoft 365 or Office 365 organization. For more information about AADConnect, see Integrating your on-premises identities with Azure Active Directory.

Exchange ActiveSync clients will be seamlessly redirected to Microsoft 365 or Office 365 when a user's mailbox is moved to Exchange Online. To support this, ActiveSync clients need to support HTTP 451 redirect. When a client is redirected, the profile on the device is updated with the URL of the Exchange Online service. This means the client will no longer attempt to contact the on-premises Exchange server when trying to find the mailbox.

Messaging policy and compliance

There are several new and updated message policy and compliance features in Exchange 2016.

Data loss prevention

To comply with business standards and industry regulations, organizations need to protect sensitive information and prevent its inadvertent disclosure. Examples of sensitive information that you might want to prevent from leaking outside your organization include credit card numbers, social security numbers, health records, or other personally identifiable information (PII). With a DLP policy and mail flow rules (also known as transport rules) in Exchange 2016, you can now identify, monitor, and protect 80 different types of sensitive information with new conditions and actions:

- With the new condition **Any attachment has these properties, including any of these words**, a mail flow rule can match messages where the specified property of the attached Office document contains specified words. This condition makes it easy to integrate your Exchange mail flow rules and DLP policies with SharePoint, Windows Server 2012 R2 File Classification Infrastructure (FCI), or a third-party classification system.
- With the new action **Notify the recipient with a message**, a mail flow rule can send a notification to the recipient with the text you specify. For example, you can inform the recipient that the message was rejected by a mail flow rule, or that it was marked as spam and will be delivered to their Junk Email folder.
- The action **Generate incident report and send it to** has been updated to enable the notification of multiple recipients by allowing a group address to be configured as the recipient.

To learn more about DLP, see Data loss prevention in Exchange 2016.

In-place Archiving, retention, and eDiscovery

Exchange 2016 includes the following improvements to In-Place Archiving, retention, and eDiscovery to help your organization meet its compliance needs:

- **Public folder support for In-Place eDiscovery and In-Place Hold**: Exchange 2016 integrates public folders into the In-Place eDiscovery and Hold workflow. You can use In-Place eDiscovery to search public folders in your organization, and you can put an In-Place Hold on public folders. And similar to placing a mailbox on hold, you can place a query-based and a time-based hold on public folders. Currently, you can only search and place a hold on all public folders. In later releases, you'll be able to choose specific public folders to search and place on hold. For more information, see Search and place a hold on public folders using In-Place eDiscovery.

- **Compliance Search**: Compliance Search is a new eDiscovery search tool in Exchange 2016 with new and improved scaling and performance capabilities. You can use it to search very large numbers of mailboxes in a single search. In fact, there's no limit on the number of mailboxes that can be included in a single search, so you can search all mailboxes in your organization at once. There's also no limit on the number of searches that can run at the same time. For In-Place eDiscovery in Exchange 2016, the limits are the same as in Exchange 2013: you can search up to 10,000 mailboxes in a single search and your organization can run a maximum of two In-Place eDiscovery searches at the same time.

  In Exchange 2016, Compliance Search is only available by using the Exchange Management Shell. For information about using the Compliance Search cmdlets, see the following topics:

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

  > **NOTE**
  >
  > To have access to the Compliance Search cmdlets, an administrator or eDiscovery manager must be assigned the Mailbox Search management role or be a member of the Discovery Management role group.

For more information, see Messaging policy and compliance in Exchange 2016.

Improved performance and scalability

In Exchange 2016, the search architecture has been redesigned. Previously, search was a synchronous operation that was not very fault-tolerant. The new architecture is asynchronous and decentralized. It distributes the work across multiple servers and keeps retrying if any servers are too busy. This means that we can return results more reliability, and faster.

Another advantage of the new architecture is that search scalability is improved. The number of mailboxes you can search at once using the console has increased from 5k to 10k for both mailboxes and archive mailboxes, allowing you to search a total of 20k mailboxes at the same time.

What's new when upgrading from Exchange 2010 to Exchange 2016 RTM?

Exchange admin center

Exchange 2016 architecture

Setup

Microsoft 365 or Office 365 hybrid

Messaging policy and compliance

Antimalware protection

Mail flow and the transport pipeline

Sharing and collaboration

Integration with SharePoint and Skype for Business

Clients

Batch mailbox moves

High availability and site resilience

Exchange workload management

Exchange admin center

Exchange 2016 provides a single unified management console that allows for ease of use and is optimized for management of on-premises, online, or hybrid deployments. The *Exchange admin center* (EAC) in Exchange 2016 replaces the Exchange Management Console (EMC) and the Exchange Control Panel (ECP) that were used in Exchange 2010 (but the name of the EAC virtual directory is still "ECP"). Some EAC features include:

- **List view**: The list view in EAC has been designed to remove key limitations that existed in ECP. ECP was limited to displaying up to 500 objects and, if you wanted to view objects that weren't listed in the details pane, you needed to use searching and filtering to find those specific objects. In Exchange 2016, the viewable limit from within the EAC list view is approximately 20,000 objects. After the EAC returns the results, the EAC client performs the searching and sorting, which greatly increases the performance compared to the ECP in Exchange 2010. In addition, paging has been added so that you can page to the results. You can also configure page size and export to a .csv file.

- **Add/Remove columns to the Recipient list view**: You can choose which columns to view, and with local cookies, you can save your custom list views per machine that you use to access the EAC.

- **Secure the ECP virtual directory**: You can control access to the EAC from inside and outside your corporate network without affecting user access to their Outlook on the web options. For more information, see Turn off access to the Exchange admin center.

- **Tool consolidation**: The functionality of these management tools has been integrated into the EAC:

  - The Public Folder administration console.

  - The Role Based Access Control (RBAC) User Editor in the Exchange Toolbox.

  - The Call Statistics and User Call Logs tools for Unified Messaging.

- **Notifications**: In Exchange 2016, the EAC now has a Notification viewer so that you can view notifications and alerts for:

  - Exchange certificates that are installed on any Exchange 2016 server in your organization that are expired or within 30 days of expiring.

  - Mailbox moves and migrations (Mailbox Replication Service or MRS tasks). Notifications are displayed when migrations are started, in process, and completed.

    You can also configure email addresses to receive these notifications in the EAC, or all events by using the **Set-Notification** cmdlet in the Exchange Management Shell.

- **Groups enhancements**: By default, up to 500 recipients are returned when you open the **Select Members** window, however, you can choose to list up to 10,000 recipients by clicking **Get All Results** beneath the recipient list. We now support browsing more than 500 recipients by using the scroll bar, and we've also added enhanced search features so you can filter recipients in the recipient list. You can filter by:

  - City
  - Company
  - Country/region
  - Department
  - Office
  - Title

- **Delivery reports**: Administrators can use the EAC to track delivery information for email messages sent to or received by any user in the organization. You just select a mailbox, and then search for messages sent to or received by a different user. You can filter the search by words in the subject line. The results track a message through the delivery process, and indicate whether the message was successfully delivered, is pending delivery, or wasn't delivered. For more information, see Track messages with delivery reports.

- **Certificate management**: Administrators can use the EAC to manage Exchange certificates on multiple servers from a central location, which helps to minimize the amount of interaction that's required to manage Exchange certificates. For more information about certificate management procedures in Exchange 2016, see Certificate procedures in Exchange 2016.

For more information about the EAC, see Exchange admin center in Exchange 2016.

Exchange 2016 architecture

Today, CPU horsepower is significantly less expensive and is no longer a constraining factor. With that constraint lifted, the primary design goal for Exchange 2016 is for simplicity of scale, hardware utilization, and failure isolation. With Exchange 2016, we reduced the number of server roles to two: the Mailbox and Edge Transport server roles.

The Mailbox server in Exchange 2016 includes all of the server components from the Mailbox, Client Access, Hub Transport, and Unified Messaging server roles in Exchange 2010:

- Client Access services provide authentication, limited redirection, and proxy services. Client Access services don't do any data rendering and offer all the usual client access protocols: HTTP, POP and IMAP, and SMTP.

- Mailbox services include: the backend client access protocols, Transport service, Mailbox databases, and Unified Messaging. The Mailbox server handles all activity for the active mailboxes on that server.

Like previous versions of Exchange. the Edge Transport role is typically deployed in your perimeter network, outside your internal Active Directory forest, and is designed to minimize the attack surface of your Exchange deployment. By handling all Internet-facing mail flow, it also adds additional layers of message protection and security against viruses and spam, and can apply mail flow rules (also known as transport rules) to control message flow.

The Exchange 2016 architecture provides the following benefits:

- **Version upgrade flexibility**: No more rigid upgrade requirements. Mailbox servers can be upgraded independently and in any order in relation to other Mailbox servers.

- **Session indifference**: With Exchange 2010, session affinity to the Client Access server role was required for several protocols. In Exchange 2016, the client access and mailbox components reside on the same Mailbox server. No session affinity is required between Mailbox servers, Edge Transport servers, or mail servers on the Internet. This allows inbound client connections to Mailbox servers to be balanced using techniques provided by load-balancing technology like least connection or round-robin.

- **Deployment simplicity**: With an Exchange 2010 site-resilient design, you needed up to eight different namespaces: two Internet Protocol namespaces, two for Outlook Web App fallback, one for Autodiscover, two for RPC Client Access, and one for SMTP. With Exchange 2016, the most organizations only need two namespaces for coexistence with Exchange 2010: one for client protocols and one for Autodiscover. Depending on how you configure your mail routing, you might also need an additional namespace for SMTP routing.

For more information about the Exchange 2016 architecture, see Exchange 2016 architecture.

Setup

Setup has been completely rewritten so that installing Exchange 2016 and making sure you've got the latest product rollups and security fixes is easier than ever. Here are some of the improvements we've made:

- **Improved readiness checks**: Readiness checks make sure that your computer and your organization are ready for Exchange 2016. After you've provided the necessary information about your installation to Setup, the readiness checks are run before installation begins. The new readiness check engine now runs through all checks before reporting back to you on what actions need to be performed before Setup can continue, and it does so faster than ever. As with previous versions of Exchange, you can tell Setup to install the Windows features that are required by Setup so you don't have to install them manually.

- **Simplified and modern wizard**: We've removed all the steps in the Setup wizard that aren't absolutely required for you to install Exchange. What's left is an easy-to-follow wizard that takes you through the installation process one step at a time.

For more information, see Planning and deployment.

Microsoft 365 or Office 365 hybrid

The Hybrid Configuration Wizard (HCW) that was included with Exchange 2013 is moving to become a cloud-based application. When you choose to configure a hybrid deployment in Exchange 2016, you'll be prompted to download and install the wizard as a small app. The wizard will function the same in previous versions of Exchange, with a few new benefits:

- The wizard can be updated quickly to support changes in the Microsoft 365 or Office 365 service.

- The wizard can be updated to account for issues that are encountered when customers try to configure a hybrid deployment.

- Improved troubleshooting and diagnostics to help you resolve issues that you encounter.

- The same wizard will be used by everyone who's configuring a hybrid deployment with Exchange 2013 or Exchange 2016.

In addition to Hybrid Configuration Wizard improvements, multi-forest hybrid deployments are being simplified with Azure Active Directory Connect (AADConnect). AADConnect introduces management agents that will make it significantly easier to synchronize multiple on-premises Active Directory forests with a single Microsoft 365 or Office 365 organization.

Exchange ActiveSync clients will be seamlessly redirected to Microsoft 365 or Office 365 when a user's mailbox is moved to Exchange Online. To support this, ActiveSync clients need to support HTTP 451 redirect. When a client is redirected, the profile on the device is updated with the URL of the Exchange Online service. This means the client will no longer attempt to contact the on-premises Exchange server when trying to find the mailbox.

Messaging policy and compliance

There are several new and updated message policy and compliance features in Exchange 2016.

Data loss prevention

Data loss prevention (DLP) capabilities help you protect your sensitive data and inform users of internal compliance policies. DLP can also help keep your organization safe from users who might mistakenly send sensitive information to unauthorized people. DLP helps you identify, monitor, and protect sensitive data through deep content analysis. Exchange 2016 offers built-in DLP policies based on regulatory standards such as personally identifiable information (PII) and payment card industry data security standards (PCI), and is extensible to support other policies important to your business. With a DLP policy in Exchange 2016, you can now identify, monitor, and protect 80 different types of sensitive information. For more information, see Sensitive information types in Exchange 2016. Additionally, the new policy tips in Outlook 2016 inform users about policy violations before sensitive data is sent.

To learn more, see Data loss prevention in Exchange 2016

Mail flow rules (transport rules)

You can use Exchange mail flow rules (also known as transport rules) to look for specific conditions in messages that pass through your organization and take action on them. For example, your organization might require that certain types of messages are blocked or rejected in order to meet legal or compliance requirements, or to implement specific business needs. Mail flow rules are similar to the Inbox rules that are available in Outlook. The main difference between mail flow rules and Inbox rules is that mail flow rules take action on messages while they're in transit as opposed to after the message is delivered. Mail flow rules also contain a richer set of conditions, exceptions, and actions, which gives you the flexibility to implement many types of messaging policies.

These features are new to mail flow rules in Exchange 2016:

- Exchange mail flow rules can now identify 80 different types of sensitive information, including 30 new sensitive information types focusing on identifiers from South America, Europe, and Asia. These 80 built-in types are included, but you can also develop your own type from scratch. For more information on these sensitive information types, see Sensitive information types in Exchange 2016.

- With the new condition **Any attachment has these properties, including any of these words**, a mail flow rule can match messages where the specified property of the attached Office document contains specified words. This condition makes it easy to integrate your Exchange mail flow rules and DLP policies with SharePoint, Windows Server 2012 R2 File Classification Infrastructure (FCI), or a third-party classification system.

- With the new action **Notify the recipient with a message**, a mail flow rule can send a notification to the recipient with the text you specify. For example, you can inform the recipient that the message was rejected by a mail flow rule, or that it was marked as spam and will be delivered to their Junk Email folder.

- The action **Generate incident report and send it to** has been updated to enable the notification of multiple recipients by allowing a group address to be configured as the recipient

- Additional mail flow rules predicates and actions.

For more information, see Mail flow rules in Exchange 2016.

Azure Rights Management connector

The Azure Rights Management connector (also known as the Microsoft Rights Management connector or RMS connector) is an optional application that helps you enhance data protection for your Exchange 2016 server by connecting to the cloud-based Azure Rights Management service (also known as Microsoft Rights Management or Azure RMS). Once you install the RMS connector, it provides continuous data protection throughout the life span of the information and because these services are customizable, you can define the level of protection you need. For example, you can limit email message access to specific users or set view-only rights for certain messages.

For more information, see Deploying the Azure Rights Management connector.

In-place Archiving, retention, and eDiscovery

Exchange 2016 includes the following improvements to In-Place Archiving, retention, and eDiscovery to help your organization meet its compliance needs:

- **In-Place Hold**: In-Place Hold is a unified hold model that allows you to meet legal hold requirements in the following scenarios:
  - Preserve the results of the query (query-based hold), which allows for scoped immutability across mailboxes.
  - Place a time-based hold to meet retention requirements (for example, retain all items in a mailbox for seven years, a scenario that required the use of Single Item Recovery/Deleted Item Retention in Exchange 2010).
  - Place a mailbox on indefinite hold (similar to litigation hold in Exchange 2010).
  - Place a user on multiple holds to meet different case requirements.
- **In-Place eDiscovery**: In-Place eDiscovery allows authorized users to search mailbox data across all mailboxes and In-Place Archives in an Exchange 2016 organization and copy messages to a discovery mailbox for review. In Exchange 2016, In-Place eDiscovery allows discovery managers to perform more efficient searches and hold.
  - **Federated search** allows you to search and preserve data across multiple data repositories. With Exchange 2016, you can perform in-place eDiscovery searches across Exchange, SharePoint, and Skype for Business. You can use the eDiscovery Center in SharePoint 2013 to perform In-Place eDiscovery search and hold.
  - **Query-based In-Place Hold** allows you to save the results of the query, which allows for scoped immutability across mailboxes.
  - **Export search results** Discovery Managers can export mailbox content to a .pst file from the SharePoint 2013 eDiscovery Console. Mailbox export request cmdlets are no longer required to export a mailbox to a .pst file.
  - **Keyword statistics**: Search statistics are offered on a per search term basis. This enables a Discovery Manager to quickly make intelligent decisions about how to further refine the search query to provide better results. eDiscovery search results are sorted by relevance.
  - **KQL syntax**: Discovery Managers can use Keyword Query Language (KQL) syntax in search queries. KQL is similar to the Advanced Query Syntax (AQS), that was used for discovery searches in Exchange 2010.
  - **In-Place eDiscovery and Hold wizard**: Discovery Managers can use the In-Place eDiscovery and Hold wizard to perform eDiscovery and hold operations.

    > **NOTE**
    >
    > If SharePoint 2013 isn't available, a subset of the eDiscovery functionality is available in the Exchange admin center.

  - **Public folder support for In-Place eDiscovery and In-Place Hold**: Exchange 2016 has integrated public folders into the In-Place eDiscovery and Hold workflow. You can use In-Place eDiscovery to search public folders in your organization, and you can put a In-Place Hold on public folders. And similar to placing a mailbox on hold, you can place a query-based and a time-based hold on public folders. Currently, you can only search and place a hold on all public folders. In later releases, you'll be able to choose specific public folders to search and place on hold. For more information, see Search and place a hold on public folders using In-Place eDiscovery.
  - **Compliance Search**: Compliance Search is a new eDiscovery search tool in Exchange 2016 with

  new and improved scaling and performance capabilities. You can use it to search very large numbers of mailboxes in a single search. In fact, there's no limit on the number of mailboxes that can be included in a single search, so you can search all mailboxes in your organization at once. There's also no limit on the number of searches that can run at the same time. For In-Place eDiscovery in Exchange 2016, the limits are the same as in Exchange 2013: you can search up to 10,000 mailboxes in a single search and your organization can run a maximum of two In-Place eDiscovery searches at the same time.

  In Exchange 2016, Compliance Search is only available by using the Exchange Management Shell. For information about using the Compliance Search cmdlets, see the following topics:

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

  > **NOTE**
  >
  > To have access to the Compliance Search cmdlets, an administrator or eDiscovery manager must be assigned the Mailbox Search management role or be a member of the Discovery Management role group.

- **Search across primary and archive mailboxes in Outlook on the web**: Users can search across their primary and archive mailboxes in Outlook on the web. Two separate searches are no longer necessary.

- **Archive Skype for Business content**: Exchange 2016 supports archiving of Skype for Business content in a user's mailbox. You can place Skype for Business content on hold using In-Place Hold and use In-Place eDiscovery to search Skype for Business content archived in Exchange.

For more information, see Messaging policy and compliance in Exchange 2016.

Auditing

Exchange 2016 includes the following improvements to auditing:

- **Auditing reports**: The EAC includes auditing functionality so that you can run reports or export entries from the mailbox audit log and the administrator audit log. The mailbox audit log records whenever a mailbox is accessed by someone other than the person who owns the mailbox. This can help you determine who has accessed a mailbox and what they have done. The administrator audit log records any action, based on an Exchange Management Shell cmdlet, performed by an administrator. This can help you troubleshoot configuration issues or identify the cause of problems related to security or compliance. For more information, see Exchange Auditing Reports.

- **Viewing the administrator audit log**: Instead of exporting the administrator audit log, which can take up to 24 hours to receive in an email message, you can view administrator audit log entries in the EAC. To do this, go to **Compliance Management** > **Auditing** and click **View the administrator audit log**. Up to 1000 entries will be displayed on multiple pages. To narrow the search, you can specify a date range.

  As an additional improvement, you can also export the audit log data in a format that's common to both Exchange 2016 and SharePoint Server 2016. This makes it easier to integrate with third-party tools to view the data and create richer reports.

  For more information, see View the Administrator Audit Log.

Antimalware protection

The built-in malware filtering capabilities of Exchange 2016 helps protect your network from malicious software that's transferred by email messages. All messages sent or received by your Exchange 2016 Mailbox server are scanned for malware (viruses and spyware) by using the built-in Malware Agent. If malware is detected, the message is deleted. Notifications may also be sent to senders or administrators when an infected message is deleted and not delivered. You can also choose to replace infected attachments with either default or custom messages that notify the recipients of the malware detection.

For more information about antimalware protection, see Antimalware protection in Exchange 2016.

Mail flow and the transport pipeline

How messages flow through an organization and what happens to them has changed significantly in Exchange 2016. Following is a brief overview of the changes:

- **Transport pipeline**: The transport pipeline in Exchange 2016 is now made up of several different services: the Front End Transport service, the Transport service, and the Mailbox Transport service. For more information, see Mail flow and the transport pipeline.

- **Routing**: Mail routing in Exchange 2016 recognizes DAG boundaries as well as Active Directory site boundaries. Also, mail routing has been improved to queue messages more directly for internal recipients. For more information, see Mail routing.

- **Connectors**: The default maximum message size for a Send connector or a Receive connector has increased from 10MB to 25MB. For more information, see Connector limits.

  You can configure Send connectors to route outbound mail through the Front End transport service on Mailbox servers. For more information, see Configure Send connectors to proxy outbound mail.

Recipients

Administrators can now use the EAC to create a *group naming policy*, which lets you standardize and manage the names of distribution groups that are created by users in your organization. You can automatically add a prefix or suffix to the name of the distribution group when it's created, and you can block specific words from being used in group names. For more information, see Create a Distribution Group Naming Policy.

For more information about recipients in Exchange 2016, see Recipients.

Sharing and collaboration

Exchange 2016 includes the following enhancements for sharing and collaboration:

- **Public folders**: Public folders now take advantage of the existing high availability and storage technologies of the mailbox store. The public folder architecture uses specially designed mailboxes to store both the hierarchy and the public folder content. This new design also means that there is no longer a public folder database. Public folder replication now uses the continuous replication model. High availability for the hierarchy and content mailboxes is provided by the database availability group (DAG). With this design, we're moving away from a multi-master replication model to a single-master replication model. For more information, see Public folders.

- **Shared mailboxes**: In previous versions of Exchange, creating a shared mailbox was a multi-step process in which you had to use the Exchange Management Shell to set the delegate permissions. Now you can create a shared mailbox in one step via the Exchange admin center (EAC). In the EAC, go to **Recipients** > **Shared** to create a shared mailbox. Shared mailboxes are a recipient type so you can easily search for your shared mailboxes in either the EAC or by using the Exchange Management Shell. For more information, see Shared mailboxes.

Integration with SharePoint and Skype for Business

Exchange 2016 offers greater integration with SharePoint and Skype for Business. Benefits of this enhanced integration include:

- Skype for Business Server 2015 can archive content in Exchange 2016 and use Exchange 2016 as a contact store.
- Discovery Managers can perform In-Place eDiscovery and Hold searches across SharePoint, Exchange, and Skype for Business data.

For more information, see Plan Exchange 2016 integration with SharePoint and Skype for Business.

Clients

Outlook on the web (formerly Outlook Web App)

Outlook Web App is now known as Outlook on the web, which continues to let users access their Exchange mailbox from almost any web browser.

> **NOTE**
>
> Supported Web browsers for Outlook on the web in Exchange 2016 are Microsoft Edge, Internet Explorer 11, and the most recent versions of Mozilla Firefox, Google Chrome, and Apple Safari.

In Exchange 2016, the former Outlook Web App user interface is updated and optimized for tablets and smart phones, in addition to desktop and laptop computers. New Exchange 2016 features include:

- **Platform-specific experiences for phones** for both iOS and Android.
- **Premium Android experience** using Chrome on devices running Android version 4.2 or later.
- **Apps for Outlook** which allow users and administrators to extend the capabilities of Outlook on the web.
- **Email improvements**, including a new single-line view of the Inbox with an optimized reading pane, archiving, emojis, and the ability to undo mailbox actions like deleting a message or moving a message.
- **Contact linking** the ability for users to add contacts from their LinkedIn accounts.
- **Calendar** has an updated look and new features, including email reminders for Calendar events, ability to propose a new time in meeting invitations, improved search, and birthday calendars.
- **Search suggestions and refiners** for an improved search experience that helps users find the information they want, faster. Search suggestions try to anticipate what the user's looking for and returns results that might be what the user is looking for. Search refiners will help a user more easily find the information they're looking for by providing contextually-aware filters. Filters might include date ranges, related senders, and so on.
- **Themes** Exchange 2016 provides over 50 built-in themes.
- **Options** for individual mailboxes have been overhauled.
- **Link preview** which enables users to paste a link into messages, and Outlook on the web automatically generates a rich preview to give recipients a peek into the contents of the destination. This works with video links as well.
- **Inline video** player saves the user time by keeping them in the context of their conversations. An inline preview of a video automatically appears after inserting a video URL.
- **Link preview** which enables users to paste a link into messages, and Outlook on the web automatically generates a rich preview to give recipients a peek into the contents of the destination. This works with video links as well.

- **Pins and Flags** which allow users to keep essential emails at the top of their inbox (Pins) and mark others for follow-up (Flags). Pins are now folder specific, great for anyone who uses folders to organize their email. Quickly find and manage flagged items with inbox filters or the new Task module, accessible from the app launcher.

- **Performance improvements** in a number of areas across Outlook on the web, including creating calendar events, composing, loading messages in the reading pane, popouts, search, startup, and switching folders.

- **New Outlook on the web action pane** that allows you to quickly click those actions you most commonly use such as New, Reply all, and Delete. A few new actions have been added as well including Archive, Sweep, and Undo.

Offline Outlook on the web

Internet Explorer 11 and Windows Store apps using JavaScript support the Application Cache API (or AppCache) as defined in the HTML5 specification, which allows you to create offline web applications. AppCache enables webpages to cache (or save) resources locally, including images, script libraries, style sheets, and so on. In addition, AppCache allows URLs to be served from cached content using standard Uniform Resource Identifier (URI) notation. The following is a list of the browsers that support AppCache:

- Microsoft Edge

- Internet Explorer 11 or later versions

- Google Chrome 44 or later versions

- Mozilla Firefox 39 or later versions

- Apple Safari 8 or later (only on OS X/iOS) versions

MAPI over HTTP

MAPI over HTTP is now the default protocol that Outlook uses to communicate with Exchange. MAPI over HTTP improves the reliability and stability of the Outlook and Exchange connections by moving the transport layer to the industry-standard HTTP model. This allows a higher level of visibility of transport errors and enhanced recoverability. Additional functionality includes support for an explicit pause-and-resume function, which enables supported clients to change networks or resume from hibernation while maintaining the same server context.

Note: MAPI over HTTP isn't enabled in organizations where the following conditions are both true:

- You're installing Exchange 2016 in an organization that already has Exchange 2013 servers installed.

- MAPI over HTTP wasn't enabled in Exchange 2013.

While MAPI over HTTP is now the default communication protocol between Outlook and Exchange, clients that don't support it will fall back to Outlook Anywhere (RPC over HTTP). RPC (RPC over TCP) is no longer supported.

For more information, see MAPI over HTTP in Exchange 2016.

Document collaboration

Exchange 2016, along with SharePoint Server 2016, enables Outlook on the web users to link to and share documents that are stored in OneDrive for Business in an on-premises SharePoint server instead of attaching files to messages. Users in an on-premises environment can collaborate on files in the same manner that's used in Microsoft 365 or Office 365.

For more information about SharePoint Server 2016, see New and improved features in SharePoint Server 2016.

When an Exchange 2016 user receives a Word, Excel, or PowerPoint file in an email attachment, and the file is stored in OneDrive for Business or on-premises SharePoint, the user will now have the option of viewing and editing that file in Outlook on the web alongside the message. To do this, you'll need a separate computer in your on-premises organization that's running Office Online Server. For more information, see Install Office Online

Server in an Exchange 2016 organization.

Exchange 2016 also brings the following improvements to document collaboration:

- Saving files to OneDrive for Business.
- Uploading a file to OneDrive for Business.
- Most Recently Used lists populated with both local and online files.

Batch mailbox moves

Exchange 2016 makes use of batch moves. The move architecture is built on top of MRS (Mailbox Replication service) moves with enhanced management capability. The batch move architecture features the following enhancements:

- Ability to move multiple mailboxes in large batches.
- Email notification during move with reporting.
- Automatic retry and automatic prioritization of moves.
- Primary and personal archive mailboxes can be moved together or separately.
- Option for manual move request finalization, which allows you to review a move before you complete it.
- Periodic incremental syncs to migrate the changes.

For more information, see Manage on-premises mailbox moves in Exchange 2016.

High availability and site resilience

The high availability model of the mailbox component has not changed significantly since Exchange 2010. The unit of high availability is still the database availability group (DAG). The DAG still uses Windows Server failover clustering. Continuous replication still supports both file mode and block mode replication. However, there have been some improvements. Failover times have been reduced as a result of transaction log code improvements and deeper checkpoint on the passive databases. The Exchange Store service has been re-written in managed code. Now, each database runs under its own process, which isolates store issues to a single database.

Exchange 2016 uses DAGs and mailbox database copies, along with other features such as single item recovery, retention policies, and lagged database copies, to provide high availability, site resilience, and Exchange native data protection. The high availability platform, the Exchange Information Store and the Extensible Storage Engine (ESE), have all been enhanced to provide greater availability, easier management, and to reduce costs. These enhancements include:

- **Managed availability**: With managed availability, internal monitoring and recovery-oriented features are tightly integrated to help prevent failures, proactively restore services, and initiate server failovers automatically or alert administrators to take action. The focus is on monitoring and managing the end user experience rather than just server and component uptime to help keep the service continuously available.
- **Managed Store**: See the Managed Store section.
- **Support for multiple databases per disk**: Exchange 2016 includes enhancements that enable you to support multiple databases (mixtures of active and passive copies) on the same disk, thereby leveraging larger disks in terms of capacity and IOPS as efficiently as possible.
- **Automatic reseed**: Enables you to quickly restore database redundancy after disk failure. If a disk fails, the database copy stored on that disk is copied from the active database copy to a spare disk on the same server. If multiple database copies were stored on the failed disk, they can all be automatically re-seeded on a spare disk. This enables faster reseeds, as the active databases are likely to be on multiple servers and the data is copied in parallel.

- **Automatic recovery from storage failures**: This feature continues the innovation that was introduced in Exchange 2010 to allow the system to recover from failures that affect resiliency or redundancy. In addition to the Exchange 2010 bugcheck behaviors, Exchange 2016 includes additional recovery behaviors for long I/O times, excessive memory consumption by `MSExchangeRepl.exe`, and severe cases where the system is in such a bad state that threads can't be scheduled.

- **Lagged copy enhancements**: Lagged copies can now use automatic log play down to care for themselves (to a certain extent). Lagged copies will automatically play down log files in a variety of situations, such as single page restore and low disk space scenarios. If the system detects that page patching is required for a lagged copy, the logs will be automatically replayed into the lagged copy. Lagged copies will also invoke this auto replay feature when a low disk space threshold has been reached, and when the lagged copy has been detected as the only available copy for a specific period of time. In addition, lagged copies can leverage Safety Net, making recovery or activation much easier. *Safety Net* is improved functionality in Exchange 2016 based on the transport dumpster of Exchange 2010.

- **Single copy alert enhancements**: The single copy alert that was introduced in Exchange 2010 is no longer a separate scheduled script. It's now integrated into the managed availability components within the system and is a native function within Exchange.

- **DAG network auto-configuration**: DAGs networks can be automatically configured by the system based on configuration settings. In addition to manual configuration options, DAGs can also distinguish between MAPI and Replication networks and configure DAG networks automatically.

For more information about these features, see High availability and site resilience and Changes to high availability and site resilience over previous versions.

Managed Store

In Exchange 2016, *Managed Store* is the name of the Information Store processes, `Microsoft.Exchange.Store.Service.exe` and `Microsoft.Exchange.Store.Worker.exe`. The new Managed Store is written in C# and is tightly integrated with the Microsoft Exchange Replication service (`MSExchangeRepl.exe`) to provide higher availability through improved resiliency. In addition, the Managed Store allows more granular management of resource consumption, and has improved diagnostics for faster root cause analysis.

The Managed Store works with the Microsoft Exchange Replication service to manage mailbox databases, which continue to use the Extensible Storage Engine (ESE) database engine. Exchange 2016 includes significant changes to the mailbox database schema that provide many optimizations over previous versions of Exchange. The Microsoft Exchange Replication service is also responsible for all service availability related to Mailbox servers. These architectural changes enable faster database failover and better physical disk failure handling.

The Managed Store uses the same search platform as SharePoint Server 2016 to provide more robust indexing and searching when compared to Microsoft Search engine that was used in previous versions of Exchange.

For more information, see High availability and site resilience.

Exchange workload management

An Exchange workload is an Exchange server feature, protocol, or service that has been explicitly defined for the purposes of Exchange system resource management. Each Exchange workload consumes system resources such as CPU, mailbox database operations, or Active Directory requests to execute user requests or run background work. Examples of Exchange workloads include Outlook on the web, Exchange ActiveSync, mailbox migration, and mailbox assistants.

There are two ways to manage Exchange workloads in Exchange 2016:

- **Monitor the health of system resources**: Managing workloads based on the health of system resources.

- **Control how resources are consumed by individual users**: Controlling how resources are consumed by individual users was possible in Exchange 2010 (user throttling), and this capability has been expanded for Exchange 2016.

For more information about these features, see User workload management in Exchange 2016.

What's discontinued in Exchange Server

8/3/2020 • 7 minutes to read • Edit Online

This topic discusses the components, features, and functionality that's been removed, discontinued, or replaced in Exchange 2019.

Discontinued features from Exchange 2016 to Exchange 2019

This section lists the Exchange 2016 features that are no longer available in Exchange 2019.

Architecture

FEATURE	COMMENTS AND MITIGATION
Unified Messaging (UM)	Unified Messaging has been removed from Exchange 2019. We recommend that Exchange 2019 organizations transition to Skype for Business Cloud Voice Mail.

Discontinued features from Exchange 2013 to Exchange 2019

This section lists the Exchange 2013 features that are no longer available in Exchange 2019.

Architecture

FEATURE	COMMENTS AND MITIGATION
Unified Messaging (UM)	Unified Messaging has been removed from Exchange 2019. We recommend that Exchange 2019 organizations transition to Skype for Business Cloud Voice Mail.
Client Access server role	The Client Access server role has been replaced by Client Access services that run on the Mailbox server role. The Mailbox server role now performs all functionality that was previously included with the Client Access server role. For more information about the new Mailbox server role, see Exchange Server architecture.
MAPI/CDO library	The MAPI/CDO library has been replaced by Exchange Web Services (EWS), Exchange ActiveSync (EAS), and Representational State Transfer (REST)* APIs. If an application uses the MAPI/CDO library, it needs to move to EWS, EAS, or the REST APIs to communicate with Exchange 2019.

* REST APIs will be included in a future release of Exchange 2019.

De-emphasized features in Exchange 2019

The following features are being de-emphasized in Exchange 2019 and may not be included in future versions of Exchange.

- Third-party replication APIs
- RPC over HTTP

- Database availability group (DAG) support for failover cluster administrative access points

This topic discusses the components, features, or functionality that have been removed, discontinued, or replaced in Exchange 2016.

Discontinued features from Exchange 2013 to Exchange 2016

This section lists the Exchange 2013 features that are no longer available in Exchange 2016.

Architecture

FEATURE	COMMENTS AND MITIGATION
Client Access server role	The Client Access server role has been replaced by Client Access services that run on the Mailbox server role. The Mailbox server role now performs all functionality that was previously included with the Client Access server role. For more information about the new Mailbox server role, see Exchange Server architecture.
MAPI/CDO library	The MAPI/CDO library has been replaced by Exchange Web Services (EWS), Exchange ActiveSync (EAS), and Representational State Transfer (REST)* APIs. If an application uses the MAPI/CDO library, it needs to move to EWS, EAS, or the REST APIs to communicate with Exchange 2016.

* REST APIs will be included in a future release of Exchange 2016.

De-emphasized features in Exchange 2016

The following features are being de-emphasized in Exchange 2016 and may not be included in future versions of Exchange.

- Third-party replication APIs
- RPC over HTTP
- Database Availability Group support for failover cluster administrative access points

Discontinued features from Exchange 2010 to Exchange 2016

This section lists the Exchange 2010 features that are no longer available in Exchange 2016.

Architecture

FEATURE	COMMENTS AND MITIGATION
Hub Transport server role	The Hub Transport server role has been replaced by Transport services which run on the Mailbox server role. The Mailbox server role includes the Microsoft Exchange Transport, Microsoft Exchange Mailbox Transport Delivery, the Microsoft Exchange Mailbox Transport Submission, and the Microsoft Exchange Frontend Transport service. For more information, see Mail flow and the transport pipeline.

FEATURE	COMMENTS AND MITIGATION
Unified Messaging server role	The Unified Messaging server role has been replaced by Unified Messaging services which run on the Mailbox and Client Access server roles. The Mailbox server role includes the Microsoft Exchange Unified Messaging service and the Client Access server role includes the Microsoft Exchange Unified Messaging Call Router service. For more information, see Voice Architecture Changes.
MAPI/CDO library	The MAPI/CDO library has been replaced by Exchange Web Services (EWS), Exchange ActiveSync (EAS), and Representational State Transfer (REST)* APIs. If an application uses the MAPI/CDO library, it needs to move to EWS, EAS, or the REST APIs to communicate with Exchange 2016.

* REST APIs will be included in a future release of Exchange 2016.

Management interfaces

FEATURE	COMMENTS AND MITIGATION
Exchange Management Console and Exchange Control Panel	The Exchange Management Console and the Exchange Control Panel have been replaced by the Exchange admin center (EAC). EAC uses the same virtual directory (/ecp) as the Exchange Control Panel. For more information, see Exchange admin center in Exchange Server.

Client access

FEATURE	COMMENTS AND MITIGATION
Outlook 2003 is not supported	To connect Microsoft Outlook to Exchange 2016, the use of the Autodiscover service is required. However, Microsoft Outlook 2003 doesn't support the use of the Autodiscover service.
RPC/TCP access for Outlook clients	In Exchange 2016, Microsoft Outlook clients can connect using Outlook Anywhere (RPC/HTTP) or MAPI over HTTP Outlook 2013 Service Pack 1 and later. If you have Outlook clients in your organization, using Outlook Anywhere and/or MAPI over HTTP is required. For more information, see Outlook Anywhere and MAPI over HTTP in Exchange Server.

Outlook Web App and Outlook

FEATURE	COMMENTS AND MITIGATION
Spell check	Outlook Web App no longer has built-in spell check services. Instead, it uses the spell check features in your Web browsers.
Customizable filters	Outlook Web App no longer has customizable filtered views and no longer supports saving filtered views to Favorites. Customizable filters have been replaced by fixed filters that can be used to view all messages, unread messages, messages sent to the user, or flagged messages.

FEATURE	COMMENTS AND MITIGATION
Message flags	The ability to set a custom date on a message flag isn't available in Outlook Web App. You can use Outlook to set custom dates.
Chat contact list	The chat contact list that appeared in the folder list in Outlook Web App for Exchange 2010 is no longer available.
Search folders	The ability for users to use Search folders isn't currently available in Outlook Web App.
Web Parts	Outlook on the web no longer includes support for Web Parts. Customers will need to develop replacement functionality to meet this need in their environments.

Mail flow

FEATURE	COMMENTS AND MITIGATION
Linked connectors	The ability to link a Send connector to a Receive connector has been removed. Specifically, the *LinkedReceiveConnector* parameter has been removed from New-SendConnector and Set-SendConnector.

Antispam and antimalware

FEATURE	COMMENTS AND MITIGATION
Antispam agent management in the EMC	In Exchange 2010, when you enabled the antispam agents on a Hub Transport server, you could manage the antispam agents in the Exchange Management Console (EMC). In Exchange 2016, when you enable the antispam agents on a Mailbox server, you can't manage the agents using the EAC. You can only use the Exchange Management Shell. For information about how to enable the antispam agents on a Mailbox server, see Enable antispam functionality on Mailbox servers.
Connection Filtering agent on Hub Transport servers	In Exchange 2010, when you enabled the antispam agents on a Hub Transport server, the Attachment Filter agent was the only antispam agent that wasn't available. In Exchange 2016, when you enable the antispam agents on a Mailbox server, the Attachment Filter agent and the Connection Filtering agent aren't available. The Connection Filtering agent provides IP Allow List and IP Block List capabilities. For information about how to enable the antispam agents on a Mailbox server, see Enable antispam functionality on Mailbox servers. **Note**: The only way to enable the Connection Filtering agent is to install an Edge Transport server in the perimeter network. For more information, see Edge Transport servers.

Messaging policy and compliance

FEATURE	COMMENTS AND MITIGATION

FEATURE	COMMENTS AND MITIGATION
Managed Folders	In Exchange 2010, you use managed folders for messaging retention management (MRM). In Exchange 2016, managed folders aren't supported. You must use retention policies for MRM. **Note**: Cmdlets related to managed folders are still available. You can create managed folders, managed content settings and managed folder mailbox policies, and apply a managed folder mailbox policy to a user, but the MRM assistant skips processing of mailboxes that have a managed folder mailbox policy applied.
Port Managed Folder wizard	In Exchange 2010, you use the Port Managed Folder wizard to create retention tags based on managed folder and managed content settings. In Exchange 2016, the Exchange admin center doesn't include this functionality. You can use the **New-RetentionPolicyTag** cmdlet with the *ManagedFolderToUpgrade* parameter to create a retention tag based on a managed folder.

Unified Messaging and voice mail

FEATURE	COMMENTS AND MITIGATION
Directory lookups using Automatic Speech Recognition (ASR)	In Exchange 2010, Outlook Voice Access users can use speech inputs using Automatic Speech Recognition (ASR) to search for users listed in the directory. Speech inputs could be also used in Outlook Voice Access to navigate menus, messages, and other options. However, even if an Outlook Voice Access user is able to use speech inputs, they have to use the telephone key pad to enter their PIN, and navigate personal options. In Exchange 2016, authenticated and non-authenticated Outlook Voice Access users can't search for users in the directory using speech inputs or ASR in any language. However, callers that call into an auto attendant can use speech inputs in multiple languages to navigate auto attendant menus and search for users in the directory.

Mailbox database copies

FEATURE	COMMENTS AND MITIGATION
Update-MailboxDatabaseCopy Update Mailbox Database Copy wizard	Content index catalog seeding is no longer possible over the replication network; it can only be done over a MAPI network. This is true even when you use the `-Network` parameter in the **Update-MailboxDatabaseCopy** cmdlet.

Updates for Exchange Server

8/3/2020 • 2 minutes to read • Edit Online

Exchange follows a quarterly delivery model to release Cumulative Updates (CUs) that address issues reported by customers. CUs sometimes also add new features and functionality.

Critical product updates are packages that address a Microsoft-released security bulletin or that contain a change in time zone definitions. When in Mainstream Support, critical product updates are released as needed on a monthly basis for the most recently released CU and for the immediately previous CU. When in Extended Support, critical product updates are released as needed on a monthly basis for only the most recently released CU.

Each CU is a full installation of Exchange that includes all updates and changes from previous CUs. When installing a new Exchange server using the latest released CU, you don't need to install Exchange RTM or any previously released CU.

VERSION	BLOG POST	
Exchange 2019 CU6	Released: June 2020 Quarterly Exchange Updates	
Exchange 2019 CU5	Released: March 2020 Quarterly Exchange Updates	
Exchange 2019 CU4	Released: December 2019 Quarterly Exchange Updates	
Exchange 2019 CU3	Released: September 2019 Quarterly Exchange Updates	
Exchange 2019 CU2	Released: June 2019 Quarterly Exchange Updates	
Exchange 2019 CU1	Released: February 2019 Quarterly Exchange Updates	
Exchange 2019 RTM	Exchange Server 2019 Now Available	

For information about the new features you'll get when you upgrade to Exchange 2019 from previous versions of Exchange, see What's new in Exchange Server.

To get the latest version of Exchange 2016, download and install Cumulative Update 17 for Exchange Server 2016. Because each CU is a full installation of Exchange that includes updates and changes from all previous CUs, you don't need to install any previous CUs or Exchange 2016 RTM first.

The following table contains links to Exchange Team blog posts ("What's New" information) for this and other Exchange 2016 CUs.

VERSION	BLOG POST
Exchange 2016 CU17	Released: June 2020 Quarterly Exchange Updates

VERSION	BLOG POST
Exchange 2016 CU16	Released: March 2020 Quarterly Exchange Updates
Exchange 2016 CU15	Released: December 2019 Quarterly Exchange Updates
Exchange 2016 CU14	Released: September 2019 Quarterly Exchange Updates
Exchange 2016 CU13	Released: June 2019 Quarterly Exchange Updates
Exchange 2016 CU12	Released: February 2019 Quarterly Exchange Updates
Exchange 2016 CU11	Released: October 2018 Quarterly Exchange Updates
Exchange 2016 CU10	Released: June 2018 Quarterly Exchange Updates
Exchange 2016 CU9	Released: March 2018 Quarterly Exchange Updates
Exchange 2016 CU8	Released: December 2017 Quarterly Exchange Updates
Exchange 2016 CU7	Released: September 2017 Quarterly Exchange Updates
Exchange 2016 CU6	Released: June 2017 Quarterly Exchange Updates
Exchange 2016 CU5	Released: March 2017 Quarterly Exchange Updates
Exchange 2016 CU4	Released: December 2016 Quarterly Exchange Updates
Exchange 2016 CU3	Released: September 2016 Quarterly Exchange Updates
Exchange 2016 CU2	Released: June 2016 Quarterly Exchange Updates
Exchange 2016 CU1	Released: March 2016 Quarterly Exchange Updates
Exchange 2016 RTM	Exchange Server 2016: Forged in the cloud. Now available on-premises

For information about the new features you'll get when you upgrade to Exchange 2016 from previous versions of Exchange, see What's new in Exchange Server.

- To upgrade to the latest CU after you've downloaded it, see Upgrade Exchange to the latest Cumulative Update.
- For downloads and updates for other versions of Exchange, see Exchange Server build numbers and release dates.

Exchange Server build numbers and release dates

8/3/2020 • 13 minutes to read • Edit Online

You can use the information in this topic to verify the version of Exchange that is running in your organization.

This topic is organized in sections that correspond to the major releases of Exchange. Each section lists build numbers for each Service Pack (SP), Cumulative Update (CU), or Update Rollup (RU) of the specific Exchange release.

Download links for the latest CU, RU, and SP for Exchange Server 2019, Exchange Server 2016, Exchange Server 2013, Exchange Server 2010, and Exchange Server 2007 are included.

> **NOTE**
>
> In the following sections, RTM stands for release to manufacturing (the first version of the product).

Exchange Server 2019

The table in this section provides build numbers and general release dates for each version of Microsoft Exchange Server 2019.

To view the build number of an Exchange 2019 server, run the following command in the Exchange Management Shell.

```
Get-ExchangeServer | Format-List Name,Edition,AdminDisplayVersion
```

PRODUCT NAME	RELEASE DATE	BUILD NUMBER (SHORT FORMAT)	BUILD NUMBER (LONG FORMAT)
Exchange Server 2019 CU6	June 16, 2020	15.2.659.4	15.02.0659.004
Exchange Server 2019 CU5	March 17, 2020	15.2.595.3	15.02.0595.003
Exchange Server 2019 CU4	December 17, 2019	15.2.529.5	15.02.0529.005
Exchange Server 2019 CU3	September 17, 2019	15.2.464.5	15.02.0464.005
Exchange Server 2019 CU2	June 18, 2019	15.2.397.3	15.02.0397.003
Exchange Server 2019 CU1	February 12, 2019	15.2.330.5	15.02.0330.005
Exchange Server 2019 RTM	October 22, 2018	15.2.221.12	15.02.0221.012
Exchange Server 2019 Preview	July 24, 2018	15.2.196.0	15.02.0196.000

Exchange Server 2016

The table in this section provides build numbers and general release dates for each version of Microsoft Exchange

Server 2016.

To view the build number of an Exchange 2016 server, run the following command in the Exchange Management Shell.

```
Get-ExchangeServer | Format-List Name,Edition,AdminDisplayVersion
```

PRODUCT NAME	RELEASE DATE	BUILD NUMBER (SHORT FORMAT)	BUILD NUMBER (LONG FORMAT)
Exchange Server 2016 CU17	June 16, 2020	15.1.2044.4	15.01.2044.004
Exchange Server 2016 CU16	March 17, 2020	15.1.1979.3	15.01.1979.003
Exchange Server 2016 CU15	December 17, 2019	15.1.1913.5	15.01.1913.005
Exchange Server 2016 CU14	September 17, 2019	15.1.1847.3	15.01.1847.003
Exchange Server 2016 CU13	June 18, 2019	15.1.1779.2	15.01.1779.002
Exchange Server 2016 CU12	February 12, 2019	15.1.1713.5	15.01.1713.005
Exchange Server 2016 CU11	October 16, 2018	15.1.1591.10	15.01.1591.010
Exchange Server 2016 CU10	June 19, 2018	15.1.1531.3	15.01.1531.003
Exchange Server 2016 CU9	March 20, 2018	15.1.1466.3	15.01.1466.003
Exchange Server 2016 CU8	December 19, 2017	15.1.1415.2	15.01.1415.002
Exchange Server 2016 CU7	September 19, 2017	15.1.1261.35	15.01.1261.035
Exchange Server 2016 CU6	June 27, 2017	15.1.1034.26	15.01.1034.026
Exchange Server 2016 CU5	March 21, 2017	15.1.845.34	15.01.0845.034
Exchange Server 2016 CU4	December 13, 2016	15.1.669.32	15.01.0669.032
Exchange Server 2016 CU3	September 20, 2016	15.1.544.27	15.01.0544.027
Exchange Server 2016 CU2	June 21, 2016	15.1.466.34	15.01.0466.034
Exchange Server 2016 CU1	March 15, 2016	15.1.396.30	15.01.0396.030
Exchange Server 2016 RTM	October 1, 2015	15.1.225.42	15.01.0225.042
Exchange Server 2016 Preview	July 22, 2015	15.1.225.16	15.01.0225.016

Exchange Server 2013

The table in this section provides build numbers and general release dates for each version of Microsoft Exchange Server 2013.

To view the build number of an Exchange 2013 server, run the following command in the Exchange Management Shell.

```
Get-ExchangeServer | Format-List Name,Edition,AdminDisplayVersion
```

PRODUCT NAME	RELEASE DATE	BUILD NUMBER (SHORT FORMAT)	BUILD NUMBER (LONG FORMAT)
Exchange Server 2013 CU23	June 18, 2019	15.0.1497.2	15.00.1497.002
Exchange Server 2013 CU22	February 12, 2019	15.0.1473.3	15.00.1473.003
Exchange Server 2013 CU21	June 19, 2018	15.0.1395.4	15.00.1395.004
Exchange Server 2013 CU20	March 20, 2018	15.0.1367.3	15.00.1367.003
Exchange Server 2013 CU19	December 19, 2017	15.0.1365.1	15.00.1365.001
Exchange Server 2013 CU18	September 19, 2017	15.0.1347.2	15.00.1347.002
Exchange Server 2013 CU17	June 27, 2017	15.0.1320.4	15.00.1320.004
Exchange Server 2013 CU16	March 21, 2017	15.0.1293.2	15.00.1293.002
Exchange Server 2013 CU15	December 13, 2016	15.0.1263.5	15.00.1263.005
Exchange Server 2013 CU14	September 20, 2016	15.0.1236.3	15.00.1236.003
Exchange Server 2013 CU13	June 21, 2016	15.0.1210.3	15.00.1210.003
Exchange Server 2013 CU12	March 15, 2016	15.0.1178.4	15.00.1178.004
Exchange Server 2013 CU11	December 15, 2015	15.0.1156.6	15.00.1156.006
Exchange Server 2013 CU10	September 15, 2015	15.0.1130.7	15.00.1130.007
Exchange Server 2013 CU9	June 17, 2015	15.0.1104.5	15.00.1104.005
Exchange Server 2013 CU8	March 17, 2015	15.0.1076.9	15.00.1076.009
Exchange Server 2013 CU7	December 9, 2014	15.0.1044.25	15.00.1044.025
Exchange Server 2013 CU6	August 26, 2014	15.0.995.29	15.00.0995.029
Exchange Server 2013 CU5	May 27, 2014	15.0.913.22	15.00.0913.022
Exchange Server 2013 SP1	February 25, 2014	15.0.847.32	15.00.0847.032
Exchange Server 2013 CU3	November 25, 2013	15.0.775.38	15.00.0775.038
Exchange Server 2013 CU2	July 9, 2013	15.0.712.24	15.00.0712.024

PRODUCT NAME	RELEASE DATE	BUILD NUMBER (SHORT FORMAT)	BUILD NUMBER (LONG FORMAT)
Exchange Server 2013 CU1	April 2, 2013	15.0.620.29	15.00.0620.029
Exchange Server 2013 RTM	December 3, 2012	15.0.516.32	15.00.0516.032

Exchange Server 2010

The tables in this section provide build numbers and general release dates for each version of Microsoft Exchange Server 2010.

To view the build number of an Exchange 2010 server, run the following command in the Exchange Management Shell:

```
Get-Command ExSetup | ForEach {$_.FileVersionInfo}
```

Exchange Server 2010 SP3 build numbers

PRODUCT NAME	RELEASE DATE	BUILD NUMBER (SHORT FORMAT)	BUILD NUMBER (LONG FORMAT)
Update Rollup 30 for Exchange Server 2010 SP3	February 11, 2020	14.3.496.0	14.03.0496.000
Update Rollup 29 for Exchange Server 2010 SP3	July 9, 2019	14.3.468.0	14.03.0468.000
Update Rollup 28 for Exchange Server 2010 SP3	June 7, 2019	14.3.461.1	14.03.0461.001
Update Rollup 27 for Exchange Server 2010 SP3	April 9, 2019	14.3.452.0	14.03.0452.000
Update Rollup 26 for Exchange Server 2010 SP3	February 12, 2019	14.3.442.0	14.03.0442.000
Update Rollup 25 for Exchange Server 2010 SP3	January 8, 2019	14.3.435.0	14.03.0435.000
Update Rollup 24 for Exchange Server 2010 SP3	September 5, 2018	14.3.419.0	14.03.0419.000
Update Rollup 23 for Exchange Server 2010 SP3	August 13, 2018	14.3.417.1	14.03.0417.001
Update Rollup 22 for Exchange Server 2010 SP3	June 19, 2018	14.3.411.0	14.03.0411.000
Update Rollup 21 for Exchange Server 2010 SP3	May 7, 2018	14.3.399.2	14.03.0399.002
Update Rollup 20 for Exchange Server 2010 SP3	March 5, 2018	14.3.389.1	14.03.0389.001

PRODUCT NAME	RELEASE DATE	BUILD NUMBER (SHORT FORMAT)	BUILD NUMBER (LONG FORMAT)
Update Rollup 19 for Exchange Server 2010 SP3	December 19, 2017	14.3.382.0	14.03.0382.000
Update Rollup 18 for Exchange Server 2010 SP3	July 11, 2017	14.3.361.1	14.03.0361.001
Update Rollup 17 for Exchange Server 2010 SP3	March 21, 2017	14.3.352.0	14.03.0352.000
Update Rollup 16 for Exchange Server 2010 SP3	December 13, 2016	14.3.336.0	14.03.0336.000
Update Rollup 15 for Exchange Server 2010 SP3	September 20, 2016	14.3.319.2	14.03.0319.002
Update Rollup 14 for Exchange Server 2010 SP3	June 21, 2016	14.3.301.0	14.03.0301.000
Update Rollup 13 for Exchange Server 2010 SP3	March 15, 2016	14.3.294.0	14.03.0294.000
Update Rollup 12 for Exchange Server 2010 SP3	December 15, 2015	14.3.279.2	14.03.0279.002
Update Rollup 11 for Exchange Server 2010 SP3	September 15, 2015	14.3.266.2	14.03.0266.002
Update Rollup 10 for Exchange Server 2010 SP3	June 17, 2015	14.3.248.2	14.03.0248.002
Update Rollup 9 for Exchange Server 2010 SP3	March 17, 2015	14.3.235.1	14.03.0235.001
Update Rollup 8 v2 for Exchange Server 2010 SP3	December 12, 2014	14.3.224.2	14.03.0224.002
Update Rollup 8 v1 for Exchange Server 2010 SP3 (recalled)	December 9, 2014	14.3.224.1	14.03.0224.001
Update Rollup 7 for Exchange Server 2010 SP3	August 26, 2014	14.3.210.2	14.03.0210.002
Update Rollup 6 for Exchange Server 2010 SP3	May 27, 2014	14.3.195.1	14.03.0195.001
Update Rollup 5 for Exchange Server 2010 SP3	February 24, 2014	14.3.181.6	14.03.0181.006
Update Rollup 4 for Exchange Server 2010 SP3	December 9, 2013	14.3.174.1	14.03.0174.001

PRODUCT NAME	RELEASE DATE	BUILD NUMBER (SHORT FORMAT)	BUILD NUMBER (LONG FORMAT)
Update Rollup 3 for Exchange Server 2010 SP3	November 25, 2013	14.3.169.1	14.03.0169.001
Update Rollup 2 for Exchange Server 2010 SP3	August 8, 2013	14.3.158.1	14.03.0158.001
Update Rollup 1 for Exchange Server 2010 SP3	May 29, 2013	14.3.146.0	14.03.0146.000
Exchange Server 2010 SP3	February 12, 2013	14.3.123.4	14.03.0123.004

Build numbers for previous releases of Exchange Server 2010

PRODUCT NAME	RELEASE DATE	BUILD NUMBER (SHORT FORMAT)	BUILD NUMBER (LONG FORMAT)
Update Rollup 8 for Exchange Server 2010 SP2	December 9, 2013	14.2.390.3	14.02.0390.003
Update Rollup 7 for Exchange Server 2010 SP2	August 3, 2013	14.2.375.0	14.02.0375.000
Update Rollup 6 Exchange Server 2010 SP2	February 12, 2013	14.2.342.3	14.02.0342.003
Update Rollup 5 v2 for Exchange Server 2010 SP2	December 10, 2012	14.2.328.10	14.02.0328.010
Update Rollup 5 for Exchange Server 2010 SP2	November 13, 2012	14.3.328.5	14.03.0328.005
Update Rollup 4 v2 for Exchange Server 2010 SP2	October 9, 2012	14.2.318.4	14.02.0318.004
Update Rollup 4 for Exchange Server 2010 SP2	August 13, 2012	14.2.318.2	14.02.0318.002
Update Rollup 3 for Exchange Server 2010 SP2	May 29, 2012	14.2.309.2	14.02.0309.002
Update Rollup 2 for Exchange Server 2010 SP2	April 16, 2012	14.2.298.4	14.02.0298.004
Update Rollup 1 for Exchange Server 2010 SP2	February 13, 2012	14.2.283.3	14.02.0283.003
Exchange Server 2010 SP2	December 4, 2011	14.2.247.5	14.02.0247.005
Update Rollup 8 for Exchange Server 2010 SP1	December 10, 2012	14.1.438.0	14.01.0438.000

PRODUCT NAME	RELEASE DATE	BUILD NUMBER (SHORT FORMAT)	BUILD NUMBER (LONG FORMAT)
Update Rollup 7 v3 for Exchange Server 2010 SP1	November 13, 2012	14.1.421.3	14.01.0421.003
Update Rollup 7 v2 for Exchange Server 2010 SP1	October 10, 2012	14.1.421.2	14.01.0421.002
Update Rollup 7 for Exchange Server 2010 SP1	August 8, 2012	14.1.421.0	14.01.0421.000
Update Rollup 6 for Exchange Server 2010 SP1	October 27, 2011	14.1.355.2	14.01.0355.002
Update Rollup 5 for Exchange Server 2010 SP1	August 23, 2011	14.1.339.1	14.01.0339.001
Update Rollup 4 for Exchange Server 2010 SP1	July 27, 2011	14.1.323.6	14.01.0323.006
Update Rollup 3 for Exchange Server 2010 SP1	April 6, 2011	14.1.289.7	14.01.0289.007
Update Rollup 2 for Exchange Server 2010 SP1	December 9, 2010	14.1.270.1	14.01.0270.001
Update Rollup 1 for Exchange Server 2010 SP1	October 4, 2010	14.1.255.2	14.01.0255.002
Exchange Server 2010 SP1	August 23, 2010	14.1.218.15	14.01.0218.015
Update Rollup 5 for Exchange Server 2010	December 13, 2010	14.0.726.0	14.00.0726.000
Update Rollup 4 for Exchange Server 2010	June 10, 2010	14.0.702.1	14.00.0702.001
Update Rollup 3 for Exchange Server 2010	April 13, 2010	14.0.694.0	14.00.0694.000
Update Rollup 2 for Exchange Server 2010	March 4, 2010	14.0.689.0	14.00.0689.000
Update Rollup 1 for Exchange Server 2010	December 9, 2009	14.0.682.1	14.00.0682.001
Exchange Server 2010 RTM	November 9, 2009	14.0.639.21	14.00.0639.021

Exchange Server 2007

The tables in this section provide build numbers and general release dates for each version of Microsoft Exchange Server 2007.

NOTE

The version information for Exchange Server 2007 SP1 is displayed correctly in the Exchange Management Console, in the Exchange Management Shell, and in the **About Exchange Server 2007 Help** dialog box. However, after you apply Exchange 2007 SP1 to an Edge Transport server that's running the RTM version of Exchange 2007, the version information for the Edge Transport server isn't updated in the Exchange Management Console unless the Edge Transport server is resubscribed to the Active Directory site. This is because the Edge Transport server doesn't directly update Active Directory by using any configuration information. Instead, the version information for Edge Transport servers is recorded in Active Directory during the creation of an Edge Subscription.

To view the build number of an Exchange 2007 server, run the following command in the Exchange Management Shell:

```
Get-Command ExSetup | ForEach {$_.FileVersionInfo}
```

PRODUCT NAME	RELEASE DATE	BUILD NUMBER (SHORT FORMAT)	BUILD NUMBER (LONG FORMAT)
Update Rollup 23 for Exchange Server 2007 SP3	March 21, 2017	8.3.517.0	8.03.0517.000
Update Rollup 22 for Exchange Server 2007 SP3	December 13, 2016	8.3.502.0	8.03.0502.000
Update Rollup 21 for Exchange Server 2007 SP3	September 20, 2016	8.3.485.1	8.03.0485.001
Update Rollup 20 for Exchange Server 2007 SP3	June 21, 2016	8.3.468.0	8.03.0468.000
Update Rollup 19 forExchange Server 2007 SP3	March 15, 2016	8.3.459.0	8.03.0459.000
Update Rollup 18 forExchange Server 2007 SP3	December, 2015	8.3.445.0	8.03.0445.000
Update Rollup 17 forExchange Server 2007 SP3	June 17, 2015	8.3.417.1	8.03.0417.001
Update Rollup 16 for Exchange Server 2007 SP3	March 17, 2015	8.3.406.0	8.03.0406.000
Update Rollup 15 for Exchange Server 2007 SP3	December 9, 2014	8.3.389.2	8.03.0389.002
Update Rollup 14 for Exchange Server 2007 SP3	August 26, 2014	8.3.379.2	8.03.0379.002
Update Rollup 13 for Exchange Server 2007 SP3	February 24, 2014	8.3.348.2	8.03.0348.002

PRODUCT NAME	RELEASE DATE	BUILD NUMBER (SHORT FORMAT)	BUILD NUMBER (LONG FORMAT)
Update Rollup 12 for Exchange Server 2007 SP3	December 9, 2013	8.3.342.4	8.03.0342.004
Update Rollup 11 for Exchange Server 2007 SP3	August 13, 2013	8.3.327.1	8.03.0327.001
Update Rollup 10 for Exchange Server 2007 SP3	February 11, 2013	8.3.298.3	8.03.0298.003
Update Rollup 9 for Exchange Server 2007 SP3	December 10, 2012	8.3.297.2	8.03.0297.002
Update Rollup 8-v3 for Exchange Server 2007 SP3	November 13, 2012	8.3.279.6	8.03.0279.006
Update Rollup 8-v2 for Exchange Server 2007 SP3	October 9, 2012	8.3.279.5	8.03.0279.005
Update Rollup 8 for Exchange Server 2007 SP3	August 13, 2012	8.3.279.3	8.03.0279.003
Update Rollup 7 for Exchange Server 2007 SP3	April 16, 2012	8.3.264.0	8.03.0264.000
Update Rollup 6 for Exchange Server 2007 SP3	January 26, 2012	8.3.245.2	8.03.0245.002
Update Rollup 5 for Exchange Server 2007 SP3	September 21, 2011	8.3.213.1	8.03.0213.001
Update Rollup 4 for Exchange Server 2007 SP3	May 28, 2011	8.3.192.1	8.03.0192.001
Update Rollup 3-v2 for Exchange Server 2007 SP3	March 30, 2011	8.3.159.2	8.03.0159.002
Update Rollup 2 for Exchange Server 2007 SP3	December 10, 2010	8.3.137.3	8.03.0137.003
Update Rollup 1 for Exchange Server 2007 SP3	September 9, 2010	8.3.106.2	8.03.0106.002
Exchange Server 2007 SP3	June 7, 2010	8.3.83.6	8.03.0083.006

Build numbers for previous releases of Exchange Server 2007

PRODUCT NAME	RELEASE DATE	BUILD NUMBER (SHORT FORMAT)	BUILD NUMBER (LONG FORMAT)
Update Rollup 5 for Exchange Server 2007 SP2	December 7, 2010	8.2.305.3	8.02.0305.003

PRODUCT NAME	RELEASE DATE	BUILD NUMBER (SHORT FORMAT)	BUILD NUMBER (LONG FORMAT)
Update Rollup 4 for Exchange Server 2007 SP2	April 9, 2010	8.2.254.0	8.02.0254.000
Update Rollup 3 for Exchange Server 2007 SP2	March 17, 2010	8.2.247.2	8.02.0247.002
Update Rollup 2 for Exchange Server 2007 SP2	January 22, 2010	8.2.234.1	8.02.0234.001
Update Rollup 1 for Exchange Server 2007 SP2	November 19, 2009	8.2.217.3	8.02.0217.003
Exchange Server 2007 SP2	August 24, 2009	8.2.176.2	8.02.0176.002
Update Rollup 10 for Exchange Server 2007 SP1	April 13, 2010	8.1.436.0	8.01.0436.000
Update Rollup 9 for Exchange Server 2007 SP1	July 16, 2009	8.1.393.1	8.01.0393.001
Update Rollup 8 for Exchange Server 2007 SP1	May 19, 2009	8.1.375.2	8.01.0375.002
Update Rollup 7 for Exchange Server 2007 SP1	March 18, 2009	8.1.359.2	8.01.0359.002
Update Rollup 6 for Exchange Server 2007 SP1	February 10, 2009	8.1.340.1	8.01.0340.001
Update Rollup 5 for Exchange Server 2007 SP1	November 20, 2008	8.1.336.1	8.01.0336.01
Update Rollup 4 for Exchange Server 2007 SP1	October 7, 2008	8.1.311.3	8.01.0311.003
Update Rollup 3 for Exchange Server 2007 SP1	July 8, 2008	8.1.291.2	8.01.0291.002
Update Rollup 2 for Exchange Server 2007 SP1	May 9, 2008	8.1.278.2	8.01.0278.002
Update Rollup 1 for Exchange Server 2007 SP1	February 28, 2008	8.1.263.1	8.01.0263.001
Exchange Server 2007 SP1	November 29, 2007	8.1.240.6	8.01.0240.006
Update Rollup 7 for Exchange Server 2007	July 8, 2008	8.0.813.0	8.00.0813.000

PRODUCT NAME	RELEASE DATE	BUILD NUMBER (SHORT FORMAT)	BUILD NUMBER (LONG FORMAT)
Update Rollup 6 for Exchange Server 2007	February 21, 2008	8.0.783.2	8.00.0783.002
Update Rollup 5 for Exchange Server 2007	October 25, 2007	8.0.754.0	8.00.0754.000
Update Rollup 4 for Exchange Server 2007	August 23, 2007	8.0.744.0	8.00.0744.000
Update Rollup 3 for Exchange Server 2007	June 28, 2007	8.0.730.1	8.00.0730.001
Update Rollup 2 for Exchange Server 2007	May 8, 2007	8.0.711.2	8.00.0711.002
Update Rollup 1 for Exchange Server 2007	April 17, 2007	8.0.708.3	8.00.0708.003
Exchange Server 2007 RTM	March 8, 2007	8.0.685.25	8.00.0685.025

Exchange Server 2003

The following table lists the build numbers and general release dates for each version of Microsoft Exchange Server 2003. To view the build number of Exchange Server 2003, open the **Properties** dialog box of the server object.

PRODUCT NAME	RELEASE DATE	BUILD NUMBER
Exchange Server 2003 post-SP2	August 2008	6.5.7654.4
Exchange Server 2003 post-SP2	March 2008	6.5.7653.33
Exchange Server 2003 SP2	October 19, 2005	6.5.7683
Exchange Server 2003 SP1	May25, 2004	6.5.7226
Exchange Server 2003	September 28, 2003	6.5.6944

Exchange 2000 Server

The following table lists the build numbers and general release dates for each version of Microsoft Exchange 2000 Server. To view the build number of Exchange 2000 Server, open the **Properties** dialog box of the server object.

PRODUCT NAME	RELEASE DATE	BUILD NUMBER
Exchange 2000 Server post-SP3	August 2008	6.0.6620.7
Exchange 2000 Server post-SP3	March 2008	6.0.6620.5
Exchange 2000 Server post-SP3	August 2004	6.0.6603

PRODUCT NAME	RELEASE DATE	BUILD NUMBER
Exchange 2000 Server post-SP3	April 2004	6.0.6556
Exchange 2000 Server post-SP3	September 2003	6.0.6487
Exchange 2000 Server SP3	July 18, 2002	6.0.6249
Exchange 2000 Server SP2	November 29, 2001	6.0.5762
Exchange 2000 Server SP1	June 21, 2001	6.0.4712
Exchange 2000 Server	November 29, 2000	6.0.4417

Exchange Server 5.5

The following table lists the build numbers and general release dates for each version of Microsoft Exchange Server version 5.5.

PRODUCT NAME	RELEASE DATE	BUILD NUMBER
Exchange Server version 5.5 SP4	November 1, 2000	5.5.2653
Exchange Server version 5.5 SP3	September 9, 1999	5.5.2650
Exchange Server version 5.5 SP2	December 23, 1998	5.5.2448
Exchange Server version 5.5 SP1	August 5, 1998	5.5.2232
Exchange Server version 5.5	February 3, 1998	5.5.1960

Exchange Server 5.0

PRODUCT NAME	RELEASE DATE	BUILD NUMBER
Exchange Server 5.0 SP2	February 19, 1998	5.0.1460
Exchange Server 5.0 SP1	June 18, 1997	5.0.1458
Exchange Server 5.0	May 23, 1997	5.0.1457

Exchange Server 4.0

The following table lists the build numbers and general release dates for each version of Microsoft Exchange Server 4.0.

PRODUCT NAME	RELEASE DATE	BUILD NUMBER
Exchange Server 4.0 SP5	May 5, 1998	4.0.996
Exchange Server 4.0 SP4	March 28, 1997	4.0.995

PRODUCT NAME	RELEASE DATE	BUILD NUMBER
Exchange Server 4.0 SP3	October 29, 1996	4.0.994
Exchange Server 4.0 SP2	July 19, 1996	4.0.993
Exchange Server 4.0 SP1	May 1, 1996	4.0.838
Exchange Server 4.0 Standard Edition	June 11, 1996	4.0.837

Release notes for Exchange Server

8/3/2020 • 5 minutes to read • Edit Online

TIP

Looking for the Exchange 2013 release notes? See Release notes for Exchange 2013.

Welcome to Microsoft Exchange Server 2019! This topic contains important information that you need to know to successfully deploy Exchange 2019. Please read this topic completely before beginning your deployment.

Known issues in Exchange Server 2019

When you attempt to uninstall Exchange Server from Windows 2019 Server Core using the Exchange Setup Wizard, the operation will fail. The wizard attempts to launch the Windows Control Panel to uninstall Exchange, but the Control Panel does not exist in Windows Server Core. To uninstall Exchange from Windows Server Core, run the following Setup command from the command line:

~~~~
```PowerShell
Setup.exe /IAcceptExchangeServerLicenseTerms /mode:Uninstall
```
~~~~

This issue will be resolved in a future CU update for Exchange Server 2019.

Welcome to Microsoft Exchange Server 2016! This topic contains important information that you need to know to successfully deploy Exchange 2016. Please read this topic completely before beginning your deployment.

Setup

- **Installing Exchange using Delegate Admin permissions causes Setup to fail**: When a user who is a member of only the Delegated Setup role group attempts to install Exchange on a pre-provisioned server, Setup will fail. This happens because the Delegated Setup group lacks the permissions required to create and configure certain objects in Active Directory.

  To work around this issue, do one of the following:

  - Add the user installing Exchange to the Domain Admins Active Directory security group.

  - Install Exchange using a user that is a member of the Organization Management role group.

Mailbox

- **Moving mailboxes from earlier versions of Exchange to Exchange 2016 CU5 or later can fail**: When you attempt to move a mailbox from an earlier version of Exchange to Exchange CU5 or later using a migration batch request, the move might fail. This can happen if the migration system mailbox isn't located on an Exchange 2016 server with CU5 or later installed.

  Before you can move mailboxes to Exchange 2016 CU5 or later using a migration batch request, you need to move the migration mailbox to an Exchange server running CU5 or later using the following steps.

  1. Open the Exchange Management Shell on your Exchange 2016 Mailbox server.

  2. Run the following command to get a list of mailbox databases that are located on your Exchange

2016 servers. Copy the name of the mailbox database where you want to move the migration mailbox to the clipboard.

```
Get-MailboxDatabase | Where {$_.AdminDisplayVersion -Like "*15.1*"} | Format-Table Name,
ServerName
```

3. Run the following command to move the migration mailbox to your Exchange 2016 server. Paste the mailbox database name you copied in the previous step after *TargetDatabase*.

```
New-MoveRequest "Migration.8f3e7716-2011-43e4-96b1-aba62d229136" -TargetDatabase "<mailbox
database name>"
```

- **Mailbox servers running different versions of Exchange can be added to the same database availability group**: The **Add-DatabaseAvailabilityGroupServer** cmdlet and the Exchange admin center incorrectly allow an Exchange 2013 server to be added to an Exchange 2016-based database availability group (DAG), and vice versa. Exchange supports adding only Mailbox servers running the same version (Exchange 2013 versus Exchange 2016, for example) to a DAG. Additionally, the Exchange admin center displays both Exchange 2013 and Exchange 2016 servers in the list of servers available to add to a DAG. This could allow an administrator to inadvertently add a server running an incompatible version of Exchange to a DAG (for example, adding an Exchange 2013 server to an Exchange 2016-based DAG).

  There is currently no workaround for this issue. Administrators must be diligent when adding a Mailbox server to a DAG. Add only Exchange 2013 servers to Exchange 2013-based DAGs, and only Exchange 2016 servers to Exchange 2016-based DAGs. You can differentiate each version of Exchange by looking at the **Version** column in the list of servers in the Exchange admin center. The following are the server versions for Exchange 2013 and Exchange 2016:

  - **Exchange 2013** 15.0 (Build xxx.xx)
  - **Exchange 2016** 15.1 (Build xxx.xx)

- **Can't connect to archive mailbox when using MAPI over HTTP**: In Exchange 2016, MAPI over HTTP can be enabled per-mailbox. An issue exists that prevents users from accessing their archive mailbox, if one is configured, when the following are true:

  - MAPI over HTTP is enabled on the user's mailbox.
  - MAPI over HTTP is disabled at the organization level.

    When these conditions are true, the user won't be able to open their archive mailbox and they'll get the error **The set of folders cannot be opened. The attempt to log on to Microsoft Exchange has failed.**

    To work around the issue, do one of the following

  - Open the archive mailbox using Outlook on the web.
  - Disable MAPI over HTTP on the mailbox by running the following command.

  ```
  Set-CasMailbox <email address> -MapiHttpEnabled $False
  ```

- **Notifications Broker service stops after 30 seconds** When you start your Exchange server, you might notice the **Notifications Broker** service start and then stop after approximately 30 seconds. If you attempt to start the service manually, it will successfully start and then stop, again after approximately 30 seconds. No errors or warnings are included in the Event log.

This behavior is expected in on-premises deployments of Exchange 2016. The **Notifications Broker** service performs a configuration check on each time the server starts. If there is nothing for the **Notifications Broker** service to do, it stops automatically until the next time the server is restarted.

Mail flow

- **Edge Transport servers can reject mail sent to valid recipients** Exchange 2016 Edge Transport servers may reject messages sent to valid internal recipients when the following are true:

  - Exchange 2016 Cumulative Update 1 (CU1) is installed on the server.

  - Recipient validation is enabled on the server.

    When an Edge Transport rejects a message because of this issue, the sender will receive a non-delivery report (NDR) with the status code **5.1.10**, and the error **Recipient not found by SMTP address lookup**. The recipient won't receive the message.

    To work around the issue, do **one** of the following:

  - Disable recipient validation on the affected Edge Transport server(s) by running the following command.

    ```
    Set-RecipientFilterConfig -RecipientValidationEnabled $False
    ```

  - Disable the recipient validation cache on the affected Edge Transport server(s) by running the following command.

    ```
    Get-TransportService | Set-TransportService -RecipientValidationCacheEnabled $False
    ```

    Caution

    Disabling the recipient validation cache causes Exchange to verify that recipients on inbound messages are valid by querying the local instance of Active Directory Lightweight Directory Services. This can significantly increase the resources Exchange needs to process messages. Before you disable the recipient validation cache, verify that your server has sufficient capacity to handle the additional demand.

- Configure your firewall or external mail exchanger (MX) DNS record to send mail to an Edge Transport server that doesn't have Exchange 2016 Cumulative Update 1 installed. You might need to configure your firewall to allow TCP port 25 to connect to the new Internet-facing server.

- Configure your firewall or external MX DNS record to send mail to an Exchange 2016 Mailbox server. You might need to configure your firewall to allow TCP port 25 to connect to the new Internet-facing server.

Exchange architecture

8/3/2020 • 6 minutes to read • Edit Online

Exchange use a single building block architecture that provides email services for deployments at all sizes, from small organizations to the largest multi-national corporations. This architecture is describe in the following diagram.

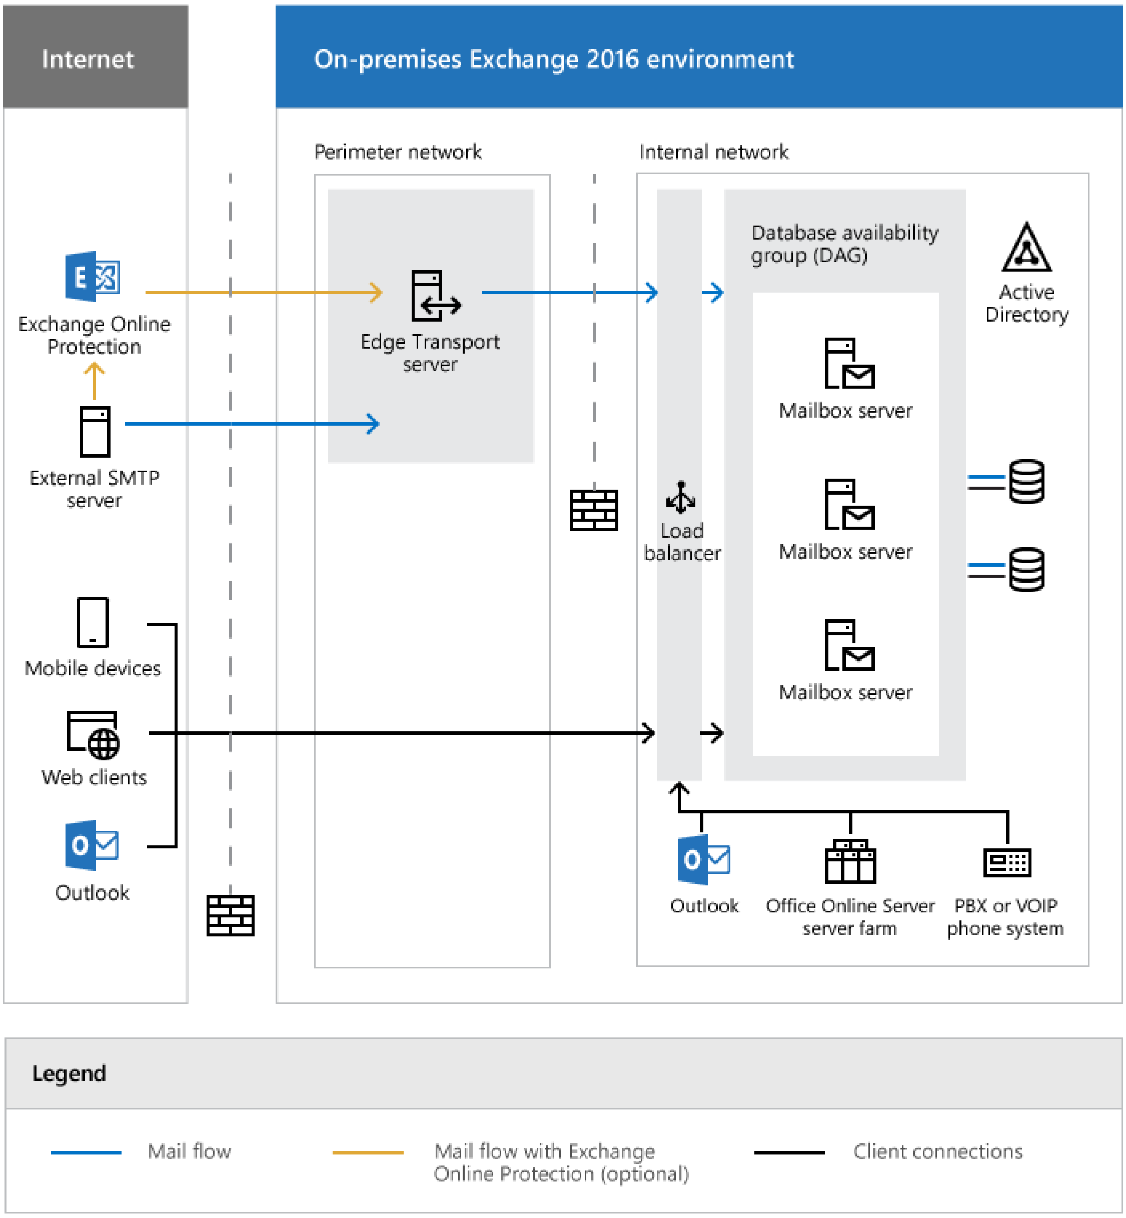

Individual components are described in the following sections.

Server communication architecture

Communication between Exchange servers and past and future versions of Exchange occurs at the protocol layer. Cross-layer communication isn't allowed. This communication architecture is summarized as "every server is an island". This architecture has the following benefits:

- Reduced inter-server communications.

- Version-aware communications.
- Isolated failures.
- Integrated design inside each server.

Protocol layer communication between Exchange servers is shown in the following diagram.

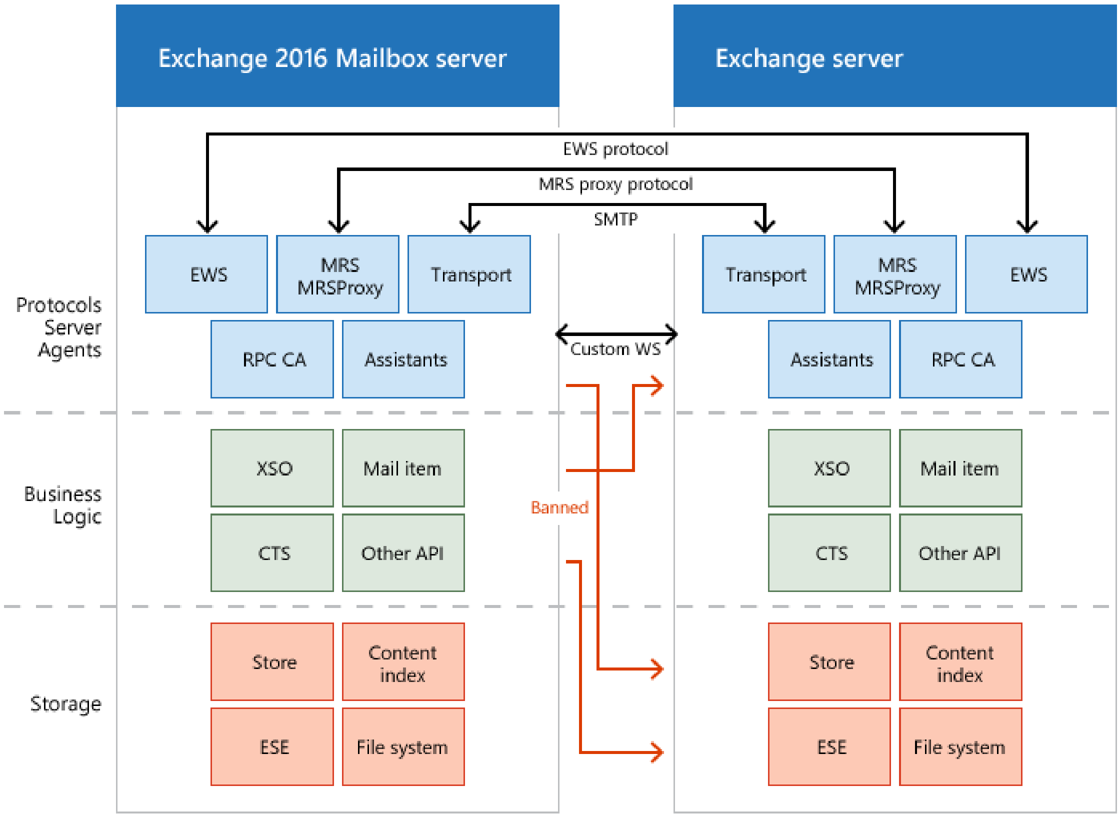

Server role architecture

Exchange uses Mailbox servers and Edge Transport servers. These server roles are described in the following sections.

Mailbox servers

- Mailbox servers contain the transport services that are used to route mail. For more information, see Mail flow and the transport pipeline
- Mailbox servers contain mailbox databases that process, render, and store data. For more information, see Manage mailbox databases in Exchange Server.
- Mailbox servers contain the Client Access services that accept client connections for all protocols. These frontend services are responsible for routing or *proxying* connections to the corresponding backend services on a Mailbox server. Clients don't connect directly to the backend services. For more information, see the Client Access protocol architecture section later in this topic.
- In Exchange 2016, Mailbox servers contain the Unified Messaging (UM) services that provide voice mail and other telephony features to mailboxes.

> **NOTE**
>
> Unified Messaging is not available in Exchange 2019.

- You manage Mailbox servers by using the Exchange admin center (EAC) and the Exchange Management Shell. For more information, see Exchange admin center in Exchange Server and Exchange Server PowerShell (Exchange Management Shell).

Edge Transport servers

- Edge Transport servers handle all external mail flow for the Exchange organization.

- Edge Transport servers are typically installed in the perimeter network, and are subscribed to the internal Exchange organization. The EdgeSync synchronization process makes recipient and other configuration information available to the Edge Transport server as mail enters and leaves the Exchange organization.

- Edge Transport servers provide antispam and mail flow rules as mail enters and leaves your Exchange organization. For more information, see Antispam protection in Exchange Server

- You manage Edge Transport servers by using the Exchange Management Shell. For more information, see Exchange Server PowerShell (Exchange Management Shell).

For more information about Edge Transport servers, see Edge Transport servers.

High availability architecture

The high availability features in Exchange Server are described in the following sections.

Mailbox high availability

A database availability group (DAG) is the fundamental element of the high availability and site resilience framework that's built into Exchange Server. A DAG is a group of Mailbox servers that host a set of databases and provides automatic, database-level recovery from database, network, and server failures. And DAGs in Exchange 2016 or later have been improved compared to Exchange 2013. For more information about DAGs, see Database availability groups.

Transport high availability

- The Transport service makes redundant copies of all messages in transit. This feature is known as *shadow redundancy*.

- The transport service makes redundant copies of all delivered messages. This feature is known as *Safety Net*.

- In Exchange Server, a DAG represents a transport high availability boundary. You can achieve site resilience by spanning a DAG across multiple Active Directory sites.

- In Exchange Server, transport high availability is more than a best effort for message redundancy, because redundancy doesn't depend on supported features of the sending mail server. Therefore, you can say that Exchange Server attempts to guarantee message redundancy by keeping multiple copies of messages during and after delivery.

For more information, see Transport high availability.

Client Access protocol architecture

The Client Access services on Exchange Mailbox servers are responsible for accepting all forms of client connections. The Client Access (frontend) services proxy these connections to the backend services on the destination Mailbox server (the local server or a remote Mailbox server that holds the active copy of the user's mailbox). Clients don't directly connect to the backend services. This communication is shown in the following diagram.

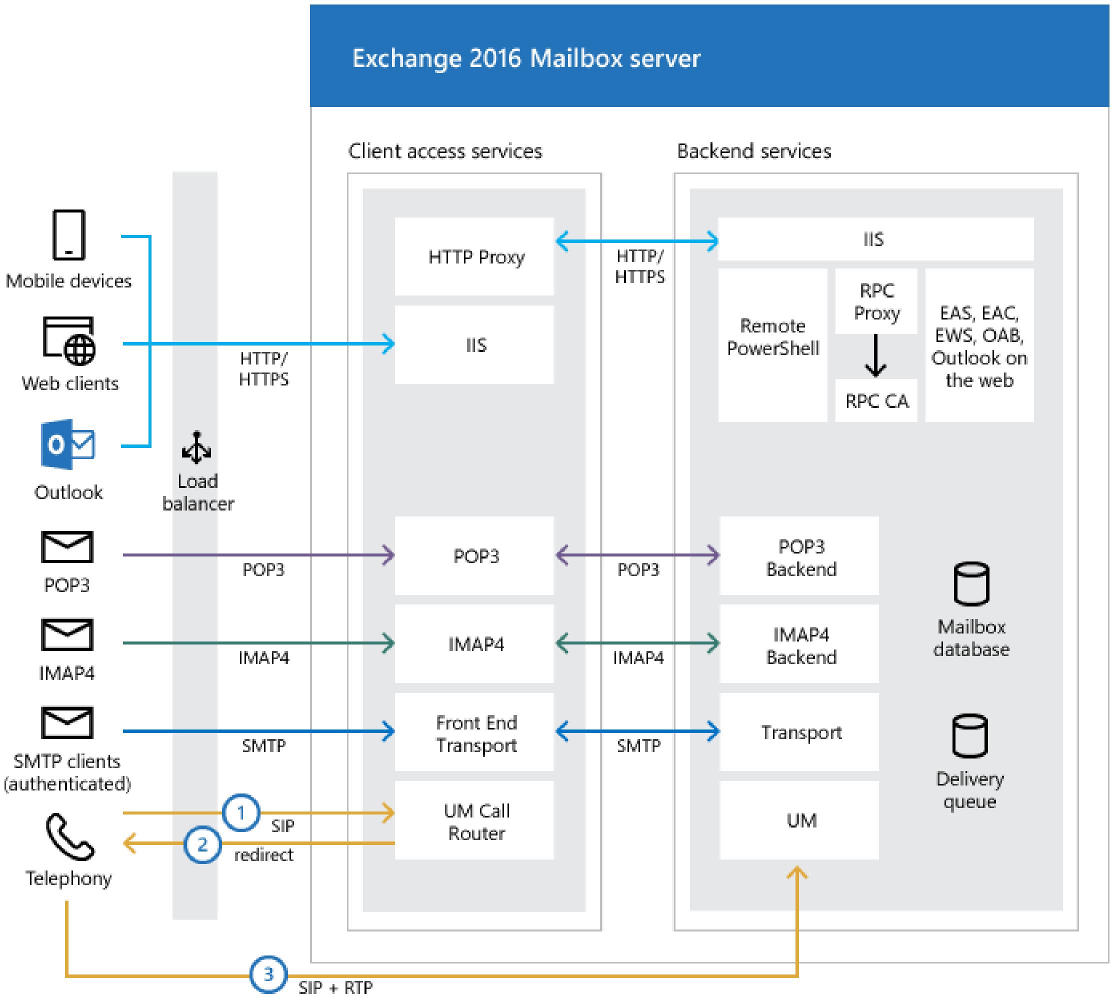

The protocol that's used by a client determines the protocol that's used to proxy the request to the backend services on the destination Mailbox server. For example, if the client connected using HTTP, the Mailbox server uses HTTP to proxy the request to the destination Mailbox server (secured via SSL using a self-signed certificate). If the client used IMAP or POP, then the protocol that's used is IMAP or POP.

In Exchange 2016, telephony requests are different than other client connections. Instead of proxying the request, the Mailbox server *redirects* the request to the Mailbox server that holds the active copy of the user's mailbox. Telephony devices are required to establish their SIP and RTP sessions directly with the Unified Messaging services on the destination Exchange 2016 Mailbox server.

NOTE

Unified Messaging is not available in Exchange 2019.

Exchange architecture changes

- **Server role consolidation**: In Exchange 2013 or earlier, you could install the Client Access server role and the Mailbox server role on separate computers. In Exchange 2016 or later, the Client Access server role is automatically installed as part of the Mailbox server role, and the Client Access server role isn't available as a separate installation option. This change reflects the philosophy of Exchange server role co-location that's been a recommended best practice since Exchange 2010. A multi-role Exchange server architecture gives you the following tangible benefits:
  - All Exchange servers in your environment (with the likely exception of any Edge Transport servers) can be exactly the same: the same hardware, the same configuration, etc. This uniformity simplifies hardware purchasing, and also maintenance and management of the Exchange servers.

  - You'll likely need fewer physical Exchange servers. This results in lower ongoing maintenance costs, fewer Exchange server licenses, and reduced rack, floor space, and power requirements.

  - Scalability is improved, because you're distributing the workload across a greater number of physical machines. During a failure, the load on the remaining Exchange multi-role servers increases only incrementally, which ensures the other functions on the Exchange servers aren't adversely affected.

  - Resiliency is improved, because a multi-role Exchange server can survive a greater number of Client Access role (or service) failures and still provide service.

- **Search improvements**: The local search instance is now able to read data from the local mailbox database copy. As a result, passive search instances no longer need to coordinate with their active counterparts to perform index updates, and bandwidth requirements between the active copy and a passive copy have been reduced by 40% compared to previous versions of Exchange. Also, search is now able to perform multiple asynchronous disk reads prior to a user completing a search term. This populates the cache with relevant information, and provides sub-second search query latency for online clients like Outlook on the web.

- **Office Online Server Preview for Outlook on the web document preview**: In Exchange 2013 or earlier, Outlook Web App included WebReady Document Viewing for the built-in preview of Office and PDF documents. In Exchange 2016 or later, Outlook on the web uses Office Online Server Preview to provide rich preview and editing capabilities for documents. While this provides a consistent document experience with other products like SharePoint and Skype for Business, it does require you to deploy Office Online Server Preview in your on-premises environment if you don't already have it. For more information, see Install Office Online Server in an Exchange organization.

- **MAPI over HTTP is the default for Outlook connections**: MAPI over HTTP was introduced in Exchange 2013 Service Pack 1, and offers improvements over the traditional Outlook Anywhere (RPC over HTTP) connection method. In Exchange 2016 or later, MAPI over HTTP is enabled by default, and offers additional controls, such as the ability to enable or disable MAPI over HTTP per user, and whether to advertise it to external clients. For more information, see MAPI over HTTP in Exchange Server.

Mailbox servers

8/3/2020 • 2 minutes to read • Edit Online

In Microsoft Exchange Server 2010, the Mailbox server role hosted both mailbox and public folder databases and also provided email message storage. In Exchange Server 2016 and Exchange Server 2019, the Mailbox server role contains transport services for routing mail, mailbox databases, Client access services to accept client connections, and (in Exchange 2016 only) Unified Messaging (UM) components.

For more information about Exchange mailbox servers and how they complement other server roles, see Exchange architecture.

Managing Mailbox servers

The following articles contain procedures for common Mailbox server management tasks:

- Manage mailbox databases in Exchange Server
- Manage on-premises mailbox moves in Exchange Server
- Prepare mailboxes for cross-forest move requests
- Enable the MRS Proxy endpoint for remote moves

Manage mailbox databases in Exchange Server

8/3/2020 • 8 minutes to read • Edit Online

A mailbox database is a unit of granularity where mailboxes are created and stored. A mailbox database is stored as an Exchange database (.edb) file. In Exchange 2016 and 2019, each mailbox database has its own properties that you can configure.

What do you need to know before you begin?

- Estimated time to complete each procedure: 10 minutes
- To open the Exchange admin center (EAC), see Exchange admin center in Exchange Server. To open the Exchange Management Shell, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Mailbox databases" entry in the Recipients Permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Create a mailbox database

Use the EAC to create a mailbox database

1. From the Exchange admin center (EAC), navigate to **Servers**.
2. Select **Databases**, and then click the + symbol to create a database.
3. Use the new database wizard to create your database.

Use the Exchange Management Shell to create a mailbox database

For an example of how to create a mailbox database, see Example 1 in New-MailboxDatabase.

How do you know this worked?

To verify that you have successfully created a database, do the following:

- From the EAC, verify that the mailbox database you created is listed on the **Databases** page.
- From the Exchange Management Shell, verify that the database was created on server Mailbox01 by running the following command.

```
Get-MailboxDatabase -Server "Mailbox01"
```

Get mailbox database properties

For detailed syntax and parameter information, see Get-MailboxDatabase.

Use the Exchange Management Shell to get mailbox database properties

For an example of how to get mailbox database properties, see Example 3 in Get-MailboxDatabase.

How do you know this worked?

To verify that you have successfully retrieved your mailbox database information, do the following:

- From the Exchange Management Shell, verify that all your mailbox database information is represented correctly.

Set mailbox database properties

Use the EAC to set mailbox database properties

1. From the EAC, navigate to **Servers**.

2. Select **Databases**, and then click to select the mailbox database you want to configure.

3. Click **Edit** ✎ to configure the attributes of a mailbox database.

4. Use the **General** tab to view status about the mailbox database, including the mailbox database path, last backup, and mailbox database status:

   - **Database path**: This read-only field displays the full path to the Exchange database (.edb) file for the selected mailbox database. To view the entire path, you may have to click the path and use the Right Arrow key. You can't use this field to change the path. To change the location of the database files, use the Move-DatabasePath cmdlet.

   - **Last full backup**: This read-only field displays the date and time of the last complete backup of the mailbox database.

   - **Last incremental backup**: This read-only field displays the date and time of the last incremental backup of the mailbox database.

   - **Status**: This read-only field displays whether the mailbox database is mounted or dismounted.

   - **Mounted on server**: This read-only field displays which server the database is mounted on.

   - **Master**: This read-only field displays the master server for the mailbox database. The Mailbox server that hosts the active copy of a database is referred to as the mailbox database master.

   - **Master type**: This read-only field displays the type of mailbox database master.

   - **Modified**: This read-only field displays the date and time the database was last modified.

   - **Servers hosting a copy of this database**: This read-only field displays the other servers that have a copy of this database.

5. Use the **Maintenance** tab to configure mailbox database settings, including specifying a journal recipient, setting a maintenance schedule, and mounting the database at startup:

   - **Journal Recipient**: Click **Browse** to specify a recipient to enable journaling on this mailbox database. Remove the recipient listed to disable journaling.

   - **Maintenance schedule**: Use this list to select one of the preset maintenance schedules. You can also configure a custom schedule. To configure a custom schedule, click **Customize**.

   - **Enable background database maintenance (24 x 7 ESE scanning)**: Select this check box to enable online database scanning, which runs continuously in the background. Online database scanning performs a checksum calculation of the database and performs operations that allow Exchange to scan for lost space on the database and recover it. If you select this check box, Exchange scans the database no more than one time per day and will issue a warning event if it can't finish scanning the database in a seven-day period.

- **Don't mount this database at startup**: Select this check box to prevent Exchange from mounting this mailbox database when it starts.
- **This database can be overwritten by a restore**: Select this check box to allow the mailbox database to be overwritten during a restore process.
- **Enable circular logging**: Select this check box to enable circular logging.

6. Use the **Limits** tab to specify the storage limits, the warning message interval, and the deletion settings for a mailbox database:

   - **Issue warning at (GB)**: Select this check box to automatically warn mailbox users that their mailbox is approaching its storage limit. To specify the storage limit, select the check box, and then specify in gigabytes (GB) how much content can be stored in the mailbox before a warning email message is sent to the mailbox users. You can enter a value from 0 through 2,097,151 megabytes (MB) (2.0 terabytes).
   - **Prohibit send at (GB)**: Select this check box to prevent users from sending new email messages after the size of their mailbox reaches the specified limit. To specify this limit, select the check box, and then type the size of the mailbox in GB at which you want to prohibit the sending of new email messages and notify the user. You can enter a value from 0 through 2,097,151 MB (2.0 terabytes).
   - **Prohibit send and receive at (GB)**: Select this check box to prevent users from sending and receiving email messages after their mailbox size reaches the specified limit. To specify this limit, select the check box, and then type the size of the mailbox in GB at which you want to prohibit the sending and receiving of email messages and notify the user. You can enter a value from 0 through 2,097,151 MB (2.0 terabytes).
   - **Keep deleted items for (days)**: Select this check box to set the number of days that deleted items are retained in a mailbox. You can enter a value from 0 through 24,855 days.
   - **Keep deleted mailboxes for (days)**: Select this check box to set the number of days that deleted mailboxes are retained. You can enter a value from 0 through 24,855 days.
   - **Don't permanently delete items until the database has been backed up**: Select this check box to prevent mailboxes and email messages from being deleted until after the mailbox database has been backed up.

7. Use the **Client Settings** tab to select the offline address book (OAB) for the mailbox:

   - **Offline address book**: To select an offline address book, click **Browse**, and then select the offline address book.

Use the Exchange Management Shell to set mailbox database properties

For an example of how to set mailbox database properties, see Example 1 in Set-MailboxDatabase.

How do you know this worked?

To verify that you have successfully set the attributes, do the following:

- Verify that your changes are saved in the EAC.
- From the Exchange Management Shell, run the following command to retrieve mailbox database properties.

  ```
  Get-MailboxDatabase -Identity MailboxDatabase01 -Status | Format-List
  ```

Move a mailbox database path

For detailed syntax and parameter information, see Move-DatabasePath.

Use the Exchange Management Shell to move a mailbox database path

For an example of how to set mailbox database properties, see Example 1 in Move-DatabasePath.

How do you know this worked?

To verify that you have successfully moved the database path, do the following:

1. From the EAC, select **Servers** > **Databases**, and then click to select the appropriate mailbox.

2. Click the **pen** symbol and verify that the database path is correct.

Mount a mailbox database

For detailed syntax and parameter information, see Mount-Database.

Use the Exchange Management Shell to mount a mailbox database

For an example of how to mount a mailbox database, see Example 1 in Mount-Database.

How do you know this worked?

To verify that you have successfully mounted the mailbox database, do the following.

- From the Exchange Management Shell, run the following command to retrieve mailbox database properties for all mailbox databases.

```
Get-MailboxDatabase -IncludePreExchange2013
```

Dismount a mailbox database

For detailed syntax and parameter information, see Dismount-Database.

Use the Exchange Management Shell to dismount a mailbox database

For an example of how to dismount a mailbox database, see Example 1 in Dismount-Database.

How do you know this worked?

To verify that you have successfully dismounted the database, do the following:

1. From EAC, select **Servers** > **Databases**, and then click to select the appropriate mailbox.

2. Click the **pen** symbol, and verify that the database status is **Dismounted**.

Remove a mailbox database

Use the EAC to remove a mailbox database

1. From the EAC, select **Servers** > **Databases**, and then click to select the appropriate mailbox.

2. Click **Delete** 🗑 to remove the mailbox database.

Use the Exchange Management Shell to remove a mailbox database

For detailed syntax and parameter information, see Remove-MailboxDatabase.

1. Run the following command to remove the mailbox database MyDatabase.

```
Remove-MailboxDatabase -Identity "MyDatabase"
```

2. When you're prompted about whether you're sure that you want to perform the action, type **Y**.

3. When the dialog box appears stating that the database was removed successfully, note the location of the Exchange database (.edb) file. If you want to remove this file from the hard drive, you must remove it manually.

How do you know this worked?

To verify that you have successfully removed the mailbox database, do the following:

- From the EAC, select **Servers** > **Databases**.

- Verify that the mailbox database has been removed.

Manage on-premises mailbox moves in Exchange Server

8/3/2020 • 12 minutes to read • Edit Online

In Exchange Server, users' primary mailboxes and archive mailboxes can reside on different databases. A *move request* is the process of moving a mailbox from one mailbox database to another. A *local move request* is a mailbox move that occurs within a single Active Directory forest (as opposed to a remote move request that occurs between Active Directory forests). You use the procedures in this topic for local move requests of primary mailboxes, archive mailboxes, or both in on-premises. Using the move request functionality, you can move the primary mailbox and the associated archive to the same database or to separate ones.

The following two services process your move request to move mailboxes:

- Exchange Mailbox Replication service (MRS)
- Exchange Mailbox Replication Proxy

The procedures in this topic will help you with on-premises mailbox moves. You can use the Exchange Management Shell and the Exchange admin center (EAC) to move mailboxes in your on-premises organization.

For more information about the Mailbox replication service and proxy, see Learn more about MRS Proxy. For more information about mailbox moves, see Mailbox moves in Exchange Server.

What do you need to know before you begin?

- Estimated time to complete each procedure: 20 minutes
- For more information about accessing and using the EAC, see Exchange admin center in Exchange Server. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Mailbox Move and Migration Permissions " entry in Recipients Permissions.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Create local move requests

You can create local move requests for:

- A single mailbox.
- Multiple mailboxes (also known as a *batch move request*).
- Multiple mailboxes that you specify in a comma-separated value (CSV) file (also known as a *migration batch*).

When you create local move requests in the EAC (for a single mailbox, multiple mailboxes, or multiple mailboxes specified in a CSV file), the request is visible to the **Get-MigrationBatch** cmdlet in the Exchange Management Shell. When the request has been completed (automatically or manually), the results for each individual mailbox are visible to the **Get-MoveRequest** cmdlet.

To create new local move requests in the Exchange Management Shell, you only use the **New-MigrationBatch** cmdlet for migration batches (the mailboxes are specified in a CSV file). To create local move requests that don't use a CSV file (individual mailboxes or batch move requests), you need to use the **New-MoveRequest** cmdlet, and these requests aren't visible to the **Get-MigrationBatch** batch cmdlet (or related ***-MigrationBatch*** cmdlets).

Use the EAC to create a local move request

1. In the EAC, go to **Recipients** > **Migration** > click **Add** +, and then select **Move to a different database**.

2. The **New local mailbox move** wizard opens. On the **Select users** page, configure one of these options:

   - **Select the users that you want to move**: Select one or more users:

     Note: Even if you're only interested in moving a user's archive mailbox, you select the user's primary mailbox.

     - Click **Add** +. In the **Select Mailbox** dialog box that appears, select one ore more mailboxes. When you're finished, click **OK**.

     - To remove mailboxes from the list, select the mailbox, and then click **Remove** −.

   - **Specify the users with a CSV file**: Click **Browse** and go to the location of the comma-separated value (CSV) file that specifies the mailboxes to move. For more information about the CSV file requirements for local move requests, see CSV Files for Mailbox Migration.

   - **Allow unknown columns in the CSV file**:

     - If you leave this check box unselected, move will ignore (silently skip) unknown columns in the CSV file (including optional columns with misspelled column headers). All unknown columns are treated like extra columns that aren't used.

     - If you select this check box, the migration fails if there are any unknown columns in the CSV file. This setting protects against spelling errors in the required and optional column headers, but the CSV file can't contain any unrelated columns.

   When you're finished, click **Next**.

3. On the **Move configuration** page, configure these settings:

   - **New migration batch name**: Enter a descriptive name for the mailbox move operation.

   - **Archive**: Select one of these options:

     - **Move the primary mailbox and the archive mailbox if one exists**

     - **Move primary mailbox only, without moving archive mailbox**

     - **Move archive mailbox only, without moving primary mailbox**

   - **Target database**: This setting affects moves for primary mailboxes.

     - To specify the database for the primary mailbox, click **Browse**. In the **Select Mailbox Database** dialog box that appears, select the database.

     - If you don't specify a database, the automatic distribution logic in Exchange will randomly select a database in the Active Directory site.

   - **Target archive database**: This setting affects moves for archive mailboxes.

     - To specify the database for the archive mailbox, click **Browse**. In the **Select Mailbox Database** dialog box that appears, select the database.

     - If you don't specify a database, the archive mailbox is moved to the same location as the primary mailbox.

   - **Bad item limit**: Specifies the maximum number of corrupted items that are allowed in the mailbox before the request fails. The default value in the EAC is 10. Don't specify a value greater than 50 here. If you want to set the limit to 51 or higher, use the *BadItemLimit* parameter and the *AcceptLargeDataLoss* switch in the Exchange Management Shell.

   When you're finished, click **Next**.

4. On the **Start the batch** page, configure these settings:

   - **After the batch is complete, a report will be sent to the following recipients.**: The default value is the account that you're using to move the mailboxes. Click **Browse** to add or remove recipients. When you're finished, click **OK**.

   - **Please select the preferred option to start the batch**: Select one of these options:

   - **Manually start the batch later**

   - **Automatically start the batch**: This is the default value.

   - **Please select the preferred option to complete the batch**: Select one of these options:

   - **Manually complete the batch**

   - **Automatically complete the migration batch**: This is the default value.

   When you're finished, click **New**.

Use the Exchange Management Shell to create a local move request for individual or multiple mailboxes

A local move request for an individual mailbox uses the **New-MailboxMove** cmdlet. But, a local move request for multiple mailboxes that doesn't specify the mailboxes in a CSV file also uses the **New-MailboxMove** cmdlet. A local move request for multiple mailboxes that doesn't use a CSV file is also known as a *batch move request*.

To create a local move request for an individual mailbox, use this syntax:

```
New-MoveRequest "<DescriptiveName>"] -Identity <MailboxIdentity> [<-ArchiveOnly | -PrimaryOnly>] [-
TargetDatabase <DatabaseIdentity>] [-ArchiveTargetDatabase<DatabaseIdentity>] [-Priority <PriorityValue>] [-
BadItemLimit <Value>] [-AcceptLargeDataLoss]
```

This example creates a new local move request with these settings:

- **Mailbox**: The primary mailbox and archive mailbox (if it exists) for Angela Gruber (agruber@contoso.com). If you only want to move the primary mailbox, use the *PrimaryOnly* switch. If you only want to move the archive mailbox, use the *ArchiveOnly* switch.

- **Target database for the primary mailbox**: MBX DB02. If we don't use the *TargetDatabase* parameter, the automatic distribution logic in Exchange will randomly select a database in the Active Directory site.

- **Target database for the archive mailbox**: MBX DB03. If we don't use the *ArchiveTargetDatabase* parameter or the *PrimaryOnly* switch, the archive mailbox database will be moved to the same database as the primary mailbox.

  If we use the *ArchiveOnly* switch without using the *ArchiveTargetDatabase* parameter, the automatic distribution logic in Exchange will randomly select a database in the Active Directory site.

- **Priority**: `Normal`, because we aren't using the *Priority* parameter.

- **Bad item limit**: 10 (the default value in the Exchange Management Shell is 0). Because the value is less than 51, we don't need to use the `AcceptLargeDataLoss` switch.

```
New-MoveRequest -Identity agruber@contoso.com -TargetDatabase "MBX 02" -ArchiveTargetDatabase "MBX 03" -
BadItemLimit 10
```

This example uses similar settings, but only moves Angela's primary mailbox.

```
New-MoveRequest -Identity agruber@contoso.com -PrimaryOnly-TargetDatabase "MBX 02" -BadItemLimit 10
```

This example uses similar settings, but only moves Angela's archive mailbox.

```
New-MoveRequest -Identity agruber@contoso.com -ArchiveOnly -ArchiveTargetDatabase "MBX 03" -BadItemLimit 10
```

For detailed syntax and parameter information, see New-MoveRequest.

A batch move request uses virtually the same syntax as a move request for an individual mailbox. The main differences are:

- You don't use the *Identity* parameter to specify the mailbox. Instead, you use the **Get-Mailbox** or **Get-User** cmdlets to generate the list of mailboxes that you want to move, and you pipeline the results to the **New-MoveRequest** cmdlet.

- You name the batch move with the *BatchName* parameter.

This example creates a batch move request with these settings:

- **Mailboxes to move**: All mailbox on the database named MBX DB01.

- **Batch name**: MBX DB01 to MBX DB02.

- **Target database**: MBX DB02. If we didn't use the *TargetDatabase* parameter, the automatic distribution logic in Exchange would randomly select databases in the Active Directory site.

- **Target database for archive mailboxes**: MBX DB02. Because we aren't using the *ArchiveTargetDatabase* parameter or the *PrimaryOnly* switch, the archive mailbox database is moved to the same database as the primary mailbox.

  If we use the *ArchiveOnly* switch without using the *ArchiveTargetDatabase* parameter, the automatic distribution logic in Exchange will randomly select databases in the Active Directory site.

- **Priority**: `High`

- **Bad item limit**: 51 (the default value in the Exchange Management Shell is 0), so we also need to use the *AcceptLargeDataLoss* switch.

```
Get-Mailbox -Database "MBX DB01" | New-MoveRequest -BatchName "MBX DB01 to MBX DB02" -TargetDatabase "MBX
DB02" -Priority High -BadItemLimit 51 -AcceptLargeDataLoss
```

For detailed syntax and parameter information, see New-MoveRequest.

Use the Exchange Management Shell to create a local move request from a CSV file

A local move request for mailboxes that are specified in a CSV file is known as a *migration batch*, and uses the **New-MigrationBatch** cmdlet.

For more information about the CSV file requirements for local move requests, see CSV Files for Mailbox Migration.

NOTE

All mailboxes that are specified in the CSV file will be migrated, even if they are outside of the RBAC scope (for example, an OU) that gives the admin permissions to migrate mailboxes.

To create a migration batch, use this syntax:

```
New-MigrationBatch -Local [-AutoStart] [-AutoComplete] -Name "<MigrationBatchName>" -CSVData ([Byte[]](Get-Content -Encoding Byte -Path "<PathAndFileName>" -ReadCount 0)) [<-ArchiveOnly | -PrimaryOnly>] [-TargetDatabases "<MailboxDatabase1>","<MailboxDatabase1>"... [-TargetArchiveDatabases "<MailboxDatabase1>","<MailboxDatabase1>"...] [-Priority <PriorityValue>] [-BadItemLimit <Value>] [-AcceptLargeDataLoss]
```

This example creates a migration batch with these settings:

- **CSV file that specifies the mailboxes to move**: C:\Users\Administrator\Desktop\LocalMove 01.csv. If you only want to move the primary mailbox, use the *PrimaryOnly* switch, or the **MailboxType** value `PrimaryOnly` in the CSV file. If you only want to move the archive mailbox, use the *ArchiveOnly* switch, or the **MailboxType** value `ArchiveOnly` in the CSV file.

- **Batch name**: LocalMove 01.

- **Target database**: MBX DB02. If we don't use the *TargetDatabase* parameter, and the primary mailbox databases aren't specified in the CSV file, the automatic distribution logic in Exchange randomly selects databases in the Active Directory site.

- **Target database for archive mailboxes**: MBX DB02. Because we aren't using the *ArchiveTargetDatabase* parameter (in the command or the CSV file), the archive mailbox database is moved to the same database as the primary mailbox.

  If we use the *ArchiveOnly* switch (in the command or CSV file) without using the *ArchiveTargetDatabase* parameter (in the command or CSV file), the automatic distribution logic in Exchange will randomly select databases in the Active Directory site.

- **When to start the migration**: Immediately, because we're using the *AutoStart* switch. If we don't use this switch, we need to use the **Start-MigrationBatch** cmdlet to start the migration batch after it's created.

- **When to complete the migration**: After the mailboxes complete their initial synchronization, because we're using the *AutoComplete* switch. If we don't use this switch, we need to use the **Complete-MigrationBatch** cmdlet to start the migration batch after it's created

- **Priority**: `Normal`, because we aren't using the *Priority* parameter.

- **Bad item limit**: 10 (the default value in the Exchange Management Shell is 0). Because the value is less than 51, we don't need to use the `AcceptLargeDataLoss` switch.

```
New-MigrationBatch -Local -AutoStart -AutoComplete -Name "LocalMove 01" -CSVData ([System.IO.File]::ReadAllBytes("C:\Users\Administrator\Desktop\LocalMove 01.csv")) -TargetDatabases "MBX DB02" -BadItemLimit 10
```

How do you know this worked?

To verify that you've successfully created a local move request, do any of these steps:

- In the EAC, go to **Recipients** > **Migration** and verify the status of the move request (note that you might need to click **Refresh** ⟳). You can select the move request, and see more information in the details pane, or by clicking **Edit** ✎.
- In the EAC, go to **Recipients** > **Migration** and click **Status For All Batches**.
- Check the notification message. The sender is Microsoft Outlook. When the move request is complete, you'll get a message with the subject `Migration batch <MigrationBatchName> has completed successfully`.
- In the EAC, click the notification viewer 🔔 to view the status of the request.
- In the Exchange Management Shell, replace *<MailboxIdentity>* with the name, email address, or alias of the mailbox, and run this command to verify the basic property values:

  ```
  Get-MoveRequest -Identity <MailboxIdentity> | Format-List DisplayName,Alias,Status,*database*
  ```

- In the Exchange Management Shell, replace *<BatchName>* with the batch name value of the move request, and run this command to verify the basic property values:

  ```
  Get-MoveRequest -BatchName <BatchName> | Format-List DisplayName,Alias,Status,*database*
  ```

  Note: If you created the move request in the EAC, the batch name value is `MigrationService:<BatchNameValueFromTheEAC>`.

- If you created the move request in the EAC, replace *<BatchName>* with the batch name value you specified, and run this command in the Exchange Management Shell to verify summary information about all mailboxes in the move:

  ```
  Get-MigrationUserStatistics -BatchId <BatchName>
  ```

- If you created the move request in the EAC, replace *<EmailAddress>* with the email address of the moved mailbox, and run this command to see detailed information about the specified mailbox:

  ```
  Get-MigrationUserStatistics -Identity <EmailAddress> | Format-List
  ```

For more information, see Get-MigrationUserStatistics.

Display migration batches

For an example of how to use the Exchange Management Shell to display a migration batch, see Example 2 in Get-MigrationBatch.

Create a cross-forest move using a .csv batch file

This example configures the migration endpoint, and then creates a cross-forest batch move from the source forest to the target forest using a .csv file.

```
New-MigrationEndpoint -Name Fabrikam -ExchangeRemote -Autodiscover -EmailAddress tonysmith@fabrikam.com -
Credentials (Get-Credential fabrikam\tonysmith)
$csvData=[System.IO.File]::ReadAllBytes("C:\Users\Administrator\Desktop\batch.csv")
New-MigrationBatch -CSVData $csvData -Timezone "Pacific Standard Time" -Name FabrikamMerger -SourceEndpoint
Fabrikam -TargetDeliveryDomain "mail.contoso.com"
```

For more information about preparing your forest for cross-forest moves, see the following topics:

- Prepare mailboxes for cross-forest move requests
- Prepare Mailboxes for Cross-Forest Moves Using Sample Code
- Prepare mailboxes for cross-forest moves using the Exchange Management Shell

For detailed syntax and parameter information, see New-MigrationBatch and New-MoveRequest.

How do you know this worked?

To verify that you have successfully completed your migration, do the following:

- From the Exchange Management Shell, run the following command to retrieve mailbox move information.

  ```
  Get-MigrationUserStatistics -Identity BatchName -Status | Format-List
  ```

For more information, see Get-MigrationUserStatistics.

Prepare mailboxes for cross-forest move requests

8/3/2020 • 11 minutes to read • Edit Online

Mailbox moves and mailbox migrations in Exchange 2016 and Exchange 2019 from one forest to another require that you prepare the destination forest, which is made easier by Exchange tools and cmdlets. Exchange 2016 supports mailbox moves and migrations using the Exchange Management Shell, specifically the **New-MoveRequest** and **New-MigrationBatch** cmdlets. You can also move the mailbox in the Exchange admin center (EAC).

To move an Exchange mailbox from a source forest to the target Exchange 2016 or Exchange 2019 target forest, the target forest needs to contain a valid mail user (also known as a mail-enabled user) with a specified set of Active Directory attributes.

- In Exchange 2016, you can move an Exchange 2010, Exchange 2013, or Exchange 2016 mailbox from a source Exchange forest to a target Exchange 2016 forest. If there's at least one Exchange 2016 Mailbox server in the target forest, the forest is considered an Exchange 2016 forest.

- In Exchange 2019, you can move an Exchange 2013, Exchange 2016, or Exchange 2019 mailbox from a source Exchange forest to a target Exchange 2019 forest. If there's at least one Exchange 2019 Mailbox server in the target forest, the forest is considered an Exchange 2019 forest.

To prepare for the mailbox move, you need to create mail users (also known as mail-enabled users) with the required Active Directory attributes in the target forest. There are two recommended approaches for creating mail users with the necessary attributes:

- If you deployed Identity Lifecycle Manager (ILM) for cross-forest global address list (GAL) synchronization, we recommend that you use Microsoft Identity Manager 2016 Service Pack 1. We've created sample code that you can use to learn how to customize ILM to synchronize the source mailbox user and target mail user.

  For more information, including how to download the sample code, see Prepare mailboxes for cross-forest moves using sample code.

- If you created the target mail user using an Active Directory tool other than ILM or Microsoft Identity Integration Server (MIIS), use the **Update-Recipient** cmdlet with the *Identity* parameter to generate the **LegacyExchangeDN** attribute for the target mail user. We've created a sample PowerShell script that reads from and writes to Active Directory and calls the **Update-Recipient** cmdlet.

  For more information about using the sample script, see Prepare mailboxes for cross-forest moves using the Exchange Management Shell.

After creating the target mail user, you can then run the **New-MoveRequest** or the **New-MigrationBatch** cmdlets to move the mailbox to the target Exchange 2016 or Exchange 2019 forest.

For more information about remote move requests, see the following topics:

- New-MigrationBatch

- New-MoveRequest

The remainder of this topic describes the mail user Active Directory attributes that are required for a mailbox move. These attributes are configured for you when you use either the code or the script to prepare for the mailbox move. However, you can manually copy these attributes using an Active Directory editor.

Active Directory user attributes required for a mailbox move

To support a remote mailbox move, the mail user object in the target Exchange forest must have the Active Directory attributes that are described in this section:

- Mandatory attributes
- Optional attributes
- Linked attributes
- Linked user attributes
- Resource mailbox attributes
- Additional attributes

Mandatory attributes

The following table lists the minimum set of attributes that need to be configured in ILM on the target mail user for the **New-MoveRequest** cmdlet to function correctly.

Mail user attributes

ACTIVE DIRECTORY ATTRIBUTE	ACTION
displayName	Copy the corresponding attribute of the source mailbox or generate a new value.
Mail	Directly copy the corresponding attribute of the source mailbox.
mailNickname	Copy the corresponding attribute of the source mailbox or generate a new value.
msExchArchiveGUID and msExchArchiveName	Directly copy the corresponding attribute of the source mailbox.
msExchMailboxGUID	Directly copy the corresponding attribute of the source mailbox.
msExchRecipientDisplayType	-2147483642 decimal (equivalent to 0x80000006 hex).
msExchRecipientTypeDetails	128 decimal (0x80 hex).
msExchUserCulture	Directly copy the corresponding attribute of the source mailbox.
msExchVersion	44220983382016 (decimal).
cn	Copy the corresponding attribute of the source mailbox or generate a new value.

ACTIVE DIRECTORY ATTRIBUTE	ACTION
proxyAddresses	Copy source mailbox's **proxyAddresses** attribute. Additionally, copy source mailbox's **LegacyExchangeDN** as an X500 address in the **proxyAddresses** attribute of the target mail user. **Note**: The **proxyAddresses** of the source mailbox user must contain an SMTP address that matches the authoritative domain of the target forest. This allows the **New-MoveRequest** cmdlet to correctly select the **targetAddress** of the source mail-enabled user (converted from the source mailbox user after the mailbox move request is complete) to ensure that mail routing is still functional.
sAMAccountName	Copy the corresponding attribute of the source mailbox or generate a new value. Ensure that the value is unique within the target forest domain that the target mail user belongs to.
targetAddress	Set to an SMTP address in the **proxyAddresses** attribute of the source mailbox. This SMTP address must belong to the authoritative domain of the source forest.
userAccountControl	Constant: 514 (equivalent to 0x202, ACCOUNTDISABLE | NORMAL_ACCOUNT).
userPrincipalName	Copy the corresponding attribute of the source mailbox or generate a new value. Because the mail user is logon disabled, this **userPrincipalName** isn't used.

Optional attributes

The following attributes aren't required for the **New-MoveRequest** cmdlet to function correctly; however, synchronizing them provides a better end-to-end user experience after moving the mailbox. Because the GAL in the target forest displays this target mail user, you should set the following GAL-related attributes.

GAL-related attributes

MAIL USER'S ACTIVE DIRECTORY ATTRIBUTE	ACTION
c	Directly copy the corresponding attribute of the source mailbox.
co	Directly copy the corresponding attribute of the source mailbox.
countryCode	Directly copy the corresponding attribute of the source mailbox.
company	Directly copy the corresponding attribute of the source mailbox.
department	Directly copy the corresponding attribute of the source mailbox.
facsimileTelephoneNumber	Directly copy the corresponding attribute of the source mailbox.

MAIL USER'S ACTIVE DIRECTORY ATTRIBUTE	ACTION
givenName	Directly copy the corresponding attribute of the source mailbox.
homePhone	Directly copy the corresponding attribute of the source mailbox.
info	Directly copy the corresponding attribute of the source mailbox.
initials	Directly copy the corresponding attribute of the source mailbox.
l	Directly copy the corresponding attribute of the source mailbox.
mobile	Directly copy the corresponding attribute of the source mailbox.
msExchAssistantName	Directly copy the corresponding attribute of the source mailbox.
msExchHideFromAddressLists	Directly copy the corresponding attribute of the source mailbox.
otherHomePhone	Directly copy the corresponding attribute of the source mailbox.
otherTelephone	Directly copy the corresponding attribute of the source mailbox.
pager	Directly copy the corresponding attribute of the source mailbox.
physicalDeliveryOfficeName	Directly copy the corresponding attribute of the source mailbox.
postalCode	Directly copy the corresponding attribute of the source mailbox.
sn	Directly copy the corresponding attribute of the source mailbox.
st	Directly copy the corresponding attribute of the source mailbox.
streetAddress	Directly copy the corresponding attribute of the source mailbox.
telephoneAssistant	Directly copy the corresponding attribute of the source mailbox.
telephoneNumber	Directly copy the corresponding attribute of the source mailbox.

MAIL USER'S ACTIVE DIRECTORY ATTRIBUTE	ACTION
title	Directly copy the corresponding attribute of the source mailbox.

Linked attributes

A *linked attribute* is an Active Directory attribute that references other Active Directory objects in the local forest. You can't directly copy the linked attribute values from a mailbox in the source forest to a mail user in the target forest. Instead, you do the following steps:

1. Find the Active Directory objects in the source forest that the source mailbox attribute refers to.
2. Find the corresponding Active Directory objects in the target forest.
3. Set the target mail user's attribute to refer to the Active Directory objects in the target forest.

Linked attributes

MAIL USER'S ACTIVE DIRECTORY ATTRIBUTE	ACTION
altRecipient	Correspond to the source mailbox's **altRecipient** attribute.
deliverAndRedirect	Directly copy the corresponding attribute of the source mailbox. This attribute is a Boolean value that should be set along with **altRecipient**.
Manager (and its backlinks)	Correspond to the source mailbox's manager attribute.
MemberOf (backlinks)	This is the backlink of group member attribute.
publicDelegates (and its backlinks)	Correspond to the source mailbox's **publicDelegates** attribute.

Linked user attributes

If you want to move a mailbox to an Exchange resource forest, the mailbox in the resource forest is considered a *linked mailbox*. In this scenario, you need to create a linked mail user in the (target) resource forest. To create a linked mail user, you need to set the attributes shown in the following table.

Linked mail user attributes

ACTIVE DIRECTORY ATTRIBUTE	ACTION
msExchMasterAccountHistory	Directly copy the corresponding attribute of the source mailbox.
msExchMasterAccountSid	If the source mailbox has **msExchMasterAccountSid**, copy it. Otherwise, copy the source mailbox's **objectSid**.
msExchRecipientDisplayType	Constant:-1073741818 decimal (equivalent to `*unsigned* 0xC0000006`).

> **NOTE**
>
> A linked mailbox can only be created if there's a forest trust between the source forest and target forest.

If the source object is disabled and the **msExchMasterAccountSid** attribute is set to self (resource mailbox, shared mailbox), don't stamp anything on the target user.

If the source object is disabled and the **msExchMasterAccountSid** attribute isn't set, the mailbox is invalid.

If the source object is enabled and the **msExchMasterAccountSid** attribute is set, the mailbox is invalid.

Resource mailbox attributes

If you want to move a resource mailbox to an Exchange forest, you need to set the attributes shown in the following table on the target mail user.

Resource mailbox attributes

MAIL USER'S ACTIVE DIRECTORY ATTRIBUTE	ACTION
msExchRecipientDisplayType	If the source mailbox is a conference room: Constant: -2147481850 decimal (equivalent to `*unsigned* 0x80000706`). If the source mailbox is an equipment mailbox: Constant: -2147481594 decimal (equivalent to `*unsigned* 0x80000806`).
msExchResourceCapacity	Directly copy the corresponding attribute of the source mailbox.
msExchResourceDisplay	Directly copy the corresponding attribute of the source mailbox.
msExchResourceMetaData	Directly copy the corresponding attribute of the source mailbox.
msExchResourceSearchProperties	Directly copy the corresponding attribute of the source mailbox.

Additional attributes

Resource mailbox attributes

MAIL USER'S ACTIVE DIRECTORY ATTRIBUTES	DESCRIPTION
comment	Directly copy the corresponding attribute of the source mailbox.
deletedItemFlags	Directly copy the corresponding attribute of the source mailbox.
delivContLength	Directly copy the corresponding attribute of the source mailbox.
departmentNumber	Directly copy the corresponding attribute of the source mailbox.
description	Directly copy the corresponding attribute of the source mailbox.
division	Directly copy the corresponding attribute of the source mailbox.

MAIL USER'S ACTIVE DIRECTORY ATTRIBUTES	DESCRIPTION
employeeID	Directly copy the corresponding attribute of the source mailbox.
employeeNumber	Directly copy the corresponding attribute of the source mailbox.
employeeType	Directly copy the corresponding attribute of the source mailbox.
extensionAttribute1-15	Directly copy the corresponding attribute of the source mailbox.
homePostalAddress	Directly copy the corresponding attribute of the source mailbox.
internationalISDNNumber	Directly copy the corresponding attribute of the source mailbox.
ipPhone	Directly copy the corresponding attribute of the source mailbox.
language	Directly copy the corresponding attribute of the source mailbox.
lmPwdHistory	Directly copy the corresponding attribute of the source mailbox.
localeID	Directly copy the corresponding attribute of the source mailbox.
mAPIRecipient	Directly copy the corresponding attribute of the source mailbox.
middleName	Directly copy the corresponding attribute of the source mailbox.
msDS-PhoneticCompanyName	Directly copy the corresponding attribute of the source mailbox.
msDS-PhoneticDepartment	Directly copy the corresponding attribute of the source mailbox.
msDS-PhoneticDisplayName	Directly copy the corresponding attribute of the source mailbox.
msDS-PhoneticFirstName	Directly copy the corresponding attribute of the source mailbox.
msDS-PhoneticLastName	Directly copy the corresponding attribute of the source mailbox.
msExchBlockedSendersHash	Directly copy the corresponding attribute of the source mailbox.

MAIL USER'S ACTIVE DIRECTORY ATTRIBUTES	DESCRIPTION
msExchELCExpirySuspensionEnd	Directly copy the corresponding attribute of the source mailbox.
msExchELCExpirySuspensionStart	Directly copy the corresponding attribute of the source mailbox.
msExchELCMailboxFlags	Directly copy the corresponding attribute of the source mailbox.
msExchExternalOOFOptions	Directly copy the corresponding attribute of the source mailbox.
msExchMessageHygieneFlags	Directly copy the corresponding attribute of the source mailbox.
msExchMessageHygieneSCLDeleteThreshold	Directly copy the corresponding attribute of the source mailbox.
msExchMessageHygieneSCLJunkThreshold	Directly copy the corresponding attribute of the source mailbox.
msExchMessageHygieneSCLQuarantineThreshold	Directly copy the corresponding attribute of the source mailbox.
msExchMessageHygieneSCLRejectThreshold	Directly copy the corresponding attribute of the source mailbox.
msExchMDBRulesQuota	Directly copy the corresponding attribute of the source mailbox.
msExchPoliciesExcluded	Directly copy the corresponding attribute of the source mailbox.
msExchSafeRecipientsHash	Directly copy the corresponding attribute of the source mailbox.
msExchSafeSendersHash	Directly copy the corresponding attribute of the source mailbox.
msExchUMSpokenName	Directly copy the corresponding attribute of the source mailbox.
otherFacsimileTelephoneNumber	Directly copy the corresponding attribute of the source mailbox.
otherIpPhone	Directly copy the corresponding attribute of the source mailbox.
otherMobile	Directly copy the corresponding attribute of the source mailbox.
otherPager	Directly copy the corresponding attribute of the source mailbox.

MAIL USER'S ACTIVE DIRECTORY ATTRIBUTES	DESCRIPTION
preferredDeliveryMethod	Directly copy the corresponding attribute of the source mailbox.
personalPager	Directly copy the corresponding attribute of the source mailbox.
personalTitle	Directly copy the corresponding attribute of the source mailbox.
photo	Directly copy the corresponding attribute of the source mailbox.
pOPCharacterSet	Directly copy the corresponding attribute of the source mailbox.
pOPContentFormat	Directly copy the corresponding attribute of the source mailbox.
postalAddress	Directly copy the corresponding attribute of the source mailbox.
postOfficeBox	Directly copy the corresponding attribute of the source mailbox.
primaryInternationalISDNNumber	Directly copy the corresponding attribute of the source mailbox.
primaryTelexNumber	Directly copy the corresponding attribute of the source mailbox.
showInAdvancedViewOnly	Directly copy the corresponding attribute of the source mailbox.
street	Directly copy the corresponding attribute of the source mailbox.
terminalServer	Directly copy the corresponding attribute of the source mailbox.
textEncodedORAddress	Directly copy the corresponding attribute of the source mailbox.
thumbnailLogo	Directly copy the corresponding attribute of the source mailbox.
thumbnailPhoto	Directly copy the corresponding attribute of the source mailbox.
url	Directly copy the corresponding attribute of the source mailbox.
userCert	Directly copy the corresponding attribute of the source mailbox.

MAIL USER'S ACTIVE DIRECTORY ATTRIBUTES	DESCRIPTION
userCertificate	Directly copy the corresponding attribute of the source mailbox.
userSMIMECertificate	Directly copy the corresponding attribute of the source mailbox.
wWWHomePage	Directly copy the corresponding attribute of the source mailbox.

Prepare mailboxes for cross-forest moves using the Exchange Management Shell

8/3/2020 • 8 minutes to read • Edit Online

Exchange Server supports mailbox moves and migrations using the Exchange Management Shell **New-MoveRequest** and **New-MigrationBatch** cmdlets. You can also move the mailbox in the Exchange admin center (EAC).

- In Exchange 2016, you can move an Exchange 2010, Exchange 2013, or Exchange 2016 mailbox from a source Exchange forest to a target Exchange 2016 forest.
- In Exchange 2019, you can move an Exchange 2013, Exchange 2016, or Exchange 2019 mailbox from a source Exchange forest to a target Exchange 2019 forest.

To run the **New-MoveRequest** and **New-MigrationBatch** cmdlets, a mail user must exist in the target Exchange forest, and the mail user must have a minimum set of required Active Directory attributes.

The sample Exchange PowerShell script described in this topic supports this task by synchronizing mailbox users from an Exchange source forest to Exchange target forests as mail users (also known as mail-enabled users). The script copies the Active Directory attributes of the mailbox users in the source forest to the target forest, and then uses the **Update-Recipient** cmdlet to turn the target objects into mail users.

For more information about using and writing scripts, see About Scripts. For more information about preparing for cross-forest moves, see Prepare mailboxes for cross-forest move requests.

Looking for other management tasks related to remote move requests? Check out Manage on-premises mailbox moves in Exchange Server.

What do you need to know before you begin?

- Locate the Prepare-MoveRequest.ps1 script in %ExchangeInstallPath%Scripts. By default, %ExcangeInstallPath% is C:\Program Files\Microsoft\Exchange Server\V15\ (note the trailing '\').
- To run the sample script, you need the following:
  - An Exchange source forest (where the mailbox currently resides).
    - For Exchange 2016 target forests, the source mailbox can be in Exchange 2010, Exchange 2013, or Exchange 2016.
    - For Exchange 2019 target forests, the source mailbox can be in Exchange 2013, Exchange 2016, or Exchange 2019.
  - A target forest with Exchange 2016 or Exchange 2019 installed (where the mailbox will be moved to).

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Prepare-MoveRequest.ps1 script to prepare mailboxes for

cross-forest moves

Run the script from the Exchange Management Shell on a Mailbox server in the target Exchange 2016 or Exchange 2019 forest. The script copies the mailbox attributes from the source forest.

To assign a specific authentication credential for the remote forest domain controller, you must first run the Windows PowerShell **Get-Credential** cmdlet and store the user input in a temporary variable. When you run the **Get-Credential** cmdlet, the cmdlet asks for the user name and password of the account used during authentication with the remote forest domain controller. You can then use the temporary variable in the Prepare-MoveRequest.ps1 script. For more information about the **Get-Credential** cmdlet, see Get-Credential.

> **NOTE**
>
> Make sure that you use two separate credentials for the local forest and the remote forest when calling this script.

1. Run the following commands to get the local forest and remote forest credentials.

   ```
   $LocalCredentials = Get-Credential
   ```

   ```
   $RemoteCredentials = Get-Credential
   ```

2. Run the following commands to pass the credential information to the *LocalForestCredential* and *RemoteForestCredential* parameters in the Prepare-MoveRequest.ps1 script.

   ```
   Prepare-MoveRequest.ps1 -Identity JohnSmith@Fabrikan.com -RemoteForestDomainController
   DC001.Fabrikam.com -RemoteForestCredential $RemoteCredentials -LocalForestDomainController
   DC001.Contoso.com -LocalForestCredential $LocalCredentials
   ```

Parameter set of the script

The following table describes the parameter set for the script.

PARAMETER	REQUIRED	DESCRIPTION
Identity	Required	The *Identity* parameter uniquely identifies a mailbox in the source forest. Identity can be any of the following values: Common name (CN), Alias, **proxyAddress** property, **objectGuid** property, or **DisplayName** property
RemoteForestCredential	Required	The *RemoteForestCredential* parameter specifies the administrator who has permissions to copy data from the source forest Active Directory.
RemoteForestDomainController	Required	The *RemoteForestDomainController* parameter specifies a domain controller in the source forest where the mailbox resides.

PARAMETER	REQUIRED	DESCRIPTION
DisableEmailAddressPolicy	Optional	The *DisableEmailAddressPolicy* parameter specifies whether the Email Address Policy (EAP) should be disabled when creating a **MailUser** object in the target forest. When you specify this parameter, the EAP in the target forest won't be applied. **Note**: When you specify this parameter, the **MailUser** object won't have e-mail address mapping in the local forest domain stamped. This is usually stamped by the EAP.
LinkedMailUser	Optional	The *LinkedMailUser* switch specifies whether to create a linked MailUser in the local forest for the mailbox user in the remote forest. If the switch is provided, the script creates a target **MailUser** object linked to the source mailbox. If the switch is omitted, the script creates a regular target **MailUser** object.
LocalForestCredential	Optional	The *LocalForestCredential* parameter specifies the administrator with permissions to write data to the target forest Active Directory. We recommend that you explicitly specify this parameter to avoid Active Directory permission issues. If the remote forest and the local forest have a trusted relationship configured, don't use a user account from the remote forest as the local forest credential, even though the remote user account may have permission to modify Active Directory in the local forest.
LocalForestDomainController	Optional	The *LocalForestDomainController* parameter specifies a domain controller in the target forest where the mail user will be created. We recommend that you specify this parameter to avoid possible domain controller replication delay issues in the local forest that could occur if a random domain controller is selected.

PARAMETER	REQUIRED	DESCRIPTION
MailboxDeliveryDomain	Optional	The *MailboxDeliveryDomain* parameter specifies an authoritative domain of the source forest so that the script can select the correct source mailbox user's **proxyAddress** property as the target mail user's **targetAddress** property. By default, the primary SMTP address of the source mailbox user is set as the **targetAddress** property of the target mail user.
OverWriteLocalObject	Optional	The *OverWriteLocalObject* parameter is used for users created by the Active Directory Migration Tool. The properties are copied from the existing mail contact to the newly created mail user. However, after this copy, the script also copies the properties from the source forest user to the newly created mail user.
TargetMailUserOU	Optional	The *TargetMailuserOU* parameter specifies the organizational unit (OU) under which the target mail user will be created.
UseLocalObject	Optional	The *UseLocalObject* parameter specifies whether to convert the existing local object to the required target mail user if the script detects an object in the local forest that conflicts with the to-be-created mail user.

Examples

This section contains several examples of how you can use the Prepare-MoveRequest.ps1 script.

Example: Single linked mail user

This example provisions a single linked mail user in the local forest, when there is forest trust between the remote forest and local forest.

1. Run the following commands to get the local forest and remote forest credentials.

   ```
   $LocalCredentials = Get-Credential
   ```

   ```
   $RemoteCredentials = Get-Credential
   ```

2. Run the following command to pass the credential information to the *LocalForestCredential* and *RemoteForestCredential* parameters in the Prepare-MoveRequest.ps1 script.

   ```
   Prepare-MoveRequest.ps1 -Identity JamesAlvord@Contoso.com -RemoteForestDomainController
   DC001.Fabrikam.com -RemoteForestCredential $RemoteCredentials -LocalForestDomainController
   DC001.Contoso.com -LocalForestCredential $LocalCredentials -LinkedMailUser
   ```

Example: Pipelining

This example supports pipelining if you supply a list of mailbox identities.

1. Run the following command.

   ```
   $UserCredentials = Get-Credential
   ```

2. Run the following command to pass the credential information to the *RemoteForestCredential* parameter in the Prepare-MoveRequest.ps1 script.

   ```
   "IanP@Contoso.com", "JoeAn@Contoso.com" | Prepare-MoveRequest.ps1 -RemoteForestDomainController
   DC001.Fabrikam.com -RemoteForestCredential $UserCredentials
   ```

Example: Use a .csv file to bulk-create mail users

You can generate a .csv file containing a list of mailbox identities from the source forest, which allows you to pipe the content of this file into the script to bulk-create the target mail users.

For example, the content of the .csv file can be:

```
Identity
Ian@contoso.com
John@contoso.com
Cindy@contoso.com
```

This example calls a .csv file to bulk create the target mail users.

1. Run the following command to get the remote forest credentials.

   ```
   $UserCredentials = Get-Credential
   ```

2. Run the following command to pass the credential information to the *RemoteForestCredential* parameter in the Prepare-MoveRequest.ps1 script.

   ```
   Import-Csv Test.csv | Prepare-MoveRequest.ps1 -RemoteForestDomainController DC001.Fabrikam.com -
   RemoteForestCredential $UserCredentials
   ```

Script behavior per target object

This section describes how the script performs in relation to several scenarios for target objects.

Duplicate target mail-enabled object

When the script attempts to create a target mail user from the source mailbox user, and it detects a duplicate local mail-enabled object, it uses the following logic:

- If the source mailbox user's **masterAccountSid** attribute equals any target object's **objectSid** or **masterAccountSid** attribute:
  - If the target object isn't mail-enabled, the script returns an error because the script doesn't support converting an object that isn't mail-enabled to a mail user.
  - If the target object is mail-enabled, the target object is a duplicate.
- If an address in the source mailbox user's **proxyAddress** properties (smtp/x500 only) equals an address in a target object's **proxyAddress** properties (smtp/x500 only), the target object is a duplicate.

The script prompts the user about the duplicate objects.

If the target mail-enabled object is a mail user or mail contact, which is most likely created by a cross-forest global address list (GAL) synchronization deployment, you can run the script again with the *UseLocalObject* parameter to use the target mail-enabled object for mailbox migration.

Mail user

If the target object is a mail user, the script copies the following attributes from the source mailbox user to the target mail user:

- **msExchMailboxGUID**
- **msExchArchiveGUID**
- **msExchArchiveName**

If the *LinkedMailUser* parameter is set, the script copies the source **objectSid** / **masterAccountSid** attribute.

Mail contact

If the target object is a mail contact, the script deletes the existing contact and copies all its attributes to a new mail user. The script also copies the following attributes from the source mailbox user:

- **msExchMailboxGUID**
- **msExchArchiveGUID**
- **msExchArchiveName**
- **sAMAccountName**
- **userAccountControl** (set to 514; equivalent to `0x202, ACCOUNTDISABLE | NORMAL_ACCOUNT`)
- **userPrincipalName**

If the *LinkedMailUser* parameter is set, the script copies the source **objectSid** / **masterAccountSid** attribute.

LegacyExchangeDN attribute

When the **Update-Recipient** cmdlet is called to convert the target object into a mail user, a new **LegacyExchangeDN** attribute is generated for the target mail user. The script copies the **LegacyExchangeDN** attribute of the target mail user as an x500 address to the **proxyAddress** properties of the source mailbox user.

Enable the MRS Proxy endpoint for remote moves

8/3/2020 • 3 minutes to read • Edit Online

The Mailbox Replication service (MRS) has a proxy endpoint that's required for cross-forest mailbox moves and remote move migrations between your on-premises Exchange organization and Microsoft 365 or Office 365. You enable the MRS proxy endpoint in the Exchange Web Services (EWS) virtual directory settings in the Client Access (frontend) services on Exchange 2016 or Exchange 2019 Mailbox servers.

Where you enable the MRS Proxy endpoint depends on the type and direction of the mailbox move:

- **Cross-forest enterprise moves**: For cross-forest moves that are initiated from the target forest (known as a *pull* move type), you need to enable the MRS Proxy endpoint on Mailbox servers in the source forest. For cross-forest moves that are initiated from the source forest (known as a *push* move type), you need to enable the MRS Proxy endpoint on Mailbox servers in the target forest.

- **Remote move migrations between an on-premises Exchange organization and Microsoft 365 or Office 365**. For both onboarding and offboarding remote move migrations, you need to enable the MRS Proxy endpoint on Mailbox servers in your on-premises Exchange organization.

Note: If you use theExchange admin center (EAC) to move mailboxes, cross-forest moves and onboarding remote move migrations are pull move types, because you initiate the request from the target environment. Offboarding remote move migrations are push move types because you initiate the request from the source environment.

What do you need to know before you begin?

- Estimated time to complete: 2 minutes per server.

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Exchange Web Services permissions" section in the Clients and mobile devices permissions topic.

- If you've deployed multiple Mailbox servers in your Exchange organization, you should enable the MRS Proxy endpoint in the Client Access services on each Mailbox server. If you add additional Mailbox servers, be sure to enable the MRS Proxy endpoint on the new servers. Cross-forest moves and remote move migrations can fail if the MRS Proxy endpoint isn't enabled on all Mailbox servers.

- If you don't perform cross-forest moves or remote move migrations, keep MRS Proxy endpoints disabled in the Client Access services on Mailbox servers to reduce the attack surface of your organization.

- Exchange Online requires Windows authentication for the MRS proxy endpoint in the Exchange Web Services (EWS) virtual directory in the Client Access (front end) services.

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to enable the MRS Proxy endpoint

1. In the EAC, go to **Servers** > **Virtual Directories**.

2. Select the EWS virtual directory that you want to configure.

   - You can use the **Select server** drop-down list to filter the Exchange servers by name.

   - To only display EWS virtual directories, select **EWS** in the **Select type** drop-down list.

   After you've selected the EWS virtual directory that you want to configure, click **Edit** ✎.

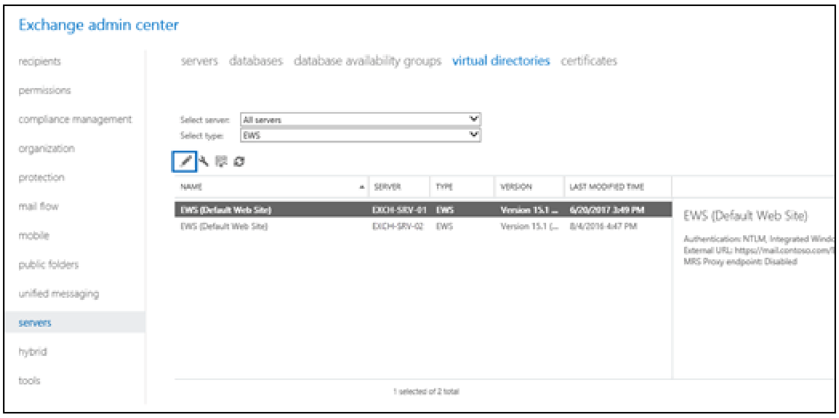

Select server: All servers
Select type: EWS

NAME	SERVER	TYPE	VERSION	LAST MODIFIED TIME
EWS (Default Web Site)	EXCH-SRV-01	EWS	Version 15.1 ...	6/20/2017 3:49 PM
EWS (Default Web Site)	EXCH-SRV-02	EWS	Version 15.1 (...	8/4/2016 4:47 PM

1 selected of 2 total

EWS (Default Web Site)

Authentication: NTLM, Integrated Windo
External URL: https://mail.contoso.com/
MRS Proxy endpoint: Disabled

3. On the properties page that opens, on the **General** tab, select the **Enable MRS Proxy endpoint** check box, and then click **Save**.

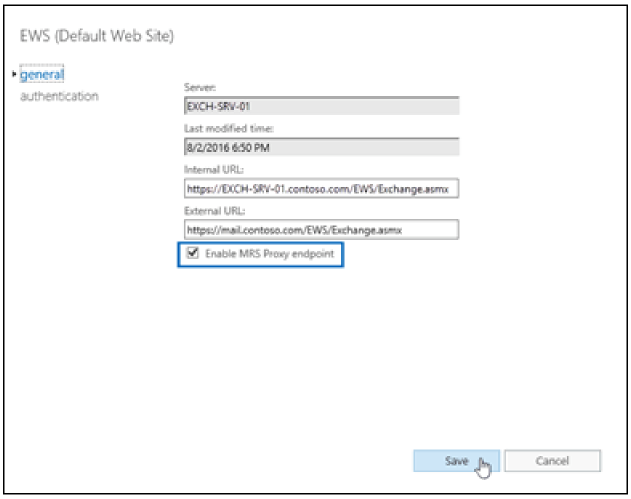

EWS (Default Web Site)

Server:
EXCH-SRV-01

Last modified time:
8/2/2016 6:50 PM

Internal URL:
https://EXCH-SRV-01.contoso.com/EWS/Exchange.asmx

External URL:
https://mail.contoso.com/EWS/Exchange.asmx

☑ Enable MRS Proxy endpoint

Use the Exchange Management Shell to enable the MRS Proxy endpoint

To enable the MRS Proxy endpoint, use this syntax:

```
Set-WebServicesVirtualDirectory -Identity "[<Server>\]EWS (Default Web Site)" -MRSProxyEnabled $true
```

This example enables the MRS Proxy endpoint in Client Access services on the Mailbox server named EXCH-SRV-01.

```
Set-WebServicesVirtualDirectory -Identity "EXCH-SRV-01\EWS (Default Web Site)" -MRSProxyEnabled $true
```

This example enables the MRS Proxy endpoint in Client Access services on all Mailbox servers in your Exchange organization.

```
Get-WebServicesVirtualDirectory | Set-WebServicesVirtualDirectory -MRSProxyEnabled $true
```

For detailed syntax and parameter information, see Set-WebServicesVirtualDirectory.

How do you know this worked?

To verify that you've successfully enabled the MRS Proxy endpoint, do any of these steps:

- In the EAC, go to **Servers** > **Virtual Directories** > select the EWS virtual directory, and verify in the details pane that the MRS Proxy endpoint is enabled.

  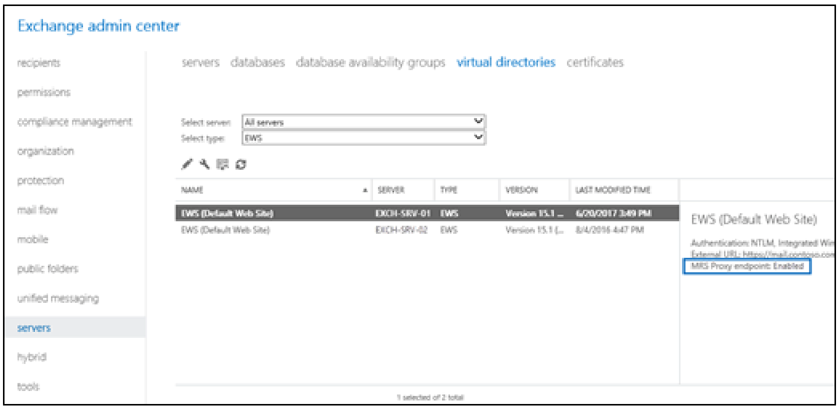

  Select server: All servers
  Select type: EWS

NAME	SERVER	TYPE	VERSION	LAST MODIFIED TIME
EWS (Default Web Site)	EXCH-SRV-01	EWS	Version 15.1 ...	6/20/2017 3:49 PM
EWS (Default Web Site)	EXCH-SRV-02	EWS	Version 15.1 (...	8/4/2016 4:47 PM

  EWS (Default Web Site)
  Authentication: NTLM, Integrated Win
  External URL: https://mail.contoso.com
  MRS Proxy endpoint: Enabled

  1 selected of 2 total

- Run this command in the Exchange Management Shell, and verify that the **MRSProxyEnabled** property for the EWS virtual directory has the value `True`:

  ```
  Get-WebServicesVirtualDirectory | Format-Table -Auto Identity,MRSProxyEnabled
  ```

- Use the **Test-MigrationServerAvailability** cmdlet in the Exchange Management Shell to test communication with the remote servers that hosts the mailboxes that you want to move (or the servers in your on-premises Exchange organization for offboarding remote move migrations from Microsoft 365 or Office 365).

  Replace *<EmailAddress>* with the email address of one of the mailboxes that you want to move, and run this command in the Exchange Management Shell:

  ```
  Test-MigrationServerAvailability -ExchangeRemoteMove -Autodiscover -EmailAddress <EmailAddress> -
  Credentials (Get-Credential)
  ```

  To run this command successfully, the MRS Proxy endpoint must be enabled.

  For detailed syntax and parameter information, see Test-MigrationServerAvailability.

Recreate missing arbitration mailboxes

8/3/2020 • 5 minutes to read • Edit Online

Exchange Server contains five special system mailboxes (seven in Exchange 2016 CU8 and later) known as *arbitration mailboxes*. Arbitration mailboxes are used for storing different types of system data and for managing messaging approval workflow. The following table lists each type of arbitration mailbox and their responsibilities.

ARBITRATION MAILBOX NAME	DISPLAY NAME	PERSISTED CAPABILITIES	FUNCTION
FederatedEmail.4c1f4d8b-8179-4148-93bf-00a95fa1e042	Microsoft Exchange Federation Mailbox	none	This mailbox stores data used to maintain federation between different Exchange organizations. This includes Rights Management Services, cross-premises mail-flow monitoring probes and responses, notifications, online archives, messaging records management, and cross-premises free/busy information.
Migration.8f3e7716-2011-43e4-96b1-aba62d229136	Microsoft Exchange Migration	Management	Stores data for the Exchange migration service to use when moving mailboxes in batches.
SystemMailbox{1f05a927-XXXX-XXXX-XXXX-XXXXXXXXXXXX} (for example, SystemMailbox{1f05a927-9350-4efe-a823-5529c2d64109}; most of the mailbox name is unique to your organization)	Microsoft Exchange Approval Assistant	none	This mailbox is provisioned for use by the Exchange approval framework for recipient moderation and auto group approval requests.
SystemMailbox{bb558c35-97f1-4cb9-8ff7-d53741dc928c}	Microsoft Exchange	ClientExtensions GMGen MailRouting MessageTracking OABGen PstProvider UMGrammar UMGrammarReady (Exchange 2016 only)	This is known as an organization mailbox. It is used for creating offline address books (OABs). To load-balance OAB generation across your organization, including across geographically separate sites, you can create additional organization mailboxes.

ARBITRATION MAILBOX NAME	DISPLAY NAME	PERSISTED CAPABILITIES	FUNCTION
SystemMailbox{e0dc1c29-89c3-4034-b678-e6c29d823ed9}	Microsoft Exchange	UMDataStorage	Discovery system mailbox. Provisioned for use by the e-Discovery feature, which is used by compliance officers to locate messages that match specified selection criteria. This mailbox is also used by Unified Messaging in Exchange 2016 for storing UM console attending files and other information.
SystemMailbox{D0E409A0-AF9B-4720-92FE-AAC869B0D201} (Exchange 2016 CU8 and later)	Microsoft Exchange	none	
SystemMailbox{2CE34405-31BE-455D-89D7-A7C7DA7A0DAA} (Exchange 2016 CU8 and later)	Microsoft Exchange	none	

If you need to re-create one of more of these arbitration mailboxes, see the instructions that follow.

What do you need to know before you begin?

- Estimated time to complete: 10 minutes per procedure.
- You need to be assigned permissions before you can perform these procedures. To see what permissions you need, see the "Recipient Provisioning Permissions" section in the Recipients Permissions topic.
- To run `Setup.exe /PrepareAD`, your account needs to be a member of the Enterprise Admins security group.
- The computer that you use to run `Setup.exe /PrepareAD` requires access to Setup.exe in the Exchange installation files:
  1. Use your most recently downloaded copy of the Exchange ISO image file, or download an updated copy from Updates for Exchange Server.
  2. In File Explorer, right-click on the Exchange ISO image file and then select **Mount**. Note the virtual DVD drive letter that's assigned.
  3. Open a Windows Command Prompt window. For example:
     - Press the Windows key + 'R' to open the **Run** dialog, type cmd.exe, and then press **OK**.
     - Press **Start**. In the **Search** box, type **Command Prompt**, then in the list of results, select **Command Prompt**.
- For more information about opening the Exchange Management Shell, see Open the Exchange Management Shell.
- For more information about running Exchange Setup in unattended mode, see Use unattended mode in Exchange Setup.

Re-create an arbitration mailbox

Use the following instructions to re-create a particular type of arbitration mailbox.

Re-create the Microsoft Exchange Federation Mailbox

To re-create the arbitration mailbox FederatedEmail.4c1f4d8b-8179-4148-93bf-00a95fa1e042, run the following commands:

1. If the mailbox is missing, run the following command from a Windows Command Prompt window:

   ```
   <Virtual DVD drive letter>:\Setup.exe /IAcceptExchangeServerLicenseTerms /PrepareAD
   ```

   For example:

   ```
   E:\Setup.exe /IAcceptExchangeServerLicenseTerms /PrepareAD
   ```

2. In the Exchange Management Shell, run the following command:

   ```
   Enable-Mailbox -Identity "FederatedEmail.4c1f4d8b-8179-4148-93bf-00a95fa1e042" -Arbitration
   ```

Re-create the Microsoft Exchange Migration mailbox

To re-create the arbitration mailbox Migration.8f3e7716-2011-43e4-96b1-aba62d229136, run the following commands:

1. If the mailbox is missing, run the following command from a Windows Command Prompt window:

   ```
   <Virtual DVD drive letter>:\Setup.exe /IAcceptExchangeServerLicenseTerms /PrepareAD
   ```

   For example:

   ```
   E:\Setup.exe /IAcceptExchangeServerLicenseTerms /PrepareAD
   ```

2. In the Exchange Management shell, run the following command:

   ```
   Enable-Mailbox -Identity "Migration.8f3e7716-2011-43e4-96b1-aba62d229136" -Arbitration
   ```

3. In the Exchange Management Shell, set the Persisted Capabilities (msExchCapabilityIdentifiers) for the mailbox by running the following command:

   ```
   Set-Mailbox -Identity "Migration.8f3e7716-2011-43e4-96b1-aba62d229136" -Arbitration -Management $true -
   Force
   ```

Re-create the Microsoft Exchange Approval Assistant mailbox

To re-create the arbitration mailbox SystemMailbox{1f05a927-XXXX-XXXX-XXXX-XXXXXXXXXXXX}, run the following commands:

1. If the mailbox is missing, run the following command from a Windows Command Prompt window:

   ```
   <Virtual DVD drive letter>:\Setup.exe /IAcceptExchangeServerLicenseTerms /PrepareAD
   ```

For example:

```
E:\Setup.exe /IAcceptExchangeServerLicenseTerms /PrepareAD
```

2. In the Exchange Management Shell, run the following command:

```
Get-User -ResultSize Unlimited | where {$_.Name -like "SystemMailbox{1f05a927*"} | Enable-Mailbox -
Arbitration
```

Re-create the Microsoft Exchange organization mailbox for OABs

To re-create the arbitration mailbox SystemMailbox{bb558c35-97f1-4cb9-8ff7-d53741dc928c}, run the following commands:

1. If the mailbox is missing, run the following command from a Windows Command Prompt window:

```
<Virtual DVD drive letter>:\Setup.exe /IAcceptExchangeServerLicenseTerms /PrepareAD
```

For example:

```
E:\Setup.exe /IAcceptExchangeServerLicenseTerms /PrepareAD
```

2. In the Exchange Management Shell, run the following command:

```
Enable-Mailbox -Identity "SystemMailbox{bb558c35-97f1-4cb9-8ff7-d53741dc928c}" -Arbitration
```

3. In the Exchange Management Shell, set the Persisted Capabilities (msExchCapabilityIdentifiers) for the mailbox by running the following command:

```
Get-Mailbox "SystemMailbox{bb558c35-97f1-4cb9-8ff7-d53741dc928c}" -Arbitration | Set-Mailbox -
Arbitration -UMGrammar $true -OABGen $true -GMGen $true -ClientExtensions $true -MessageTracking $true -
PstProvider $true -MaxSendSize 1GB -Force
```

4. In the Exchange Management Shell, add the required capabilities to the mailbox by running the following commands:

```
$OABMBX = Get-Mailbox "SystemMailbox{bb558c35-97f1-4cb9-8ff7-d53741dc928c}" -Arbitration; Set-ADUser
$OABMBX.SamAccountName -Add @{"msExchCapabilityIdentifiers"="40","42","43","44","47","51","52","46"}
```

Re-create the Microsoft Exchange Discovery system mailbox

To re-create the arbitration mailbox SystemMailbox{e0dc1c29-89c3-4034-b678-e6c29d823ed9}, run the following commands:

1. If the mailbox is missing, run the following command from a Windows Command Prompt window:

```
<Virtual DVD drive letter>:\Setup.exe /IAcceptExchangeServerLicenseTerms /PrepareAD
```

For example:

```
E:\Setup.exe /IAcceptExchangeServerLicenseTerms /PrepareAD
```

2. In the Exchange Management shell, run the following command:

   ```
   Enable-Mailbox -Identity "SystemMailbox{e0dc1c29-89c3-4034-b678-e6c29d823ed9}" -Arbitration
   ```

3. In the Exchange Management Shell, set the Persisted Capabilities (msExchCapabilityIdentifiers) for the mailbox by running the following command:

   ```
   Set-Mailbox -Identity "SystemMailbox{e0dc1c29-89c3-4034-b678-e6c29d823ed9}" -Arbitration -UMDataStorage
   $true -Force
   ```

Re-create the Microsoft Exchange 2016 CU8 and later system mailboxes

To re-create the arbitration mailbox SystemMailbox{D0E409A0-AF9B-4720-92FE-AAC869B0D201} and SystemMailbox{2CE34405-31BE-455D-89D7-A7C7DA7A0DAA}, run the following commands:

1. If the mailboxes are missing, run the following command from a Windows Command Prompt window:

   ```
   <Virtual DVD drive letter>:\Setup.exe /IAcceptExchangeServerLicenseTerms /PrepareAD
   ```

   For example:

   ```
   E:\Setup.exe /IAcceptExchangeServerLicenseTerms /PrepareAD
   ```

2. In the Exchange Management shell, run the following command:

   ```
   Enable-Mailbox -Identity "SystemMailbox{D0E409A0-AF9B-4720-92FE-AAC869B0D201}" -Arbitration
   Enable-Mailbox -Identity "SystemMailbox{2CE34405-31BE-455D-89D7-A7C7DA7A0DAA}" -Arbitration
   ```

3. In the Exchange Management Shell, set the Persisted Capabilities (msExchCapabilityIdentifiers) for the mailbox by running the following command:

   ```
   $ShardMBX = Get-Mailbox -Identity "SystemMailbox{2CE34405-31BE-455D-89D7-A7C7DA7A0DAA}" -Arbitration
   Set-Mailbox -Identity "SystemMailbox{2CE34405-31BE-455D-89D7-A7C7DA7A0DAA}" -Arbitration
   Set-ADUser $ShardMBX.SamAccountName -Add @{"msExchCapabilityIdentifiers"="66"}
   Set-ADUser $ShardMBX.SamAccountName -Add @{"msExchMessageHygieneSCLDeleteThreshold"="9"}
   Set-ADUser $ShardMBX.SamAccountName -Add @{"msExchMessageHygieneSCLJunkThreshold"="4"}
   Set-ADUser $ShardMBX.SamAccountName -Add @{"msExchMessageHygieneSCLQuarantineThreshold"="9"}
   Set-ADUser $ShardMBX.SamAccountName -Add @{"msExchMessageHygieneSCLRejectThreshold"="7"}
   ```

How do you know this worked?

To verify that you've successfully re-created the arbitration mailbox, set the search scope to search the entire Active Directory forest, an then use the **Get-Mailbox** cmdlet with the *Arbitration* switch to retrieve system mailboxes.

```
Set-ADServerSettings -ViewEntireForest $true; Get-Mailbox -Arbitration | Format-Table Name,DisplayName
```

View the results of the command to verify that appropriate system mailbox, either by Name or Display Name from the above table, has been re-created.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Managed Store in Exchange Server

8/3/2020 • 5 minutes to read • Edit Online

The Managed Store is the name for the Information Store (also known as the Store) processes in Exchange Server 2016 and Exchange Server 2019. Introduced in Exchange Server 2013, the Managed Store uses a controller/worker process model that provides storage process isolation and faster database failover. The Managed Store also uses a static database caching mechanism that replaces the dynamic buffer algorithm in previous versions of Exchange.

The multi-process model that's used by the Managed Store consists of the following processes on the Mailbox server:

- A single store service controller process for the whole Exchange server (Microsoft.Exchange.Store.Service.exe, also known as MSExchangeIS).
- Ane worker process for each mounted database (Microsoft.Exchange.Store.Worker.exe). When a database is mounted, a new worker process is instantiated that services only that database. When a database is dismounted, the worker process for that database is terminated.

For example, if you have 40 mailbox databases mounted on a Mailbox server, there will be 41 processes running for the Managed Store: one for each database, and one for the store service process controller. The store process controller monitors the health of all store worker processes on the server. A forcible or unexpected termination the Microsoft.Exchange.Store.Service.exe causes an immediate failover of all active database copies on the server.

The Managed Store is also tightly integrated with the Microsoft Exchange Replication service (MSExchangeRepl.exe) and Active Manager. The controller process, worker processes, and Replication service work together to provide greater availability and reliability as described in the following list:

- **Microsoft Exchange Replication service process (MSExchangeRepl.exe)**:
  - Responsible for issuing mount and dismount operations to the Store.
  - Initiates recovery action on storage or database failures reported by the Store, the Extensible Storage Engine (ESE), and Managed Availability responders.
  - Detects unexpected database failures.
  - Provides the administrative interface for management tasks.
- **Store service process/controller (Microsoft.Exchange.Store.Service.exe)**:
  - Manages each worker process lifetime based on the mount and dismount operations received from the Replication service.
  - Handles incoming requests from the Windows Service Control Manager.
  - Logs failure items when store worker process problems detected (for example, hang or unexpected exit).
  - Terminates store worker processes in response failover event.
- **Store worker process (Microsoft.Exchange.Store.Worker.exe)**
  - Responsible for executing RPC operations for mailboxes on a database.
  - RPC endpoint instance within worker process is the database GUID.
  - Provides database cache for a database.

Static database caching algorithm

The Managed Store uses a very simple and straightforward algorithm for determining database cache as compared to *dynamic buffer allocation* that was used in the previous versions of Exchange. The memory that's allocated for each database cache (that is, each store worker process) is based on number of local database copies and configured value of the *MaximumActiveDatabases* parameter on the **Set-MailboxServer** cmdlet (the default value is $null or blank). If the value of *MaximumActiveDatabases* is greater than number of current database copies, then the cache calculation is based on the number of database copies.

The static algorithm allocates memory for the ESE cache of each store worker process based on the amount of physical RAM that's installed in the server. This is referred to the *Max Cache Target* of the database. 25% of total server memory is allocated to the ESE cache, and is referred to as the *Server Cache Size Target*.

> **NOTE**
>
> You can override the Server Cache Size Target, and therefore the amount of memory allocated to the Store for ESE cache by using `msExchESEParamCacheSizeMax` attribute of the *InformationStore* object in Active Directory (the value configured is the number of 32 KB pages to allocate across all store processes).

A static amount of this cache is allocated to active and passive copies. The store worker process will be allocated the Max Cache Target only when servicing an active database copy. Passive database copies are allocated 20 percent of the Max Cache Target. The remainder is reserved by the Store, and allocated to the worker process if the database transitions from passive to active.

Max Cache Target is calculated only at Store startup. Therefore, if you add or remove databases or database copies, you must restart the Store controller service (MSExchangeIS) so that the cache can be adjusted accordingly. If the service is not restarted, new databases will have a smaller cache size target than databases that existed before the last service startup. In this scenario, the sum of database cache size targets will likely exceed the Server Cache Size Target until MSExchangeIS is restarted.

Example database cache calculations

Here are example database caching calculations that are based on a Mailbox server's memory and database configuration.

Example 1

Mailbox server configuration:

- 48 GB of memory
- Two active databases and two passive databases
- *MaximumActiveDatabases* parameter: not configured

The amount of database cache is 3 GB for each active database copy worker process and 0.6 GB for each passive database copy worker process. Here's how these values are calculated:

- **Server Cache Size Target**: 25% of the amount of memory: 48 GB * 0.25 = **12 GB**.
- **Database Max Cache Target**: Divide the Server Cache Size Target by the total number of active and passive databases: 12 GB / 4 databases = **3 GB**.
- **Memory used for passive database copies**: 20% of the Database Max Cache Target: 3 GB * 0.20 = **0.6 GB**.

Of the 12 GB of memory that's assigned to the Server Cache Size Target:

- 7.2 GB will be in use by database worker processes.
- 4.8 GB will be reserved by the Information Store for the two passive database copies in case they become active copies. If this happens, they will use their Max Cache Target of 3 GB.

Example 2

Mailbox server configuration:

- 48 GB of memory
- Two active databases and two passive databases
- *MaximumActiveDatabases* parameter: 2

The amount of database cache is 5 GB for each active database copy worker process and 0.2 GB for each passive database copy worker process. Here's how these values are calculated:

- **Server Cache Size Target**: 25% of the amount of memory: 48 GB * 0.25 = **12 GB**.
- **Database Max Cache Target**: Divide the Server Cache Size Target by the sum of:
  - The number of active databases
  - 20% of the number of passive databases

  12 GB / (2A + (2P * 0.20)) = **5 GB**
- **Memory used for passive database copies**: 20% of the Database Max Cache Target: 5 GB * 0.20 = **1 GB**.

Out of the 12 GB of memory assigned to the Server Cache Size Target:

- 12 GB will be in use by database worker processes
- No memory will be reserved by the Information Store for the two passive database copies because they cannot become active copies (*MaximumActiveDatabases* is configured with a value of 2, and there are already 2 active database copies on the server).

Managed Store Limits in Exchange Server

8/3/2020 • 5 minutes to read • Edit Online

The Managed Store in Exchange Server 2016 and Exchange Server 2019 is the name for the Information Store (also known as the Store) processes that manages mailbox databases. The Managed Store has connection and usage limits that prevent a single application or a single user from using all of the available connections, which could result in downtime. This topic describes the limits and how you can change them.

For more information about the Managed Store, see Managed Store in Exchange Server.

> **NOTE**
>
> Connections by administrator accounts have maximum session limits of 64000.
>
> Exchange Online limits (including Managed Store limits) are described in the Exchange Online Limits.

Terminology

Knowledge of the following terms will help you understand the types of connections referenced in this topic.

- **Sessions**

  Sessions represent the connections used by services and client applications (for example, Microsoft Outlook) to connect to the Managed Store. Services and clients can have multiple sessions at a particular time. The terms *connections* and *sessions* can be used interchangeably.

- **Threads**

  Threads represent concurrently executing requests to the Managed Store. For example, if a user opens a folder in Outlook, Outlook executes a request to the Managed Store on behalf of the user. That execution of the request is a single thread.

  For all clients, the maximum number of threads **per mailbox database** is 50. The exception is the Availability service, which has a maximum limit of 16 per user.

Session limits

Session limits are based on connections per mailbox database on the server.

The types of connection limits are:

- **Max sessions per process**: The maximum number of sessions that an Exchange service can have open at one time on a mailbox database.

- **Max user sessions per process**: The maximum number of sessions for a specific protocol for a single user.

The types of client connections to the Managed Store and the limits based on those connections are described in the following table.

CLIENT TYPE	MAX SESSIONS PER MAILBOX DATABASE	DEFAULT NUMBER OF USER SECTIONS PER MAILBOX DATABASE
Admin	10000	n/a
Availability service	10000	16
Content indexing	10000	n/a
Exchange ActiveSync	n/a	16
Exchange Web Services	n/a	16
Management	n/a	16
MAPI on the Middle Tier (MoMT)	n/a	32
MSExchangeMailboxAssistants: Events	10000	n/a
MSExchangeMailboxAssistants: Timed	10000	n/a
MSExchange Remote Procedure Call	n/a	16
Outlook on the web (formerly known as Outlook Web App)	n/a	16
POP3 and IMAP4	n/a	16
Transport	10000	n/a
Unified Messaging (Exchange 2016 only)	n/a	16
Others	n/a	16

Use the following procedure to modify the default session limits.

Notes:

- When you modify a session limit, you need to modify that limit on all Mailbox servers within a database availability group (DAG). If you don't make the same changes on all servers, the results will be inconsistent.
- To increase a session limit in the Client Access (frontend) services, you need to use the Set-ThrottlingPolicy cmdlet in the Exchange Management Shell.

WARNING

Incorrectly editing the registry can cause serious problems that may require you to reinstall your operating system. Problems resulting from editing the registry incorrectly may not be able to be resolved. Before editing the registry, back up any valuable data.

1. Open the Registry Editor. For example, press Windows key + R, and then run **regedit**.
2. Go to the following location in the registry:

\\HKEY_LOCAL_MACHINE
\SYSTEM\CurrentControlSet\Services\MSExchangeIS\ParametersSystem.

3. Select the **ParametersSystem** subkey, click **Edit** > **New**, and then select **DWORD (32-bit) Value**.

   The new value is created as **New Value #1** in the right pane.

4. Rename the new key to one of the following values, and then press Enter:

   - **Maximum Allowed Sessions Per User**: This limit specifies the maximum allowable sessions per user.
   - **Maximum Allowed Service Sessions Per User**: This limit specifies the maximum allowed service sessions per user.
   - **Maximum Allowed Exchange Sessions Per Service**: This limit specifies the maximum allowed Exchange sessions per service. The default value is 10,000.

5. Select the new key, and then click **Edit** > **Modify**.

6. In the dialog that opens, switch the **Base** value to **Decimal** and enter the new session limit in the **Value data** field.

   When you're finished, click **OK**.

Open item limits

Open item limits are limits placed on the number of items that can be opened by a single mailbox in a single session. However, a user can have multiple sessions opened simultaneously. For example, if a user has two sessions opened, the user could open 1,000 folders.

The open item limits are described in the following table

ITEM TYPE	REGISTRY OBJECT TYPE	MAX OPENED PER SESSION
ACL View	objtACLView	500
Attachment	objtAttachment	500
Attachment View	objtAttachmentView	500
CStream	objtCStream	Not applicable
Folder	objtFolder	500
Folder View	objtFolderView	500
FX Destination Stream	objtFXDstStrm	500
FX Source Stream	objtFXSrcStrm	500
Message	objtMessage	250
Message View	objtMessageView	500
Notification	objtNotify	500,000

ITEM TYPE	REGISTRY OBJECT TYPE	MAX OPENED PER SESSION
Rule View	objtRulesView	Not applicable
Stream	objtStream	250

You can limit the maximum number of resources that a MAPI client (for example, Outlook) can use simultaneously.

Note: When you modify a session limit, you need to modify that limit on all Mailbox servers within a database availability group (DAG). If you don't make the same changes on all servers, the results will be inconsistent.

> **WARNING**
> Incorrectly editing the registry can cause serious problems that may require you to reinstall your operating system. Problems resulting from editing the registry incorrectly may not be able to be resolved. Before editing the registry, back up any valuable data.

1. Open the Registry Editor. For example, press Windows key + R, and then run **regedit**.

2. Go to the following location in the registry:

   \\HKEY_LOCAL_MACHINE \SYSTEM\CurrentControlSet\Services\MSExchangeIS\ParametersSystem

3. Select the **ParametersSystem** subkey, click **Edit** > **New**, and then select **Key**.

   The new value is created as **New Key #1** in the left pane.

4. Rename the new key to **MaxObjsPerMapiSession**, and then press Enter.

5. Select the **MaxObjsPerMapiSession** subkey, click **Edit** > **New**, and then select **DWORD (32-bit) Value**.

   The new key is created as **New Value #1** in the right pane.

6. Rename the key to match one of the **Registry object type** values in the table. For example, to modify the number of messages that can be opened, enter *objtMessage* and then press Enter.

7. Select the new key, and then click **Edit** > **Modify**.

8. In the dialog that opens, switch the **Base** value to **Decimal** and enter the new limit in the **Value data** field. For example, enter **350** to increase the value for *objtMessage*.

   When you're finished, click **OK**.

9. Restart the Microsoft Exchange Information Store service by running the following command in Windows PowerShell or the Exchange Management Shell:

   ```
   Restart-Service MSExchangeIS
   ```

Item size limits

Item size limits are the limits placed on items within a user's mailbox. You configure these limits by using the *MaxSendSize* and *MaxReceiveSize* parameters on the Set-Mailbox cmdlet in the Exchange Management Shell.

ITEM TYPE	LIMIT
Message (saved)	Maximum size of the SendLimit, ReceiveLimit
Message (sent)	Maximum size of the SendLimit

Edge Transport servers

8/3/2020 • 3 minutes to read • Edit Online

Edge Transport servers handle all inbound and outbound Internet mail flow by providing mail relay and smart host services for your Exchange organization. Agents running on the Edge Transport server provide additional layers of message protection and security. These agents provide protection against spam and apply mail flow rules (also known as transport rules) to control mail flow. All of these features work together to help minimize the exposure of your internal Exchange to threats on the Internet.

Because the Edge Transport server is installed in the perimeter network, it's never a member of your organization's internal Active Directory forest and doesn't have access to Active Directory information. However, the Edge Transport server requires data that resides in Active Directory: for example, connector information for mail flow and recipient information for antispam recipient lookup tasks. This data is synchronized to the Edge Transport server by the Microsoft Exchange EdgeSync service (EdgeSync). EdgeSync is a collection of processes run on an Exchange 2016 or Exchange 2019 Mailbox server to establish one-way replication of recipient and configuration information from Active Directory to the Active Directory Lightweight Directory Services (AD LDS) instance on the Edge Transport server. EdgeSync copies only the information that's required for the Edge Transport server to perform antispam configuration tasks and to enable end-to-end mail flow. EdgeSync performs scheduled updates so the information in AD LDS remains current. For more information about Edge Subscriptions and EdgeSync, see Edge Subscriptions.

You can install more than one Edge Transport server in the perimeter network. Deploying more than one Edge Transport server provides redundancy and failover capabilities for your inbound message flow. You can load balance the SMTP traffic to your organization among Edge Transport servers by defining more than one MX record with the same priority value for your mail domain. You can achieve consistency in the configuration among multiple Edge Transport servers by using cloned configuration scripts.

The Edge Transport server role lets you manage the following message-processing scenarios.

Internet mail flow

Edge Transport servers accept messages coming into the Exchange organization from the Internet. After the messages are processed by the Edge Transport server, mail is routed to an internal Exchange Mailbox server; first to the Front End Transport service, and then to the Transport service.

All messages sent to the Internet from inside the organization are routed to Edge Transport servers after the messages are processed by the Transport service on the Exchange Mailbox server. You can configure the Edge Transport server to use DNS to resolve MX resource records for external SMTP domains, or you can configure the Edge Transport server to forward messages to a smart host for DNS resolution.

Antispam protection

In Exchange Server, antispam features provide services to block unsolicited commercial email (spam) at the network perimeter.

Spammers use a variety of techniques to send spam into your organization. Edge Transport servers help prevent users from ever receiving spam by providing a collection of agents that work together to provide different layers of spam filtering and protection. Establishing tarpitting intervals on connectors makes email harvesting attempts ineffective.

Mail flow rules on Edge Transport servers

Mail flow rules on Edge Transport servers are used to control the flow of messages sent to or received from the internet. Mail flow rules are configured on each Edge Transport server to help protect corporate network resources and data by applying an action to messages meeting specified conditions. Mail flow rule conditions are based on data, such as specific words or text patterns in the message subject, body, header, or from address; the spam confidence level (SCL); or the attachment type. Actions determine how the message is processed when a specified condition is true. Possible actions include quarantining a message, dropping or rejecting a message, appending additional recipients, or logging an event. Optional exceptions exempt particular messages from having an action applied.

Address rewriting

Address rewriting presents a consistent email address appearance to external recipients. You configure address rewriting on Edge Transport servers to modify the SMTP addresses on inbound and outbound messages. Address rewriting is especially useful for newly merged organizations that want to present a consistent email address appearance.

Edge Subscriptions

8/3/2020 • 15 minutes to read • Edit Online

Edge Subscriptions are used to populate the Active Directory Lightweight Directory Services (AD LDS) instance on the Edge Transport server with Active Directory data. Although creating an Edge Subscription is optional, subscribing an Edge Transport server to the Exchange organization provides a simpler management experience and enhances antispam features. You need to create an Edge Subscription if you plan to use recipient lookup or safelist aggregation, or if you plan to help secure SMTP communications with partner domains by using Mutual Transport Layer Security (MTLS).

Edge Subscription process

An Edge Transport server doesn't have direct access to Active Directory. The configuration and recipient information the Edge Transport server uses to process messages is stored locally in AD LDS. Creating an Edge Subscription establishes secure, automatic replication of information from Active Directory to AD LDS. The Edge Subscription process provisions the credentials used to establish a secure LDAP connection between the internal Exchange Mailbox servers and a subscribed Edge Transport server. The Microsoft Exchange EdgeSync service (EdgeSync) that runs on Mailbox servers performs periodic one-way synchronization to transfer up-to-date data to AD LDS. This reduces the administration tasks you perform in the perimeter network by letting you configure the Mailbox server and then synchronize that information to the Edge Transport server.

You subscribe an Edge Transport server to the Active Directory site that contains the Mailbox servers responsible for transferring messages to and from your Edge Transport servers. The Edge Subscription process creates an Active Directory site membership affiliation for the Edge Transport server. The site affiliation enables Mailbox servers in the Exchange organization to relay messages to the Edge Transport server for delivery to the Internet without having to configure explicit Send connectors.

One or more Edge Transport servers can be subscribed to a single Active Directory site. However, an Edge Transport server can't be subscribed to more than one Active Directory site. If you have more than one Edge Transport server deployed, each server can be subscribed to a different Active Directory site. Each Edge Transport server requires an individual Edge Subscription.

To deploy an Edge Transport server and subscribe it to an Active Directory site, follow these steps:

1. Install the Edge Transport server role.

2. Prepare for the Edge Subscription:

   - License the Edge Transport server.

   - Open ports in the firewall for mail flow and EdgeSync synchronization.

   - Verify that the Mailbox servers and the Edge Transport server can locate one another using DNS name resolution.

   - On the Mailbox Server, configure the transport settings to be replicated to the Edge Transport server.

3. On the Edge Transport server, create and export an Edge Subscription file by running the **New-EdgeSubscription** cmdlet.

4. Copy the Edge Subscription file to a Mailbox server or a file share that's accessible from the Active Directory site containing your Mailbox servers.

5. Import the Edge Subscription file to the Active Directory site by running the **New-EdgeSubscription**

cmdlet on the Mailbox server.

Prepare for the Edge Subscription

Before you can subscribe your Edge Transport server to your Exchange organization, you need to make sure your infrastructure and your Mailbox servers are prepared for the EdgeSync synchronization. To prepare for EdgeSync, you need to:

- **License the Edge Transport server**: The licensing information for the Edge Transport server is captured when the Edge Subscription is created. Subscribed Edge Transport servers need to be subscribed to the Exchange organization after the license key has been applied on the Edge Transport server. If the license key is applied on the Edge Transport server after you perform the Edge Subscription process, licensing information will not be updated in the Exchange organization, and you will need to resubscribe the Edge Transport server.

- **Verify that the required ports are open in the firewall**: The following ports are used by subscribed Edge Transport servers:

  - **SMTP**: Port 25/TCP must be open for inbound and outbound mail flow between the Internet and the Edge Transport server, and between the Edge Transport server and the internal Exchange organization.

  - **Secure LDAP**: Non-standard port 50636/TCP is used for directory synchronization from Mailbox servers to AD LDS on the Edge Transport server. This port is required for successful EdgeSync synchronization.

  > **NOTE**
  >
  > Port 50389/TCP is used locally by LDAP to bind to the AD LDS instance. This port doesn't have to be open on the firewall; it's used locally on the Edge Transport server.

  If your environment requires specific ports, you can modify the ports used by AD LDS using the `ConfigureAdam.ps1` script provided with Exchange. Modify the ports before you create the Edge Subscription. If you modify the ports after you create the Edge Subscription, you need to remove the Edge Subscription and create another one.

- **Verify that DNS host name resolution is successful from the Edge Transport server to the Mailbox servers and from the Mailbox servers to the Edge Transport server**

- **Configure the following transport settings for propagation to the Edge Transport server**

  - **Internal SMTP servers**: Use the *InternalSMTPServers* parameter on the **Set-TransportConfig** cmdlet to specify a list of internal SMTP server IP addresses or IP address ranges to be ignored by the Sender ID and Connection Filtering agents on the Edge Transport server.

  - **Accepted domains**: Configure all authoritative domains, internal relay domains, and external relay domains.

  - **Remote domains**: Configure the settings for the default remote domain object (used for recipients in all remote domains), and configure remote domain objects as required for recipients in specific remote domains.

Create and export an Edge Subscription file on the Edge Transport server

When you create an Edge Subscription file by running the **New-EdgeSubscription** cmdlet on the Edge Transport server, the following actions occur:

- An AD LDS account called the EdgeSync bootstrap replication account (ESBRA) is created. These ESBRA credentials are used to authenticate the first EdgeSync connection to the Edge Transport server. This account is configured to expire 24 hours after being created. Therefore, you need to complete the five-step subscription process described in the previous section within 24 hours. If the ESBRA expires before the Edge Subscription process is complete, you will need to run the **New-EdgeSubscription** cmdlet again to create a new Edge Subscription file.

- The ESBRA credentials are retrieved from AD LDS and written to the Edge Subscription file. The public key for the Edge Transport server's self-signed certificate is also exported to the Edge Subscription file. The credentials written to the Edge Subscription file are specific to the server that exported the file.

- Any previously created configuration objects on the Edge Transport server that will now be replicated to AD LDS from Active Directory are deleted from AD LDS, and the Exchange Management Shell cmdlets used to configure those objects are disabled. However, you can still use the **Get-*** cmdlets to view those objects. Running the **New-EdgeSubscription** cmdlet disables the following cmdlets on the Edge Transport server:

    - **Set-SendConnector**
    - **New-SendConnector**
    - **Remove-SendConnector**
    - **New-AcceptedDomain**
    - **Set-AcceptedDomain**
    - **Remove-AcceptedDomain**
    - **New-RemoteDomain**
    - **Set-RemoteDomain**
    - **Remove-RemoteDomain**

This example creates and exports the Edge Subscription file on the Edge Transport server.

```
New-EdgeSubscription -FileName "C:\Data\EdgeSubscriptionInfo.xml"
```

> **NOTE**
>
> When you run the **New-EdgeSubscription** cmdlet on the Edge Transport server, you receive a prompt to acknowledge the commands that will be disabled and the configuration that will be overwritten on the Edge Transport server. To bypass this confirmation, you need to use the *Force* parameter. This parameter is useful when you script the **New-EdgeSubscription** cmdlet. You can also use the *Force* parameter to overwrite an existing file when you resubscribe an Edge Transport server.

Import the Edge Subscription file on a Mailbox server

When you import the Edge Subscription file to the Active Directory site by running the **New-EdgeSubscription** cmdlet on a Mailbox server, the following actions occur:

- The Edge Subscription is created, joining the Edge Transport server to the Exchange organization. EdgeSync will propagate configuration data to this Edge Transport Server, creating an Edge configuration object in Active Directory.

- Each Mailbox server in the Active Directory site receives notification from Active Directory that a new Edge Transport server has been subscribed. The Mailbox server retrieves the ESBRA from the Edge Subscription file. The Mailbox server then encrypts the ESBRA by using the public key of the Edge Transport server's self-signed certificate. The encrypted credentials are then written to the Edge configuration object.

- Each Mailbox server also encrypts the ESBRA using its own public key and then stores the credentials in its own configuration object.

- EdgeSync replication accounts (ESRAs) are created in Active Directory for each Edge Transport-Mailbox server pair. Each Mailbox server stores its ESRA credentials as an attribute of the Mailbox server configuration object.
- Send connectors are automatically created to relay messages outbound from the Edge Transport server to the Internet, and inbound from the Edge Transport server to the Exchange organization. For more information, see the Send connectors created automatically by the Edge Subscription section in this topic.
- The Microsoft Exchange EdgeSync service that runs on Mailbox servers uses the ESBRA credentials to establish a secure LDAP connection between a Mailbox server and the Edge Transport server, and performs the initial replication of data. The following data is replicated to AD LDS:
  - Topology data
  - Configuration data
  - Recipient data
  - ESRA credentials
- The Microsoft Exchange Credential Service that runs on the Edge Transport server installs the ESRA credentials. These credentials are used to authenticate and secure later synchronization connections.
- The EdgeSync synchronization schedule is established.
- The Microsoft Exchange EdgeSync service running on the Mailbox servers in the subscribed Active Directory site then performs one-way replication of data from Active Directory to AD LDS on a regular schedule. You can also use the **Start-EdgeSynchronization** cmdlet to override the EdgeSync synchronization schedule and immediately start synchronization.

This example subscribes an Edge Transport server to the specified site and automatically creates the Internet Send connector and the Send connector from the Edge Transport server to the Mailbox servers.

```
New-EdgeSubscription -FileData ([byte[]]$(Get-Content -Path "C:\Data\EdgeSubscriptionInfo.xml" -Encoding Byte
-ReadCount 0)) -Site "Default-First-Site-Name"
```

> **NOTE**
>
> The default values of the *CreateInternetSendConnector* and *CreateInboundSendConnector* parameters are both `$true`, so you don't need to use them in this command.

Send connectors created automatically by the Edge Subscription

By default, when you import the Edge Subscription file to a Mailbox server, the Send connectors required to enable end-to-end mail flow between the Internet and the Exchange organization are created automatically, and any existing Send connectors on the Edge Transport server are deleted.

The Edge Subscription creates the following Send connectors:

- A Send connector named EdgeSync - Inbound to <*Site Name*> that's configured to relay messages from the Edge Transport server to the Exchange organization.
- A Send connector named EdgeSync - <*Site Name*> to Internet that's configured to relay messages from the Exchange organization to the Internet.

Also, subscribing an Edge Transport server to the Exchange organization allows the Mailbox servers in the subscribed Active Directory site to use the invisible and implicit intra-organization Send connector to relay messages to the Edge Transport server.

Inbound Send connector to receive messages from the Internet

When you run the **New-EdgeSubscription** cmdlet on the Mailbox server, the *CreateInboundSendConnector* parameter is set to the value `$true`. This creates the Send connector needed to send messages from the Edge Transport server to the Exchange organization. The following table shows the configuration of this Send connector.

Automatic inbound Send connector configuration

PROPERTY	VALUE
Name	EdgeSync - Inbound to < *Site Name*>
AddressSpaces	`SMTP:--;1` The `--` value in the address space represents all authoritative and internal relay accepted domains for the Exchange organization. Any messages the Edge Transport server receives for these accepted domains are routed to this Send connector and relayed to the smart hosts.
SourceTransportServers	< *Edge Subscription name*>
Enabled	True
DNSRoutingEnabled	False
SmartHosts	`--` The `--` value in the list of smart hosts represents all Mailbox servers in the subscribed Active Directory site. Any Mailbox servers you add to the subscribed Active Directory site after you establish the Edge Subscription don't participate in the EdgeSync synchronization process. However, they are automatically added to the list of smart hosts for the automatically created inbound Send connector. If more than one Mailbox server is located in the subscribed Active Directory site, inbound connections will be load balanced across the smart hosts.

You can't modify the address space or list of smart hosts at creation time for the automatically created inbound Send connector. However, you can set the *CreateInboundSendConnector* parameter to the value `$false` when you create an Edge Subscription. This allows you to manually configure a Send connector from the Edge Transport server to the Exchange organization.

Outbound Send connector to send messages to the Internet

When you run the **New-EdgeSubscription** cmdlet on the Mailbox server, the *CreateInternetSendConnector* parameter is set to the value `$true`. This creates the Send connector needed to send messages from the Exchange organization to the Internet. The following table shows the default configuration of this Send connector.

Automatic Internet Send connector configuration

PROPERTY	VALUE
Name	EdgeSync - < *Site Name*> to Internet
AddressSpaces	`SMTP:*;100`

PROPERTY	VALUE
SourceTransportServers	< *Edge Subscription name*> The name of the Edge Subscription is the same as the name of the subscribed Edge Transport server.
Enabled	True
DNSRoutingEnabled	True
DomainSecureEnabled	True

If more than one Edge Transport server is subscribed to the same Active Directory site, no additional Send connectors to the Internet are created. Instead, all Edge Subscriptions are added to the same Send connector as the source server. This load balances outbound connections to the Internet across the subscribed Edge Transport servers.

The outbound Send connector is configured to send email messages from the Exchange organization to all remote SMTP domains, using DNS routing to resolve domain names to MX resource records.

Details about the EdgeSync service

After you subscribe an Edge Transport server to an Active Directory site, EdgeSync will replicate configuration and recipient data to the Edge Transport servers. The service replicates the following data from Active Directory to AD LDS:

- Send connector configuration
- Accepted domains
- Remote domains
- Safe Senders Lists
- Blocked Senders Lists
- Recipients
- List of send and receive domains used in domain secure communications with partners
- List of SMTP servers listed as internal in your organization's transport configuration
- List of Mailbox servers in the subscribed Active Directory site

EdgeSync uses a mutually authenticated and authorized secure LDAP channel to transfer data from the Mailbox server to the Edge Transport server.

To replicate data to AD LDS, the Mailbox server binds to a global catalog server to retrieve updated data. EdgeSync initiates a secure LDAP session between a Mailbox server and the subscribed Edge Transport server over the non-standard TCP port 50636.

When you first subscribe an Edge Transport server to an Active Directory site, the initial replication that populates AD LDS with data from Active Directory can take five minutes or more, depending on the quantity of data in the directory service. After initial replication, EdgeSync only synchronizes new and changed objects, and removes any deleted objects.

Synchronization schedule

Different types of data synchronize on different schedules. The EdgeSync synchronization schedule specifies the maximum interval between EdgeSync synchronizations. EdgeSync synchronization occurs at the following intervals:

- Configuration data: 3 minutes.
- Recipient data: 5 minutes.
- Topology data: 5 minutes

If you want to change these intervals, use the **Set-EdgeSyncServiceConfig** cmdlet. Using the **Start-EdgeSynchronization** cmdlet on the Mailbox server to force Edge Subscription synchronization overrides the timer for the next scheduled EdgeSync synchronization, and starts EdgeSync immediately.

Selection of Mailbox servers

Each subscribed Edge Transport server is associated with a particular Active Directory site. If more than one Mailbox server exists in the site, any of these Mailbox servers can replicate data to the subscribed Edge Transport servers. To avoid contention among Mailbox servers when synchronizing, the preferred Mailbox server is selected as follows:

1. The first Mailbox server in the Active Directory site to perform a topology scan and discover the new Edge Subscription performs the initial replication. Because this discovery is based on the timing of the topology scan, any Mailbox server in the site may perform the initial replication.

2. The Mailbox server performing the initial replication establishes an EdgeSync lease option and sets a lock on the Edge Subscription. The lease option establishes that particular Mailbox server as the preferred server providing synchronization services to that Edge Transport server. The lock prevents EdgeSync running on another Mailbox server from taking over the lease option.

3. The EdgeSync lease option lasts for one hour. During that hour, no other EdgeSync service can take over the option unless a manual synchronization is started before the end of the hour. If the preferred Mailbox server isn't available to provide EdgeSync service at the time manual synchronization is started, after a five-minute wait, the lock is released and another EdgeSync service can take over the lease option and perform synchronization.

4. Unless manual synchronization is started, synchronization occurs based on the EdgeSync synchronization schedule. If the preferred server isn't available when a scheduled synchronization occurs, after a five-minute wait, the lock is released and another EdgeSync service can take over the lease option and perform synchronization.

This method of locking and leasing prevents more than one instance of EdgeSync from pushing data to the same Edge Transport server at the same time.

Notes:

- In Exchange 2016 organizations, if you also have Exchange 2010 Hub Transport servers in the subscribed Active Directory site, Exchange 2016 Mailbox servers will always take precedence and perform the replication.

- When you subscribe an Edge Transport server to an Active Directory site, all Mailbox servers installed in that Active Directory site at that time can participate in the EdgeSync synchronization process. If one of those servers is removed, the EdgeSync service that's running on the remaining Mailbox servers will continue the data synchronization process. However, if you later install new Mailbox servers in the Active Directory site, they won't automatically participate in EdgeSync synchronization. If you want to enable those new Mailbox servers to participate in EdgeSync synchronization, you will need to subscribe the Edge Transport server again.

The following table lists the EdgeSync properties related to locking and leasing. You can use the **Set-EdgeSyncServiceConfig** cmdlet to configure these properties.

EdgeSync lease properties

PARAMETER	DEFAULT VALUE	DESCRIPTION
LockDuration	`00:05:00` (5 minutes)	This setting determines how long a particular EdgeSync service will acquire a lock. If the EdgeSync service on the Mailbox server that's holding this lock doesn't respond, after five minutes the EdgeSync service on another Mailbox server will take over the lease. Forcing immediate EdgeSync synchronization doesn't override this setting.
OptionDuration	`01:00:00` (1 hour)	This setting determines how long an EdgeSync service can declare a lease option on an Edge Transport server. If the EdgeSync service holding the lease is unavailable and doesn't restart during this option period, no other Exchange EdgeSync service will take over the lease option unless you force EdgeSync synchronization.
LockRenewalDuration	`00:01:00` (1 minute)	This setting determines how frequently the lock field is updated when an EdgeSync service has acquired a lock to an Edge Transport server.

Procedures for Edge Subscriptions

8/3/2020 • 6 minutes to read • Edit Online

After you've subscribed an Edge Transport server to an Active Directory site in your Exchange organization as described in Edge Subscriptions, you might need to perform maintenance tasks on the Edge Subscription. These tasks are described in this topic.

What do you need to know before you begin?

- Estimated time to complete each procedure: 10 minutes.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "EdgeSync" entry and the "Edge Transport servers" section in the Mail flow permissions topic.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Remove an Edge Subscription

You may occasionally want to remove an Edge Subscription from the Exchange organization or from both the Exchange organization and the Edge Transport server. If you plan to later resubscribe the Edge Transport server to the Exchange organization, don't remove the Edge Subscription from the Edge Transport server. When you remove the Edge Subscription from an Edge Transport server, all replicated data is deleted from AD LDS. This can take a long time if you have lots of recipient data.

To completely remove an Edge Subscription, you need to run this procedure on the Edge Transport server you wish to remove and on an Exchange 2016 or Exchange 2019 Mailbox server in the Active Directory site where the Edge Transport server is subscribed.

After you remove the Edge Subscription, synchronization of information from AD LDS stops. All accounts stored in AD LDS are removed, and the Edge Transport server is removed from the source server list of any Send connector. You will no longer be able to use Edge Transport server features that rely on Active Directory data.

1. To remove the Edge Subscription from the Edge Transport server, use the following syntax.

   ```
   Remove-EdgeSubscription <EdgeTransportServerIdentity>
   ```

   For example, to remove the Edge Subscription on the Edge Transport server named Edge01, run the following command.

   ```
   Remove-EdgeSubscription Edge01
   ```

2. To remove the Edge Subscription from the Mailbox server, use the following syntax.

   ```
   Remove-EdgeSubscription <EdgeTransportServerIdentity>
   ```

For example, to remove the Edge Subscription for the Edge Transport server named Edge01 on a Mailbox server in the subscribed Active Directory site, run the following command.

```
Remove-EdgeSubscription Edge01
```

You will need to remove the Edge Subscription if:

- You no longer want the Edge Transport server to participate in EdgeSync synchronization. You will need to remove the Edge Subscription from both the Edge Transport server and from the Exchange organization.
- An Edge Transport server is being decommissioned. In this scenario, you only need to remove the Edge Subscription from the Exchange organization. If you uninstall the Edge Transport server role from the computer, the AD LDS instance and all Active Directory data stored in AD LDS will also be removed.
- You want to change the Active Directory site association for the Edge Subscription. You will only need to remove the Edge Subscription from the Exchange organization. After the Edge Subscription is removed from the Exchange organization, you can resubscribe the Edge Transport server to a different Active Directory site.

When you remove an Edge Subscription from the Exchange organization:

- Synchronization of information from Active Directory to AD LDS stops.
- The ESRA accounts are removed from both Active Directory and AD LDS.
- The Edge Transport server is removed from the *SourceTransportServers* property of any Send connector.
- The automatic inbound Send connector from the Edge Transport server to the Exchange organization is removed from AD LDS.

When you remove the Edge Subscription from an Edge Transport server:

- You can no longer use Edge Transport server features that rely on Active Directory data.
- Replicated data is removed from AD LDS.
- Tasks that were disabled when the Edge Subscription was created are re-enabled to allow for local configuration.

Resubscribe an Edge Transport server

Occasionally you may have to resubscribe an Edge Transport server to an Active Directory site. When the Edge Subscription is re-created, new credentials are generated and you need to follow the complete Edge Subscription process. You will need to resubscribe an Edge Transport server if:

- You add new Mailbox servers in the subscribed Active Directory site, and you want the new Mailbox server to participate in EdgeSync synchronization.
- You applied the license key for the Edge Transport server after creating the Edge Subscription. Licensing information for the Edge Transport server is captured when the Edge Subscription is created. Subscribed Edge Transport servers only appear as licensed if they are subscribed to the Exchange organization after the license key has already been applied on the Edge Transport server. If the license key is applied on the Edge Transport server after you perform the Edge Subscription process, the licensing information won't be updated in the Exchange organization, and you will need to resubscribe the Edge Transport server.
- The ESRA credentials are compromised.

IMPORTANT

To resubscribe an Edge Transport server, export a new Edge Subscription file on the Edge Transport server and then import the XML file on a Mailbox server. You will need to resubscribe the Edge Transport server to the same Active Directory site where it was originally subscribed. You don't need to first remove the original Edge Subscription; the resubscription process will overwrite the existing Edge Subscription.

Add or Remove a Mailbox server

If you add a Mailbox server to an Active Directory site that already has an Edge Transport server subscribed, the new Mailbox server doesn't automatically participate in EdgeSync synchronization. To enable a newly deployed Mailbox server to participate in EdgeSync synchronization, you need to resubscribe each Edge Transport server to the Active Directory site.

Removing a Mailbox server from an Active Directory site where an Edge Transport server is subscribed won't affect EdgeSync synchronization unless that Mailbox server is the only Mailbox server in that site. If you remove all Mailbox servers from the Active Directory site where an Edge Transport server is subscribed, that site's subscribed Edge Transport servers are orphaned.

Run EdgeSync manually

You may want to manually run EdgeSync if you've made significant changes to the configuration or recipients in Active Directory and want your changes synchronized immediately. You can run a full synchronization, or only synchronize changes made since the last replication.

A manual EdgeSync resets the EdgeSync synchronization schedule. The next automatic synchronization is based on when you ran the manual synchronization.

To manually run EdgeSync, use the following syntax.

```
Start-EdgeSynchronization [-Server <MailboxServerIdentity>] [-TargetServer <EdgeTransportServerIdentity> [-ForceFullSync]
```

The following example starts EdgeSync with the following options:

- The synchronization is initiated from the Exchange Mailbox server named Mailbox01.
- All Edge Transport servers are synchronized.
- Only the changes since the last replication are synchronized.

```
Start-EdgeSynchronization -Server Mailbox01
```

This example starts EdgeSync with the following options:

- The synchronization is initiated from the local Mailbox server.
- Only the Edge Transport server named Edge03 is synchronized.
- All recipient and configuration data are fully synchronized.

```
Start-EdgeSynchronization -TargetServer Edge03 -ForceFullSync
```

Verify EdgeSync results

You can use the **Test-EdgeSynchronization** cmdlet to verify that the Edge synchronization is working. This cmdlet reports synchronization status of subscribed Edge Transport servers.

The output of this cmdlet lets you view objects that have not been synchronized to the Edge Transport server. The task compares data stored in Active Directory against data stored in AD LDS and reports any data inconsistencies.

You can use the *ExcludeRecipientTest* parameter on the **Test-EdgeSynchronization** cmdlet to exclude validation of recipient data synchronization. If you include this parameter, only the synchronization of configuration objects is validated. Validating recipient data will take longer than validating only configuration data.

Verify EdgeSync results for a single recipient

To verify EdgeSync results for a single recipient, use the following syntax on a Mailbox server in the subscribed Active Directory site.

```
Test-EdgeSynchronization -VerifyRecipient <emailaddress>
```

This example verifies EdgeSync results for the user kate@contoso.com.

```
Test-EdgeSynchronization -VerifyRecipient kate@contoso.com
```

Configure internet mail flow through Edge Transport servers without using EdgeSync

8/3/2020 • 10 minutes to read • Edit Online

We recommend that you use the Edge Subscription process to establish mail flow between your Exchange organization and an Edge Transport server as described in Edge Subscriptions. However, certain situations may prevent you from subscribing the Edge Transport server to your Exchange organization. To manually establish mail flow between your Exchange organization and an unsubscribed Edge Transport server, you need to manually create and/or modify the following Send connectors and Receive connectors:

On the Edge Transport server:

- Create a dedicated Send connector to only send messages to the internet.
- Create a dedicated Send connector to only send messages to Mailbox servers in the Exchange organization.[1]
- Create a dedicated Receive connector to only receive messages from Mailbox servers in the Exchange organization[2]
- Modify the default Receive connector to only accept messages only from the internet.

On a Mailbox server:

- Create a dedicated Send connector to relay outgoing messages to the Edge Transport server

[1]The Send connector that's created by an EdgeSync subscription for delivering email into the Exchange organization is configured to use Exchange Server (GSSAPI) authentication. The EdgeSync subscription identifies the Edge Transport server as an Exchange server to the internal Active Directory forest, which also allows Exchange Server authentication. By definition, there is no EdgeSync subscription in this scenario, so you'll need to improvise:

- You can configure Basic authentication over TLS to provide authentication and encryption for email traffic between the Edge Transport server and the internal Exchange organization. This method has the following issues:
  - You need to configure an Active Directory account that belongs to the Exchange Servers universal security group for authentication on the Send connector that relays messages from the Edge Transport server to the internal Exchange organization. Be sure to safeguard the account credentials, and you can configure the account to allow logon only to specific computers. You also need a local account on the Edge Transport server for authentication on the Send connector that relays messages from the internal Exchange organization to the Edge Transport server.
  - Messages coming from these Send connectors will be seen as authenticated SMTP by the destination Mailbox server. This means the default Receive connector named Client Frontend <ServerName> in the Front End Transport service will accept the messages on port 587, and the messages are accepted in the backend Transport service using the default Receive connector named Client Proxy <ServerName> on port 465.
  - To provide encryption, you need to use a certificate. The self-signed certificate on the Edge Transport server won't be recognized by the internal Exchange Organization (again, the EdgeSync subscription usually takes care of this). You'll need to manually import the self-signed certificate on each Mailbox or use a certificate from a trusted third-party certification authority.
- If you don't want the messages coming from the Edge Transport server to be identified as authenticated

SMTP and therefore using the corresponding client Receive connectors, you can use Externally Secured as the authentication method, which means email traffic between the Edge Transport server and the internal Exchange organization isn't authenticated or encrypted *by Exchange*. If you use this method, you **must** configure and use an external encryption method (for example, IPsec or a VPN).

[2]Instead of a dedicated Receive connector, you can configure and use the default Receive connector on the Edge Transport server for both incoming internet messages and incoming messages from internal Mailbox servers (an EdgeSync subscription uses this Receive connector for both connections).

For more information about Send connectors, see Send connectors. For more information about Receive connectors, see Receive connectors.

Before you begin

- Estimated time to complete this task: 30 minutes.

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Send connectors" entry, the "Send connectors - Edge Transport" entry, and the "Receive connectors - Edge Transport" entry in the Mail flow permissions topic.

- On Edge Transport servers, you can only use the Exchange Management Shell to create Send connectors and Receive connectors. On Mailbox servers, you can use the Exchange admin center (EAC) or the Exchange Management Shell to create Send connectors.

  To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.

  For information about opening and using the EAC, see Exchange admin center in Exchange Server.

- The basic configuration of an Edge Transport server in the perimeter network must allow for resolving public domains for internet email and internal host names for internal email. There are different ways to do this, but you can configure the network adapter that's connected to the external (public) network segment to use a public DNS server, and configure the network adapter that's connected to the internal (private) network segment to use a DNS server in the perimeter network.

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Edge Transport Server Procedures

Step 1: Create a dedicated Send connector to only send messages to the internet

This Send connector requires the following configuration:

- **Name**: To Internet (or any descriptive name)

- **Usage type**: Internet

- **Address spaces**: "*" (all domains)

- **Network settings**: Use DNS MX records to route mail automatically. Depending on your network configuration, you can also route mail through a smart host. The smart host then routes mail to the internet.

To create a Send connector that's configured to send messages to the internet, run this command:

```
New-SendConnector -Name "To Internet" -AddressSpaces * -Usage Internet -DNSRoutingEnabled $true
```

For detailed syntax and parameter information, see New-SendConnector.

Step 2: Create a dedicated Send connector to only send messages to the Exchange organization

This Send connector requires the following configuration:

- **Name**: To Internal Org (or any descriptive name)
- **Usage type**: Internal
- **Address spaces**: `--` (indicates all accepted domains for the Exchange organization)
- DNS routing disabled (smart host routing enabled)
- **Smart hosts**: FQDN of one or more Mailbox servers as smart hosts. For example, mbxserver01.contoso.com and mbxserver02.contoso.com.
- **Smart host authentication methods**: Basic authentication over TLS
- **Smart host authentication credentials**: Credentials for the user account in the internal domain that's a member of the Exchange Servers universal security group. You need to use the **Get-Credential** cmdlet to store the credentials. Use the format *<Domain>*\ *<UserName>* or the user principal name (UPN; for example, chris@contoso.com) to enter the username.

To create a Send connector configured to send messages to the Exchange organization, replace the smart host values with the Mailbox servers in your organization, and run this command:

```
New-SendConnector -Name "To Internal Org" -Usage Internal -AddressSpaces "--" -DNSRoutingEnabled $false -
SmartHosts mbxserver01.contoso.com,mbxserver02.contoso.com -SmartHostAuthMechanism BasicAuthRequireTLS -
AuthenticationCredential (Get-Credential)
```

For detailed syntax and parameter information, see New-SendConnector.

Step 3: Modify the default Receive connector to only accept messages from the internet

Make the following configuration changes to the default Receive connector:

- Modify the name to indicate that the connector will be used solely to receive email from the internet (the default name is Default internal receive connector *<ServerName>*).
- Change the network bindings to accept messages only from the network adapter that is accessible from the internet (for example, 10.1.1.1 and the standard SMTP TCP port value of 25).

To modify the default Receive connector to only accept messages from the internet, replace *< ServerName>* and bindings ith the name of your Edge Transport server and external network adapter configuration, and run this command:

```
Set-ReceiveConnector -Identity "Default internal Receive connector ServerName>" -Name "From Internet" -Bindings
10.1.1.1:25
```

For detailed syntax and parameter information, see Set-ReceiveConnector.

Step 4: Create a dedicated Receive connector to only accept messages from the Exchange organization

This Receive connector requires the following configuration:

- **Name**: From Internal Org (or any descriptive name)

- **Usage type**: Internal
- **Local network bindings**: Internal network-facing network adapter (for example, 10.1.1.2 and the standard SMTP TCP port value of 25).
- **Remote network settings**: IP address of one or more Mailbox servers in the Exchange organization. For example, 192.168.5.10 and 192.168.5.20.
- **Authentication methods**: TLS, Basic authentication, Basic authentication over TLS, and Exchange Server authentication.

To create a Receive connector configured to only accept messages from the Exchange organization, replace the bindings and remote IP ranges with your values, and run this command.

```
New-ReceiveConnector -Name "From Internal Org" -Usage Internal -AuthMechanism
TLS,BasicAuth,BasicAuthRequireTLS,ExchangeServer -Bindings 10.1.1.2:25 -RemoteIPRanges
192.168.5.10,192.168.5.20
```

For detailed syntax and parameter information, see New-ReceiveConnector.

How do you know this worked?

To verify that you have successfully configured the required Send connectors and Receive connectors on the Edge Transport server, run this command on the Edge Transport server and verify the property values:

```
Get-SendConnector | Format-List
Name,Usage,AddressSpaces,SourceTransportServers,DSNRoutingEnabled,SmartHosts,SmartHostAuthMechanism; Get-
ReceiveConnector | Format-List Name,Usage,AuthMechanism,Bindings,RemoteIPRanges
```

Mailbox server procedures

You don't need to modify the default Receive connectors on Mailbox servers. For more information about default Receive connectors on Mailbox servers, see Default Receive connectors created during setup.

Step 5: Create a dedicated Send connector to send outgoing messages to the Edge Transport server

This Send connector requires the following configuration:

- **Name**: To Edge (or any descriptive name)
- **Usage type**: Internal
- **Address spaces**: "*" (all external domains)
- DNS routing disabled (smart host routing enabled)
- **Smart hosts**: IP address or FQDN of the Edge Transport server. For example, edge01.contoso.net.
- **Source servers**: FQDN of one or more Mailbox servers. For example, mbxserver01.contoso.com and mbxserver02.contoso.com.
- **Smart host authentication methods**: Basic authentication over TLS.
- **Smart host authentication credentials**: Credentials for the user account on the Edge Transport server.

Use the EAC to create a Send connector to send outgoing messages to the Edge Transport server

1. In the EAC, go to **Mail flow** > **Send connectors**, and then click **Add** +. This starts the **New Send connector** wizard.
2. On the first page, configure these settings:

   - **Name**: Enter To Edge.

   - **Type**: Select **Internal**.

   Click **Next**.

3. On the next page, select **Route mail through smart hosts**, and then click **Add** +. In the **Add smart host** dialog box that appears, identify the Edge Transport server by using one of these values:

   - **IP address**: For example, 10.1.1.2.

   - **Fully qualified domain name (FQDN)**: For example, edge01.contoso.net. Note that the source Mailbox servers for the Send connector must be able to resolve the Edge Transport server in DNS by using this FQDN. If they can't, use the IP address instead.

   Click **Save**.

4. On the next page, in the **Smart host authentication** section, select **Basic authentication**, and then configure these additional settings:

   - Select **Offer basic authentication only after starting TLS**

   - In the **User name** and **Password** fields, enter the credentials for the local user account on the Edge Transport server.

   Click **Next**.

5. On the next page, in the **Address space** section, click **Add** +. In the **Add domain** dialog box that appears, enter the following information:

   - **Type**: Verify SMTP is selected.

   - **Fully Qualified Domain Name (FQDN)**: Enter an asterisk (*) to indicate the Send connector is used for all external domains.

   - **Cost**: Verify 1 is entered. A lower value indicates a more preferred route.

   Click **Save**.

6. Back on the previous page, the **Scoped send connector** setting is important if your organization has Exchange servers installed in multiple Active Directory sites:

   - If you don't select **Scoped send connector**, the connector is usable by all transport servers (Exchange 2019 Mailbox servers, Exchange 2016 Mailbox servers, Exchange 2013 Mailbox servers, and Exchange 2010 Hub Transport servers) in the entire Active Directory forest. This is the default value.

   - If you select **Scoped send connector**, the connector is only usable by other transport servers in the same Active Directory site.

   Click **Next**.

7. On the next page, in the **Source server** section, click **Add** +. In the **Select a Server** dialog box that appears, select one or more Mailbox servers that you want to use to send outgoing mail through the Edge Transport server. Select a Mailbox server and click **Add ->** (repeat as many times a necessary), click **OK**, and then click **Finish**.

Use the Exchange Management Shell to create a Send connector to send outgoing messages to the Edge Transport server

To create a Send connector to send outgoing messages to the Edge Transport server, replace the smart hosts and source Mailbox servers with your values, and run this command:

```
New-SendConnector -Name "To Edge" -Usage Internal -AddressSpaces * -DNSRoutingEnabled $false -SmartHosts
edge01.contoso.com -SourceTransportServers mbxserver01.contoso.com,mbxserver02.contoso.com -
SmartHostAuthMechanism BasicAuthRequireTLS -AuthenticationCredential (Get-Credential)
```

For detailed syntax and parameter information, see New-SendConnector.

How do you know this worked?

To verify that you've successfully created a Send connector to send outgoing messages to the Edge Transport server, use either of these steps:

- In the EAC, go to **Mail flow** > **Send connectors**, select the Send connector named To Edge > click **Edit** ✎, and verify the property values.

- In the Exchange Management Shell, run this command on a Mailbox server to verify the property values:

  ```
  Get-SendConnector -Identity "To Edge" | Format-List
  Usage,AddressSpaces,DSNRoutingEnabled,SmartHosts,SourceTransportServers,SmartHostAuthMechanism
  ```

Address rewriting on Edge Transport servers

8/3/2020 • 9 minutes to read • Edit Online

Address rewriting in Exchange Server modifies the email addresses of senders and recipients in messages that enter or leave your organization through an Edge Transport server. Two transport agents on the Edge Transport server provide the rewriting functionality: the Address Rewriting Inbound Agent and the Address Rewriting Outbound Agent. The primary reason for address rewriting on outbound messages is to present a single, consistent email domain to external recipients. The primary reason for address rewriting on inbound messages is to deliver messages to the correct recipient.

The *address rewrite entry*, which you create, specifies the internal addresses (the email addresses you want to change) and the external addresses (the final email addresses you want). You can specify whether email addresses are rewritten in inbound and outbound messages, or in outbound messages only. You can create address writing entries for a single user (chris@contoso.com to support@contoso.com), a single domain (contoso.com to fabrikam.com), or for multiple subdomains with exceptions (*.fabrikam.com to contoso.com, except legal.fabrikam.com).

> **IMPORTANT**
> Regardless of how you plan to use address rewriting, you need to verify that the resulting email addresses are unique in your organization so you don't end up with duplicates. Address rewriting doesn't verify the uniqueness of a rewritten email address.

To configure address rewriting, see Address rewriting procedures on Edge Transport servers.

Scenarios for address rewriting

The following scenarios are examples of how you can use address rewriting:

- **Group consolidation**: Some organizations segment their internal businesses into separate domains that are based on business or technical requirements. This configuration can cause email messages to appear as if they come from separate groups or even separate organizations.

  The following example shows how an organization, Contoso, Ltd., can hide its internal subdomains from external recipients:

  - Outbound messages from the northamerica.contoso.com, europe.contoso.com, and asia.contoso.com domains are rewritten so they appear to originate from a single contoso.com domain. All messages are rewritten as they pass through Edge Transport servers that provide SMTP connectivity between the whole organization and the Internet.

  - Inbound messages to contoso.com recipients are relayed by the Edge Transport server to a Mailbox server. The message is delivered to the correct recipient based on the proxy address that's configured on the recipient's mailbox.

- **Mergers and acquisitions**: An acquired company might continue to run as a separate business, but you can use address rewriting to make the two organizations appear as if they're one integrated organization.

  The following example shows how Contoso, Ltd. can hide the email domain of the newly acquired company, Fourth Coffee:

  - Contoso, Ltd. wants all outbound messages from Fourth Coffee's Exchange organization to appear as if they originate from contoso.com. All messages from both organizations are sent through the Edge Transport servers at Contoso, Ltd., where email messages are rewritten from *user*@fourthcoffee.com to *user*@contoso.com.

    - Inbound messages to *user*@contoso.com are rewritten and routed to *user*@fourthcoffee.com mailboxes. Inbound messages that are sent to *user*@fourthcoffee.com are routed directly to Fourth Coffee's email servers.

- **Partners**: Many organizations use external partners to provide services for their customers, other organizations, or their own organization. To avoid confusion, the organization might replace the email domain of the partner organization with its own email domain.

  The following example shows how Contoso, Ltd. can hide a partner's email domain:

    - Contoso, Ltd. provides support for the larger Wingtip Toys organization. Wingtip Toys wants a unified email experience for its customers, and it requires all messages from support personnel at Contoso, Ltd. to appear as if they were sent from Wingtip Toys. All outbound messages that relate to Wingtip Toys are sent through their Edge Transport servers, and all contoso.com email addresses are rewritten to wingtiptoys.com email addresses.

    - Inbound messages for support@wingtiptoys.com are accepted by Wingtip Toy's Edge Transport servers, rewritten, and then routed to the support@contoso.com email address.

Message properties modified by address rewriting

A standard SMTP email message consists of a *message envelope* and message content. The message envelope contains information that's required for transmitting and delivering the message between SMTP messaging servers. The message content contains message header fields (collectively called the *message header*) and the message body. The message envelope is described in RFC 2821, and the message header is described in RFC 2822.

When a sender composes an email message and submits it for delivery, the message contains the basic information that's required to comply with SMTP standards, such as a sender, a recipient, the date and time that the message was composed, an optional subject line, and an optional message body. This information is contained in the message itself and, by definition, in the message header.

The sender's mail server generates a message envelope for the message by using the sender's and recipient's information found in the message header. It then transmits the message to the Internet for delivery to the recipient's messaging server. Recipients never see the message envelope because it's generated by the message transmission process, and it isn't actually part of the message.

Address rewriting changes an email address by rewriting specific fields in the message header or message envelope. Address rewriting changes several fields in outbound messages, but only one field in inbound email messages. The following table shows which SMTP header fields are rewritten in outbound and inbound messages.

Message fields rewritten on outbound and inbound messages

FIELD NAME	LOCATION	OUTBOUND MESSAGES	INBOUND MESSAGES
MAIL FROM	Message envelope	Rewritten	Not rewritten
RCPT TO	Message envelope	Not rewritten	Rewritten
To	Message header	Not Rewritten	Rewritten
Cc	Message header	Not Rewritten	Rewritten

FIELD NAME	LOCATION	OUTBOUND MESSAGES	INBOUND MESSAGES
From	Message header	Rewritten	Not rewritten
Sender	Message header	Rewritten	Not rewritten
Reply-To	Message header	Rewritten	Not rewritten
Return-Receipt-To	Message header	Rewritten	Not rewritten
Disposition-Notification-To	Message header	Rewritten	Not rewritten
Resent-From	Message header	Rewritten	Not rewritten
Resent-Sender	Message header	Rewritten	Not rewritten

What address rewriting doesn't change

Address rewriting doesn't modify any message header fields that would break SMTP functionality. For example, modifying certain header fields can affect routing loop detection, invalidate the digital signature, or make a rights-protected message unreadable. Therefore, the following header fields aren't modified by address rewriting.

- **Return-Path**
- **Received**
- **Message-ID**
- **X-MS-TNEF-Correlator**
- **Content-Type Boundary=string**
- Header fields located inside MIME body parts

Address rewriting ignores domains that aren't controlled by the Exchange organization. In other words, the domain needs to be configured as an authoritative accepted domain in the Exchange organization. Rewriting non-authoritative domains would cause an uncontrollable form of message relay.

Address rewriting also doesn't modify the header fields of messages that are embedded in another message. Senders and recipients expect embedded messages to remain intact and be delivered without modification, as long as the messages don't trigger mail flow rules (also known as transport rules) that are implemented between the sender and recipient.

Considerations for outbound-only address rewriting

Outbound-only address rewriting on an Edge Transport server modifies the sender's email address as messages leave the Exchange organization. You can configure outbound-only address rewriting for a single user (chris@contoso.com to support@contoso.com), or for a single domain (contoso.com to fabrikam.com). You are required to configure outbound-only address rewriting for multiple subdomains (*.fabrikam.com to contoso.com).

The rewritten email address needs to be configured as a proxy address on the affected recipients. For example, if laura@sales.contoso.com is rewritten to laura@contoso.com, the proxy address laura@contoso.com needs to be configured on Laura's mailbox. This allows replies and inbound messages to be delivered correctly.

Considerations for inbound and outbound address rewriting

Inbound and outbound, or *bidirectional* address rewriting on an Edge Transport server modifies the sender's email address as messages leave the Exchange organization, and the recipient's email address as messages enter the Exchange organization.

You can configure bidirectional address rewriting for a single user (chris@contoso.com to support@contoso.com), and a single domain (contoso.com to fabrikam.com). You can't configure bidirectional address rewriting for multiple subdomains (*.fabrikam.com to contoso.com).

Considerations for rewriting email addresses in multiple domains

When you flatten multiple internal domains or subdomains into a single external domain, you need to consider the following factors:

- **Verify unique aliases**: All email aliases (the part to the left of the @ sign) need to be unique across all subdomains. For example, if there is a joe@sales.contoso.com, there can't be a joe@marketing.contoso.com because the rewritten email address for both users would be joe@contoso.com.
- **Add proxy addresses**: The rewritten email address needs to be configured as a proxy address on all affected senders in the affected domains. For example, if joe@sales.contoso.com is rewritten to joe@contoso.com, you need to add the proxy address joe@contoso.com to Joe's mailbox. This allows replies and inbound messages to be delivered correctly.
- **Mail contacts for non-Exchange organizations**: If you're rewriting email addresses from a non-Exchange email system, you need to create mail contacts in Exchange to represent the users in the non-Exchange email system. These email contacts need to contain the original email addresses and the rewritten email addresses. For example, if joe@unix.contoso.com is rewritten to joe@contoso.com, you need to create a mail contact with joe@unix.contoso.com as the external email address and joe@contoso.com as a proxy address.

Verify unique aliases

When you rewrite email addresses in multiple subdomains, you need to make sure that all email aliases are unique across all your subdomains. For example, consider the following configuration:

The following users are in the subdomains sales.contoso.com, marketing.contoso.com, and research.contoso.com:

- maria@sales.contoso.com
- chris@sales.contoso.com
- david@marketing.contoso.com
- brian@marketing.contoso.com
- chris@research.contoso.com
- adam@research.contoso.com

Suppose you want to rewrite the subdomains sales.contoso.com, marketing.contoso.com, and research.contoso.com into the single domain contoso.com.

When the email addresses in each subdomain are rewritten, a conflict occurs between chris@sales.contoso.com and chris@research.contoso.com, because both email addresses are rewritten to chris@contoso.com. To resolve this issue, you need to change the email address of one of the affected recipients. For example, you can change chris@research.contoso.com to christopher@research.contoso.com so the email address is rewritten to christopher@contoso.com.

Priority of address rewrite entries

If a user's email address matches multiple address rewrite entries, the email address is only rewritten once based on the closest match. The following list describes the order of precedence of address rewrite entries from highest priority to lowest priority:

1. **Individual email addresses**: An address rewrite entry is configured to rewrite the email address of john@contoso.com to support@contoso.com.

2. **Domain or subdomain mapping**: An address rewrite entry is configured to rewrite all contoso.com email addresses to northwindtraders.com or all sales.contoso.com email addresses to contoso.com.

3. **Domain flattening**: An address rewrite entry is configured to rewrite *.contoso.com email addresses to contoso.com.

For example, consider an Edge Transport server where the following outbound address rewrite entries are configured:

- *.contoso.com email addresses are rewritten to contoso.com

- japan.sales.contoso.com email addresses are rewritten to contoso.jp

If masato@japan.sales.contoso.com sends an email message, the address is rewritten to masato@contoso.jp, because that entry most closely matches the sender's email address.

Digitally signed, encrypted, and rights-protected messages

Address rewriting shouldn't affect most signed, encrypted, or rights-protected messages. If address rewriting were to invalidate or otherwise change the security status of these types of messages in any way, address rewriting isn't applied.

The following values can be rewritten because the information isn't part of message signing, encryption, or rights protection:

- Fields in the message envelope.

- Top-level message body headers.

The following values aren't rewritten because the information is part of message signing, encryption, or rights protection:

- Header fields located inside MIME body parts that may be signed.

- The boundary string parameter of the MIME content type.

Address rewriting procedures on Edge Transport servers

8/3/2020 • 8 minutes to read • Edit Online

You can create address rewrite entries on Edge Transport servers that apply to a single recipient, a specific domain or subdomain, or multiple subdomains. Address rewriting can be outbound only, or inbound and outbound (bidirectional). When you create address rewrite entries, remember the following:

- Verify that the resulting email addresses are unique in your organization.
- Only literal strings are supported in the email address values.
- The wildcard character (*) is supported only in the internal address (the addresses you want to change). Valid syntax for using the wildcard character is ***.contoso.com**. The values ***contoso.com** or **sales.*.com** are not allowed.
- When you use the wildcard character, you need to configure the address rewriting as outbound only (you need to set the *OutboundOnly* parameter to the value `$true`), and outbound only address rewriting requires that you configure the rewritten email address as a proxy address on the affected recipients.
- By default, address rewriting is bidirectional for a single recipient, or for a specific domain or subdomain (the default value for the *OutboundOnly* parameter is `$false`).

For more information about address rewriting, see Address rewriting on Edge Transport servers.

What do you need to know before you begin?

- Estimated time to complete each procedure: 10 minutes.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Edge Transport servers" section in the Mail flow permissions topic.
- You can only use PowerShell to perform this procedure. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- Be careful when you configure address rewriting. Any changes that you make are immediately applied when you run the command. Consider running the command with the *WhatIf* parameter. For more information about the *WhatIf* parameter, see WhatIf and Confirm.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to enable or disable address rewriting

To completely enable or disable address rewriting, you enable or disable the address rewriting agents. By default, the address rewriting agents on an Edge Transport server are enabled.

To disable address rewriting, run the following command:

```
Disable-TransportAgent "Address Rewriting Inbound Agent"; Disable-TransportAgent "Address Rewriting Outbound
Agent"
```

To enable address rewriting, run the following command:

```
Enable-TransportAgent "Address Rewriting Inbound Agent"; Enable-TransportAgent "Address Rewriting Outbound
Agent"
```

How do you know this worked?

To verify that you have successfully enabled or disabled address rewriting, run the following command to verify the **Enabled** property value:

```
Get-TransportAgent "Address Rewriting *"
```

Use the Exchange Management Shell to view address rewrite entries

To view a summary list of all address rewrite entries, run the following command.

```
Get-AddressRewriteEntry
```

To view details of an address rewrite entry, use the following syntax.

```
Get-AddressRewriteEntry <AddressRewriteEntryIdentity> | Format-List
```

The following example displays the details of the address rewrite entry named Rewrite Contoso.com to Northwindtraders.com:

```
Get-AddressRewriteEntry "Rewrite Contoso.com to Northwindtraders.com" | Format-List
```

For more information, see Get-AddressRewriteEntry.

Use the Exchange Management Shell to create address rewrite entries

Rewrite the email address for a single recipient

To rewrite the email address for a single recipient, use the following syntax:

```
New-AddressRewriteEntry -Name "<Descriptive Name>" -InternalAddress <internal email address> -ExternalAddress
<external email address> [-OutboundOnly <$true | $false>]
```

This example rewrites the email address of all messages entering and leaving the Exchange organization for joe@contoso.com. Outbound messages are rewritten so they appear to come from support@nortwindtraders.com. Inbound messages sent to support@northwindtraders.com are rewritten to joe@contoso.com for delivery to the recipient (the *OutboundOnly* parameter is `$false` by default).

```
New-AddressRewriteEntry -Name "joe@contoso.com to support@northwindtraders.com" -InternalAddress
joe@contoso.com -ExternalAddress support@northwindtraders.com
```

Rewrite email addresses in a single domain or subdomain

To rewrite the email addresses in a single domain or subdomain, use the following syntax:

```
New-AddressRewriteEntry -Name "<Descriptive Name>" -InternalAddress <domain or subdomain> -ExternalAddress
<domain> [-OutboundOnly <$true | $false>]
```

This example rewrites the email addresses of all messages entering and leaving the Exchange organization for the contoso.com domain. Outbound messages are rewritten so they appear to come from the fabrikam.com domain. Inbound messages sent to fabrikam.com email addresses are rewritten to contoso.com for delivery to the recipients (the *OutboundOnly* parameter is `$false` by default).

```
New-AddressRewriteEntry -Name "Contoso to Fabrikam" -InternalAddress contoso.com -ExternalAddress fabrikam.com
```

This example rewrites the email addresses of all messages leaving the Exchange organization for the sales.contoso.com subdomain. Outbound messages are rewritten so they appear to come from the contoso.com domain. Inbound messages sent to contoso.com email addresses aren't rewritten.

```
New-AddressRewriteEntry -Name "sales.contoso.com to contoso.com" -InternalAddress sales.contoso.com -
ExternalAddress contoso.com -OutboundOnly $true
```

Rewrite email addresses in multiple subdomains

To rewrite the email addresses in a domain and all subdomains, use the following syntax.

```
New-AddressRewriteEntry -Name "<Descriptive Name>" -InternalAddress *.<domain> -ExternalAddress <domain> -
OutboundOnly $true [-ExceptionList <domain1,domain2...>]
```

This example rewrites the email addresses of all messages leaving the Exchange organization for the contoso.com domain and all subdomains. Outbound messages are rewritten so they appear to come from the contoso.com domain. Inbound messages sent to contoso.com recipients can't be rewritten, because a wildcard is used in the *InternalAddress* parameter.

```
New-AddressRewriteEntry -Name "Rewrite all contoso.com subdomains" -InternalAddress *.contoso.com -
ExternalAddress contoso.com -OutboundOnly $true
```

This example is just like the previous example, except now messages sent from the legal.contoso.com and corp.contoso.com subdomains are never rewritten:

```
New-AddressRewriteEntry -Name "Rewrite all contoso.com subdomains except legal.contoso.com and
corp.contoso.com" -InternalAddress *.contoso.com -ExternalAddress contoso.com -OutboundOnly $true -
ExceptionList legal.contoso.com,corp.contoso.com
```

For more information, see New-AddressRewriteEntry.

How do you know this worked?

To verify that you have successfully created address rewrite entries, do the following:

1. Replace *<AddressRewriteEntryIdentity>* with the name of the address rewrite entry, and run the following

command to verify the property values:

```
Get-AddressRewriteEntry <AddressRewriteEntryIdentity> | Format-List
```

2. From a mailbox that's affected by the address rewrite entry, send a test message to an external mailbox. Verify the test message appears to originate from the rewritten email address.

3. Reply to the test message from the external mailbox. Verify the original mailbox receives the reply.

Use the Exchange Management Shell to modify address rewrite entries

The configuration options that are available when you modify an existing address rewrite entry are identical to the configuration options when you create a new address rewrite entry.

Modify an address rewrite entry for a single recipient

To modify an address rewrite entry that rewrites the email address of a single recipient, use the following syntax:

```
Set-AddressRewriteEntry <AddressRewriteEntryIdentity> [-Name "<Descriptive Name>"] [-InternalAddress <internal
email address>] [-ExternalAddress <external email address>] [-OutboundOnly <$true | $false>]
```

This example modifies the following properties of the address rewrite entry named "joe@contoso.com to support@nortwindtraders.com":

- Changes the external address to support@northwindtraders.net.

- Changes the name of the address rewrite entry to "joe@contoso.com to support@northwindtraders.net".

- Changes the value of *OutboundOnly* to `$true`. Note that this change requires you to configure support@northwindtraders.net as a proxy address on Joe's mailbox.

```
Set-AddressRewriteEntry "joe@contoso.com to support@nortwindtraders.com" -Name "joe@contoso.com to
support@northwindtraders.net" -ExternalAddress support@northwindtraders.net -OutboundOnly $true
```

Modify an address rewrite entry for a single domain or subdomain

To modify an address rewrite entry that rewrites the email addresses from a single domain or subdomain, use the following syntax.

```
Set-AddressRewriteEntry <AddressRewriteEntryIdentity> [-Name "<Descriptive Name>"] [-InternalAddress <domain
or subdomain>] [-ExternalAddress <domain>] [-OutboundOnly <$true | $false>]
```

This example changes the internal address value of the address rewrite entry named "Northwind Traders to Contoso".

```
Set-AddressRewriteEntry "Northwindtraders to Contoso" -InternalAddress northwindtraders.net
```

Modify an address rewrite entry for multiple subdomains

To modify an address rewrite entry that rewrites the email addresses in a domain and all subdomains, use the following syntax.

```
Set-AddressRewriteEntry <AddressRewriteEntryIdentity> [-Name "<Descriptive Name>"] [-InternalAddress *.
<domain>] [-ExternalAddress <domain>] [-ExceptionList <list of domains>]
```

To replace the existing exception list values of an address rewrite entry, use the following syntax:

```
Set-AddressRewriteEntry <AddressRewriteEntryIdentity> -ExceptionList <domain1,domain2,...>
```

This example replaces the existing exception list for the address rewrite entry named Contoso to Northwind Traders with the values marketing.contoso.com and legal.contoso.com:

```
Set-AddressRewriteEntry "Contoso to Northwind Traders" -ExceptionList sales.contoso.com,legal.contoso.com
```

To add or remove exception list values without affecting other exception list entries, use the following syntax:

```
Set-AddressRewriteEntry <AddressRewriteEntryIdentity> -ExceptionList @{Add="<domain1>","<domain2>"...;
Remove="<domain3>","<domain4>"...}
```

This example adds finanace.contoso.com and removes marketing.contoso.com from the exception list of the address rewrite entry named Contoso to Northwind Traders:

```
Set-AddressRewriteEntry "Contoso to Northwind Traders" -ExceptionList @{Add="finanace.contoso.com";
Remove="marketing.contoso.com"}
```

For more information, see Set-AddressRewriteEntry.

How do you know this worked?

To verify that you have successfully modified an address rewrite entry, do the following:

1. Replace *<AddressRewriteEntryIdentity>* with the name of the address rewrite entry, and run the following command to verify the property values:

   ```
   Get-AddressRewriteEntry <AddressRewriteEntryIdentity> | Format-List
   ```

2. From a mailbox that's affected by the address rewrite entry, send a test message to an external mailbox. Verify the test message appears to originate from the rewritten email address.

3. From the external mailbox, reply to the test message. Verify the original mailbox receives the reply.

Use the Exchange Management Shell to remove address rewrite entries

To remove a single address rewrite entry, use the following syntax:

```
Remove-AddressRewriteEntry <AddressRewriteEntryIdentity>
```

This example removes the address rewrite entry named "Contoso.com to Northwindtraders.com":

```
Remove-AddressRewriteEntry "Contoso.com to Northwindtraders.com"
```

To remove multiple address rewrite entries, use the following syntax:

```
Get-AddressRewriteEntry [<search criteria>] | Remove-AddressRewriteEntry [-WhatIf]
```

This example removes all address rewrite entries:

```
Get-AddressRewriteEntry | Remove-AddressRewriteEntry
```

This example simulates the removal of address rewrite entries that contain the text "to contoso.com" in the name. The *WhatIf* switch allows you to preview the result without committing any changes.

```
Get-AddressRewriteEntry "*to contoso.com" | Remove-AddressRewriteEntry -WhatIf
```

If you're satisfied with the result, run the command again without the *WhatIf* switch to remove the address rewrite entries.

```
Get-AddressRewriteEntry "*to contoso.com" | Remove-AddressRewriteEntry
```

For more information, see Remove-AddressRewriteEntry.

How do you know this worked?

To verify that you have successfully removed an address rewrite entry, do the following:

1. Run the command `Get-AddressRewriteEntry`, and verify that the address rewrite entries you removed aren't listed.

2. From a mailbox that was affected by the address rewrite entry, send a test message to an external mailbox. Verify the test message is no longer affected by the removed address rewrite entry.

3. From the external mailbox, reply to the test message. Verify the original mailbox receives the reply and that the message is unaffected by the removed address rewrite entry.

Import address rewrite entries on Edge Transport servers

8/3/2020 • 4 minutes to read • Edit Online

You can bulk-create or import address rewriting information into an Edge Transport server by using a comma-separated value (CSV) file. The following list describes common scenarios that require you to do this:

- You are replacing an address rewriting solution with an Edge Transport server.
- You enter into an agreement with a third-party solution provider that requires you to rewrite their email addresses.
- You acquire another organization, and you need to temporarily rewrite the email addresses in the acquired organization.

You can use a spreadsheet application like Microsoft Excel to create the CSV file. Format the file as described in this topic and save it as a .csv file.

The first row, or *header row*, of the CSV file lists the names of the parameters. Each parameter is separated by a comma. The required and optional parameters are described in the following table.

PARAMETER	REQUIRED OR OPTIONAL	DESCRIPTION
Name	Required	A unique, descriptive name for the address rewrites entry.
InternalAddress	Required	The address you want to change. You can use the following values: • A single email address (chris@contoso.com) • A single domain or subdomain (contoso.com or sales.contoso.com) • A domain and all subdomains (*.contoso.com)
ExternalAddress	Required	The final email address you want. You can use the following values: • A single email address if you specified a single email address for *InternalAddress* • A single domain or subdomain for all other values of *InternalAddress*
ExceptionList	Optional	Available only when you're rewriting email addresses in a domain and all subdomains (*.contoso.com). Specifies one or more subdomains you want to exclude from address rewriting. Enclose the value in double quotation marks, and separate multiple values by commas. For example, `"marketing.contoso.com"` or `"marketing.contoso.com,legal.contoso.com"` .

PARAMETER	REQUIRED OR OPTIONAL	DESCRIPTION
OutboundOnly	Optional	`False` means that addresses are written on inbound and outbound mail. `True` means that addresses are rewritten on outbound mail only, and you need to manually configure the rewritten email address as a proxy address on the affected recipients. The default value is `False`, but you need to set it to `True` if *InternalAddress* contains the wildcard character (*.contoso.com). The *OutboundOnly* parameter value in the CSV file is `True` or `False`, not `$True` or `$False`.

Each row under the header row represents an individual address rewrite entry. The values in each row need to be in the same order as the parameter names in the header row. Each value is separated by a comma.

What do you need to know before you begin?

- Estimated time to complete this task: 15 minutes
- Make sure you understand the ramifications of address rewriting. For example, the rewritten email address need to be unique in your Exchange organization, and you might need to configure proxy addresses on the affected recipients. For more information, see Address rewriting on Edge Transport servers and Address rewriting procedures on Edge Transport servers.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Address Rewriting agent" entry in the Mail flow permissions topic.
- If you have more than one Edge Transport server, we recommend that you use the procedures in this topic to import the address rewrite entries into a single Edge Transport server and then clone the configuration of that Edge Transport server to the other Edge Transport servers in your organization. For more information about how to clone an Edge Transport server, see Using Edge Transport Server Cloned Configuration.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Step 1: Create the CSV file

When you create the CSV file, consider the following items:

- If you specify values for optional parameters in the CSV file, every row must include a value in that column. If you want to create multiple address rewrite entries where some entries have optional parameters and some entries do not, you need to separate those address rewrite entries into different CSV files, and then import each CSV file separately.
- If the CSV file contains non-ASCII characters, be sure to save the CSV file with UTF-8 encoding or other Unicode encoding. Saving the CSV file with UTF-8 encoding or other Unicode encoding might be easier when the system locale of the computer matches the language that's used in the CSV file.

The following example shows how a CSV file can be populated with the optional *ExceptionList* and *OutboundOnly* parameters included:

```
Name,InternalAddress,ExternalAddress,ExceptionList,OutboundOnly
"Wingtip
UK",*.wingtiptoys.co.uk,tailspintoys.com,"legal.wingtiptoys.co.uk,finance.wingtiptoys.co.uk,support.wingtiptoys
.co.uk",True
"Wingtip
USA",*.wingtiptoys.com,tailspintoys.com,"legal.wingtiptoys.com,finance.wingtiptoys.com,support.wingtiptoys.com,
corp.wingtiptoys.com",True
"Wingtip
Canada",*.wingtiptoys.ca,tailspintoys.com,"legal.wingtiptoys.ca,finance.wingtiptoys.ca,support.wingtiptoys.ca",
True
```

Step 2: Import the CSV file

To import the CSV file, use the following syntax:

```
Import-Csv <FileNameAndPath> | ForEach {New-AddressRewriteEntry -Name $_.Name -InternalAddress
$_.InternalAddress -ExternalAddress $_.ExternalAddress -OutboundOnly ([Bool]::Parse($_.OutboundOnly)) -
ExceptionList $_.ExceptionList}
```

This example imports the address rewrite entries from C:\My Documents\ImportAddressRewriteEntries.csv.

```
Import-Csv "C:\My Documents\ImportAddressRewriteEntries.csv" | ForEach {New-AddressRewriteEntry -Name $_.Name -
InternalAddress $_.InternalAddress -ExternalAddress $_.ExternalAddress -OutboundOnly
([Bool]::Parse($_.OutboundOnly)) -ExceptionList $_.ExceptionList}
```

How do you know this step worked?

To verify that you have successfully imported address rewrite entries from a CSV file, use either of the following procedures:

- To see all address rewrite entries, run the following command:

  ```
  Get-AddressRewriteEntry
  ```

- To see details about a specific address rewrite entry, replace *<AddressRewriteIdentity>* with the name of the address rewrite entry, and run the following command:

  ```
  Get-AddressRewriteEntry "<AddressRewriteIdentity>" | Format-List
  ```

Client Access services

8/3/2020 • 3 minutes to read • Edit Online

In Exchange Server, the Client Access services on Mailbox servers provide authentication and proxy services for internal and external client connections. The Client Access services are stateless, so data isn't queued or stored in them. In Exchange Server, the Client Access services are part of the Mailbox server, so you can't configure a standalone Client Access server like you could in previous versions of Exchange. For more information, see Client access protocol architecture.

Client connectivity in Exchange 2016 and Exchange 2019 is similar to Exchange 2013, but different from Exchange 2010:

- Outlook clients use MAPI over HTTP or Outlook Anywhere (RPC over HTTP). In Exchange 2016 and Exchange 2019, MAPI over HTTP is enabled by default.
- Exchange 2016 and Exchange 2019 require fewer namespaces for site-resilient solutions than Exchange 2010. For more information, see Namespace Planning in Exchange 2016.

Client Access services functionality

The Client Access services in Exchange Server function much like a front door, admitting all client connection requests and routing them to the correct mailbox database. The Client Access services provide network security such as Transport Layer Security (TLS) encryption, and manage client connections through redirection and proxying. The Client Access services authenticate client connections and typically proxy the connection request to the Mailbox server that holds the active copy of the user's mailbox. In some cases, the Client Access services might redirect the request to the Client Access services on another Exchange server, either in a different location or on a more recent version of Exchange.

The Client Access services have the following features:

- **Stateless services**: In previous versions of Exchange, many of the Client Access protocols required session affinity. For example, Outlook Web App in Exchange 2010 required that all requests from a particular client be handled by a specific Client Access server within a load balanced array of Client Access servers. In Exchange 2016 and Exchange 2019, the Client Access services are stateless. In other words, because all processing for the mailbox happens in the backend services on the Mailbox server, it doesn't matter which instance of the Client Access service in an array of Client Access services receives each individual client request. This means that session affinity is no longer required at the load balancer level. This allows inbound connections to Client Access services to be balanced by using simple load balancing techniques such as DNS round-robin. It also allows hardware load balancing devices to support significantly more concurrent connections. For more information, see Load Balancing in Exchange 2016.
- **Connection pooling**: The Client Access services handle client authentication and send the **AuthN** data to the backend services on the Mailbox server. The account that's used by the Client Access services to connect to the backend services on Mailbox servers is a privileged account that's a member of the Exchange Servers group. This allows the Client Access services to pool connections to the backend services on Mailbox servers effectively. An array of Client Access services can handle millions of client connections from the Internet, but far fewer connections are used to proxy the requests to the backend services on Mailbox servers than in Exchange 2010. This improves processing efficiency and end-to-end latency.

NOTE

Don't confuse an array of Client Access services with an RPC Client Access Server array that was used for RPC over TCP client connections in Exchange 2010. In Exchange 2016 and Exchange 2019, an array of Client Access services simply indicates a group of load-balanced Client Access services on Exchange 2016 or Exchange 2019 servers.

Management tasks in the Client Access services

- **Digital certificates**: Although Exchange Server uses self-signed certificates to encrypt and authenticate connections between Exchange servers, you need to install and configure certificates to encrypt client connections. For more information, see Digital certificates and encryption in Exchange Server.
- **Kerberos authentication for load-balanced Client Access services**: For more information, see Configure Kerberos authentication for load-balanced Client Access services.

Exchange admin center in Exchange Server

8/3/2020 • 11 minutes to read • Edit Online

The Exchange admin center (EAC) is the web-based management console in Exchange Server that's optimized for on-premises, online, and hybrid Exchange deployments. The EAC was introduced in Exchange Server 2013, and replaces the Exchange Management Console (EMC) and the Exchange Control Panel (ECP), which were the two management interfaces in Exchange Server 2010.

Looking for the Exchange Online version of this topic? See Exchange admin center in Exchange Online.

Looking for the Exchange Online Protection version of this topic? See Exchange admin center in Exchange Online Protection.

Accessing the EAC

The URL of the EAC is controlled by the Internet Information Services (IIS) virtual directory named ECP in the Client Access (frontend) services on the Mailbox server. Yes, the virtual directory is named ECP, not EAC.

- **Internal URL**: By default, this value contains the fully-qualified domain name (FQDN) of the Exchange server in the format `https://<ServerFQDN>/ecp`. For example, `https://mailbox01.contoso.com/ecp`. To access the EAC in a web browser on the Exchange server itself, you can use the value `https://localhost/ecp`.
- **External URL**: By default, this value is unconfigured. Before you can connect to the EAC from the Internet, you need to configure the following settings:
  - The external URL value on the ECP virtual directory. For more information, see Step 4: Configure external URLs in Configure mail flow and client access on Exchange servers.
  - A corresponding record in your public DNS.
  - A TLS certificate that contains or matches the host name entry. Very likely, this will be a subject alternative name (SAN) certificate or a wildcard certificate, because most of the client services are all available under the same website on the Exchange server. For more information, see Certificate requirements for Exchange services.

    After you configure the settings, a common external URL value for the EAC would resemble `https://mail.contoso.com/ecp`.

    Note: External users who connect to Outlook on the web (formerly known as Outlook Web App) also need access to the EAC to access their own **Options** page. You can disable external administrator access to the EAC while still allowing users to access their **Options** page in Outlook on the web. For more information, see Turn off access to the Exchange admin center.

The easiest way to find the internal and external URL values for the EAC (without using **Servers** > **Virtual directories** in the EAC itself) is by using the **Get-EcpVirtualDirectory** cmdlet in the Exchange Management Shell. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.

These examples show you how to find the internal and external URL values for the EAC virtual directories in your organization:

- To find the values on all Exchange servers in your organization, run the following command:

```
Get-EcpVirtualDirectory | Format-List Server,Name,*Url
```

- To find the values on the server named Mailbox01, run the following command:

```
Get-EcpVirtualDirectory | Format-List Name,*Url
```

- To find the value for the virtual directory named "ecp (Default Web Site)" on the server named Mailbox01, run the following command.

```
Get-EcpVirtualDirectory -Identity "Mailbox01\ecp (Default Web Site)" | Format-List *Url
```

For more information, see Get-EcpVirtualDirectory.

In Exchange 2016, if you're in a coexistence environment with Exchange 2010, the location of your mailbox controls the default behavior for opening the EAC or ECP:

- If your mailbox is located on the Exchange 2010 Mailbox server, you get the Exchange 2010 ECP by default. You can access the EAC by adding the Exchange version to the URL (which is 15 for both Exchange 2013 and Exchange 2016). For example, to access the EAC through the Client Access (frontend) services on the Mailbox server named Mailbox01, use the following URL: `https://Mailbox01/ecp/?ExchClientVer=15` .
- If your mailbox is located on an Exchange 2016 Mailbox server, and you want to access the ECP on the Exchange 2010 Client Access server named CAS01, use the following URL: `https://CAS01/ecp/?ExchClientVer=14` .

Common user interface elements in the EAC

The section describes the user interface elements that are common across the EAC.

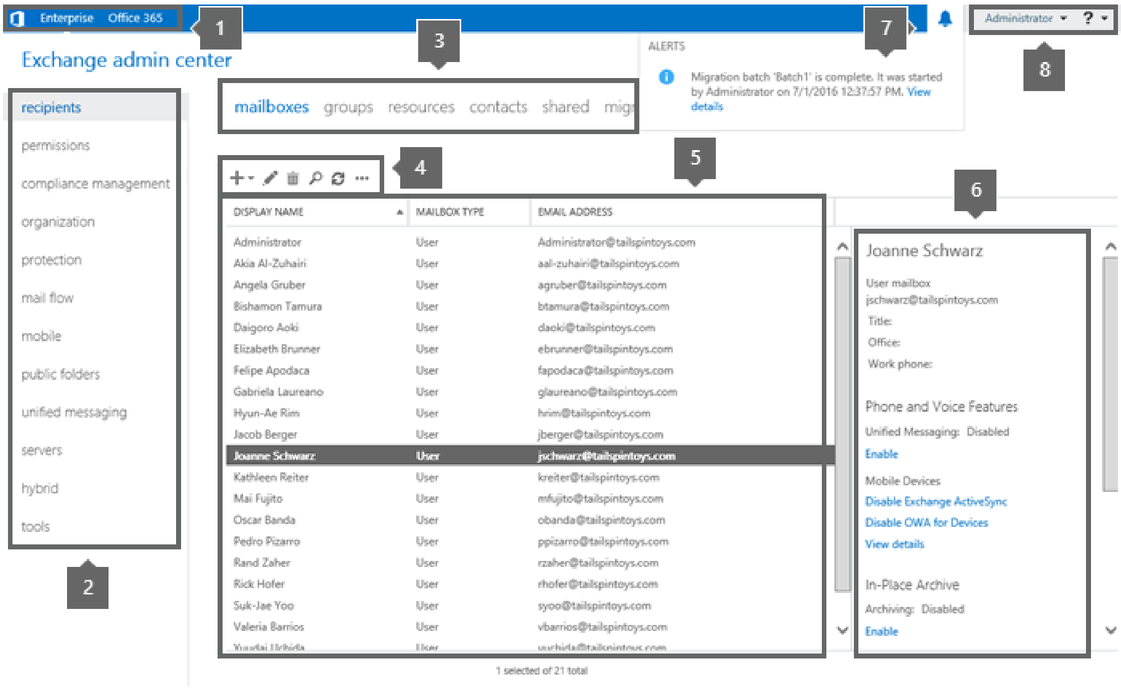

ALERTS

Migration batch 'Batch1' is complete. It was started by Administrator on 7/1/2016 12:37:57 PM. View details

DISPLAY NAME	MAILBOX TYPE	EMAIL ADDRESS
Administrator	User	Administrator@tailspintoys.com
Akia Al-Zuhairi	User	aal-zuhairi@tailspintoys.com
Angela Gruber	User	agruber@tailspintoys.com
Bishamon Tamura	User	btamura@tailspintoys.com
Daigoro Aoki	User	daoki@tailspintoys.com
Elizabeth Brunner	User	ebrunner@tailspintoys.com
Felipe Apodaca	User	fapodaca@tailspintoys.com
Gabriela Laureano	User	glaureano@tailspintoys.com
Hyun-Ae Rim	User	hrim@tailspintoys.com
Jacob Berger	User	jberger@tailspintoys.com
Joanne Schwarz	User	jschwarz@tailspintoys.com
Kathleen Reiter	User	kreiter@tailspintoys.com
Mai Fujito	User	mfujito@tailspintoys.com
Oscar Banda	User	obanda@tailspintoys.com
Pedro Pizarro	User	ppizarro@tailspintoys.com
Rand Zaher	User	rzaher@tailspintoys.com
Rick Hofer	User	rhofer@tailspintoys.com
Suk-Jae Yoo	User	syoo@tailspintoys.com
Valeria Barrios	User	vbarrios@tailspintoys.com
Yuudai Uchida	User	yuchida@tailspintoys.com

1 selected of 21 total

Joanne Schwarz

User mailbox
jschwarz@tailspintoys.com
Title:
Office:
Work phone:

Phone and Voice Features

Unified Messaging: Disabled
Enable

Mobile Devices
Disable Exchange ActiveSync
Disable OWA for Devices
View details

In-Place Archive

Archiving: Disabled
Enable

1: Cross-premises navigation

The cross-premises navigation allows you to easily switch between your Exchange Online and on-premises Exchange deployments. If you don't have an Exchange Online organization, the **Office 365** link takes you to a page that compares plans and pricing for Microsoft 365 and Office 365 services.

2: Feature pane

The feature pane is the first level of navigation for most of the tasks that you'll perform in the EAC, and is organized by the following feature areas:

- **Recipients**: Manage mailboxes, groups, resource mailboxes (room and equipment mailboxes), contacts, shared mailboxes, and mailbox migrations and moves. For more information, see the following topics:
  - Create user mailboxes in Exchange Server and Manage user mailboxes
  - Manage distribution groups and Manage dynamic distribution groups
  - Create and manage room mailboxes
  - Manage mail contacts and Manage mail users
  - Create shared mailboxes in the Exchange admin center
- **Permissions**: Manage role-based access control (RBAC) administrator roles, user roles, and Outlook on the web policies. For more information, see the following topics.
  - Manage role groups , Manage role group members, and Manage role assignment policies.
  - View or configure Outlook on the web mailbox policy properties
- **Compliance management**: This is where you'll manage In-Place eDiscovery, In-Place Hold, auditing (mailbox audit logging and administrator audit logging), data loss prevention (DLP), retention policies, retention tags, and journal rules. For more information, see the following topics:
  - In-Place eDiscovery in Exchange Server and In-Place Hold and Litigation Hold in Exchange Server
  - Mailbox audit logging in Exchange Server and Administrator audit logging in Exchange Server
  - Data loss prevention in Exchange Server
  - Retention policies and Retention tags.
  - Journaling in Exchange Server
- **Organization**: Manage federated sharing, Outlook Apps, and address lists. For more information, see the following topics:
  - Sharing
  - Install or remove add-ins for Outlook for your Exchange 2013 organization
  - Address lists in Exchange Server
- **Protection**: Manage antimalware protection for your organization. For more information, see Antimalware protection in Exchange Server.
- **Mail flow**: Manage mail flow rules (also known as transport rules), delivery reports, accepted domains, remote domains, email address policies, Receive connectors, and Send connectors. For more information, see the following topics:
  - Mail flow rules in Exchange Server
  - Track messages with delivery reports
  - Address lists in Exchange Server
  - Accepted domains in Exchange Server

    - Remote Domains
    - Email address policies in Exchange Server
    - Receive connectors
    - Send connectors

- **Mobile**: Manage the mobile devices that you allow to connect to your organization. You can manage mobile device access and mobile device mailbox policies. For more information, see the following topics:
    - Mobile devices
    - Mobile device mailbox policies

- **Public folders**: Manage public folders and public folder mailboxes. For more information, see Public folders.

- **Unified Messaging**: Manage UM dial plans and UM IP gateways. (UM is not available in Exchange 2019.) For more information, see the following topics:
    - UM Dial Plans
    - UM IP Gateways

- **Servers**: View and manage server-specific settings, databases, database availability groups (DAGs), virtual directories, and certificates. For more information, see the following topics:
    - POP3 and IMAP4 in Exchange Server
    - Configure the Startup Mode on a Client Access Server and Configure the Startup Mode on a Mailbox Server
    - Message retry, resubmit, and expiration intervals
    - Configure message tracking , Configure connectivity logging in Exchange Server, and Protocol logging
    - Manage Outlook Anywhere
    - Manage mailbox database copies
    - Manage database availability groups
    - Virtual Directory Management
    - Certificate procedures in Exchange Server

- **Hybrid**: Set up and configure a Hybrid organization.

3: Tabs

The **Setup** tab allows you to run the Hybrid Configuration Wizard or modify the settings of your existing hybrid deployment.

4: Toolbar

When you click most tabs, you'll see a toolbar. The toolbar has icons that perform specific actions. The following table describes the most common icons and their actions. To see the action that's associated with an icon (the icon's title), simply hover over the icon.

ICON	NAME	ACTION
+	Add, New	Create a new object. Some of these icons have an associated down arrow you can click to show additional objects you can create. For example, in **Recipients** > **Mailboxes**, clicking the down arrow displays **User mailbox** and **Linked mailbox** as additional options.
✏	Edit	Edit an object.
🗑	Delete	Delete an object. Some delete icons have a down arrow you can click to show additional options.
🔍	Search	Open a search box so you can enter text for an object that you want to find you want to find in a long list of objects.
🔄	Refresh	Refresh the list view.
•••	More options	View more actions you can perform for that tab's objects. For example, in **Recipients** > **Mailboxes** clicking this icon shows the following options: **Disable**, **Add/Remove columns**, **Export data to a CSV file**, **Connect a mailbox**, and **Advanced search**.
↑ ↓	Up arrow and down arrow	Move an object up or down in the list, when the order is important. For example, in **Mail flow** > **Email address policies** click the up arrow to move the policy higher in the list, which increases the priority of the policy by specifying which policy is applied first. You can also use these arrows to navigate the public folder hierarchy and to move rules up or down in the list view.
⎘	Copy	Copy an object so you can make changes to it without changing the original object. For example, in **Permissions** > **Admin roles**, select a role from the list view, and then click this icon to create a new role group based on an existing one.

ICON	NAME	ACTION
—	Remove	Remove an item from a list. For example, in the **Public Folder Permissions** dialog box, you can remove users from the list of users allowed to access the public folder by selecting the user and clicking this icon.

5: List view

Tabs that contain many objects display those objects in a list view. The viewable limit in the EAC list view is approximately 20,000 objects. Paging is included so you can skip to the results that you want to see. In the **Recipients** list view, you can also configure page size and export the data to a CSV file.

6: Details pane

When you select an object from the list view, more information about that object is displayed in the details pane. For some object types, the details pane includes quick management tasks. For example, if you navigate to **Recipients** > **Mailboxes** and select a mailbox from the list view, the details pane (among other options) displays an option to enable or disable the archive for that mailbox.

Some object types also allow you to bulk edit multiple objects in the details pane. You can select multiple objects in the list view by selecting an object, holding the Shift key, and selecting an object farther down in the list, or by holding down the CTRL key as you select each object. If bulk edit is available for the object types that you selected, you'll see the available options in the details pane. For example, at **Recipients** > **Mailboxes**, when you select multiple mailboxes of the same type, the title of the details pane changes to **Bulk Edit**, and you can update contact and organization information, custom attributes, mailbox quotas, Outlook on the web settings, and more.

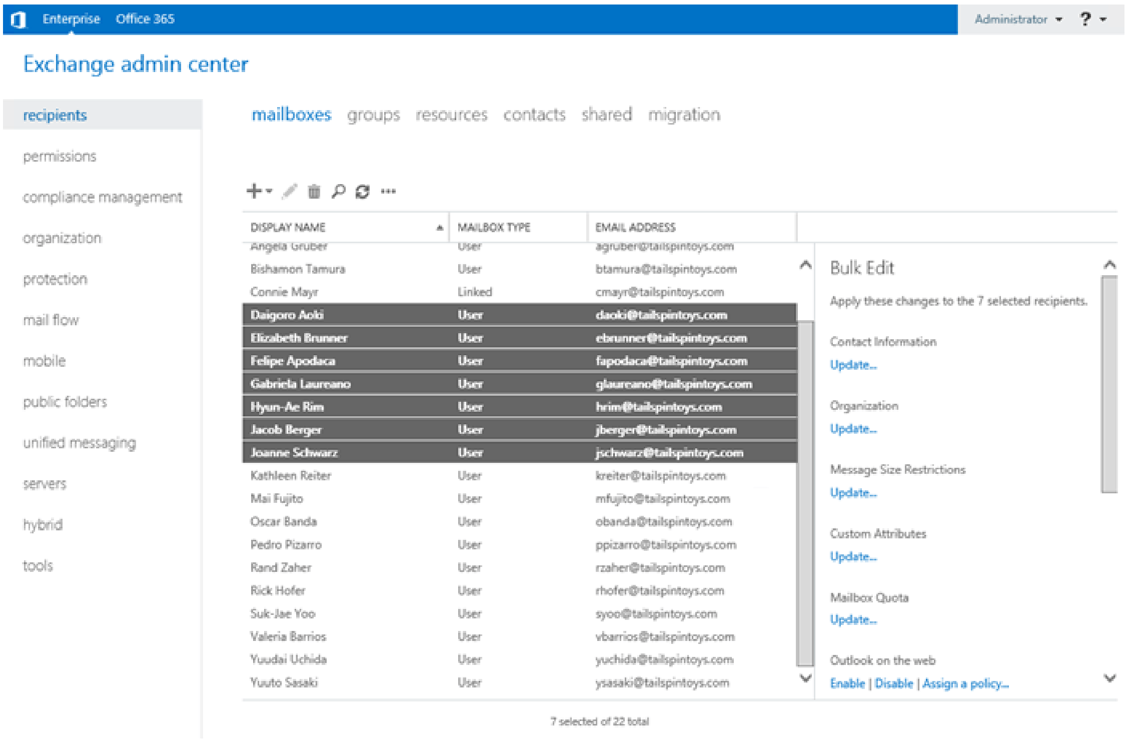

DISPLAY NAME	MAILBOX TYPE	EMAIL ADDRESS
Angela Gruber	User	agruber@tailspintoys.com
Bishamon Tamura	User	btamura@tailspintoys.com
Connie Mayr	Linked	cmayr@tailspintoys.com
Daigoro Aoki	User	daoki@tailspintoys.com
Elizabeth Brunner	User	ebrunner@tailspintoys.com
Felipe Apodaca	User	fapodaca@tailspintoys.com
Gabriela Laureano	User	glaureano@tailspintoys.com
Hyun-Ae Rim	User	hrim@tailspintoys.com
Jacob Berger	User	jberger@tailspintoys.com
Joanne Schwarz	User	jschwarz@tailspintoys.com
Kathleen Reiter	User	kreiter@tailspintoys.com
Mai Fujito	User	mfujito@tailspintoys.com
Oscar Banda	User	obanda@tailspintoys.com
Pedro Pizarro	User	ppizarro@tailspintoys.com
Rand Zaher	User	rzaher@tailspintoys.com
Rick Hofer	User	rhofer@tailspintoys.com
Suk-Jae Yoo	User	syoo@tailspintoys.com
Valeria Barrios	User	vbarrios@tailspintoys.com
Yuudai Uchida	User	yuchida@tailspintoys.com
Yuuto Sasaki	User	ysasaki@tailspintoys.com

7 selected of 22 total

Bulk Edit

Apply these changes to the 7 selected recipients.

Contact Information
Update...

Organization
Update...

Message Size Restrictions
Update...

Custom Attributes
Update...

Mailbox Quota
Update...

Outlook on the web
Enable | Disable | Assign a policy...

7: Notifications

The EAC includes a notification viewer that displays information about:

- Expiring and expired certificates.
- The status of mailbox moves and migrations (also known as Mailbox Replication Service tasks or

MRS tasks). You can also use the notification viewer to opt-in to receive email notifications about these tasks.

- Exporting mailbox content to .pst files.

To show or hide the notification viewer, click the icon (🔔).

Notifications are alerts that are sent to the arbitration mailbox named `FederatedEmail.4c1f4d8b-8179-4148-93bf-00a95fa1e042`. The EAC checks this mailbox for alerts every 30 seconds. Notifications remain in the arbitration mailbox until they are removed by the component that sent them, or until they expire (they should be removed by the Managed Folder Assistant after 30 days).

You can also use the **Get-Notification** cmdlet in the Exchange Management Shell to view more details about notifications, and the **Set-Notification** cmdlet to request notification emails for future alerts.

8: Me tile and Help

The *Me tile* allows you to sign out of the EAC and sign in as a different user by clicking on the drop-down menu that's next to your account name.

Click the help icon (**?**) to view the help content for the tab that you're currently on. If you click on the drop-down menu that's next to the help icon, you can perform the following additional actions:

- **Disable Help bubble**: The Help bubble displays contextual help for fields when you create or edit objects in the EAC. From here, you can globally turn off or turn on the Help bubble for all fields in the EAC.

- **Performance console**: The Performance console displays many counters that relate to the performance of the EAC.

- **Copyright** and **Privacy**: Click these links to read the copyright and privacy information for Exchange Server.

Supported browsers

The levels of support for operating system and browser combinations that you can use to access the EAC are described in the following tables.

Notes:

- The levels of support for the EAC are:

  - **Supported**: All functionality and features are supported and have been fully tested.

  - **Unsupported**: The browser and operating system combination isn't supported, **or** hasn't been tested. For more information about supported versions of Internet Explorer on Windows, see Internet Explorer Support Announcement.

  - **n/a**: The browser and operating system combination isn't possible. For example, an older browser on a newer operating system, or vice-versa.

- Operating system and browser combinations that aren't listed are unsupported. This includes iOS and Android.

- Third-party plug-ins might cause issues with the EAC for supported browsers.

Client operating systems

WEB BROWSER	WINDOWS 7	WINDOWS 8.1	WINDOWS 10	MAC OS X	LINUX
Internet Explorer 9	Unsupported	n/a	n/a	n/a	n/a
Internet Explorer 10	Unsupported	n/a	n/a	n/a	n/a
Internet Explorer 11	Supported	Supported	Supported	n/a	n/a
Microsoft Edge	n/a	n/a	Supported	n/a	n/a
Mozilla Firefox latest version or one previous	Supported	Supported	Supported	Supported	Supported
Apple Safari 6 or later versions	n/a	n/a	n/a	Supported	n/a
Google Chrome latest version or one previous	Supported	Supported	Supported	Supported	Supported

Windows Server operating systems

WEB BROWSER	WINDOWS SERVER 2008 R2	WINDOWS SERVER 2012	WINDOWS SERVER 2012 R2	WINDOWS SERVER 2016
Internet Explorer 9	Unsupported	n/a	n/a	n/a
Internet Explorer 10	Unsupported	Supported	n/a	n/a
Internet Explorer 11	Supported	n/a	Supported	Supported

Turn off access to the Exchange admin center

8/3/2020 • 8 minutes to read • Edit Online

The Exchange admin center (EAC) is the primary management interface for Exchange 2013 or later. For more information, see Exchange admin center in Exchange Server. By default, access to the EAC isn't restricted, and access to Outlook on the web (formally known as Outlook Web App) on an on an Internet-facing Exchange server also gives access to the EAC. You still need valid credentials to sign in to the EAC, but organizations may want to restrict access to the EAC for client connections from the Internet.

In Exchange Server 2019, you can use Client Access Rules to block client access to the EAC. For more information, see Client Access Rules in Exchange Server.

The EAC virtual directory is named ECP, and is managed by the *- **ECPVirtualDirectory** cmdlets. When you set the *AdminEnabled* parameter to the value `$false` on the EAC virtual directory, you disable access to the EAC for internal and external client connections, without affecting access to the **Settings** > **Options** page in Outlook on the web.

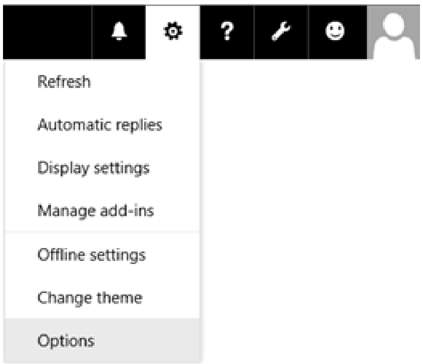

But, this configuration introduces a new problem: access to the EAC is completely disabled on the server, even for administrators on the internal network. To fix this issue, you have two choices:

- Configure a second Exchange server that's only accessible from the internal network to handle internal EAC connections.
- On the existing Exchange server, create a new Internet Information Services (IIS) web site with new virtual directories for the EAC and Outlook on the web that's only accessible from the internal network.

  Note: You need to configure the EAC **and** Outlook on the web in the new web site, because the EAC requires the Outlook on the web authentication module from the same web site.

What do you need to know before you begin?

- Estimated time to complete each procedure: 5 minutes.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Exchange admin center connectivity" entry in the Exchange infrastructure and PowerShell permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Step 1: Use the Exchange Management Shell to disable access to the EAC

Remember, this step disables access to the EAC on the server for internal and external connections, but still allows users to access their own **Settings** > **Options** page in Outlook on the web.

To disable access to the EAC on an Exchange server, use the following syntax:

```
Set-ECPVirtualDirectory -Identity "<Server>\ecp (Default Web Site)" -AdminEnabled $false
```

This example turns disables access to the EAC on the server named MBX01.

```
Set-ECPVirtualDirectory -Identity "MBX01\ecp (Default Web Site)" -AdminEnabled $false
```

How do you know this step worked?

To verify that you've disabled access to the EAC on the server, replace *<Server>* with the name of your Exchange server, and run the following command to verify the value of the **AdminEnabled** property:

```
 -Identity "MBX01\ecp (Default Web Site)" | Format-List AdminEnabled
```

When you open https://<servername>/ecp or from the internal network, your own **Settings** > **Options** page in Outlook on the web opens instead of the EAC.

Step 2: Give access to the EAC on the internal network

Choose either of the following options.

Option 1: Configure a second Exchange server that's only accessible from the internal network

The default value of the **AdminEnabled** property is `True` on the default EAC virtual directory. To confirm this value on the second server, replace *<Server>* with the name of the server, and run the following command:

```
Get-ECPVirtualDirectory -Identity "<Server>\ecp (Default Web Site)" | Format-List AdminEnabled
```

If the value is `False`, replace *<Server>* with the name of the server, and run the following command:

```
Set-ECPVirtualDirectory -Identity "<Server>\ecp (Default Web Site)" -AdminEnabled $true
```

Option 2: Create a new web site on the existing Exchange server, and configure the EAC and Outlook on the web in the new web site for the internal network

The required steps are:

1. Add a second IP address to the Exchange server.

2. Create a new web site in IIS that uses the second IP address, and assign file and folder permissions.

3. Copy the contents of the default web sites to the new web site.

4. Create new EAC and Outlook on the web virtual directories for the new web site.

5. Restart IIS for the changes to take effect.

> **IMPORTANT**
>
> When you install an Exchange Server Cumulative Update (CU), the CU won't update files in the new web site and virtual directories. After you apply the CU, you need to completely remove the new web site, virtual directories, and content in the folders and then re-create the new web site, virtual directories, and content in the folders.

Step 2a: Add a second IP address to the Exchange server

You can add a second network adapter and assign the IP address to the second network adapter, or you can assign a second IP address to the existing network adapter.

The steps to assign a second IP address to the existing network adapter are described below.

1. Open the properties of the network adapter. For example:

   a. From a Command Prompt window, the Exchange Management Shell, or the **Run** dialog, run `ncpa.cpl`.

   b. Right-click on the network adapter, and then choose **Properties**.

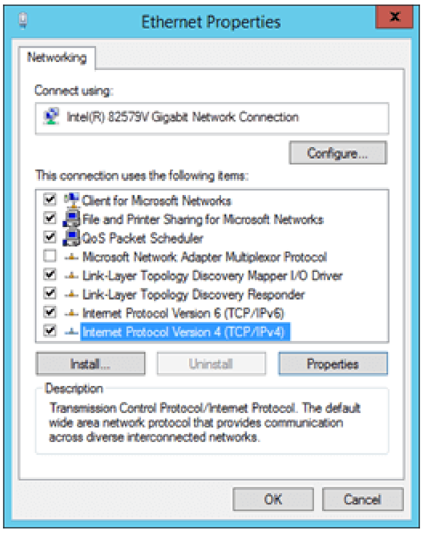

Connect using:
Intel(R) 82579V Gigabit Network Connection

This connection uses the following items:
- ☑ Client for Microsoft Networks
- ☑ File and Printer Sharing for Microsoft Networks
- ☑ QoS Packet Scheduler
- ☐ Microsoft Network Adapter Multiplexor Protocol
- ☑ Link-Layer Topology Discovery Mapper I/O Driver
- ☑ Link-Layer Topology Discovery Responder
- ☑ Internet Protocol Version 6 (TCP/IPv6)
- ☑ Internet Protocol Version 4 (TCP/IPv4)

Description
Transmission Control Protocol/Internet Protocol. The default wide area network protocol that provides communication across diverse interconnected networks.

2. In the properties of the network adapter, select **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.

3. In the **Internet Protocol Version 4 (TCP/IPv4) Properties** window that opens, on the **General** tab, click **Advanced**.

4. In the **Advanced TCP/IP Settings** window that opens, on the **IP Settings** tab, in the **IP addresses** section, click **Add** and enter the IP address.

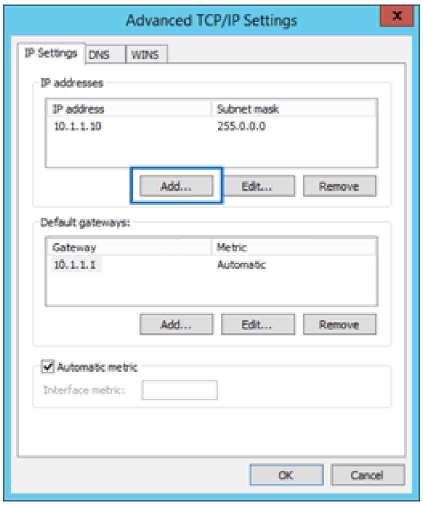

IP addresses

IP address	Subnet mask
10.1.1.10	255.0.0.0

Default gateways:

Gateway	Metric
10.1.1.1	Automatic

☑ Automatic metric

Interface metric:

Note: If you add a second network adapter, in the **Advanced TCP/IP Settings** window, on the **DNS** tab, un-check **Register this connection's address in DNS**.

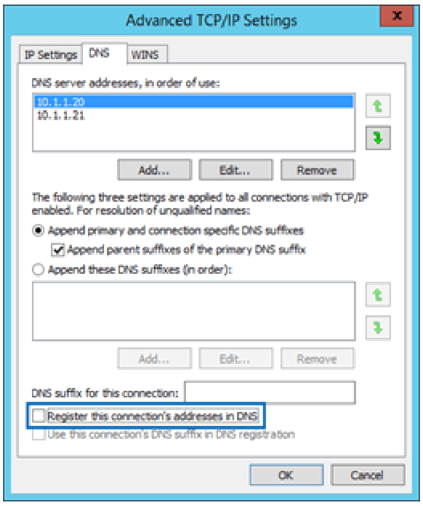

DNS server addresses, in order of use:

10.1.1.20
10.1.1.21

The following three settings are applied to all connections with TCP/IP enabled. For resolution of unqualified names:

● Append primary and connection specific DNS suffixes

☑ Append parent suffixes of the primary DNS suffix

○ Append these DNS suffixes (in order):

DNS suffix for this connection:

□ Register this connection's addresses in DNS

□ Use this connection's DNS suffix in DNS registration

Step 2b: Create a new web site in IIS that uses the second IP address, and assign file and folder permissions

1. Open IIS Manager on the Exchange server. An easy way to do this in Windows Server 2012 or later is to press Windows key + Q, type inetmgr, and select **Internet Information Services (IIS) Manager** in the results.

2. In the **Connections** pane, expand the server, select **Sites**, and in the **Actions** pane, click **Add Website**.

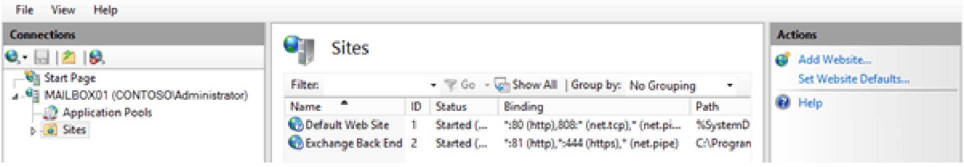

Name	ID	Status	Binding	Path
Default Web Site	1	Started (...	*:80 (http),808:* (net.tcp),* (net.pi...	%SystemD
Exchange Back End	2	Started (...	*:81 (http),*:444 (https),* (net.pipe)	C:\Progran

3. In the **Add Website** window that appears, configure the following settings:

   - **Site name**: `EAC_Secondary`
   - **Physical path**: `C:\inetpub\EAC_Secondary`
   - **Binding**
     - **Type**: https
     - **IP address**: Select the second IP address that you added in the previous step.
     - **Port**: 443
   - **SSL certificate**: Choose the certificate that you want to use (for example, the default Exchange

certificate named Microsoft Exchange).

When you're finished, click **OK**.

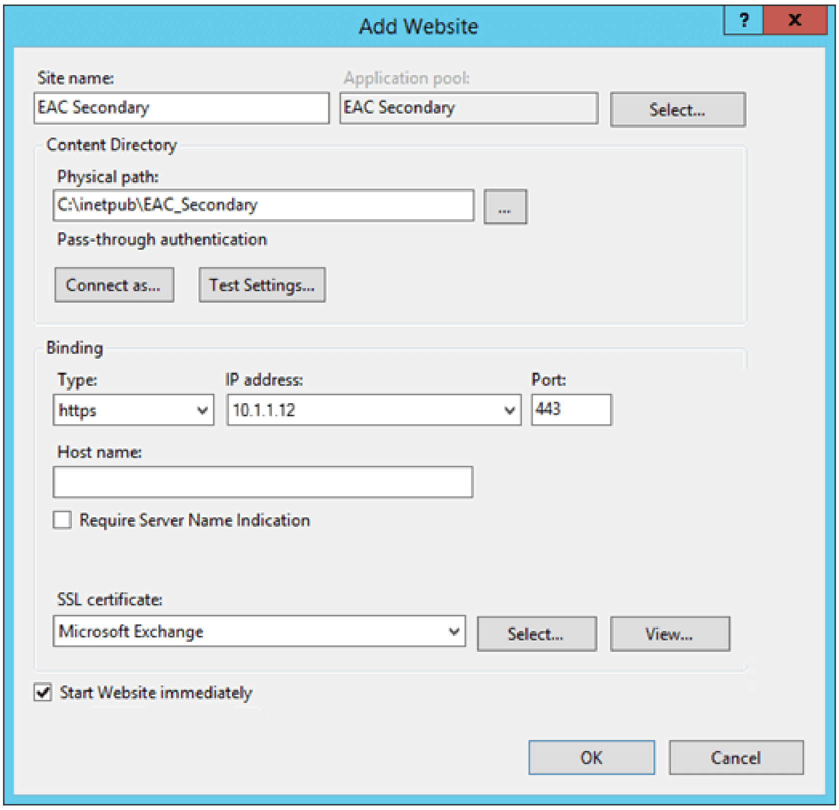

Add Website

Site name: EAC Secondary

Application pool: EAC Secondary

Content Directory

Physical path: C:\inetpub\EAC_Secondary

Pass-through authentication

Binding

Type: https

IP address: 10.1.1.12

Port: 443

Host name:

□ Require Server Name Indication

SSL certificate: Microsoft Exchange

☑ Start Website immediately

4. Create `ecp` and `owa` folders in `C:\inetpub\EAC_Secondary`.

a. In IIS Manager, select the `EAC_Secondary` web site, and in the **Actions** pane, click **Explore**.

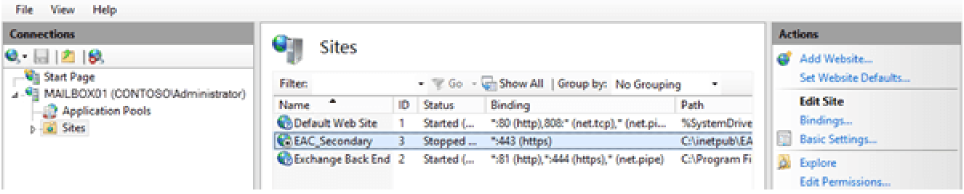

Sites

Filter: Go Show All | Group by: No Grouping

Name	ID	Status	Binding	Path
Default Web Site	1	Started (...	*:80 (http),808:* (net.tcp),* (net.pi...	%SystemDrive
EAC_Secondary	3	Stopped ...	*:443 (https)	C:\inetpub\EA
Exchange Back End	2	Started (...	*:81 (http),*:444 (https),* (net.pipe)	C:\Program Fi

b. In the File Explorer window that opens, create the following folders in `C:\inetpub\EAC_Secondary`:

- `ecp`
- `owa`

When you're finished, close File Explorer.

5. Assign **Read & Execute** permissions to the local security group named **IIS_IUSRS** on the `C:\inetpub\EAC_Secondary` folder.

a. In IIS Manger, select the `EAC_Secondary` web site, and in the **Actions** pane, click **Edit Permissions**.

b. In the **EAC_Secondary Properties** window that opens, click the **Security** tab, and then click **Edit**.

c. In the **Permissions for EAC_Secondary** window that opens, click **Add**.

d. In the **Select Users, Computers, Service Accounts or Groups** window that opens, perform the following steps:

i. Click **Locations**, and in the **Locations** dialog box that opens, select the local server, and then click **OK**.

ii. In the **Enter the object names to select** field, type IIS_IUSRS, click **Check Names**, and then click **OK**.

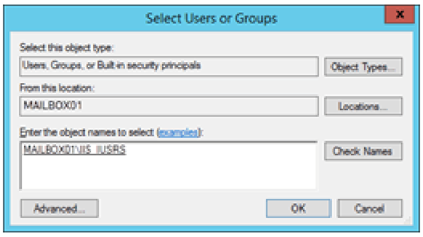

Select Users or Groups

Select this object type:
Users, Groups, or Built-in security principals

From this location:
MAILBOX01

Enter the object names to select (examples):
MAILBOX01\IIS_IUSRS

e. Back on the **Permissions for EAC_Secondary** window, select **IIS_IUSRS**, and in the **Allow** column, select **Read & Execute** (which automatically selects the **List Folder Contents** and **Read** permissions), and then click **OK** twice.

Step 2c: Copy the contents of the default web sites to the new web site

- Copy all files and folders from the Default Web Site (`C:\inetpub\wwwroot`) to `C:\inetpub\EAC_Secondary` . You can skip the following files that can't be copied:
    - `MacCertification.asmx`
    - `MobileDeviceCertification.asmx`
    - `decomission.asmx`
    - `editissuancelicense.asmx`
- Copy all files and folders from `%ExchangeInstallPath%FrontEnd\HttpProxy\ecp` to `C:\inetpub\EAC_Secondary\ecp` .
- Copy all files and folders from `%ExchangeInstallPath%FrontEnd\HttpProxy\owa` to `C:\inetpub\EAC_Secondary\owa` .

Step 2d: Use the Exchange Management Shell to create new EAC and Outlook on the web virtual directories for the new web site

To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.

Replace *<Server>* with the name of your server, and run the following commands to create the new EAC and Outlook on the web virtual directories for the new web site.

```
New-EcpVirtualDirectory -Server <Server> -Role ClientAccess -WebSiteName EAC_Secondary -Path
"C:\inetpub\EAC_Secondary\ecp"
```

```
New-OwaVirtualDirectory -Server <Server> -Role ClientAccess -WebSiteName EAC_Secondary -Path
"C:\inetpub\EAC_Secondary\owa"
```

Step 2e: Restart IIS

1. In IIS Manager, in the **Connections** pane, select the server.

2. In the **Actions** pane, click **Restart**.

Note: To restart IIS from the command line, open an elevated command prompt (a Command Prompt window that you opened by selecting **Run as administrator**) and run the following commands:

```
net stop w3svc /y
```

```
net start w3svc
```

How do you know this task worked?

To verify that you have successfully disabled access to the EAC on an Exchange server, perform the following steps:

1. Test your organization's internal and external URL for Outlook on the web. For example, if the external URL is https://mail.contoso.com/owa, and the internal URL is https://mbx01.contoso.com/owa use the following procedures to verify your configuration:

    - Verify that internal and external users can open their mailboxes by using Outlook on the web, including the **Settings** > **Options** page.

    - Verify that https://mail.contoso.com/ecp and https://mbx01.contoso.com/ecp return either of the following results:

        - **404 - website not found**

        - The user is redirected to their **Settings** > **Options** page in Outlook on the web.

2. Verify that administrators can access the EAC on the internal network based on your configuration selection:

    - **Second Exchange server**: If the second Exchange server is named MBX02, verify that https://mbx02.contoso.com/ecp opens the EAC.

    - **New EAC web site on the existing Exchange server**: If the IP address of the new EAC web site is 10.1.1.12, verify that https://10.1.1.12/ecp opens the EAC.

Autodiscover service in Exchange Server

8/3/2020 • 11 minutes to read • Edit Online

The Autodiscover service minimizes user configuration and deployment steps by providing clients access to Exchange features. For Exchange Web Services (EWS) clients, Autodiscover is typically used to find the EWS endpoint URL. However, Autodiscover can also provide information to configure clients that use other protocols. Autodiscover works for client applications that are inside or outside firewalls and in resource forest and multiple forest scenarios.

Exchange 2016 introduced changes to services that were previously handled by the multiple servers. The Mailbox server now provides Client Access services, so you can't configure a standalone Client Access server like you could in previous versions of Exchange. Autodiscover service in Exchange 2016 and Exchange 2019 is possible because:

- Exchange creates a virtual directory named `autodiscover` under the default web site in Internet Information Services (IIS).
- Active Directory stores and provides authoritative URLs for domain-joined computers.
- Client Access services on Mailbox servers provide authentication and proxy services for internal and external client connections.
- Outlook configures services with only the username and password.

NOTE

If you are a user looking for help with connecting your Outlook client to your Exchange server, see Outlook email setup.

Autodiscover services and Active Directory

Exchange stores in Active Directory the configuration of Exchange servers in the organization as well as information about your users' mailboxes. Before you install Exchange Server, you need to prepare your Active Directory forest and its domains. If you aren't familiar with Exchange forests or domains, see Step 3: Prepare Active Directory domains.

Exchange automatically creates at installation the virtual directory `autodiscover` in IIS, the frontend Client Access services web site that clients connect to. This allows Outlook to discover the Exchange mailbox settings so that users don't have to deal with manually configuring advanced settings.

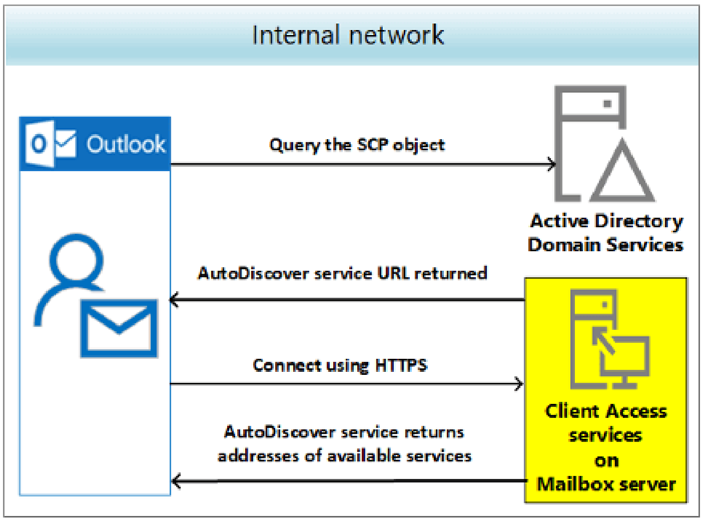
Internal network
Outlook
Query the SCP object
Active Directory Domain Services
AutoDiscover service URL returned
Connect using HTTPS
AutoDiscover service returns addresses of available services
Client Access services on Mailbox server

The SCP object is also created in Active Directory at the same time as the Autodiscover service virtual directory. The SCP stores and provides authoritative URLs of the Autodiscover service for domain-joined computers.

You need to update the SCP object to point to the Exchange server. This is necessary because Exchange servers provide additional Autodiscover information to clients to improve the discovery process. You can use the **Set-ClientAccessService** cmdlet to update the SCP object. For more information, see Set-ClientAccessService.

> **IMPORTANT**
>
> You need to be assigned permissions before you can run the **Set-ClientAccessService** cmdlet. To find the permissions required to run any cmdlet or parameter in your organization, see Find the permissions required to run any Exchange cmdlet.

Autodiscover makes it easy to retrieve the information that you need to connect to mailboxes on Exchange servers. SCP objects locate those Autodiscover servers or endpoints appropriate for the user you're retrieving settings for. And SCP objects in AD DS provide an easy way for domain-joined clients to look up Autodiscover servers.

Exchange publishes two types of SCP objects for the Autodiscover service:

- **SCP pointers**: Contains information that points to specific LDAP servers that should be used to locate Autodiscover SCP objects for the user's domain. SCP pointers are stamped with the following GUID: 67661d7F-8FC4-4fa7-BFAC-E1D7794C1F68.

- **SCP URLs**: Contains URLs for Autodiscover endpoints. SCP URLs are stamped with the following GUID: 77378F46-2C66-4aa9-A6A6-3E7A48B19596

The SCP object contains the authoritative list of Autodiscover service URLs for the forest. To learn more about locating Autodiscover service endpoints, see Generate a list of Autodiscover endpoints.

Client connectivity in Exchange 2016 and Exchange 2019 is like Exchange 2013 and differs from Exchange 2010. In Exchange 2016 and 2019, MAPI over HTTP is enabled by default, when previously Outlook clients used Outlook Anywhere (RPC over HTTP). Exchange 2016 and 2019 require fewer name spaces for site-resilient solutions than Exchange 2010, reducing to two from the previously required seven namespaces. To read more about namespace and Exchange Server, see the blog Namespace Planning in Exchange 2016.

Depending on whether you configured the Autodiscover service on a separate site, the Autodiscover service URL will be either of the following values, where `//<SMTP-address-domain>` is the primary SMTP domain address:

- `https://<SMTP-address-domain>/autodiscover/autodiscover.xml`

- `https://autodiscover.<smtp-address-domain>/autodiscover/autodiscover.xml`

For example, if the user's email address is tony@contoso.com, the primary SMTP domain address is contoso.com.

Client applications use the Autodiscover service when the application starts for the first time. For example when an Exchange Web Services (EWS) application starts for the first time, the application configures itself using the Autodiscover service. For a user's computer joined to the contoso.com domain and in the Longview regional Active Directory site, the application generates the list of these Autodiscover service endpoints:

ENDPOINT	GENERATED BY
https://longview.contoso.com/autodiscover/autodiscover.xml	SCP results
https://email.contoso.com/autodiscover/autodiscover.xml	SCP results
https://newark.contoso.com/autodiscover/autodiscover.xml	SCP results

ENDPOINT	GENERATED BY
https://contoso.com/autodiscover/autodiscover.exc	Derived from email address
https://autodiscover.contoso.com/autodiscover/autodiscover	Derived from email address

For more information about SCP objects, see Publishing with Service Connection Points.

Autodiscover in DNS

Exchange introduced namespace requirements for Autodiscover in Exchange 2010 and certificates required several of them. In a server resilience scenario, all of these elements were required:

- Primary datacenter IP namespace
- Secondary datacenter IP namespace
- Primary Outlook Web App failback namespace
- Secondary Outlook Web App failback namespace
- Transport namespace (for SMTP)
- Primary datacenter RPC Client Access namespace
- Secondary datacenter RPC Client Access namespace

Server resiliency scenarios have been improved, reducing the five namespaces to two. This is because Exchange no longer needs the RPC Client Access namespaces and Client Access services proxy requests to the Mailbox server that is hosting the active Mailbox database. A Mailbox server in one Active Directory site can proxy a session to a another Active Directory site's Mailbox server.

What this means is that unique namespaces are no longer required for *each* datacenter. For example, instead of mail.contoso.com and mail2.contoso.com, you only need a single namespace, mail.contoso.com, for the datacenter pair. Additionally, failback namespaces are no longer needed in Database Availability Groups (DAG) activation scenarios.

Autodiscover is simple to set up for your domain because it only requires that you create a CNAME resource record in your external (public) DNS. CNAME records let you hide the implementation details of your network from the clients that connect to it. Used internally in your network, CNAME records allow users to use the simpler URI mail.domain.com instead of host.examplemachinename.domain.com.

A CNAME or canonical name record is the DNS equivalent to a Windows shortcut or an Apple Mac alias. A CNAME record is an alias for an Address (A) record that maps an IP address to the target server. If, for example, your domain is contoso.com, you create a CNAME record for autodiscover.contoso.com. The name in the CNAME record must match a name in a certificate. CNAME records work only for hostnames. CNAMEs work externally, but they don't replace the URL in the browser bar. When the certificate is checked against the URL, you get a failure with a warning, but you can still access the service.

A typical CNAME record looks like this:

- Name: autodiscover
- TTL: 3600
- RR Type: CNAME
- Target: The externally accessible FQDN of the Mailbox server (for example, mail.contoso.com)

In this example, autodiscover.contoso.com resolves to mail.contoso.com. For more information, see Step 4: Configure external URLs in Configure mail flow and client access on Exchange servers.

We recommend that you create an Autodiscover CNAME record for every domain on your account, including domain aliases and accepted domains. You need to create either a CNAME or SRV record where your domain is hosted. Only then you can synchronize your offline address book, show free/busy information and enable the Out of office feature in Outlook.

Service (SRV) resource records let you specify the location of the servers for a specific service, protocol, and DNS domain. For example, if you have two Web servers in your domain, you can create SRV resource records indicating which hosts serve as Web servers. Resolvers can then retrieve all the SRV resource records for the Web servers.

A typical SRV record looks like this:

- Service: _autodiscover
- Protocol: ._tcp
- Port Number: 443
- Host: mail.contoso.com
- Priority: 0
- Weight: 0

In this example, the Outlook server namespace is mail.contoso.com.

Read more about CNAME and SRV records in the Exchange team blog, Namespace planning in Exchange 2016.

Autodiscover services in Outlook

With only the user credentials, the Outlook client can authenticate to Active Directory and search for the Autodiscover SCP objects. After the client obtains and enumerates the instances of the Autodiscover service, the client connects to the Client Access (frontend) services on the first Mailbox server in the enumerated list. The client then collects the profile information in the form of XML data that's needed to connect to the user's mailbox and available Exchange features.

> **NOTE**
>
> Depending on your DNS provider's requirements, you may need to add the fully qualified domain name (FQDN) as your hostname. In that case, if your domain is contoso.com, then your hostname would be autodiscover.contoso.com, not autodiscover.com.

You need to set up a special DNS record for your domain name that points to the server providing Autodiscover services so that Exchange accounts function correctly in Outlook. For external access, or using DNS, the client locates the Autodiscover service on the Internet by using the primary SMTP domain address from the user's email address.

The Autodiscover service uses one of these four methods to configure the email client. The first two work for small, single SMTP namespace organizations. The last two serve multiple-SMTP namespaces.

- Connect to: https://contoso.com/AutoDiscover/AutoDiscover.xml
- Connect to: https://autodiscover.contoso.com/AutoDiscover/AutoDiscover.xml
- Autodiscover redirect URL for redirection: http://autodiscover.contoso.com/autodiscover/autodiscover.xml
- Search for DNS SRV record

Some of the hostnames and URLs can be configured by using the Exchange admin center (EAC) and the Exchange Management Shell, while others require that you use PowerShell. You can learn more about that in Configure mail flow and client access.

Through the Autodiscover service, Outlook finds a new connection point made up of the user's mailbox. That is, Autodiscover uses the identification made up of a GUID, plus @, and the domain portion of the user's primary SMTP address. The Autodiscover service returns the following information to the client:

- User's display name
- Separate connection settings for internal and external connectivity
- Location of the user's mailbox (the Mailbox server that currently holds the active copy of the mailbox)
- URLs for various Outlook features that govern functionality such as free/busy information, Unified Messaging (UM) in Exchange 2016 (but not in Exchange 2019), and the offline address book (OAB)
- Outlook Anywhere server settings

You'll need to make sure that you have configured the correct external URLs for the virtual directories of the following services. The examples in the table that follows show values required for the contoso.com email domain.

SERVICE	EXCHANGE MANAGEMENT SHELL	MODIFIES
Offline Address Book	`Get-OabVirtualDirectory | Set-OabVirtualDirectory -ExternalURL https://mail.companycontoso.com/oab`	OAB virtual directories used in IIS
Exchange Web Sevices	`Get-WebServicesVirtualDirectory | Set-WebServicesVirtualDirectory -ExternalURL https://mail.companycontoso.com/ews/exchange.asmx`	Exchange Web Services virtual directories in IIS
Outlook Anywhere (RPC over HTTP)	`Get-OutlookAnywhere | Set-OutlookAnywhere -ExternalHostname mail.contoso.com -ExternalClientsRequireSsl $true`	Outlook Anywhere virtual directories in IIS

Messaging Application Programming Interface (MAPI) over HTTP (Exchange 2013 SP1 or later)	`Get-MapiVirtualDirectory | Set-MapiVirtualDirectory -ExternalURL https://mail.companycontoso.com/mapi`
	`Set-OrganizationConfig -MapiHttpEnabled $true`
		MAPI virtual directories in IIS

Click the Service name in the preceding table for more information about how to obtain or reconfigure these URLs.

When a user's Exchange information changes, Outlook uses the Autodiscover service to automatically reconfigure the user's profile. For example, if a user's mailbox is moved. or the client can't connect to the user's mailbox or to available Exchange features, Outlook will contact the Autodiscover service and automatically update the user's profile to include the information that's required to connect to the mailbox and Exchange features.

Other clients

Autodiscover service the preferred method to locate all services in Skype for Business Server 2015 . When a connection is successful, the Autodiscover service returns all the Web Services URLs for the user's home pool, including the Mobility Service (known as Mcx by the virtual directory created for the service in IIS), Lync Web App and Web scheduler URLs. However, both the internal Mobility Service URL and the external Mobility Service URL is associated with the external Web Services FQDN. Therefore, regardless of whether a mobile device is internal or external to the network, the device always connects to the Mobility Service externally through reverse proxy. The Autodiscover service also returns references to Internal/UCWA, External/UCWA and UCWA. These entries refer to the Unified Communications Web API (UCWA) web component.

Configure Autodiscover services

Autodiscover works for client applications inside and outside firewalls and in resource forest and multiple forest scenarios. For EWS clients, Autodiscover is typically used to find the EWS endpoint URL, but Autodiscover can also provide information to configure clients that use other protocols.

When you install Exchange Server, a self-signed certificate that's created and signed by the Exchange server is automatically installed on the server. However, you can also create additional self-signed certificates that you can use for other services.

Creating a certificate request is the first step in installing a new certificate on an Exchange server to configure Transport Layer Security (TLS) encryption for one or more Exchange services. You use a certificate request (also known as a certificate signing request or CSR) to obtain a certificate from a certification authority (CA). For more information, see the following topics:

- Digital certificates and encryption in Exchange Server
- Create an Exchange Server certificate request for a certification authority

> **NOTE**
>
> You can confirm your Autodiscover service by using the Microsoft Remote Connectivity Analyzer. When the connectivity is successful, also select and run the Outlook Connectivity test. If that fails, you may need to configure the external URLs in Exchange. The results from the Microsoft Remote Connectivity Analyzer should explain why connectivity failed. Generally, a connectivity failure means that you don't have the correct external URLs configured for the virtual directories of the various Outlook services.

Manage Autodiscover services

In deployments where clients connect to multiple Exchange servers, the Autodiscover SCP object is created for the (frontend) Client Access services on each Mailbox server. The SCP object contains the ServiceBindingInfo attribute with the FQDN of the Exchange server that the client connects to in the form of `https://<ExchangeServer>/autodiscover/autodiscover.xml` (for example, `https://cas01/autodiscover/autodiscover.xml)`.

You can run the Exchange ActiveSync Autodiscover and Outlook Autodiscover tests in the Microsoft Remote Connectivity Analyzer. If the user is using a local wireless network to connect to Exchange Online, the user should run both tests to make sure that the local network allows for connections to the ActiveSync endpoints.

You can get help for planning and deploying Autodiscover services as part of your Exchange deployment in Planning and deployment for Exchange Server.

Load balancing in Exchange Server

8/3/2020 • 12 minutes to read • Edit Online

Load balancing in Exchange 2016 and later build on the Microsoft high availability and network resiliency platform delivered in Exchange 2013. When this is combined with the availability of third-party load balancing solutions (both hardware and software), there are multiple options for implementing load balancing in your Exchange organization.

Exchange architecture changes introduced in Exchange 2013 brought about the Mailbox server and Client Access server roles. Compare this to Exchange 2010, where Client Access, Mailbox, Hub Transport, and Unified Messages ran on separate servers.

Using minimal server roles, Exchange 2016 and 2019 deliver:

- Simplified deployment with the Mailbox server running Client Access services and Edge Transport server roles.
- Mail flow managed in the transport pipeline, which is the collection of services, connections, queues, and components that route messages to the Transport service categoriser on the Mailbox server.
- High availability by deploying load balancers to distribute client traffic.

The HTTP protocol standard introduced with Exchange 2013 means that session affinity is no longer required in Exchange 2016 and Exchange 2019. Session affinity allows a persistent connection for messaging-enabled services so that a user doesn't have to reenter their username and password multiple times.

Previously, Exchange 2007 and Exchange 2010 supported RPC over HTTP for Outlook Anywhere. Exchange 2013 introduced MAPI over HTTP, although it wasn't enabled by default. It's now enabled by default in Exchange 2016 and Exchange 2019.

With the HTTP protocol in use, all native clients connect using HTTP and HTTPs in Exchange Server. This standard protocol removes the need for affinity, which was previously required to avoid a new prompting for user credentials whenever load balancing redirected the connection to a different server.

Server roles in Exchange Server

The reduced number of server roles for Exchange 2016 and Exchange 2019 simplifies Exchange implementation and hardware requirements. The number of server roles in Exchange 2016 and 2019 shrinks from seven to two: the Mailbox server and the Edge Transport server. The Mailbox server role includes Client Access services, while the Edge Transport server provides secure mail flow in Exchange 2016 and Exchange 2019, just as it did in earlier versions of Exchange.

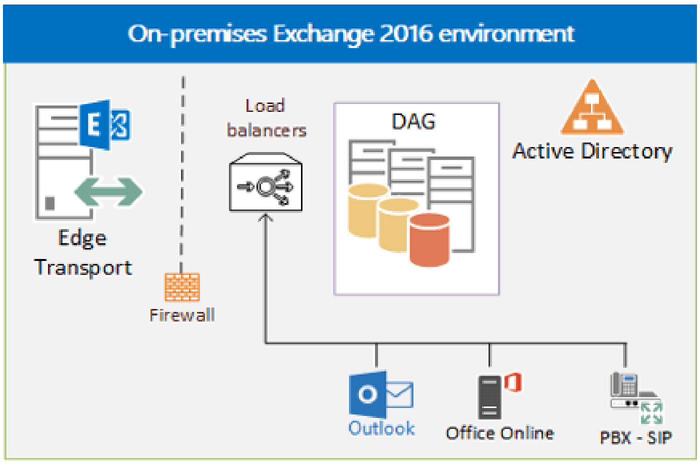

In Exchange 2013, the Client Access server role made sure that when a user attempted to access their mailbox, the server proxied the request back to the Mailbox server actively serving the user's mailbox. This meant that services such as Outlook on the web (previously known as Outlook Web App) were rendered for the user on the Mailbox itself, removing any need for affinity.

The same functionality remains in Exchange 2016 and Exchange 2019. If two Mailbox servers host different mailboxes, they can proxy traffic for each other when necessary. The Mailbox server that hosts the active copy of the mailbox serves the user accessing it, even if the user connects to a different Mailbox server.

Read more about the server role changes in Exchange Server in the topic, Exchange Server architecture.

SERVER ROLE	SERVICES
Mailbox server	Uses EdgeSync to manage one-way replication of receipt and configuration info from Active Directory to the AD LDS instance on the Edge Transport server. Copies only information needed to let Edge Transport perform antispam and enable end-to-end mail flow.
Edge Transport	Manages all inbound and outbound Internet mail flow using: • mail relay • smart hosting • agents that provide additional antispam service • agents that apply transport to control mail flow Not a member of the Active Directory forest

Although not required, the Edge Transport server sits in the perimeter network , just as in earlier Exchange versions, to provide secure inbound and outbound mail flow for your Exchange organization.

Read more about the transport service in the topic, Understanding the Transport service on Edge Transport servers.

Protocols in Exchange Server

Beginning with Exchange 2016, all native Exchange clients use the HTTP protocol to connect to a designated service, with HTTP cookies provided to the user at log in which are encrypted using the Client Access services SSL certificate. A logged in user can resume the session on a different Mailbox server running Client Access services without reauthenticating. Servers using the same SSL certificate can decrypt the client authentication cookie.

HTTP makes possible the use of service or application health checks in your Exchange network. Depending on your load balancer solution, you can implement health probes to check different components of your system.

The effect of HTTP-only access for clients is that load balancing is simpler, too. If you wanted, you could use DNS to load balance your Exchange traffic. You would simply provide the client with the IP address of every Mailbox server, and the HTTP client would handle the chores. If an Exchange server fails, the protocol attempts to connect to another server. However, there are drawbacks to load balancing to DNS, discussed in the following section *Load balancing options in Exchange Server*.

Read more about HTTP and Exchange Server in the topic MAPI over HTTP in Exchange Server.

Load balancing options in Exchange Server

In the example shown here, multiple servers configured in a database availability group (DAG) host the Mailbox servers running Client Access services. This provides high availability with a small Exchange server footprint. The client connects to the load balancer rather than directly to the Exchange servers. There is no requirement for load balancer pairs, however we recommend deploying in clusters to improve network resilience.

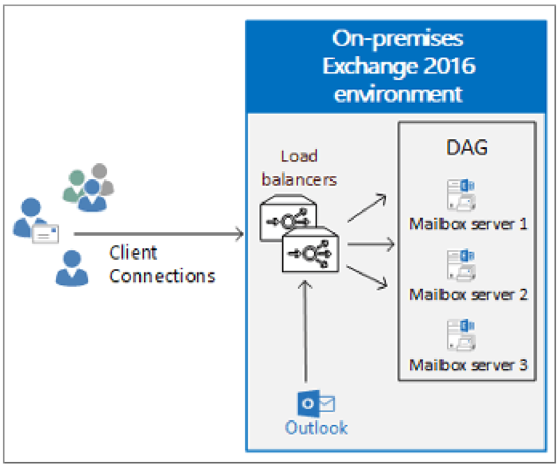

Be aware that DAGs use Microsoft Clustering Services. These services can't be enabled on the same server as Windows Network Load Balancing (NLB). Accordingly, Windows NLB is not an option when using DAGs. There are third-party software and virtual appliance solutions in this case.

Using DNS is the simplest option for load balancing your Exchange traffic . With DNS load balancing, you only have to provide your clients with the IP address of every Mailbox server. After that, DNS round robin distributes that traffic to your Mailbox servers. The HTTP client is smart enough to connect to another server should one Exchange server fail completely.

Simplicity comes at a price, however. In this case, DNS round robin isn't truly load-balancing the traffic, because there isn't a way programmatically to make sure that each server gets a fair share of the traffic. Also, there is no service level monitoring so that when a single service fails, clients are not automatically redirected to an available service. For example, if Outlook on the web is in failure mode, the clients see an error page.

DNS load balancing requires more external IP addresses when you publish externally. That means that each individual Exchange server in your organization would require an external IP address.

There are more elegant solutions to load balancing your traffic, such as hardware that uses Transport Layer 4 or Application Layer 7 to help distribute client traffic. Load balancers monitor each Exchange client-facing service, and in the event of service failure, load balancers can direct traffic to another server and take the problem server offline. Additionally, some level of load distribution makes sure that no single Mailbox server is proxying the majority of client access.

Load balancing services can use Layer 4 or Layer 7, or a combination, to manage traffic. There are benefits and drawbacks to each solution.

- Layer 4 load balancers work at the Transport layer to direct traffic without examining the contents.

  Because they don't examine the traffic contents, Layer 4 load balancers save time in transit. However, this comes with trade-offs. Layer 4 load balancers know only the IP address, protocol, and TCP port. Knowing only a single IP address, the load balancer can monitor only a single service.

  Layer 4 load balancing benefits include:

  - Requires fewer resources (no content examination).
  - Distributes traffic at the Transport layer.

    The risk with a Layer 4 solution is that if a service fails but the server is still available, clients can connect to the failed service. This means that a resilient Layer 4 implementation requires multiple IP addresses configured with separate HTTP namespaces per service, for example, owa.contoso.com, eas.contoso.com, mapi.contoso.com, which allows for service-level monitoring.

- Layer 7 load balancers work at the Application layer and can inspect the traffic content and direct it accordingly.

  Layer 7 load balancers forego the raw performance benefits of Layer 4 load balancing for the simplicity of a single namespace, for example, mail.contoso.com, and per-service monitoring. Layer 7 load balancers understand the HTTP path, such as /owa or /Microsoft-Server-ActiveSync, or /mapi, and can direct traffic to working servers based on monitoring data.

  Layer 7 load balancing benefits include:

  - Needs only a single IP address.
  - Inspects content and can direct traffic.
  - Provides notification of failed service that can be taken offline.
  - Handles load balancer SSL termination.
  - Distributes traffic at the application layer and understands the destination URL.

SSL should terminate at the load balancer as this offers a centralized place to correct SSL attacks.

The ports that need to be load balanced include some, such as those for IMAP4 or POP3, that may not even be used in your Exchange organization.

TCP PORT	ROLES	USES
25	Mailbox	Inbound SMTP
587	Mailbox	Inbound SMTP for clients
110	Mailbox	POP3 clients
143	Mailbox	IMAP4 clients
443	Mailbox	HTTPS (Outlook on the web, AutoDiscover, web services, ActiveSync, MAPI over HTTP, RPC over HTTP, OAB, EAC)
993	Mailbox	Secure IMAP4 clients
995	Mailbox	Secure POP3 clients

Load balancing deployment scenarios in Exchange Server

Exchange 2016 introduced significant flexibility for your namespace and load balancing architecture. With many options for deploying load balancing in your Exchange organization, from simple DNS to sophisticated third-party Layer 4 and Layer 7 solution, we recommend that you review them all in light of your organization's needs.

The following scenarios come with benefits and limitations, and understanding each is key to implementing the solution that best fits your Exchange organization:

- **Scenario A** Single namespace, no session affinity: Layer 4 or Layer 7
- **Scenario B** Single namespace, no session affinity: Layer 7
- **Scenario C** Single namespace with session affinity, Layer 7
- **Scenario D** Multiple namespaces and no session affinity, Layer 4

Scenario A Single namespace, no session affinity: Layer 4 or Layer 7

In this Layer 4 scenario, a single namespace, mail.contoso.com, is deployed for the HTTP protocol clients. The load balancer doesn't maintain session affinity. Because this is a layer 4 solution, the load balancer is configured to check the health of only a single virtual directory as it cannot distinguish Outlook on the web requests from RPC requests.

From the perspective of the load balancer in this example, health is per-server and not per-protocol for the designated namespace. Administrators will have to choose which virtual directory they want to target for the health probe; we recommend that you choose a heavily used virtual directory. For example, if the majority of your users utilize Outlook on the web, then choose the Outlook on the web virtual directory in the health probe.

As long as the Outlook on the web health probe response is healthy, the load balancer keeps the destination Mailbox server in the load balancing pool. However, if the Outlook on the web health probe fails for any reason, then the load balancer removes the destination Mailbox server from the load balancing pool for all requests associated with that namespace. This means that if the health probe fails, all client requests for that namespace are directed to another server, regardless of protocol.

Scenario B Single namespace, no session affinity: Layer 7

In this Layer 7 scenario, a single namespace, mail.contoso.com, is deployed for all the HTTP protocol clients. The load balancer doesn't maintain session affinity. Since the load balancer is configured for Layer 7, there is SSL termination and the load balancer knows the destination URL.

We recommend this configuration for Exchange 2016 and Exchange 2019. The load balancer is configured to check the health of the destination Mailbox servers in the load balancing pool, and a health probe is configured on each virtual directory.

For example, as long as the Outlook on the web health probe response is healthy, the load balancer will keep the destination Mailbox server in the Outlook on the web load balancing pool. However, if the Outlook on the web health probe fails for any reason, then the load balancer removes the target Mailbox server from the load balancing pool for Outlook on the web requests. In this example, health is per-protocol, which means that if the health probe fails, only the affected client protocol is directed to another server.

Scenario C Single namespace with session affinity, Layer 7

In this Layer 7 scenario, a single namespace, mail.contoso.com, is deployed for all the HTTP protocol clients. Because the load balancer is configured for Layer 7, there is SSL termination and the load balancer knows the destination URL. The load balancer is also configured to check the health of the target Mailbox servers in the load balancing pool. The health probe is configured on each virtual directory.

However, enabling session affinity decreases capacity and utilization. This is because the more involved affinity options, cookie-based load balancing or Secure Sockets Layer (SSL) session-ID, require more processing and resources. We recommend that you check with your vendor on how session affinity affects your load balancing scalability.

Just as in the previous scenario, as long as the Outlook on the web health probe response is healthy, the load balancer keeps the destination Mailbox server in the Outlook on the web load balancing pool. However, if the Outlook on the web health probe fails for any reason, then the load balancer removes the target Mailbox server from the load balancing pool for Outlook on the web requests. Here, health is per-protocol, which means that if the health probe fails, only the affected client protocol is directed to another server.

Scenario D Multiple namespaces and no session affinity

This last scenario with multiple namespaces and no session affinity offers per-protocol health checks and Layer 4 power. A unique namespace is deployed for each HTTP protocol client. For example, you would configure the HTTP protocol clients as mail.contoso.com, mapi.contoso.com, and eas.contoso.com.

This scenario provides per-protocol health checking while not requiring complex load-balancing logic. The load balancer uses Layer 4 and is not configured to maintain session affinity. The load balancer configuration checks the health of the destination Mailbox servers in the load balancing pool. In this setting, the health probes are configured to target the health of each virtual directory, as each virtual directory has a unique namespace. Because it's configured for Layer 4, the load balancer doesn't know the URL is being accessed, yet the result is as if it does know. Since health is per-protocol, if the health probe fails, only the affected client protocol is directed to another server.

Load balancing and managed availability in Exchange Server

Monitoring the available servers and services is key to high availability networks. Since some load balancing solutions have no knowledge of the target URL or the content of the request, this can introduce complexities for Exchange health probes.

Exchange 2016 and Exchange 2019 include a built-in monitoring solution, known as Managed Availability. Managed availability, also known as Active Monitoring or Local Active Monitoring, is the integration of built-in monitoring and recovery actions with the Exchange high availability platform.

Managed Availability includes an offline responder. When the offline responder is invoked, the affected protocol (or server) is removed from service.

To ensure that load balancers do not route traffic to a Mailbox server that Managed Availability has marked as offline, load balancer health probes must be configured to check `<virtualdirectory>/healthcheck.htm` , for example, https://mail.contoso.com/owa/healthcheck.htm.

Read more about managed availability in Managed availability.

Configure Kerberos authentication for load-balanced Client Access services

8/3/2020 • 13 minutes to read • Edit Online

In order for you to use Kerberos authentication with load-balanced Mailbox servers running Client Access services, you have to complete the configuration steps described in this article.

Create the alternate service account credential in Active Directory Domain Services

All Exchange servers that run Client Access services that share the same namespaces and URLs must use the same *alternate service account credential* or (ASA credential). In general, it's sufficient to have a single account for a forest for each version of Exchange.

> **IMPORTANT**
>
> Exchange 2010 and Exchange 2016 can't share the same ASA credential. If your ASA credential was created for Exchange 2010, you have to create a new one for Exchange 2016.
>
> While CNAME records are supported for shared namespaces, Microsoft recommends using A records. This ensures that the client correctly issues a Kerberos ticket request based on the shared name, and not the server FQDN.

> **NOTE**
>
> Group Managed Service Accounts (gMSA) are not supported in on-premises Exchange Server environments and thus cannot be used in this scenario.

When you set up the ASA credential, keep these guidelines in mind:

- **Account type**: We recommend that you create a computer account instead of a user account. A computer account doesn't allow interactive logon and may have simpler security policies than a user account. If you create a computer account, the password doesn't expire, but we recommend you update the password periodically anyway. You can use local group policy to specify a maximum age for the computer account and scripts to periodically delete computer accounts that do not meet current policies. Your local security policy also determines when you have to change the password. Although we recommend you use a computer account, you can create a user account.

- **Account name**: There are no requirements for the name of the account. You can use any name that conforms to your naming scheme.

- **Account group**: The account you use for the ASA credential doesn't need special security privileges. If you're using a computer account then the account needs only to be a member of the Domain Computers security group. If you're using a user account then the account needs only to be a member of the Domain Users security group.

- **Account password**: The password you provide when you create the account will be used. So when you create the account, you should use a complex password and ensure that the password conforms to your organization's password requirements.

To create the ASA credential as a computer account

1. On a domain-joined computer, run Windows PowerShell or the Exchange Management Shell.

   Use the **Import-Module** cmdlet to import the Active Directory module.

   ```
   Import-Module ActiveDirectory
   ```

2. Use the **New-ADComputer** cmdlet to create a new Active Directory computer account using this cmdlet syntax:

   ```
   New-ADComputer [-Name] <string> [-AccountPassword <SecureString>] [-AllowReversiblePasswordEncryption
   <System.Nullable[boolean]>] [-Description <string>] [-Enabled <System.Nullable[bool]>]
   ```

   Example:

   ```
   New-ADComputer -Name EXCH2016ASA -AccountPassword (Read-Host 'Enter password' -AsSecureString) -
   Description 'Alternate Service Account credentials for Exchange' -Enabled:$True -SamAccountName
   EXCH2016ASA
   ```

   Where *EXCH2016ASA* is the name of the account, the description *Alternate Service Account credentials for Exchange* is whatever you want it to be, and the value for the *SamAccountName* parameter, in this case *EXCH2016ASA*, has to be unique in your directory.

3. Use the **Set-ADComputer** cmdlet to enable the AES 256 encryption cipher support used by Kerberos using this cmdlet syntax:

   ```
   Set-ADComputer [-Name] <string> [-add @{<attributename>="<value>"]
   ```

   Example:

   ```
   Set-ADComputer EXCH2016ASA -add @{"msDS-SupportedEncryptionTypes"="28"}
   ```

   Where *EXCH2016ASA* is the name of the account and the attribute to be modified is *msDS-SupportedEncryptionTypes* with a decimal value of 28, which enables the following ciphers: RC4-HMAC, AES128-CTS-HMAC-SHA1-96, AES256-CTS-HMAC-SHA1-96.

For more information about these cmdlets, see Import-Module and New-ADComputer.

Cross-forest scenarios

If you have a cross-forest or resource-forest deployment, and you have users that are outside the Active Directory forest that contains Exchange, you must configure forest trust relationships between the forests. Also, for each forest in the deployment, you have to set up a routing rule that enables trust between all name suffixes within the forest and across forests. For more information about managing cross-forest trusts, see Configuring Partner Organizations.

Identify the Service Principal Names to associate with the ASA credential

After you create the ASA credential, you have to associate Exchange Service Principal Names (SPNs) with the ASA credential. The list of Exchange SPNs may vary with your configuration, but should include at least the following:

- **http/**: Use this SPN for Outlook Anywhere, MAPI over HTTP, Exchange Web Services, Autodiscover, and Offline Address Book.

The SPN values must match the service name on the network load balancer instead of on individual servers. To help plan which SPN values you should use, consider the following scenarios:

- Single Active Directory site
- Multiple Active Directory sites

In each of these scenarios, assume that the load-balanced, fully-qualified domain names (FQDNs) have been deployed for the internal URLs, external URLs, and the autodiscover internal URI used by members running Client Access services.

Single Active Directory site

If you have a single Active Directory site, your environment may resemble the one in the following figure:

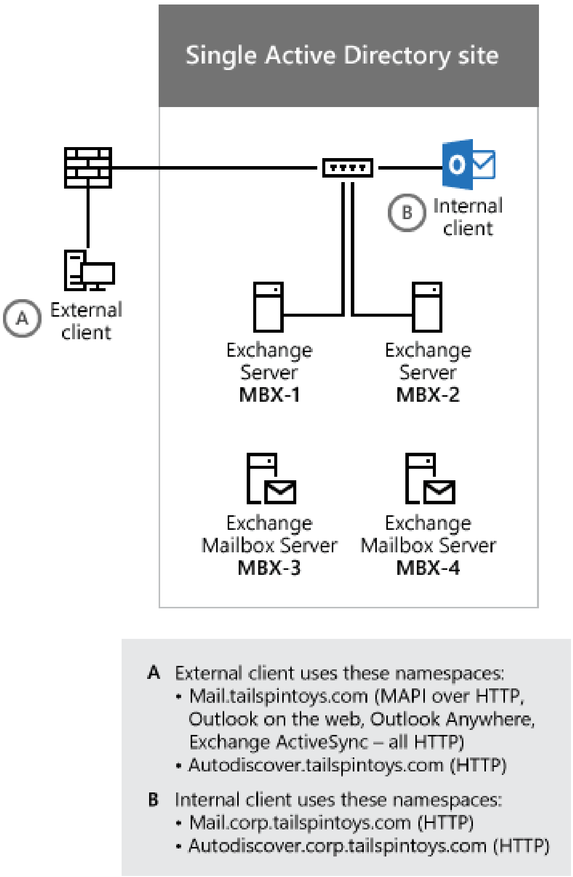
Single Active Directory site
B Internal client
A External client
Exchange Server MBX-1
Exchange Server MBX-2
Exchange Mailbox Server MBX-3
Exchange Mailbox Server MBX-4

A External client uses these namespaces:
- Mail.tailspintoys.com (MAPI over HTTP, Outlook on the web, Outlook Anywhere, Exchange ActiveSync – all HTTP)
- Autodiscover.tailspintoys.com (HTTP)

B Internal client uses these namespaces:
- Mail.corp.tailspintoys.com (HTTP)
- Autodiscover.corp.tailspintoys.com (HTTP)

Based on the FQDNs that are used by the internal Outlook clients in the preceding figure, you have to associate the following SPNs with the ASA credential:

- http/mail.corp.tailspintoys.com
- http/autodiscover.corp.tailspintoys.com

Multiple Active Directory sites

If you have multiple Active Directory sites, your environment may resemble the one in the following figure:

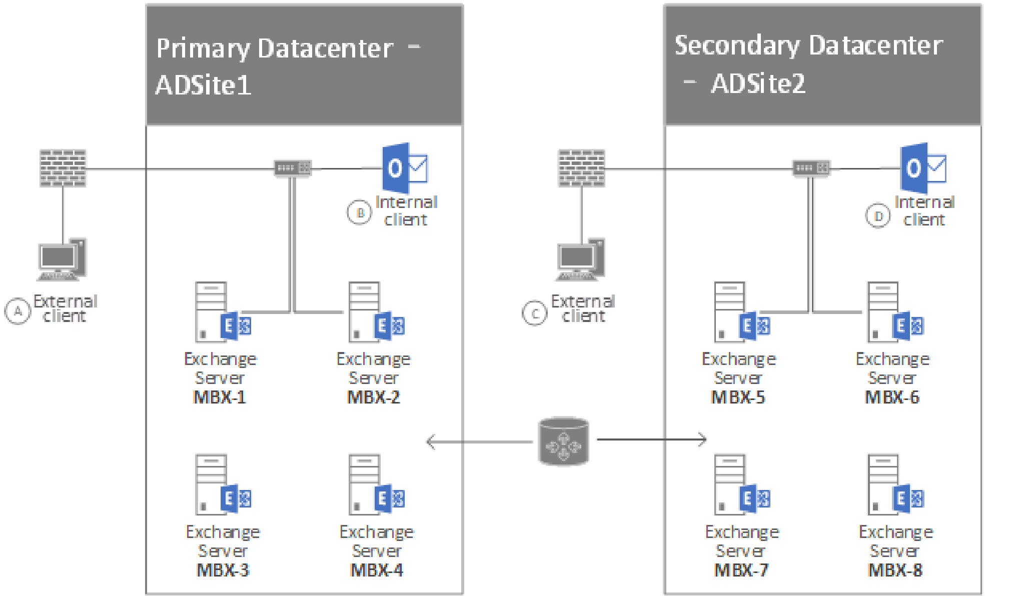

A External client uses these namespaces:
- Mail.tailspintoys.com(MAPI over HTTP, Outlook on the web, Outlook Anywhere, Exchange ActiveSync – all HTTP)
- Autodiscover.tailspintoys.com (HTTP)

B Internal client uses these namespaces:
- Mail.corp.tailspintoys.com(HTTP)
- Autodiscover.corp.tailspintoys.com(HTTP)

C External client uses these namespaces:
- Mailsdc.tailspintoys.com(MAPI over HTTP, Outlook on the web, Outlook Anywhere, Exchange ActiveSync – all HTTP)

D Internal client uses these namespaces:
- Mailsdc.corp.tailspintoys.com(HTTP)
- Autodiscoversdc.corp.tailspintoys.com(HTTP)

Based on the FQDNs that are used by the Outlook clients in the preceding figure, you would have to associate the following SPNs with the ASA credential that is used by the Mailbox servers running Client Access services in ADSite 1:

- http/mail.corp.tailspintoys.com
- http/autodiscover.corp.tailspintoys.com

You would also have to associate the following SPNs with the ASA credential that is used by the Mailbox servers running Client Access services in ADSite 2:

- http/mailsdc.corp.tailspintoys.com
- http/autodiscoversdc.corp.tailspintoys.com

Configure and then verify configuration of the ASA credential on each server running Client Access services

After you've created the account, you have to verify that the account has replicated to all AD DS domain controllers. Specifically, the account must be present on each server running Client Access services that will use the ASA credential. Next, you configure the account as the ASA credential on each server running Client Access services in your deployment.

You configure the ASA credential by using the Exchange Management Shell as described in one of these procedures:

- Deploy the ASA credential to the first Exchange server running Client Access services
- Deploy the ASA credential to subsequent Exchange servers running Client Access services

The only supported method for deploying the ASA credential is to use the RollAlternateServiceAcountPassword.ps1 script. For more information, see Using the RollAlternateserviceAccountCredential.ps1 Script in the Shell. After the script has run, we recommend that you verify that all the targeted servers have been updated correctly.

Deploy the ASA Credential to the first Exchange server running Client Access services

1. Open the Exchange Management Shell on an Exchange 2016 or Exchange 2019 server.

2. Change directories to *<Exchange 2016 installation directory>*\V15\Scripts.

3. Run the following command to deploy the ASA credential to the first Exchange 2016 or Exchange 2019 server running Client Access services:

   ```
   .\RollAlternateServiceAccountPassword.ps1 -ToSpecificServer cas-1.corp.tailspintoys.com -
   GenerateNewPasswordFor tailspin\EXCH2016ASA$
   ```

4. When you're asked if you want to change the password for the alternate service account, answer **Yes**.

The following is an example of the output that's shown when you run the RollAlternateServiceAccountPassword.ps1 script.

```
========== Starting at 01/12/2016 10:17:47 ==========
Creating a new session for implicit remoting of "Get-ExchangeServer" command...
Destination servers that will be updated:
Name                                                        PSComputerName
----                                                        --------------
cas-1                                                   cas-1.corp.tailspintoys.com
Credentials that will be pushed to every server in the specified scope (recent first):
UserName
Password
--------
--------
tailspin\EXCH2016ASA$
System.Security.SecureString
Prior to pushing new credentials, all existing credentials that are invalid or no longer work will be removed
from  the destination servers.
Pushing credentials to server mbx-1
Setting a new password on Alternate Service Account in Active Directory
Password change
Do you want to change password for tailspin\EXCH2016ASA$ in Active Directory at this time?
[Y] Yes  [N] No  [S] Suspend  [?] Help (default is "Y"): y
Preparing to update Active Directory with a new password for tailspin\EXCH2016ASA$ ...
Resetting a password in the Active Directory for tailspin\EXCH2016ASA$ ...
New password was successfully set to Active Directory.
Retrieving the current Alternate Service Account configuration from servers in scope
Alternate Service Account properties:
StructuralObjectClass QualifiedUserName Last Pwd Update       SPNs
--------------------- ----------------- --------------       ----
computer              tailspin\EXCH2016ASA$   1/12/2016 10:19:53 AM
Per-server Alternate Service Account configuration as of the time of script completion:
   Array: {mail.corp.tailspintoys.com}
Identity  AlternateServiceAccountConfiguration
--------  ------------------------------------
cas-1 Latest: 1/12/2016 10:19:22 AM, tailspin\EXCH2016ASA$
          ...
========== Finished at 01/12/2016 10:20:00 ==========
        THE SCRIPT HAS SUCCEEDED
```

Deploy the ASA credential to another Exchange server running Client Access services

1. Open the Exchange Management Shell on an Exchange 2016 or Exchange 2019 server.

2. Change directories to *<Exchange 2016 installation directory>*\V15\Scripts.

3. Run the following command to deploy the ASA credential to another Exchange 2016 or Exchange 2019 server running Client Access services:

```
.\RollAlternateServiceAccountPassword.ps1 -ToSpecificServer cas-2.corp.tailspintoys.com -CopyFrom cas-1.corp.tailspintoys.com
```

4. Repeat Step 3 for each server running Client Access services that you want to deploy the ASA credential to.

The following is an example of the output that's shown when you run the RollAlternateServiceAccountPassword.ps1 script.

```
========== Starting at 01/12/2016 10:34:35 ==========
Destination servers that will be updated:
Name                                                        PSComputerName
----                                                        --------------
cas-2                                                   cas-2.corp.tailspintoys.com
Credentials that will be pushed to every server in the specified scope (recent first):
UserName
Password
--------
--------
tailspin\EXCH2016ASA$
System.Security.SecureString
Prior to pushing new credentials, all existing credentials will be removed from the destination servers.
Pushing credentials to server mbx-2
Retrieving the current Alternate Service Account configuration from servers in scope
Alternate Service Account properties:
StructuralObjectClass QualifiedUserName Last Pwd Update       SPNs
--------------------- ----------------- ---------------       ----
computer              tailspin\EXCH2016ASA$   1/12/2016 10:19:53 AM
Per-server Alternate Service Account configuration as of the time of script completion:
   Array: cas-2.corp.tailspintoys.com
Identity  AlternateServiceAccountConfiguration
--------  ------------------------------------
cas-2 Latest: 1/12/2016 10:37:59 AM, tailspin\EXCH2016ASA$
          ...
========== Finished at 01/12/2016 10:38:13 ==========
        THE SCRIPT HAS SUCCEEDED
```

Verify the deployment of the ASA credential

- Open the Exchange Management Shell on an Exchange 2016 or Exchange 2019 server.
- Run the following command to check the settings on the server running Client Access services:

```
Get-ClientAccessServer CAS-3 -IncludeAlternateServiceAccountCredentialStatus | Format-List Name, AlternateServiceAccountConfiguration
```

- Repeat Step 2 on each server running Client Access services for which you want to verify the deployment of the ASA credential.

The following is an example of the output that's shown when you run the Get-ClientAccessServer command above and no previous ASA credential was set.

```
Name                                 : CAS-1
AlternateServiceAccountConfiguration : Latest: 1/12/2016 10:19:22 AM, tailspin\EXCH2016ASA$
                                       Previous: <Not set>
                                        ...
```

The following is an example of the output that's shown when you run the Get-ClientAccessServer command above and an ASA credential was previously set. The previous ASA credential and the date and time it was set are returned.

```
Name                                 : CAS-3
AlternateServiceAccountConfiguration : Latest: 1/12/2016 10:19:22 AM, tailspin\EXCH2016ASA$
                                       Previous: 7/15/2015 12:58:35 PM, tailspin\oldSharedServiceAccountName$
                                         ...
```

Associate Service Principal Names (SPNs) with the ASA credential

> **IMPORTANT**
>
> Don't associate SPNs with an ASA credential until you have deployed that credential to at least one Exchange Server, as described earlier in Deploy the ASA Credential to the first Exchange server running Client Access services. Otherwise, you will experience Kerberos authentication errors.

Before you associate the SPNs with the ASA credential, you have to verify that the target SPNs aren't already associated with a different account in the forest. The ASA credential must be the only account in the forest with which these SPNs are associated. You can verify that no other account in the forest is associated with the SPNs by running the **setspn** command from the command line.

Verify an SPN is not already associated with an account in a forest by running the setspn command

1. Press **Start**. In the **Search** box, type **Command Prompt**, then in the list of results, select **Command Prompt**.

2. At the command prompt, type the following command:

   ```
   setspn -F -Q <SPN>
   ```

   Where <SPN> is the SPN you want to associate with the ASA credential. For example:

   ```
   setspn -F -Q http/mail.corp.tailspintoys.com
   ```

   The command should return nothing. If it returns something, another account is already associated with the SPN. Repeat this step one time for each SPN that you want to associate with the ASA credential.

Associate an SPN with an ASA credential by using the setspn command

1. Press **Start**. In the **Search** box, type **Command Prompt**, and then select **Command Prompt** in the list of results.

2. At the command prompt, type the following command:

   ```
   setspn -S <SPN> <Account>$
   ```

   Where <SPN> is the SPN you want to associate with the ASA credential and <Account> is the account associated with the ASA credential. For example:

   ```
   setspn -S http/mail.corp.tailspintoys.com tailspin\EXCH2016ASA$
   ```

   Run this command one time for each SPN that you want to associate with the ASA credential.

Verify you associated the SPNs with the ASA credentials by using the setspn command

1. Press **Start**. In the **Search** box, type **Command Prompt**, and then select **Command Prompt** in the list of results.

2. At the command prompt, type the following command:

   ```
   setspn -L <Account>$
   ```

   Where <Account> is the account associated with the ASA credential. For example:

   ```
   setspn -L tailspin\EXCH2016ASA$
   ```

   You have to run this command only one time.

Enable Kerberos authentication for Outlook clients

1. Open the Exchange Management Shell on an Exchange 2016 or Exchange 2019 server.

2. To enable Kerberos authentication for Outlook Anywhere clients, run the following command on your Exchange 2016 or Exchange 2019 server that is running Client Access services:

   ```
   Get-OutlookAnywhere -Server CAS-1 | Set-OutlookAnywhere -InternalClientAuthenticationMethod  Negotiate
   ```

3. To enable Kerberos authentication for MAPI over HTTP clients, run the following command on your Exchange 2016 or Exchange 2019 server that is running Client Access services:

   ```
   Get-MapiVirtualDirectory -Server CAS-1 | Set-MapiVirtualDirectory -IISAuthenticationMethods
   Ntlm,Negotiate
   ```

   In hybrid environments with Exchange Online or if you use OAuth internally, run the following commands on your Exchange 2016 or Exchange 2019 server that's running Client Access services:

   ```
   $mapidir = Get-MapiVirtualDirectory -Server CAS-1
   $mapidir | Set-MapiVirtualDirectory -IISAuthenticationMethods ($mapidir.IISAuthenticationMethods
   +='Negotiate')
   ```

4. Repeat steps 2 and 3 for each Exchange 2016 or Exchange 2019 server that is running Client Access services for which you want to enable Kerberos authentication.

Verify Exchange client Kerberos authentication

After you've successfully configured Kerberos and the ASA credential, verify that clients can authenticate successfully, as described in these tasks.

Verify that the Microsoft Exchange Service Host service is running

The Microsoft Exchange Service Host service (MSExchangeServiceHost) on the server that is running Client Access services is responsible for managing the ASA credential. If MSExchangeServiceHost isn't running, Kerberos authentication isn't possible. By default, the service is configured to automatically start when the computer starts.

To verify the Microsoft Exchange Service Host service is started

1. Click **Start**, type **services.msc**, and then select **services.msc** from the list.

2. In the **Services** window, locate the **Microsoft Exchange Service Host** service in the list of services.

3. The status of the service should be **Running**. If the status is not **Running**, right-click the service, and then click **Start**.

Verify Kerberos from the server running Client Access services

When you configured the ASA credential on each server running Client Access services, you ran the **Set-**

ClientAccessServer cmdlet. After you run this cmdlet, you can use the logs to verify successful Kerberos connections.

Verify that Kerberos is working correctly by using the HttpProxy log file

1. In a text editor, browse to the folder where the HttpProxy log is stored. By default, the log is stored in the following folder:

   %ExchangeInstallPath%\Logging\HttpProxy\RpcHttp

2. Open the most recent log file, and then look for the word **Negotiate**. The line in the log file will look something like the following example:

   ```
   2014-02-19T13:30:49.219Z,e19d08f4-e04c-42da-a6be-
   b7484b396db0,15,0,775,22,,RpcHttp,mail.corp.tailspintoys.com,/rpc/rpcproxy.dll,,Negotiate,True,tailspin\
   Wendy,tailspintoys.com,MailboxGuid~ad44b1e0-e44f-4a16-9396-
   3a437f594f88,MSRPC,192.168.1.77,EXCH1,200,200,,RPC_OUT_DATA,Proxy,exch2.tailspintoys.com,15.00.0775.000,
   IntraForest,MailboxGuidWithDomain,,,,76,462,1,,1,1,,0,,0,,0,0,16272.3359,0,0,3,0,23,0,25,0,16280,1,16274
   ,16230,16233,16234,16282,?ad44b1e0-e44f-4a16-9396-3a437f594f88@tailspintoys.com:6001,,BeginRequest=2014-
   02-19T13:30:32.946Z;BeginGetRequestStream=2014-02-19T13:30:32.946Z;OnRequestStreamReady=2014-02-
   19T13:30:32.946Z;BeginGetResponse=2014-02-19T13:30:32.946Z;OnResponseReady=2014-02-
   19T13:30:32.977Z;EndGetResponse=2014-02-19T13:30:32.977Z;,PossibleException=IOException;
   ```

   If you see that the **AuthenticationType** value is **Negotiate**, the server is successfully creating Kerberos authenticated connections.

Maintain the ASA credential

If you have to refresh the password on the ASA credential periodically, use the steps for configuring the ASA credential in this article. Consider setting up a scheduled task to perform regular password maintenance. Be sure to monitor the scheduled task to ensure timely password rollovers and prevent possible authentication outages.

Turn Kerberos authentication off

To configure your servers that are running Client Access services to stop using Kerberos, disassociate or remove the SPNs from the ASA credential. If the SPNs are removed, Kerberos authentication won't be tried by your clients, and clients that are configured to use Negotiate authentication will use NTLM instead. Clients that are configured to use only Kerberos will be unable to connect. After the SPNs are removed, you should also delete the account.

To remove the ASA credential

1. Open the Exchange Management Shell on an Exchange 2016 or Exchange 2019 server, and run the following command:

   ```
   Set-ClientAccessServer CAS-1 -RemoveAlternateServiceAccountCredentials
   ```

2. Although you don't have to do this immediately, you should eventually restart all client computers to clear the Kerberos ticket cache from the computer.

Digital certificates and encryption in Exchange Server

8/3/2020 • 15 minutes to read • Edit Online

Encryption and digital certificates are important considerations in any organization. By default, Exchange Server is configured to use Transport Layer Security (TLS) to encrypt communication between internal Exchange servers, and between Exchange services on the local server. But, Exchange administrators need to consider their encryption requirements for communication with internal and external clients (computers and mobile devices), and external messaging servers.

> **NOTE**
>
> Exchange Server 2019 includes important changes to improve the security of client and server connections. The default configuration for encryption will enable TLS 1.2 only and disable support for older algorithms (namely, DES, 3DES, RC2, RC4 and MD5). It will also configure elliptic curve key exchange algorithms with priority over non-elliptic curve algorithms. In Exchange Server 2016 and later, all cryptography settings are inherited from the configuration specified in the operating system. For additional information, see Exchange Server TLS Guidance.

This topic describes the different types of certificates that are available, the default configuration for certificates in Exchange, and recommendations for additional certificates that you'll need to use with Exchange.

For the procedures that are required for certificates in Exchange Server, see Certificate procedures in Exchange Server.

Digital certificates overview

Digital certificates are electronic files that work like an online password to verify the identity of a user or a computer. They're used to create the encrypted channel that's used for client communications. A certificate is a digital statement that's issued by a certification authority (CA) that vouches for the identity of the certificate holder and enables the parties to communicate in a secure manner by using encryption.

Digital certificates provide the following services:

- **Encryption**: They help protect the data that's exchanged from theft or tampering.
- **Authentication**: They verify that their holders (people, web sites, and even network devices such as routers) are truly who or what they claim to be. Typically, the authentication is one-way, where the source verifies the identity of the target, but mutual TLS authentication is also possible.

Certificates can be issued for several uses. For example: web user authentication, web server authentication, Secure/Multipurpose Internet Mail Extensions (S/MIME), Internet Protocol security (IPsec), and code signing.

A certificate contains a public key and attaches that public key to the identity of a person, computer, or service that holds the corresponding private key. The public and private keys are used by the client and the server to encrypt data before it's transmitted. For Windows users, computers, and services, trust in the CA is established when the root certificate is defined in the trusted root certificate store, and the certificate contains a valid certification path. A certificate is considered valid if it hasn't been revoked (it isn't in the CA's certificate revocation list or CRL), or hasn't expired.

The three primary types of digital certificates are described in the following table.

TYPE	DESCRIPTION	ADVANTAGES	DISADVANTAGES
Self-signed certificate	The certificate is signed by the application that created it.	Cost (free).	The certificate isn't automatically trusted by client computers and mobile devices. The certificate needs to be manually added to the trusted root certificate store on all client computers and devices, but not all mobile devices allow changes to the trusted root certificate store. Not all services work with self-signed certificates. Difficult to establish an infrastructure for certificate lifecycle management. For example, self-signed certificates can't be revoked.
Certificate issued by an internal CA	The certificate is issued by a public key infrastructure (PKI) in your organization. An example is Active Directory Certificate Services (AD CS). For more information, see Active Directory Certificate Services Overview.	Allows organizations to issue their own certificates. Less expensive than certificates from a commercial CA.	Increased complexity to deploy and maintain the PKI. The certificate isn't automatically trusted by client computers and mobile devices. The certificate needs to be manually added to the trusted root certificate store on all client computers and devices, but not all mobile devices allow changes to the trusted root certificate store.
Certificate issued by a commercial CA	The certificate is purchased from a trusted commercial CA.	Simplified certificate deployment, because all clients, devices, and servers automatically trust the certificates.	Cost. You need to plan ahead to minimize the number of certificates that are required.

To prove that a certificate holder is who they claim to be, the certificate must accurately identify the certificate holder to other clients, devices, or servers. The three basic methods to do this are described in the following table.

METHOD	DESCRIPTION	ADVANTAGES	DISADVANTAGES
Certificate subject match	The certificate's **Subject** field contains the common name (CN) of the host. For example, the certificate that's issued to www.contoso.com can be used for the web site https://www.contoso.com.	Compatible with all clients, devices, and services. Compartmentalization. Revoking the certificate for a host doesn't affect other hosts.	Number of certificates required. You can only use the certificate for the specified host. For example, you can't use the www.contoso.com certificate for ftp.contoso.com, even when the services are installed on the same server. Complexity. On a web server, each certificate requires its own IP address binding.

METHOD	DESCRIPTION	ADVANTAGES	DISADVANTAGES
Certificate subject alternative name (SAN) match	In addition to the **Subject** field, the certificate's **Subject Alternative Name** field contains a list of multiple host names. For example: • www.contoso.com • ftp.contoso.com • ftp.eu.fabirkam.net	Convenience. You can use the same certificate for multiple hosts in multiple, separate domains. Most clients, devices, and services support SAN certificates. Auditing and security. You know exactly which hosts are capable of using the SAN certificate.	More planning required. You need to provide the list of hosts when you create the certificate. Lack of compartmentalization. You can't selectively revoke certificates for some of the specified hosts without affecting all of the hosts in the certificate.
Wildcard certificate match	The certificate's **Subject** field contains the common name as the wildcard character (*) plus a single domain or subdomain. For example, *.contoso.com or *.eu.contoso.com. The *.contoso.com wildcard certificate can be used for: • www.contoso.com • ftp.contoso.com • mail.contoso.com	Flexibility. You don't need to provide a list of hosts when you request the certificate, and you can use the certificate on any number of hosts that you may need in the future.	You can't use wildcard certificates with other top-level domains (TLDs). For example, you can't use the *.contoso.com wildcard certificate for *.contoso.net hosts. You can only use wildcard certificates for host names at the level of the wildcard. For example, you can't use the *.contoso.com certificate for www.eu.contoso.com. Or, you can't use the *.eu.contoso.com certificate for www.uk.eu.contoso.com. Older clients, devices, applications, or services might not support wildcard certificates. Wildcards aren't available with Extended Validation (EV) certificates. Careful auditing and control is required. If the wildcard certificate is compromised, it affects every host in the specified domain.

Certificates in Exchange

When you install Exchange 2016 or Exchange 2019 on a server, two self-signed certificates are created and installed by Exchange. A third self-signed certificate is created and installed by Microsoft Windows for the Web Management service in Internet Information Services (IIS). These three certificates are visible in the Exchange admin center (EAC) and the Exchange Management Shell, and are described in the following table:

NAME	COMMENTS

NAME	COMMENTS
Microsoft Exchange	This Exchange self-signed certificate has the following capabilities: • The certificate is automatically trusted by all other Exchange servers in the organization. This includes any Edge Transport servers subscribed to the Exchange organization. • The certificate is automatically enabled for all Exchange services except Unified Messaging, and is used to encrypt internal communication between Exchange servers, Exchange services on the same computer, and client connections that are proxied from the Client Access services to the backend services on Mailbox servers. (**Note**: UM is not available on Exchange 2019.) • The certificate is automatically enabled for inbound connections from external SMTP messaging servers, and outbound connections to external SMTP messaging servers. This default configuration allows Exchange to provide *opportunistic TLS* on all inbound and outbound SMTP connections. Exchange attempts to encrypt the SMTP session with an external messaging server, but if the external server doesn't support TLS encryption, the session is unencrypted. • The certificate doesn't provide encrypted communication with internal or external clients. Clients and servers don't trust the Exchange self-signed certificate, because the certificate isn't defined in their trusted root certification stores.
Microsoft Exchange Server Auth Certificate	This Exchange self-signed certificate is used for server-to-server authentication and integration by using OAuth. For more information, see Plan Exchange Server integration with SharePoint and Skype for Business.
WMSVC	This Windows self-signed certificate is used by the Web Management service in IIS to enable remote management of the web server and its associated web sites and applications. If you remove this certificate, the Web Management service will fail to start if no valid certificate is selected. Having the service in this state can prevent you from installing Exchange updates, or uninstalling Exchange from the server. For instructions on how to correct this issue, see Event ID 1007 - IIS Web Management Service Authentication

The properties of these self-signed certificates are described in the Properties of the default self-signed certificates section.

These are the key issues that you need to consider when it comes to certificates in Exchange:

- You don't need to replace the Microsoft Exchange self-signed certificate to encrypt network traffic between Exchange servers and services in your organization.
- You need additional certificates to encrypt connections to Exchange servers by internal and external clients.
- You need additional certificates to force the encryption of SMTP connections between Exchange servers and external messaging servers.

The following elements of planning and deployment for Exchange Server are important drivers for your certificate requirements:

- **Load balancing**: Do you plan to terminate the encrypted channel at load balancer or reverse proxy server, use Layer 4 or Layer 7 load balancers, and use session affinity or no session affinity? For more information, see Load Balancing in Exchange 2016.

- **Namespace planning**: What versions of Exchange are present, are you using the bound or unbound namespace model, and are you using *split-brain DNS* (configuring different IP addresses for the same host based on internal vs. external access)? For more information, see Namespace Planning in Exchange 2016.

- **Client connectivity**: What services will your clients use (web-based services, POP, IMAP, etc.) and what versions of Exchange are involved? For more information, see the following topics:

  - Client Connectivity in an Exchange 2016 Coexistence Environment with Exchange 2013

  - Client Connectivity in an Exchange 2016 Coexistence Environment with Exchange 2010

  - Client Connectivity in an Exchange 2016 Coexistence Environment with Mixed Exchange Versions

Certificate requirements for Exchange services

The Exchange services that certificates can be assigned to are described in the following table.

SERVICE	DESCRIPTION
IIS (HTTP)	By default, the following services are offered under the default website in the Client Access (frontend) services on a Mailbox server, and are used by clients to connect to Exchange: • Autodiscover • Exchange ActiveSync • Exchange admin center • Exchange Web Services • Offline address book (OAB) distribution • Outlook Anywhere (RPC over HTTP) • Outlook MAPI over HTTP • Outlook on the web • Remote PowerShell* Because you can only associate a single certificate with a website, all the DNS names that clients use to connect to these services need to be included in the certificate. You can accomplish this by using a SAN certificate or a wildcard certificate.
POP or IMAP	The certificates that are used for POP or IMAP can be different from the certificate that's used for IIS. However, to simplify administration, we recommend that you also include the host names that are used for POP or IMAP in your IIS certificate, and use the same certificate for all of these services.

SERVICE	DESCRIPTION
SMTP	SMTP connections from clients or messaging servers are accepted by one or more Receive connectors that are configured in the Front End Transport service on the Exchange server. For more information, see Receive connectors. To require TLS encryption for SMTP connections, you can use a separate certificate for each Receive connector. The certificate must include the DNS name that's used by the SMTP clients or servers to connect to the Receive connector. To simplify certificate management, consider including all DNS names for which you have to support TLS traffic in a single certificate. To require *mutual TLS authentication*, where the SMTP connections between the source and destination servers are both encrypted and authenticated, see Domain Security.
Unified Messaging (UM)	For more information, see Deploying Certificates for UM. **Note**: UM is not available in Exchange 2019.
Hybrid deployment with Microsoft 365 or Office 365	For more information, see Certificate Requirements for Hybrid Deployments.
Secure/Multipurpose Internet Mail Extensions (S/MIME)	For more information, see S/MIME for message signing and encryption.

[*] Kerberos authentication and Kerberos encryption are used for remote PowerShell access, from both the Exchange admin center and the Exchange Management Shell. Therefore, you don't need to configure your certificates for use with remote PowerShell, as long as you connect directly to an Exchange server (not to a load balanced namespace). To use remote PowerShell to connect to an Exchange server from a computer that isn't a member of the domain, or to connect from the Internet, you need to configure your certificates for use with remote PowerShell.

Best practices for Exchange certificates

Although the configuration of your organization's digital certificates will vary based on its specific needs, information about best practices has been included to help you choose the digital certificate configuration that's right for you.

- **Use as few certificates as possible**: Very likely, this means using SAN certificates or wildcard certificates. In terms of interoperability with Exchange, both are functionally equivalent. The decision on whether to use a SAN certificate vs a wildcard certificate is more about the key capabilities or limitations (real or perceived) for each type of certificate as described in the Digital certificates overview.

  For example, if all of your common names will be in the same level of contoso.com, it doesn't matter if you use a SAN certificate or a wildcard certificate. But, if need to use the certificate for autodiscover.contoso.com, autodiscover.fabrikam.com, and autodiscover.northamerica.contoso.com, you need to use a SAN certificate.

- **Use certificates from a commercial CA for client and external server connections**: Although you can configure most clients to trust any certificate or certificate issuer, it's much easier to use a certificate from a commercial CA for client connections to your Exchange servers. No configuration is required on the client to trust a certificate that's issued by a commercial CA. Many commercial CAs offer certificates that are configured specifically for Exchange. You can use the EAC or the Exchange Management Shell to generate certificate requests that work with most commercial CAs.

- **Choose the right commercial CA**: Compare certificate prices and features between CAs. For example:
  - Verify that the CA is trusted by the clients (operating systems, browsers, and mobile devices) that connect to your Exchange servers.
  - Verify that the CA supports the kind of certificate that you need. For example, not all CAs support SAN certificates, the CA might limit the number of common names that you can use in a SAN certificate, or the CA may charge extra based on the number of common names in a SAN certificate.
  - See if the CA offers a grace period during which you can add additional common names to SAN certificates after they're issued without being charged.
  - Verify that the certificate's license allows you to use the certificate on the required number of servers. Some CAs only allow you to use the certificate on one server.
- **Use the Exchange certificate wizard**: A common error when you create certificates is to forget one or more common names that are required for the services that you want to use. The certificate wizard in the Exchange admin center helps you include the correct list of common names in the certificate request. The wizard lets you specify the services that will use the certificate, and includes the common names that you need to have in the certificate for those services. Run the certificate wizard when you've deployed your initial set of Exchange 2016 or Exchange 2019 servers and determined which host names to use for the different services for your deployment.
- **Use as few host names as possible**: Minimizing the number of host names in SAN certificates reduces the complexity that's involved in certificate management. Don't feel obligated to include the host names of individual Exchange servers in SAN certificates if the intended use for the certificate doesn't require it. Typically, you only need to include the DNS names that are presented to the internal clients, external clients, or external servers that use the certificate to connect to Exchange.

  For a simple Exchange Server organization named Contoso, this is a hypothetical example of the minimum host names that would be required:

  - **mail.contoso.com**: This host name covers most connections to Exchange, including Outlook, Outlook on the web, OAB distribution, Exchange Web Services, Exchange admin center, and Exchange ActiveSync.
  - **autodiscover.contoso.com**: This specific host name is required by clients that support Autodiscover, including Outlook, Exchange ActiveSync, and Exchange Web Services clients. For more information, see Autodiscover service.

Properties of the default self-signed certificates

Some of the more interesting properties of the default self-signed certificates that are visible in the Exchange admin center and/or the Exchange Management Shell on an Exchange server are described in the following table.

	MICROSOFT EXCHANGE	MICROSOFT EXCHANGE SERVER AUTH CERTIFICATE	WMSVC
Subject	`CN=<ServerName>` (for example, `CN=Mailbox01`)	`CN=Microsoft Exchange Server Auth Certificate`	`CN=WMSvc-<ServerName>` (for example, `CN=WMSvc-Mailbox01`)

	MICROSOFT EXCHANGE	MICROSOFT EXCHANGE SERVER AUTH CERTIFICATE	WMSVC
Subject Alternative Names (CertificateDomains)	• *<ServerName>* (for example, Mailbox01) • *<ServerFQDN>* (for example, Mailbox01.contoso.com)	none	`WMSvc-<ServerName>` (for example, `WMSvc-Mailbox01`)
Has private key (HasPrivateKey)	**Yes** (True)	**Yes** (True)	**Yes** (True)
PrivateKeyExportable*	False	True	True
EnhancedKeyUsageList*	Server Authentication (1.3.6.1.5.5.7.3.1)	Server Authentication (1.3.6.1.5.5.7.3.1)	Server Authentication (1.3.6.1.5.5.7.3.1)
IISServices*	`IIS://<ServerName>/W3SVC/1, IIS://<ServerName>/W3SVC/2` (for example, `IIS://Mailbox01/W3SVC/1, IIS://Mailbox01/W3SVC/2`)	none	none
IsSelfSigned	True	True	True
Issuer	`CN=<ServerName>` (for example, `CN=Mailbox01`)	`CN=Microsoft Exchange Server Auth Certificate`	`CN=WMSvc-<ServerName>` (for example, `CN=WMSvc-Mailbox01`)
NotBefore	The date/time that Exchange was installed.	The date/time that Exchange was installed.	The date/time that the IIS Web Manager service was installed.
Expires on (NotAfter)	5 years after `NotBefore`.	5 years after `NotBefore`.	10 years after `NotBefore`.
Public key size (PublicKeySize)	2048	2048	2048
RootCAType	Registry	None	Registry
Services	IMAP, POP, IIS, SMTP	SMTP	None

*These properties aren't visible in the standard view in the Exchange Management Shell. To see them, you need to specify the property name (exact name or wildcard match) with the **Format-Table** or **Format-List** cmdlets. For example:

- `Get-ExchangeCertificate -Thumbprint <Thumbprint> | Format-List *`
- `Get-ExchangeCertificate -Thumbprint <Thumbprint> | Format-Table -Auto FriendlyName,*PrivateKey*`

For more information, see Get-ExchangeCertificate.

Further details about the default self-signed certificates that are visible in Windows Certificate Manger are described in the following table.

	MICROSOFT EXCHANGE	MICROSOFT EXCHANGE SERVER AUTH CERTIFICATE	WMSVC
Signature algorithm	sha1RSA	sha1RSA	sha1RSA
Signature hash algorithm	sha1	sha1	sha1
Key usage	Digital Signature, Key Encipherment (a0)	Digital Signature, Key Encipherment (a0)	Digital Signature, Key Encipherment (a0), Data Encipherment (b0 00 00 00)
Basic constraints	• `Subject Type=End Entity` • `Path Length Constraint=None` .	• `Subject Type=End Entity` • `Path Length Constraint=None`	n/a
Thumbprint algorithm	sha1	sha1	sha1

Typically, you don't use Windows Certificate Manger to manage Exchange certificates (use the Exchange admin center or the Exchange Management Shell). Note that the WMSVC certificate isn't an Exchange certificate.

Certificate procedures in Exchange Server

8/3/2020 • 3 minutes to read • Edit Online

Ensuring that certificates are installed and configured correctly is key to delivering a secure messaging infrastructure for the enterprise. In Exchange Server you can manage certificates in the Exchange admin center (EAC), and in the Exchange Management Shell. Certificate management in the EAC has been improved over certificate management in the Exchange Management Console in Exchange Server 2010. Specifically, certificate management in the EAC can help administrators by:

- Minimizing the number of certificates that are required.
- Minimizing the interaction that's required for certificates.
- Allowing the centralized installation and management of certificates on all Exchange servers in the organization.

For more information about certificates in Exchange, see Digital certificates and encryption in Exchange Server.

The tasks that are associated with certificate management in Exchange are described in the following table.

TASK	EAC	EXCHANGE MANAGEMENT SHELL	TOPIC	COMMENTS
Create a new self-signed certificate on an Exchange server.	**Servers** > **Certificates** > select the server > **Add** + > **Create a self-signed certificate**	**New-ExchangeCertificate**	Create a new Exchange Server self-signed certificate	You can create new self-signed certificates and configure the certificates for Exchange services in one step.
Create a new certificate request (also known as a certificate signing request or CSR) for a certification authority (CA).	**Servers** > **Certificates** > select the server > **Add** + > **Create a request for a certificate from a certification authority**	**New-ExchangeCertificate** with the *GenerateRequest* switch.	Create an Exchange Server certificate request for a certification authority	The procedures are the same for an internal CA (for example, Active Directory Certificate Services) or a commercial CA.
Complete a pending certificate request on an Exchange server.	**Servers** > **Certificates** > select the server > select the certificate request > click the **Complete** link in the details pane.	**Import-ExchangeCertificate**	Complete a pending Exchange Server certificate request	After you receive the certificate file or files from the CA, you install them on the Exchange server.

TASK	EAC	EXCHANGE MANAGEMENT SHELL	TOPIC	COMMENTS
Assign a certificate to Exchange services.	**Servers** > **Certificates** > select the server > select the certificate > **Edit** ✏ > **Services** tab.	**Enable-ExchangeCertificate**	Assign certificates to Exchange Server services	The procedures are the same for self-signed certificates, or certificates that were issued by a CA. For certificates issued by a CA, you can only assign the certificates to Exchange services after you complete the pending certificate request (install the certificate on the Exchange server).
Delete an existing certificate or certificate request from an Exchange server.	**Servers** > **Certificates** > select the server > select the certificate > **Delete** 🗑	**Remove-ExchangeCertificate**	n/a	The procedures are the same for self-signed certificates, certificate requests, or certificates issued by a CA.
Renew an existing certificate on an Exchange server.	**Servers** > **Certificates** > select the server > select the certificate > click **Renew** in the details pane.	**Get-ExchangeCertificate** and **New-ExchangeCertificate**	Renew an Exchange Server certificate	For self-signed certificates, you renew the certificate in one step. For certificates that were issued by a CA, you create a request to renew the certificate, and send the request to the CA. The notification viewer in the EAC displays a warning when a certificate on any Exchange server in your organization is about to expire.
Export an existing certificate or certificate request from an Exchange server.	**Servers** > **Certificates** > select the server > select the certificate > **More options** ••• > **Export Exchange Certificate**	**Export-ExchangeCertificate**	Export a certificate from an Exchange server	You can only export valid (unexpired) certificates where the **PrivateKeyExportable** property has the value `True`. You can only export pending certificate requests in the Exchange Management Shell. You can't import an exported pending certificate request.

TASK	EAC	EXCHANGE MANAGEMENT SHELL	TOPIC	COMMENTS
Import (install) a certificate on an Exchange server.	**Servers** > **Certificates** > select the server > **More options** ••• > **Import Exchange Certificate**	**Import-ExchangeCertificate**	Import or install a certificate on an Exchange server	Import a certificate that was exported from another server.
View existing certificates or certificate requests on an Exchange server, or view the details for a specific certificate or certificate request.	**Servers** > **Certificates** > select the server For details on a specific certificate or certificate request, select the item from the list, and then click **Edit** ✏.	**Get-ExchangeCertificate**	n/a	Some certificate properties are visible in the details pane in the EAC when you select the certificate or certificate request from the list. Some certificate properties aren't visible in the standard view in the Exchange Management Shell. To see them, you need to specify the property name (exact name or wildcard match) with the **Format-Table** or **Format-List** cmdlets. For more information, see Get-ExchangeCertificate.

Create an Exchange Server certificate request for a certification authority

8/3/2020 • 8 minutes to read • Edit Online

Creating a certificate request is the first step in installing a new certificate on an Exchange server to configure Transport Layer Security (TLS) encryption for one or more Exchange services. You use a certificate request (also known as a certificate signing request or CSR) to obtain a certificate from a certification authority (CA). The procedures are the same for obtaining certificates from an internal CA (for example, Active Directory Certificate Services), or from a commercial CA. After you create the certificate request, you send the results to the CA, and the CA uses the information to issue the actual certificate, which you install later.

You can create certificate requests in the Exchange admin center (EAC) or in the Exchange Management Shell. The **New Exchange certificate** wizard in the EAC can assist you in selecting the host names that are required in the certificate.

What do you need to know before you begin?

- Estimated time to complete: 5 minutes to complete the new certificate request. However, more time is required before the request results in a certificate. For more information, see Next steps.
- You need to plan carefully to choose the type of certificate that you want, and the host names that are required in the certificate. For more information, see Digital certificates and encryption in Exchange Server.
- Verify the certificate request requirements of the CA. Exchange generates a PKCS #10 request (.req) file that uses Base64 (default) or Distinguished Encoding Rules (DER) encoding, with an RSA public key that's 1024, 2048 (default), or 4096 bits. Note that encoding and public key options are only available in the Exchange Management Shell. For more information, see New-ExchangeCertificate.
- In the EAC, you need to store the certificate request file on a UNC path (`\\<Server>\<Share>\` or `\\<LocalServerName>\c$\`). In the Exchange Management Shell, you can specify a local path.
- To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Client Access services security" entry in the Clients and mobile devices permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to create a new certificate request

1. Open the EAC and navigate to **Servers** > **Certificates**.
2. In the **Select server** list, select the Exchange server where you want to install the certificate, and then click

Add **+**.

3. The **New Exchange certificate** wizard opens. On the **This wizard will create a new certificate or a certificate request file** page, verify that **Create a request for a certificate from a certification authority** is selected, and then click **Next**.

   Note: To create a new self-signed certificate, see Create a new Exchange Server self-signed certificate.

4. On the **Friendly name for this certificate** page, enter a descriptive name for the certificate, and then click **Next**.

5. On the **Request a wildcard certificate** page, make one of the following choices:

   - **If you want a wildcard certificate**: Select **Request a wildcard certificate**, and enter the wildcard character (*) and the domain in the **Root domain** field. For example, *.contoso.com or *.eu.contoso.com. When you're finished, click **Next**.

   - **If you want a subject alternative name (SAN) certificate**: Make no selections on this page, and click **Next**.

   - **If you want a certificate for a single host**: Make no selections on this page, and click **Next**.

6. In the **Store certificate request on this server** page, click **Browse** and select the Exchange server where you want to store the certificate request (where you want to install the certificate), click **OK**, and then click **Next**.

   Note: Steps 7 and 8 only apply to a request for a SAN certificate, or a certificate for a single host. If you selected **Request a wildcard certificate**, skip to Step 9.

7. The **Specify the domains you want to be included in your certificate** page is basically a worksheet that helps you to determine the internal and external host names that are required in the certificate for the following Exchange services:

   - Outlook on the web

   - Offline address book generation (OAB)

   - Exchange Web Services

   - Exchange ActiveSync

   - Autodiscover

   - POP

   - IMAP

   - Outlook Anywhere

     If you enter a value for each service based on the location (internal or external), the wizard determines the host names that are required in the certificate, and the information is displayed on the next page. To modify a value for a service, click **Edit** (✏) and enter the host name value that you want to use (or delete the value). When you're finished, click **Next**.

     If you've already determined the host name values that you need in the certificate, you don't need to fill out the information on this page. Instead, click **Next** to manually enter the host names on the next page.

8. The **Based on your selections, the following domains will be included in your certificate** page lists the host names that will be included in the certificate request. The host name that's used in the certificate's **Subject** field is bold, which can be hard to see if that host name is selected. You can verify the

host name entries that are required in the certificate based on the selections that you made on the previous page. Or, you can ignore the values from the last page and add, edit, or remove host name values.

- If you want a SAN certificate, the **Subject** field still requires one common name (CN) value. To select the host name for the certificate's **Subject** field, select the value and click **Set as common name** (check mark). The value should now appear bold.
- If you want a certificate for a single host name, select the other values one at a time and click **Remove** (−).

  Notes:

  - You can't delete the bold host name value that will be used for the certificate's **Subject** field. First, you need to select or add a different host name, and then click **Set as common name** (check mark).
  - The changes that you make on this page might be lost if you click the **Back** button.

9. On the **Specify information about your organization** page, enter the following values:

   - **Organization name**
   - **Department name**
   - **City/Locality**
   - **State/Province**
   - **Country/Region name**

     Note: These X.500 values are included in the certificate's **Subject** field. Although a value is required in every field before you can proceed, the CA might not care about certain fields (for example, **Department name**), while other fields are very important (for example, **Country/Region name** and **Organization name**). Check the **Subject** field requirements of your CA.

     When you're finished, click **Next**.

10. On the **Save the certificate request to the following file** page, enter the UNC path and filename for the certificate request. For example, `\\FileServer01\Data\ExchCertRequest.req`. When you're finished, click **Finish**.

The certificate request appears in the list of Exchange certificates with a status value of **Pending**. For next steps, see the Next steps section.

Use the Exchange Management Shell to create a new certificate request

To create a new certificate request for a wildcard certificate, a SAN certificate, or a certificate for a single host, use the following syntax:

```
New-ExchangeCertificate -GenerateRequest -RequestFile <FilePathOrUNCPath>\<FileName>.req [-FriendlyName
<DescriptiveName>] -SubjectName [C=<CountryOrRegion>,S=<StateOrProvince>,L=<LocalityOrCity>,O=
<Organization>,OU=<Department>],CN=<HostNameOrFQDN> [-DomainName <Host1>,<Host2>...] [-BinaryEncoded <$true |
$false>] [-KeySize <1024 | 2048 | 4096>] [-Server <ServerIdentity>]
```

This example creates a certificate request on the local Exchange server for a wildcard certificate with the following properties:

- **SubjectName**: *.contoso.com in the United States, which requires the value `C=US,CN=*.contoso.com`.

- **RequestFile**: `\\FileServer01\Data\Contoso Wildcard Cert.req`
- **FriendlyName**: Contoso.com Wildcard Cert

```
New-ExchangeCertificate -GenerateRequest -RequestFile "\\FileServer01\Data\Contoso Wildcard Cert.req" -
FriendlyName "Contoso.com Wildcard Cert" -SubjectName "C=US,CN=*.contoso.com"
```

This example creates a certificate request on the local Exchange server for a SAN certificate with the following properties:

- **SubjectName**: mail.contoso.com in the United States, which requires the value `C=US,CN=mail.contoso.com`. Note that this CN value is automatically included in the *DomainName* parameter (the **Subject Alternative Name** field).
- Additional **Subject Alternative Name** field values:
  - autodiscover.contoso.com
  - legacy.contoso.com
  - mail.contoso.net
  - autodiscover.contoso.net
  - legacy.contoso.net
- **RequestFile**: `\\FileServer01\Data\Contoso SAN Cert.req`
- **FriendlyName**: Contoso.com SAN Cert
- **DomainName**: Unquoted comma-separated list of domains

```
New-ExchangeCertificate -GenerateRequest -RequestFile "\\FileServer01\Data\Contoso SAN Cert.req" -FriendlyName
"Contoso.com SAN Cert" -SubjectName "C=US,CN=mail.contoso.com" -DomainName
autodiscover.contoso.com,legacy.contoso.com,mail.contoso.net,autodiscover.contoso.net,legacy.contoso.net
```

This example creates a request for a single subject certificate with the following properties:

- **SubjectName**: mail.contoso.com in the United States, which requires the value `C=US,CN=mail.contoso.com`.
- **RequestFile**: `\\FileServer01\Data\Mail.contoso.com Cert.req`
- **FriendlyName**: Mail.contoso.com Cert

```
New-ExchangeCertificate -GenerateRequest -RequestFile "\\FileServer01\Data\Mail.contoso.com Cert.req" -
FriendlyName "Mail.contoso.com Cert" -SubjectName "C=US,CN=mail.contoso.com"
```

Notes:

- The only required part of the X.500 *SubjectName* parameter value (the certificate's **Subject** field) is `CN=<HostNameOrFQDN>`. However, the certificate request should always include the `C=<CountryOrRegion>` value (otherwise, you might not be able to renew the certificate). Check the **Subject** field requirements of your CA.
- The *RequestFile* parameter accepts a local path or a UNC path.
- We didn't use the *BinaryEncoded* switch, so the request is Base64 encoded. The information that's displayed onscreen is also written to the file, and the contents of the file are what we need to send to the CA. If we had used the *BinaryEncoded* switch, the request would have been encoded by DER, and the certificate request

file itself is what we would need to send to the CA.

- We didn't use the *KeySize* parameter, so the certificate request has a 2048 bit RSA public key.
- For more information, see New-ExchangeCertificate.

How do you know this worked?

To verify that you have successfully created a new certificate request, perform either of the following steps:

- In the EAC at **Servers** > **Certificates**, verify the server where you stored the certificate request is selected. The request should be in the list of certificates with the **Status** value **Pending request**.
- In the Exchange Management Shell on the server where you stored the certificate request, run the following command:

  ```
  Get-ExchangeCertificate | where {$_.Status -eq "PendingRequest" -and $_.IsSelfSigned -eq $false} |
  Format-List FriendlyName,Subject,CertificateDomains,Thumbprint
  ```

Next steps

The content of a Base64 encoded certificate request file looks like this:

```
-----BEGIN NEW CERTIFICATE REQUEST-----
MIIEBjCCAu4CAQAwYzEWMBQGA1UEAwwNKi5jb250b3NvLmNvbTELMAkGA1UECwwC
SVQxEDAOBgNVBAoMB0NvbnRvc28xEDAOBgNVBAcMB1NlYXR0bGUxCzAJBgNVBAgM
AldBMQswCQYDVQQGEwJVUzCCASIwDQYJKoZIhvcNAQEBBQADggEPADCCAQoCggEB
ANZFK6JxcQMEBitJcEC82vCvr6251o28CMmrpIkl7Z0MnkCrU+BMTLBuZnIgaLvb
jlzORvH6DP/dbyR8gQEAHVrXVWdr3AJIRbqQXWwN++BM5b2O6lIrA8w41XwGNu6r
dtddi+POf8UYwot7PXw6wDsbKaTs1ePVK/0XdemdJCFIXNfCT8LY4p/KryQAyquo
XDa+Acbx7TRxG2kXNAxgPGve+mvyCyizbugXAJIz4nugJ2k/X1kGYDc7f/b80tCv
bPTcGCr09ScsbKmsQcqJ7UxiX2tScpO5AQxNxJHGL+bA6+96FBjPnFZaqPbFgI74
N6hmZdSEDgQlaGfLEGjZBGMCAwEAAaCCAVwwGgYKKwYBBAGCNw0CAzEMFgo2LjEu
NzYwMS4yMEwGCSqGSIb3DQEJDjE/MD0wDgYDVR0PAQH/BAQDAgWgMAwGA1UdEwEB
/wQCMAAwHQYDVR0OBBYEFNRw1o74zcuGyky33rl7WChgdQrlMHIGCisGAQQBgjcN
AgIxZDBiAgEBHloATQBpAGMAcgBvAHMAbwBmAHQAIABSAFMAQQAgAFMAQwBoAGEA
bgBuAGUAbAAgAEMAcgB5AHAAdABvAGcAcgBhAHAAaABpAGMAIABQAHIAbwB2AGkA
ZABlAHIDAQAwfAYJKwYBBAGCNxUUMW8wbQIBBQwrRVhIUi0zMjQ4LkVYSFItMzI0
OGRvbS5leHRlc3QubWljcm9zb2Z0LmNvbQwXRVhIUi0zMjQ4RE9NXEVYSFItMzI0
OCQMIk1pY3Jvc29mdC5FeGNoYW5nZS5TZXJ2aWNlSG9zdC5leGUwDQYJKoZIhvcN
AQEFBQADggEBAL63qVj1m2mBz53+nilnlFweOlcltXoxaF28+Kf0hrJVbH5a2Jme
tS0iKU8YXU3mZ3NnWco+5ea024f9awMIzg4z/heE5yEUFf9UtwRGSOc84r2QexPa
zT/rveTTcbliKU0EFhporl3C2uuBCdAewyLj+/k0hABH3djnmMONG6NyC5f+wMun
kkH5naiSLdsTYbq8jkWYuSqL0qdhtmauqWeAPpA0hKDkQk5eDWpOGx3mgxiaQumo
Rqw6dmQ+o8TC+lE3Tvgdfv47A84X8H7Y9h8liS4h0OfbsgEQb8LcM0YHD6yvPgcD
JCmt8A7JFHF9u6mghjiKlXaZ/i+2l10Wsu8=
-----END NEW CERTIFICATE REQUEST-----
```

You need to send this information to the CA. How you send it depends on the CA, but typically, you send the contents of the file in an email message or in the certificate request form on the CA's web site.

If the CA requires a binary certificate request that's encoded by DER (you used the **New-ExchangeCertificate** cmdlet with the *BinaryEncoded* switch), you typically send the whole certificate request file to the CA.

After you receive the certificate from the CA, you need to complete the pending certificate request. For instructions, see Complete a pending Exchange Server certificate request.

Complete a pending Exchange Server certificate request

8/3/2020 • 5 minutes to read • Edit Online

Completing a pending certificate request (also known as a certificate signing request or CSR) is the next step in configuring Transport Layer Security (TLS) encryption in Exchange Server. After you receive the certificate from the certification authority (CA), you install the certificate on the Exchange server to complete the pending certificate request.

You can complete a pending certificate request in the Exchange admin center (EAC) or in the Exchange Management Shell. The procedures are the same for completing new certificate requests or certificate renewal requests. The procedures are also the same for certificates that were issued by an internal CA (for example, Active Directory Certificate Services), or a commercial CA.

You might receive one or more of the following types of certificate files CA:

- **PKCS #12 certificate files**: These are binary certificate files that have .cer, .crt, .der, .p12, or .pfx filename extensions, and require a password when the file contains the private key or chain of trust. The CA might issue you only one binary certificate file that you need to install (protected by a password), or multiple root or intermediate binary certificate files that you also need to install.

- **PKCS #7 certificate files**: These are text certificate files that have .p7b or .p7c filename extensions. These files contain the text: `-----BEGIN CERTIFICATE-----` and `-----END CERTIFICATE-----` or `-----BEGIN PKCS7-----` and `-----END PKCS7-----`. If the CA includes a chain of certificates file with your binary certificate file, you also need to install the chain of certificates file.

What do you need to know before you begin?

- Estimated time to complete: 5 minutes.

- The procedures in this topic require you to have created a new certificate request on the Exchange server, sent the certificate request to the CA, and received the certificate from the CA. For more information, see Create an Exchange Server certificate request for a certification authority.

- In the EAC, you need to retrieve the certificate file from a UNC path (`\\<Server>\<Share>` or `\\<LocalServerName>\c$\`). In the Exchange Management Shell, you can use a local file path.

- If you renew or replace a certificate that was issued by a CA on a subscribed Edge Transport server, you need to remove the old certificate, and then delete and recreate the Edge Subscription. For more information, see Edge Subscription process.

- To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Client Access services security" entry in the Clients and mobile devices permissions topic.

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to create complete a pending certificate request

1. Open the EAC and navigate to **Servers** > **Certificates**.

2. In the **Select server** list, select the Exchange server that holds the pending certificate request.

3. A pending certificate request has the following properties:

   - In the list of certificates, the value of the **Status** field is **Pending request**.

   - When you select the certificate request from the list, there's a **Complete** link in the details pane.

   Select the pending certificate request that you want to complete, and then click **Complete** in the details pane.

4. On the **Complete pending request** page that opens, in the **File to import from** field, enter the UNC path and filename for the certificate file. For example, `\\FileServer01\Data\ContosoCert.cer`. When you're finished, click **OK**.

The certificate request becomes a certificate in the list of Exchange certificates with a **Status** value of **Valid**. For next steps, see the Next steps section.

Use the Exchange Management Shell to complete a pending certificate request

The syntax that you use to complete a pending certificate request in the Exchange Management Shell depends on the type of certificate file or files that you were issued.

To import a binary certificate file (PKCS #12 files that have .cer, .crt, .der, .p12, or .pfx filename extensions), use the following syntax:

```
Import-ExchangeCertificate -FileName "<FilePathOrUNCPath>\<FileName>" [-Password (ConvertTo-SecureString -
String '<Password> ' -AsPlainText -Force)] [-PrivateKeyExportable <$true | $false>] [-Server <ServerIdentity>]
```

This example imports the binary certificate file `\\FileServer01\Data\Contoso Cert.cer` that's protected by the password P@ssw0rd1 on the local Exchange server.

```
Import-ExchangeCertificate -FileName "\\FileServer01\Data\Contoso Cert.cer" -Password (ConvertTo-SecureString
-String 'P@ssw0rd1' -AsPlainText -Force)
```

To import a chain of certificates file (PKCS #7 text files that have .p7b or .p7c filename extensions), use the following syntax:

```
Import-ExchangeCertificate -FileData ([Byte[]](Get-Content -Encoding Byte -Path "<FilePathOrUNCPath>" -
ReadCount 0))]
```

This example imports the text certificate file `\\FileServer01\Data\Chain of Certificates.p7b` on the local Exchange server.

```
Import-ExchangeCertificate -FileData "Import-ExchangeCertificate -FileData ([Byte[]](Get-Content -Encoding
Byte -Path "\\FileServer01\Data\Chain of Certificates.p7b" -ReadCount 0))]
```

Notes:

- The *FileName* and *FileData* parameters accept local paths if the certificate file is located on the Exchange server where you're running the command, and this is the same server where you want to import the certificate. Otherwise, use a UNC path.
- If you want to be able to export the certificate from the server where you're importing it, you need to use the *PrivateKeyExportable* parameter with the value `$true`.
- For more information, see Import-ExchangeCertificate.

How do you know this worked?

To verify that you have successfully completed the certificate request and installed the certificate on the Exchange server, use either of the following procedures:

- In the EAC at **Servers** > **Certificates**, verify the server where you installed the certificate is selected. In the list of certificates, verify that the certificate has **Status** property value **Valid**.
- In the Exchange Management Shell on the server where you installed the certificate, run the following command and verify that the certificate is listed:

  ```
  Get-ExchangeCertificate | where {$_.Status -eq "Valid" -and $_.IsSelfSigned -eq $false} | Format-List
  FriendlyName,Subject,CertificateDomains,Thumbprint
  ```

Next steps

After you complete the pending certificate request by installing the certificate on the server, you need to assign the certificate to one or more Exchange services before the Exchange server is able to use the certificate for encryption. For more information, see Assign certificates to Exchange services.

Assign certificates to Exchange Server services

8/3/2020 • 4 minutes to read • Edit Online

After you install a certificate on an Exchange server, you need to assign the certificate to one or more Exchange services before the Exchange server is able to use the certificate for encryption. You can assign certificates to services in the Exchange admin center (EAC) or in the Exchange Management Shell. Once you assign a certificate to a service, you can't remove the assignment. If you no longer want to use a certificate for a specific service, you need to assign another certificate to the service, and then remove the certificate that you don't want to use.

The available Exchange services are described in the following table.

SERVICE	USES
IIS	TLS encryption for internal and external client connections that use HTTP. This includes: Autodiscover Exchange ActiveSync Exchange admin center Exchange Web Services Offline address book (OAB) distribution Outlook Anywhere (RPC over HTTP) Outlook MAPI over HTTP Outlook on the web
IMAP	TLS encryption for IMAP4 client connections. Don't assign a wildcard certificate to the IMAP4 service. Instead, use the **Set-ImapSettings** cmdlet to configure the fully qualified domain name (FQDN) that clients use to connect to the IMAP4 service.
POP	TLS encryption for POP3 client connections. Don't assign a wildcard certificate to the POP3 service. Instead, use the **Set-PopSettings** cmdlet to configure the FQDN that clients use to connect to the POP3 service.
SMTP	TLS encryption for external SMTP client and server connections. Mutual TLS authentication between Exchange and other messaging servers. When you assign a certificate to SMTP, you're prompted to replace the default Exchange self-signed certificate that's used to encrypt SMTP communication between internal Exchange servers. Typically, you don't need to replace the default SMTP certificate.
Unified Messaging (UM)	TLS encryption for client connections to the backend UM service on Exchange 2016 Mailbox servers. You can only assign a certificate to the UM service when the UM startup mode property of the service is set to TLS or Dual. If the UM startup mode is set to the default value TCP, you can't assign the certificate to the UM service. (**Note**: UM is not available in Exchange 2019). For more information, see Configure the Startup Mode on a Mailbox Server.

SERVICE	USES
Unified Messaging Call Router (UMCallRouter)	TLS encryption for client connections to the UM Call Router service in the Client Access services on Exchange 2016 Mailbox servers. You can only assign a certificate to the UM Call Router service when the UM startup mode property of the service is set to TLS or Dual. If the UM startup mode is set to the default value TCP, you can't assign the certificate to the UM Call Router service. (**Note**: UM is not available in Exchange 2019). For more information, see Configure the Startup Mode on a Client Access Server.

What do you need to know before you begin?

- Estimated time to complete: 5 minutes.

- After you do the procedures in this topic, you might need to restart Internet Information Services (IIS). In some scenarios, Exchange might continue to use the previous certificate for encrypting and decrypting the cookie that's used for Outlook on the web (formerly known as Outlook Web App) authentication. We recommend restarting IIS in environments that use Layer 4 load balancing.

- If you renew or replace a certificate that was issued by a CA on a subscribed Edge Transport server, you need to remove the old certificate, and then delete and recreate the Edge Subscription. For more information, see Edge Subscription process.

- To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Client Access services security" entry in the Clients and mobile devices permissions topic.

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to assign a certificate to Exchange services

1. Open the EAC, and navigate to **Servers** > **Certificates**.

2. In the **Select server** list, select the Exchange server that holds the certificate.

3. Select the certificate that you want to configure, and then click **Edit** ✎. The certificate needs to have the **Status** value **Valid**.

4. On the **Services** tab, in the **Specify the services you want to assign this certificate to** section, select the services. Remember, you can add services, but you can't remove them. When you're finished, click **Save**.

Use the Exchange Management Shell to assign a certificate to Exchange services

To assign a certificate to Exchange services, use the following syntax:

```
Enable-ExchangeCertificate -Thumbprint <Thumbprint> -Services <Service1>,<Service2>... [-Server
<ServerIdentity>]
```

This example assigns the certificate that has the thumbprint value `434AC224C8459924B26521298CE8834C514856AB` to the POP, IMAP, IIS, and SMTP services.

```
Enable-ExchangeCertificate -Thumbprint 434AC224C8459924B26521298CE8834C514856AB -Services POP,IMAP,IIS,SMTP
```

You can find the certificate thumbprint value by using the **Get-ExchangeCertificate** cmdlet.

How do you know this worked?

To verify that you have successfully assigned a certificate to one or more Exchange services, use either of the following procedures:

- In the EAC at **Servers** > **Certificates**, verify the server where you installed the certificate is selected. Select the certificate, and in the details pane, verify that the **Assigned to services property** contains the services that you selected.
- In the Exchange Management Shell on the server where you installed the certificate, run the following command to verify the Exchange services for the certificate:

  ```
  Get-ExchangeCertificate | Format-List FriendlyName,Subject,CertificateDomains,Thumbprint,Services
  ```

Create a new Exchange Server self-signed certificate

8/3/2020 • 5 minutes to read • Edit Online

When you install Exchange Server, a self-signed certificate that's created and signed by the Exchange server itself is automatically installed on the server. However, you can also create additional self-signed certificates that you can use.

You can create self-signed certificates certificate in the Exchange admin center (EAC) or in the Exchange Management Shell.

What do you need to know before you begin?

- Estimated time to complete: 5 minutes.
- Exchange self-signed certificates work well for encrypting communication between internal Exchange servers, but not so well for encrypting external connections, because clients, servers, and services don't automatically trust Exchange self-signed certificates. To create a certificate request (also known as a certificate signing request or CSR) for a commercial certification authority that's automatically trusted by all clients, servers, and services, see Create an Exchange Server certificate request for a certification authority.
- When you create a new self-signed certificate by using the **New-ExchangeCertificate** cmdlet, you can assign the certificate to Exchange services during the creation of the certificate. For more information about the Exchange services, see Assign certificates to Exchange Server services.
- To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Client Access services security" entry in the Clients and mobile devices permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to create a new Exchange self-signed certificate

1. Open the EAC and navigate to **Servers** > **Certificates**.
2. In the **Select server** list, select the Exchange server where you want to install the certificate, and then click **Add** ➕.
3. The **New Exchange certificate** wizard opens. On the **This wizard will create a new certificate or a certificate request file** page, select **Create a self-signed certificate**, and then click **Next**.

   Note: To create a new certificate request for a certificate authority, see Create an Exchange Server certificate request for a certification authority.
4. On the **Friendly name for this certificate** page, enter a friendly name for the certificate, and then click

Next.

5. In the **Specify the servers you want to apply this certificate to** page, click **Add** +

   On the **Select a server** page that opens, select the Exchange server where you want to install the certificate, and click **Add - >**. Repeat this step as many times as necessary. When you're finished selecting servers, click **OK**.

   When you're finished, click **Next**.

6. The **Specify the domains you want to be included in your certificate** page is basically a worksheet that helps you determine the internal and external host names that are required in the certificate for the following Exchange services:

   - Outlook on the web
   - Offline address book generation (OAB)
   - Exchange Web Services
   - Exchange ActiveSync
   - Autodiscover
   - POP
   - IMAP
   - Outlook Anywhere

     If you enter a value for each service based on the location (internal or external), the wizard determines the host names that are required in the certificate, and the information is displayed on the next page. To modify a value for a service, click **Edit** (✎) and enter the host name value that you want to use (or delete the value). When you're finished, click **Next**.

     If you've already determined the host name values that you need in the certificate, you don't need to fill out the information on this page. Instead, click **Next** to manually enter the host names on the next page.

7. The **Based on your selections, the following domains will be included in your certificate** page lists the host names that will be included in the self-signed certificate. The host name that's used in the certificate's **Subject** field is bold, which can be hard to see if that host name is selected. You can verify the host name entries that are required in the certificate based on the selections that you made on the previous page. Or, you can ignore the values from the last page and add, edit, or remove host name values.

   - If you want a SAN certificate, the **Subject** field still requires one common name (CN) value. To select the host name for the certificate's **Subject** field, select the value and click **Set as common name** (check mark). The value should now appear bold.
   - If you want a certificate for a single host name, select the other values one at a time and click **Remove** (−).

     When you're finished on this page, click **Finish**.

     Notes:

   - You can't delete the bold host name value that will be used for the certificate's **Subject** field. First, you need to select or add a different host name, and then click **Set as common name** (check mark).
   - The changes that you make on this page might be lost if you click the **Back** button.

Use the Exchange Management Shell to create a new Exchange self-signed certificate

To create a new Exchange self-signed certificate, use the following syntax:

```
New-ExchangeCertificate [-FriendlyName <DescriptiveName>] [-SubjectName [C=<CountryOrRegion>,S=<StateOrProvince>,L=<LocalityOrCity>,O=<Organization>,OU=<Department>],CN=<HostNameOrFQDN>]] [-DomainName <Host1>,<Host2>...] [-Services <None | IIS | IMAP | POP | SMTP | UM | UMCallRouter> [-PrivateKeyExportable <$true | $false>] [-Server <ServerIdentity>] -[Force]
```

This example creates a self-signed certificate on the local Exchange server with the following properties:

- **Subject**: *<ServerName>*. For example, if you run the command on the server named Mailbox01, the value is `Mailbox01`.
- **Subject alternative names**: *<ServerName>*, <Server FQDN>. For example, `Mailbox01, Mailbox01.contoso.com`.
- **Friendly name**: Microsoft Exchange
- **Services**: POP, IMAP, SMTP.

```
New-ExchangeCertificate
```

This example creates a creates a self-signed certificate on the local Exchange server with the following properties:

- **Subject**: Exchange01, which requires the value `CN=Exchange01`. Note that this value is automatically included in the *DomainName* parameter (the **Subject Alternative Name** field).
- Additional subject alternative names:
  - mail.contoso.com
  - autodiscover.contoso.com
  - Exchange01.contoso.com
  - Exchange02.contoso.com
- **Services**: SMTP, IIS
- **Friendly name**: Contoso Exchange Certificate
- The private key is exportable. This allows you to export the certificate from the server (and import it on other servers).

```
New-ExchangeCertificate -FriendlyName "Contoso Exchange Certificate" -SubjectName CN=Exchange01 -DomainName mail.contoso.com,autodiscover.contoso.com,Exchange01.contoso.com,Exchange02.contoso.com -Services SMTP,IIS -PrivateKeyExportable $true
```

Notes:

- The only required part of the X.500 *SubjectName* parameter value (the certificate's **Subject** field) is `CN=<HostNameOrFQDN>`.
- Some *Services* parameter values generate warning or confirmation messages. For more information, see Assign certificates to Exchange Server services.

- For more information, see New-ExchangeCertificate.

How do you know this worked?

To verify that you have successfully created an Exchange self-signed certificate, perform either of the following steps:

- In the EAC at **Servers** > **Certificates**, verify the server where you created the self-signed certificate is selected. The certificate should be in the list of certificates with the **Status** value **Valid**.
- In the Exchange Management Shell on the server where you created the self-signed certificate, run the following command and verify the properties:

```
Get-ExchangeCertificate | where {$_.Status -eq "Valid" -and $_.IsSelfSigned -eq $true} | Format-List
FriendlyName,Subject,CertificateDomains,Thumbprint,NotBefore,NotAfter
```

Renew an Exchange Server certificate

8/3/2020 • 6 minutes to read • Edit Online

Every certificate has a built-in expiration date. In Exchange Server, the default self-signed certificate that's installed on the Exchange server expires 5 years after Exchange was installed on the server. You can use the Exchange admin center (EAC) or the Exchange Management Shell to renew Exchange certificates. This includes Exchange self-signed certificates, and certificates that were issued by a certification authority (CA).

What do you need to know before you begin?

- Estimated time to complete: 5 minutes
- To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- For certificates that were issued by a CA, verify the certificate request requirements of the CA. Exchange generates a PKCS #10 request (.req) file that uses Base64 encoding (default) or Distinguished Encoding Rules (DER), with an RSA public key that's 1024, 2048 (default), or 4096 bits. Note that encoding and public key options are only available in the Exchange Management Shell.
- To renew a certificate that was issued by a CA, you need to renew the certificate with the same CA that issued the certificate. If you're changing CAs, or if there's a problem with the original certificate when you try to renew it, you need to create a new certificate request (also known as a certificate signing request or CSR) for a new certificate. For more information, see Create an Exchange Server certificate request for a certification authority.
- If you renew or replace a certificate that was issued by a CA on a subscribed Edge Transport server, you need to remove the old certificate, and then delete and recreate the Edge Subscription. For more information, see Edge Subscription process.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Client Access services security" entry in the Clients and mobile devices permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Renew a certificate that was issued by a certification authority

The procedures are the same for certificates that were issued by an internal CA (for example, Active Directory Certificate Services), or a commercial CA.

To renew a certificate that was issued by a CA, you create a certificate renewal request, and then you send the request to the CA. The CA then sends you the actual certificate file that you need to install on the Exchange server. The procedure is nearly identical to that of completing a new certificate request by installing the certificate on the server. For instructions, see Complete a pending Exchange Server certificate request.

Use the EAC to create a certificate renewal request for a certification authority

1. Open the EAC and navigate to **Servers** > **Certificates**.

2. In the **Select server** list, select the Exchange server that holds the certificate that you want to renew.

3. All valid certificates have a **Renew** link in the details pane that's visible when you select the certificate from the list. Select the certificate that you want to renew, and then click **Renew** in the details pane.

4. On the **Renew Exchange certificate** page that opens, in the **Save the certificate request to the following file** field, enter the UNC path and filename for the new certificate renewal request file. For example, `\\FileServer01\Data\ContosoCertRenewal.req`. When you're finished, click **OK**.

The certificate request appears in the list of Exchange certificates with a status value of **Pending**.

Use the Exchange Management Shell to create a certificate renewal request for a certification authority

To create a certificate renewal request for a certification authority on the local Exchange server, use the following syntax:

```
Get-ExchangeCertificate -Thumbprint <Thumbprint> | New-ExchangeCertificate -GenerateRequest -RequestFile
<FilePathOrUNCPath>\<FileName>.req
```

To find the thumbprint value of the certificate that you want to renew, run the following command:

```
Get-ExchangeCertificate | where {$_.Status -eq "Valid" -and $_.IsSelfSigned -eq $false} | Format-List
FriendlyName,Subject,CertificateDomains,Thumbprint,NotBefore,NotAfter
```

This example creates a certificate renewal request with the following properties:

- **Certificate to renew**: `5DB9879E38E36BCB60B761E29794392B23D1C054`

- **RequestFile**: `\\FileServer01\Data\ContosoCertRenewal.req`

```
Get-ExchangeCertificate -Thumbprint 5DB9879E38E36BCB60B761E29794392B23D1C054 | New-ExchangeCertificate -
GenerateRequest -RequestFile \\FileServer01\Data\ContosoCertRenewal.req
```

Notes:

- The *RequestFile* parameter accepts a local path or a UNC path.

- We didn't use the *BinaryEncoded* switch, so the request is Base64 encoded. The information that's displayed onscreen is also written to the file, and the contents of the file are what we need to send to the CA. If we had used the *BinaryEncoded* switch, the request would have been encoded by DER, and the certificate request file itself is what we would need to send to the CA.

- We didn't use the *KeySize* parameter, so the certificate request has a 2048 bit RSA public key.

- For more information, see Get-ExchangeCertificate and New-ExchangeCertificate.

How do you know this worked?

To verify that you have successfully created a certificate renewal request for a certification authority, perform either of the following steps:

- In the EAC at **Servers** > **Certificates**, verify the server where you stored the certificate request is selected. The request should be in the list of certificates with the **Status** value **Pending request**.

- In the Exchange Management Shell on the server where you stored the certificate request, run the following command:

```
Get-ExchangeCertificate | where {$_.Status -eq "PendingRequest" -and $_.IsSelfSigned -eq $false} |
Format-List FriendlyName,Subject,CertificateDomains,Thumbprint
```

Renew an Exchange self-signed certificate

When you renew an Exchange self-signed certificate, you're basically making a new certificate.

Use the EAC to renew an Exchange self-signed certificate

1. Open the EAC and navigate to **Servers** > **Certificates**.

2. In the **Select server** list, select the Exchange server that holds the certificate that you want to renew.

3. All valid certificates have a **Renew** link in the details pane that's visible when you select the certificate from the list. Select the certificate that you want to renew, and then click **Renew** in the details pane.

4. On the **Renew Exchange certificate** page that opens, verify the read-only list of Exchange services that the existing certificate is assigned to, and then click **OK**.

Use the Exchange Management Shell to renew an Exchange self-signed certificate

To renew a self-signed certificate, use the following syntax:

```
Get-ExchangeCertificate -Thumbprint <Thumbprint> | New-ExchangeCertificate [-Force] [-PrivateKeyExportable
<$true | $false>]
```

To find the thumbprint value of the certificate that you want to renew, run the following command:

```
Get-ExchangeCertificate | where {$_.IsSelfSigned -eq $true} | Format-List
FriendlyName,Subject,CertificateDomains,Thumbprint,NotBefore,NotAfter
```

This example renews a self-signed certificate on the local Exchange server, and uses the following settings:

- The thumbprint value of the existing self-signed certificate to renew is `BC37CBE2E59566BFF7D01FEAC9B6517841475F2D`

- The *Force* switch replaces the original self-signed certificate without a confirmation prompt.

- The private key is exportable. This allows you to export the certificate and import it on other servers.

```
Get-ExchangeCertificate -Thumbprint BC37CBE2E59566BFF7D01FEAC9B6517841475F2D | New-ExchangeCertificate -Force
-PrivateKeyExportable $true
```

How do you know this worked?

To verify that you have successfully renewed an Exchange self-signed certificate, use either of the following procedures:

- In the EAC at **Servers** > **Certificates**, verify the server where you installed the certificate is selected. In the list of certificates, verify that the certificate has **Status** property value **Valid**.

- In the Exchange Management Shell on the server where you renewed the self-signed certificate, run the following command to verify the property values:

```
Get-ExchangeCertificate | where {$_.Status -eq "Valid" -and $_.IsSelfSigned -eq $true} | Format-List
FriendlyName,Subject,CertificateDomains,Thumbprint,NotBefore,NotAfter
```

Export a certificate from an Exchange server

8/3/2020 • 2 minutes to read • Edit Online

You can export a certificate from an Exchange server as a backup or to import the certificate on other clients, devices or servers. You can export certificates in the Exchange admin center (EAC) or in the Exchange Management Shell. The resulting certificate file is a password-protected binary PKCS #12 file that contains the certificate's private key, and is suitable for importing (installing) on other servers.

What do you need to know before you begin?

- Estimated time to complete: 5 minutes.
- In the EAC, you need to export the certificate file to a UNC path (`\\<Server>\<Share>\` or `\\<LocalServerName>\c$\`). In the Exchange Management Shell, you can specify a local path.
- To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Client Access services security" entry in the Clients and mobile devices permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to export a certificate

1. Open the EAC and navigate to **Servers** > **Certificates**.
2. In the **Select server** list, select the Exchange server that contains the certificate, click **More options** •••, and select **Export Exchange certificate**.
3. On the **Export Exchange certificate** page that opens, enter the following information:
   - **File to export to**: Enter the UNC path and file name of the certificate file. For example, `\\FileServer01\Data\Fabrikam.pfx`
   - **Password**: When you export the certificate with its private key, you need to specify a password. Exporting the certificate with its private key allows you to import the certificate on other servers.

   When you're finished, click **OK**.

Use the Exchange Management Shell to export a certificate

To export a binary certificate file that you can import on other clients or servers, use the following syntax:

```
Export-ExchangeCertificate -Thumbprint <Thumbprint> -FileName "<FilePathOrUNCPath>\<FileName>.pfx" -
BinaryEncoded -Password (ConvertTo-SecureString -String '<Password> ' -AsPlainText -Force) [-Server
<ServerIdentity>]
```

This example exports a certificate from the local Exchange server to a file with the following settings:

- The certificate that has the thumbprint value `5113ae0233a72fccb75b1d0198628675333d010e` is exported to the file `C:\Data\Fabrikam.pfx`.
- The exported certificate file is encoded by DER (not Base64).
- The password for the certificate file is `P@ssw0rd1`.

```
Export-ExchangeCertificate -Thumbprint 5113ae0233a72fccb75b1d0198628675333d010e -FileName
"C:\Data\Fabrikam.pfx" -BinaryEncoded -Password (ConvertTo-SecureString -String 'P@ssw0rd1' -AsPlainText -
Force)
```

Notes:

- The *FileName* parameter accepts a local path or a UNC path.
- You can use a similar procedure to export a pending certificate request (also known as a certificate signing request or CSR). For example, if you need to resubmit the certificate request to the certification authority, and you can't find the original certificate request file. When you export a certificate request, you typically don't need to use the *Password* parameter or the *BinaryEncoded* switch, and you save the request to a .req file. Note that you can't import an exported certificate request on another server.
- For more information, see Export-ExchangeCertificate.

How do you know this worked?

To verify that you have successfully exported a certificate from an Exchange server, try importing the certificate file on another server. For more information, see Import or install a certificate on an Exchange server.

Import or install a certificate on an Exchange server

8/3/2020 • 4 minutes to read • Edit Online

To enable encryption for one or more Exchange services, the Exchange server needs to use a certificate. SMTP communication between internal Exchange servers is encrypted by the default self-signed certificate that's installed on the Exchange server. To encrypt communication with internal or external clients, servers, or services, you'll likely want to use a certificate that's automatically trusted by all clients, services and servers that connect to your Exchange organization. For more information, see Certificate requirements for Exchange services.

You can import (install) certificates on Exchange servers in the Exchange admin center (EAC) or in the Exchange Management Shell.

These are the types of certificate files that you can import on an Exchange server:

- **PKCS #12 certificate files**: These are binary certificate files that have .cer, .crt, .der, .p12, or .pfx filename extensions, and require a password when the file contains the private key or chain of trust. Examples of these types of files include:
  - Self-signed certificates that were exported from other Exchange servers by using the EAC or the **Export-ExchangeCertificate** with the *PrivateKeyExportable* parameter value `$true`. For more information, see Export a certificate from an Exchange server.
  - Certificates that were issued by a certification authority (an internal CA like Active Directory Certificate Services, or a commercial CA).
  - Certificates that were exported from other servers (for example, Skype for Business Server).
- **PKCS #7 certificate files**: These are text certificate files that have .p7b or .p7c filename extensions. These files contain the text: `-----BEGIN CERTIFICATE-----` and `-----END CERTIFICATE-----` or `-----BEGIN PKCS7-----` and `-----END PKCS7-----`. A certificate authority might include a chain of certificates file that also needs to be installed along with the actual binary certificate file.

What do you need to know before you begin?

- Estimated time to complete: 5 minutes.
- In the EAC, you need to import the certificate file from a UNC path (`\\<Server>\<Share>\` or `\\<LocalServerName>\c$\`). In the Exchange Management Shell, you can specify a local path.
- In the EAC, you can import the certificate file on multiple Exchange servers at the same time (Step 4 in the procedure).
- To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Client Access services security" entry in the Clients and mobile devices permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to import a certificate on one or more Exchange servers

1. Open the EAC and navigate to **Servers** > **Certificates**.

2. In the **Select server** list, select the Exchange server where you want to install the certificate, click **More options** •••, and select **Import Exchange certificate**.

3. The **Import Exchange certificate** wizard opens. On the **This wizard will import a certificate from a file** page, enter the following information:

   - **File to import from**: Enter the UNC path and filename of the certificate file. For example, `\\FileServer01\Data\Fabrikam.cer`

   - **Password**: If the certificate file contains the private key or chain of trust, the file is protected by a password. Enter the password here.

   When you're finished, click **Next**.

4. In the **Specify the servers you want to apply this certificate to** page, click **Add** +

   On the **Select a server** page that opens, select the Exchange server where you want to install the certificate, and click **Add ->**. Repeat this step as many times as necessary. When you're finished selecting servers, click **OK**.

   When you're finished, click **Finish**. For next steps, see the Next steps section.

Use the Exchange Management Shell to import a certificate on an Exchange server

To import a binary certificate file (PKCS #12 files that have .cer, .crt, .der, .p12, or .pfx filename extensions), use the following syntax:

```
Import-ExchangeCertificate -FileName "<FilePathOrUNCPath>\<FileName>" -Password (ConvertTo-SecureString -
String '<Password> ' -AsPlainText -Force) [-PrivateKeyExportable <$true | $false>] [-Server <ServerIdentity>]
```

This example imports the certificate file `\\FileServer01\Data\Fabrikam.pfx` that's protected by the password P@ssw0rd1 on the local Exchange server.

```
Import-ExchangeCertificate -FileName "\\FileServer01\Data\Fabrikam.pfx" -Password (ConvertTo-SecureString -
String 'P@ssw0rd1' -AsPlainText -Force)
```

To import a chain of certificates file (PKCS #7 text files that have .p7b or .p7c filename extensions) that's associated with a certificate, use the following syntax:

```
Import-ExchangeCertificate -FileData ([Byte[]](Get-Content -Encoding Byte -Path "<FilePathOrUNCPath>" -
ReadCount 0))
```

This example imports the chain of certificates file `\\FileServer01\Data\Chain of Certificates.p7b`.

```
Import-ExchangeCertificate -FileData ([Byte[]](Get-Content -Encoding Byte -Path "\\FileServer01\Data\Chain of
Certificates.p7b" -ReadCount 0))
```

Notes:

- You need to repeat this procedure on each Exchange server where you want to import the certificate (run the command on the server, or use the *Server* parameter).

- The *FileName* and *FileData* parameters accept local paths if the certificate file is located on the Exchange server where you're running the command, and this is the same server where you want to import the certificate. Otherwise, use a UNC path.

- If you want to be able to export the certificate from the server where you're importing it, you need to use the *PrivateKeyExportable* parameter with the value `$true`.

- For more information, see Import-ExchangeCertificate.

How do you know this worked?

To verify that you have successfully imported (installed) a certificate on an Exchange server, use either of the following procedures:

- In the EAC at **Servers** > **Certificates**, verify the server where you installed the certificate is selected. The certificate should be in the list of certificates with the **Status** value **Valid**.

- In the Exchange Management Shell on the server where you installed the certificate, run the following command:

  ```
  Get-ExchangeCertificate | where {$_.Status -eq "Valid"} | Format-List
  FriendlyName,Subject,CertificateDomains,Thumbprint,NotBefore,NotAfter
  ```

Next steps

After you install the certificate on the server, you need to assign the certificate to one or more Exchange services before the Exchange server is able to use the certificate for encryption. For more information, see Assign certificates to Exchange Server services.

Configure client-specific message size limits

8/3/2020 • 7 minutes to read • Edit Online

In Exchange Server, there are several different message size limits that apply to messages as they travel through your organization. For more information, see Message size and recipient limits in Exchange Server.

However, there are client-specific message size limits you can configure for Outlook on the web (fornerly known as Outlook Web App) and email clients that use Exchange ActiveSync or Exchange Web Services (EWS). If you change the Exchange organizational, connector, or user message size limits, you likely need change the limits for Outlook on the web, ActiveSync, and EWS. These limits are described in the following tables. To change the message size limit for a specific client type, you need to change **all** the values that are described in the table.

> **NOTE**
>
> For any message size limit, you need to set a value that's larger than the actual size you want enforced. This accounts for the Base64 encoding of attachments and other binary data. Base64 encoding increases the size of the message by approximately 33%, so the value you specify should be approximately 33% larger than the actual message size you want enforced. For example, if you specify a maximum message size value of 64 MB, you can expect a realistic maximum message size of approximately 48 MB.

ActiveSync

SERVICES	CONFIGURATION FILE	KEYS AND DEFAULT VALUES	SIZE
Client Access (frontend)	`%ExchangeInstallPath%FrontEn`	`maxAllowedContentLength="3000[illegible]"` (not present by default; see comments)	bytes
Client Access (frontend)	`%ExchangeInstallPath%FrontEn`	`maxRequestLength="10240"`	kilobytes
Backend	`%ExchangeInstallPath%ClientA`	`maxAllowedContentLength="3000[illegible] bytes"` (not present by default; see comments)	bytes
Backend	`%ExchangeInstallPath%ClientA`	`maxRequestLength="10240"`	kilobytes
Backend	`%ExchangeInstallPath%ClientA`	`<add key="MaxDocumentDataSize" value="10240000">`	bytes

Comments on ActiveSync limits

By default, there is no *maxAllowedContentLength* key in the `web.config` files for ActiveSync. However, the maximum message size for ActiveSync is affected by the **maxAllowedContentLength** value that is applied to all web sites on the server. The default value is 30000000 bytes. To see these values for ActiveSync on Mailbox servers in IIS Manager, perform the following steps:

1. Do one of the following steps:

   - For the Client Access (frontend) web site, open **IIS Manager**, navigate to **Sites** > **Default Web Site** and select **Microsoft-Server-ActiveSync**.

- For the backend web site, open **IIS Manager**, navigate to **Sites** > **Exchange Back End** and select **Microsoft-Server-ActiveSync**.

2. Verify the **Features View** tab is selected at the bottom, and double-click **Configuration Editor** in the **Management** section.

3. Click the drop down arrow in the **Section** field, navigate to **system.webServer** > **security** and select **requestFiltering**.

4. In the results, expand **requestLimits**, and you'll see **maxAllowedContentLength** and the default value 30000000 (bytes).

To change the **maxAllowedContentLength** value, enter a new value in bytes, and click **Apply**. You need to change the value on the Client Access web site and the back end web site.

Note: You can change the same setting in IIS manager at **Sites** > **Default Web Site** > **Microsoft-Server-ActiveSync** or **Sites** > **Exchange Back End** > **Microsoft-Server-ActiveSync** and then **Request Filtering** in the **IIS** section > **Edit Feature Settings** in the **Actions** area > **Maximum allowed content length (Bytes)** in the **Request Limits** section.

After you change the value in IIS Manager, a new *maxAllowedContentLength* key is written to the corresponding Client Access or backend web.config file that's described in the table.

Exchange Web Services

SERVICE	CONFIGURATION FILE	KEYS AND DEFAULT VALUES	SIZE
Client Access (frontend)	`%ExchangeInstallPath%FrontEnd`	`maxAllowedContentLength="67108864"`	bytes
Backend	`%ExchangeInstallPath%ClientAc`	`maxAllowedContentLength="67108864"`	bytes
Backend	`%ExchangeInstallPath%ClientAccess\exchweb\ews\web.config`	14 instances of `maxReceivedMessageSize="67108864"` (for different combinations of http/https bindings and authentication methods)	bytes

Comments on EWS limits

- In the backend `web.config` file, there are two instances of the value `maxReceivedMessageSize="1048576"` for **UMLegacyMessageEncoderSoap11Element** bindings that you don't need to modify.

- *maxRequestLength* is an ASP.NET setting that's present in both web.config files, but isn't used by EWS, so you don't need to modify it.

Outlook on the web

SERVICE	CONFIGURATION FILE	KEYS AND DEFAULT VALUES	SIZE
Client Access (frontend)	`%ExchangeInstallPath%FrontEnd`	`maxAllowedContentLength="35000000"`	bytes
Client Access (frontend)	`%ExchangeInstallPath%FrontEnd`	`maxRequestLength="35000"`	kilobytes
Backend	`%ExchangeInstallPath%ClientAc`	`maxAllowedContentLength="35000000"`	bytes

SERVICE	CONFIGURATION FILE	KEYS AND DEFAULT VALUES	SIZE
Backend	`%ExchangeInstallPath%ClientAc`	`maxRequestLength="35000"`	kilobytes
Backend	`%ExchangeInstallPath%ClientAccess\Owa\web.config`	2 instances of `maxReceivedMessageSize="35000000"` (for http and https bindings)	bytes
Backend	`%ExchangeInstallPath%ClientAccess\Owa\web.config`	2 instances of `maxStringContentLength="35000000"` (for http and https bindings)	bytes

Comments on Outlook on the web limits

- In the backend `web.config` file, there's an instance of the value `maxStringContentLength="102400"` for the **MsOnlineShellService** binding that you don't need to modify.

What do you need to know before you begin?

- Estimated time to complete: 15 minutes

- Exchange permissions don't apply to the procedures in this topic. These procedures are performed in the operating system of the Exchange server.

- Changes you save to the web.config configuration file are applied after you restart IIS.

- To allow for the 33% increase in size due to Base64 encoding, multiply your desired new maximum size value in megabytes by 4/3. To convert the value into kilobytes, multiply by 1024. To convert the value into bytes, multiply by 1048756 (1024*1024). Note that the size increase caused by Base64 encoding could be greater than 33%, and depends on several factors (for example, the attachment size, file type, compression, and the email client).

- Any customized Exchange or Internet Information Server (IIS) settings that you made in Exchange XML application configuration files on the Exchange server (for example, web.config files or the EdgeTransport.exe.config file) **will be overwritten** when you install an Exchange CU. Be sure save this information so you can easily re-apply the settings after the install. After you install the Exchange CU, you need to re-configure these settings.

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Use Notepad to configure a client-specific message size limit

1. Open the appropriate web.config files in Notepad. For example, to open the web.config files for EWS clients, run the following commands:

```
Notepad %ExchangeInstallPath%ClientAccess\exchweb\ews\web.config
```

```
Notepad %ExchangeInstallPath%FrontEnd\HttpProxy\ews\web.config
```

2. Find the relevant keys in the appropriate web.config files as described in the tables earlier in the topic. For example, for EWS clients, find the *maxAllowedContentLength* key in the Client Access and backend web.config files and all 14 instances of the value `maxReceivedMessageSize="67108864"` in the backend web.config file.

   ```
   <requestLimits maxAllowedContentLength="67108864" />
   ...maxReceivedMessageSize="67108864"...
   ```

   For example, to allow a Base64 encoded maximum message size of approximately 64 MB, change all instances of `67108864` to `89478486` (64*4/3*1048756):

   ```
   <requestLimits maxAllowedContentLength="89478486" />
   ...maxReceivedMessageSize="89478486"...
   ```

3. When you're finished, save and close the web.config files.

4. Restart IIS on the Exchange server by using either of the following methods:

   - Open IIS Manager, select the server, and in the **Actions** pane, click **Restart**.

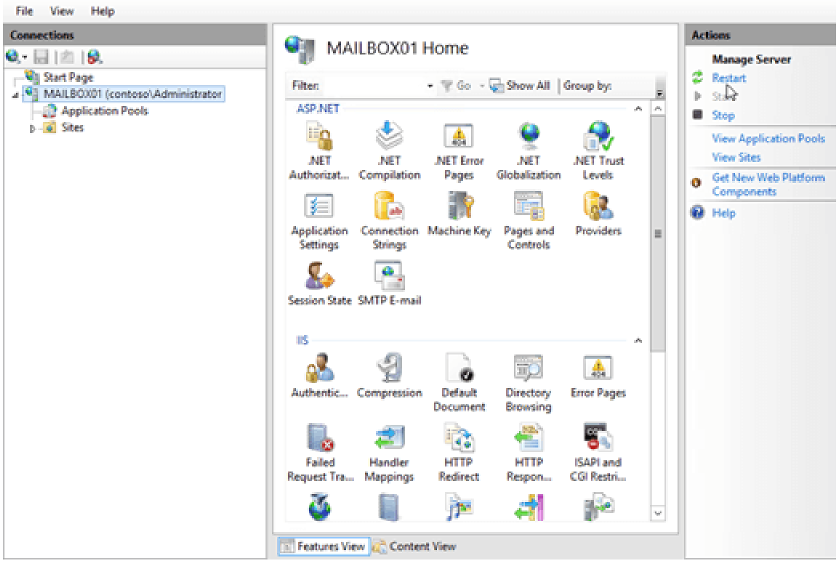

   - Run the following commands from an elevated command prompt (a Command Prompt window you open by selecting **Run as administrator**):

     ```
     net stop w3svc /y
     ```

     ```
     net start w3svc
     ```

Configure client-specific message size limits from the command line

Instead of using Notepad, you can also configure the client-specific message size limits from the command line. Open an elevated command prompt on the Exchange server (a Command Prompt window you open by selecting **Run as administrator**) and run the appropriate commands for the limits that you want to configure.

NOTE

- The size values in the commands are the default values, so you'll need to change them.
- Pay attention to whether the value is in bytes or kilobytes.

ActiveSync

```
%windir%\system32\inetsrv\appcmd.exe set config "Default Web Site/Microsoft-Server-ActiveSync/" -
section:system.webServer/security/requestFiltering /requestLimits.maxAllowedContentLength:30000000
%windir%\system32\inetsrv\appcmd.exe set config "Default Web Site/Microsoft-Server-ActiveSync/" -
section:system.web/httpRuntime /maxRequestLength:10240
%windir%\system32\inetsrv\appcmd.exe set config "Exchange Back End/Microsoft-Server-ActiveSync/" -
section:system.webServer/security/requestFiltering /requestLimits.maxAllowedContentLength:30000000
%windir%\system32\inetsrv\appcmd.exe set config "Exchange Back End/Microsoft-Server-ActiveSync/" -
section:system.web/httpRuntime /maxRequestLength:10240
%windir%\system32\inetsrv\appcmd.exe set config "Exchange Back End/Microsoft-Server-ActiveSync/" -
section:appSettings /[key='MaxDocumentDataSize'].value:10240000
```

Exchange Web Services

```
%windir%\system32\inetsrv\appcmd.exe set config "Default Web Site/ews/" -
section:system.webServer/security/requestFiltering /requestLimits.maxAllowedContentLength:67108864
%windir%\system32\inetsrv\appcmd.exe set config "Exchange Back End/ews/" -
section:system.webServer/security/requestFiltering /requestLimits.maxAllowedContentLength:67108864
%windir%\system32\inetsrv\appcmd.exe set config "Exchange Back End/ews/" -section:system.serviceModel/bindings
/customBinding.[name='EWSAnonymousHttpsBinding'].httpsTransport.maxReceivedMessageSize:67108864
%windir%\system32\inetsrv\appcmd.exe set config "Exchange Back End/ews/" -section:system.serviceModel/bindings
/customBinding.[name='EWSAnonymousHttpBinding'].httpTransport.maxReceivedMessageSize:67108864
%windir%\system32\inetsrv\appcmd.exe set config "Exchange Back End/ews/" -section:system.serviceModel/bindings
/customBinding.[name='EWSBasicHttpsBinding'].httpsTransport.maxReceivedMessageSize:67108864
%windir%\system32\inetsrv\appcmd.exe set config "Exchange Back End/ews/" -section:system.serviceModel/bindings
/customBinding.[name='EWSBasicHttpBinding'].httpTransport.maxReceivedMessageSize:67108864
%windir%\system32\inetsrv\appcmd.exe set config "Exchange Back End/ews/" -section:system.serviceModel/bindings
/customBinding.[name='EWSNegotiateHttpsBinding'].httpsTransport.maxReceivedMessageSize:67108864
%windir%\system32\inetsrv\appcmd.exe set config "Exchange Back End/ews/" -section:system.serviceModel/bindings
/customBinding.[name='EWSNegotiateHttpBinding'].httpTransport.maxReceivedMessageSize:67108864
%windir%\system32\inetsrv\appcmd.exe set config "Exchange Back End/ews/" -section:system.serviceModel/bindings
/customBinding.[name='EWSWSSecurityHttpsBinding'].httpsTransport.maxReceivedMessageSize:67108864
%windir%\system32\inetsrv\appcmd.exe set config "Exchange Back End/ews/" -section:system.serviceModel/bindings
/customBinding.[name='EWSWSSecurityHttpBinding'].httpTransport.maxReceivedMessageSize:67108864
%windir%\system32\inetsrv\appcmd.exe set config "Exchange Back End/ews/" -section:system.serviceModel/bindings
/customBinding.[name='EWSWSSecuritySymmetricKeyHttpsBinding'].httpsTransport.maxReceivedMessageSize:67108864
%windir%\system32\inetsrv\appcmd.exe set config "Exchange Back End/ews/" -section:system.serviceModel/bindings
/customBinding.[name='EWSWSSecuritySymmetricKeyHttpBinding'].httpTransport.maxReceivedMessageSize:67108864
%windir%\system32\inetsrv\appcmd.exe set config "Exchange Back End/ews/" -section:system.serviceModel/bindings
/customBinding.[name='EWSWSSecurityX509CertHttpsBinding'].httpsTransport.maxReceivedMessageSize:67108864
%windir%\system32\inetsrv\appcmd.exe set config "Exchange Back End/ews/" -section:system.serviceModel/bindings
/customBinding.[name='EWSWSSecurityX509CertHttpBinding'].httpTransport.maxReceivedMessageSize:67108864
%windir%\system32\inetsrv\appcmd.exe set config "Exchange Back End/ews/" -section:system.serviceModel/bindings
/webHttpBinding.[name='EWSStreamingNegotiateHttpsBinding'].maxReceivedMessageSize:67108864
%windir%\system32\inetsrv\appcmd.exe set config "Exchange Back End/ews/" -section:system.serviceModel/bindings
/webHttpBinding.[name='EWSStreamingNegotiateHttpBinding'].maxReceivedMessageSize:67108864
```

Outlook on the web

```
%windir%\system32\inetsrv\appcmd.exe set config "Default Web Site/owa/" -
section:system.webServer/security/requestFiltering /requestLimits.maxAllowedContentLength:35000000
%windir%\system32\inetsrv\appcmd.exe set config "Default Web Site/owa/" -section:system.web/httpRuntime
/maxRequestLength:35000
%windir%\system32\inetsrv\appcmd.exe set config "Exchange Back End/owa/" -
section:system.webServer/security/requestFiltering /requestLimits.maxAllowedContentLength:35000000
%windir%\system32\inetsrv\appcmd.exe set config "Exchange Back End/owa/" -section:system.web/httpRuntime
/maxRequestLength:35000
%windir%\system32\inetsrv\appcmd.exe set config "Exchange Back End/owa/" -section:system.serviceModel/bindings
/webHttpBinding.[name='httpsBinding'].maxReceivedMessageSize:35000000
%windir%\system32\inetsrv\appcmd.exe set config "Exchange Back End/owa/" -section:system.serviceModel/bindings
/webHttpBinding.[name='httpBinding'].maxReceivedMessageSize:35000000
%windir%\system32\inetsrv\appcmd.exe set config "Exchange Back End/owa/" -section:system.serviceModel/bindings
/webHttpBinding.[name='httpsBinding'].readerQuotas.maxStringContentLength:35000000
%windir%\system32\inetsrv\appcmd.exe set config "Exchange Back End/owa/" -section:system.serviceModel/bindings
/webHttpBinding.[name='httpBinding'].readerQuotas.maxStringContentLength:35000000
```

How do you know this worked?

To verify that you have successfully configured the client-specific message size limit, you need to send a test message to and from a mailbox by using the affected client. You can try a few smaller attachments or one large attachment so the test messages are approximately 33% less than the value you configured. For example, a configured value of 85 MB results in a realistic maximum message size of approximately 64 MB.

Availability service in Exchange Server

8/3/2020 • 2 minutes to read • Edit Online

The Availability service makes free/busy information available to Outlook and Outlook on the web (formerly known as Outlook Web App) clients. The Availability service improves information workers' calendaring and meeting scheduling experience by providing secure, consistent, and up-to-date free/busy information.

Outlook and Outlook on the web use the Availability service to perform the following tasks:

- Retrieve current free/busy information for Exchange mailboxes
- Retrieve current free/busy information from other Exchange organizations
- Retrieve published free/busy information from public folders for mailboxes on previous versions of Exchange
- View attendee working hours
- Show meeting time suggestions

How the availability service works in Exchange Server

The Availability service retrieves free/busy information directly from the target Exchange mailbox.

Outlook uses the Exchange Autodiscover service to obtain the URL of the Availability service. For more information about the Autodiscover service, see Autodiscover service.

You can use the Exchange Management Shell to configure the Availability service. You can't use the Exchange admin center (EAC) to configure the Availability service.

The Availability service API is available as a web service to let developers write third-party integration tools.

Availability service and automatic reply messages

The Availability service provides access to automatic-reply messages that users send when they are out of the office or away for an extended period of time.

Information workers use the Automatic Replies feature (formerly known as Out of Office) in Outlook and Outlook on the web to alert others when they're unavailable to respond to email messages. This functionality makes it easier to set and manage automatic reply messages for both information workers and administrators.

Methods used to retrieve free/busy information

The following table lists the methods used to retrieve free/busy information in different single-forest topologies.

CLIENT	SOURCE MAILBOX RETRIEVING FREE/BUSY INFORMATION	TARGET MAILBOX	FREE/BUSY RETRIEVAL METHOD
Outlook 2010 or later	Exchange 2010 or later	Exchange 2010 or later	The Availability service reads free/busy information from the target mailbox.

CLIENT	SOURCE MAILBOX RETRIEVING FREE/BUSY INFORMATION	TARGET MAILBOX	FREE/BUSY RETRIEVAL METHOD
Outlook on the web or Outlook Web App	Exchange 2010 or later	Exchange 2010 or later	Outlook on the web or Outlook Web App calls the Availability service API, which reads the free/busy information from the target mailbox.

Configure the Availability service for cross-forest topologies

8/3/2020 • 4 minutes to read • Edit Online

The Availability service improves information workers' free/busy information by providing secure, consistent, and up-to-date free/busy information to clients that are running Outlook. By default, this service is installed with Exchange Server. In cross-forest topologies where all connecting clients are running Outlook, the Availability service is the only method of retrieving free/busy information. You can use the Exchange Management Shell to configure the Availability service for cross-forest topologies.

> **NOTE**
>
> You can't use the Exchange admin center (EAC) to configure the Availability service for cross-forest topologies.

Using the Availability service in trusted and untrusted forests

You can use the Availability service in cross-forest topologies across trusted or untrusted forests. The type of free/busy information that's available depends on if you're using a trusted or untrusted forest.

Trusted forests: In trusted forests, you can configure the Availability service to retrieve free/busy information on a per-user basis. When the Availability service is configured to retrieve free/busy information on a per-user basis, the service can make cross-forest requests on behalf of a particular user. This allows a user in a remote forest to retrieve detailed free/busy information for someone who is not in the same forest.

Untrusted forests: In untrusted forests, you can only configure the Availability service to retrieve free/busy information on an organization-wide basis. When the Availability service makes free/busy cross-forest requests at the organizational level, free/busy information is returned for each user in the organization. In untrusted forests, it isn't possible to control the level of free/busy information that's returned on a per-user basis.

Configuring Windows for cross-forest topologies

By default, a global address list (GAL) contains mail recipients from a single forest. If you have a cross-forest environment, we recommend using Microsoft Identity Lifecycle Manager (ILM) 2007 Feature Pack 1 (FP1) to ensure that the GAL in any given forest contains mail recipients from other forests. ILM 2007 FP1 creates mail users that represent recipients from other forests, thereby allowing users to view them in the GAL and send mail. For example, users in Forest A appear as a mail user in Forest B and vice versa. Users in the target forest can then select the mail user object that represents a recipient in another forest to send mail.

To enable GAL synchronization, you create management agents that import mail-enabled users, contacts, and groups from designated Active Directory services into a centralized metadirectory. In the metadirectory, mail-enabled objects are represented as mail users. Groups are represented as contacts without any associated membership. The management agents then export these mail users to an organizational unit in the specified target forest.

What do you need to know before you begin?

- Estimated time to complete each procedure: 5 minutes.
- To open the Exchange Management Shell, see Open the Exchange Management Shell.

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Availability Service Permissions" entries in the Clients and mobile devices permissions topic.

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

- There are additional considerations when the target forest is Exchange Server 2013 or Exchange Server 2016. See Cross forest free/busy lookup fails when target forest is Exchange Server 2013 or Exchange Server 2016 for more information.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to configure per-user free/busy information in a trusted cross-forest topology

This example configures the Availability service to retrieve per-user free/busy information on a Mailbox server in the target forest.

```
Get-MailboxServer | Add-ADPermission -Accessrights Extendedright -Extendedrights "ms-Exch-EPI-Token-
Serialization" -User "<Remote Forest Domain>\Exchange servers"
```

This example defines the free/busy access method that the Availability service uses on the local Mailbox server in the source forest. The local Mailbox server is configured to access free/busy information from the forest ContosoForest.com on a per-user basis. This example uses the service account to retrieve free/busy information.

```
Add-AvailabilityAddressSpace -Forestname ContosoForest.com -AccessMethod PerUserFB -UseServiceAccount $true
```

> **NOTE**
>
> To configure bidirectional cross-forest availability, repeat these steps in the Exchange Management Shell for the target forest.

If you choose to configure cross-forest availability with trust, and also choose to use a service account (instead of specifying organization-wide or per-user credentials), you must extend permissions as shown in the example in the next section, "Use the Exchange Management Shell to configure trusted cross-forest availability with a service account." Performing that procedure in the target forest gives Mailbox servers in the source forest permission to serialize the original user context.

Use the Exchange Management Shell to configure trusted cross-forest availability with a service account

This example configures trusted cross-forest availability with a service account.

```
Get-MailboxServer | Add-ADPermission -Accessrights Extendedright -Extendedright "ms-Exch-EPI-Token-
Serialization" -User "<Remote Forest Domain>\Exchange servers"
```

For detailed information about syntax and parameters, see the following topics:

- Get-MailboxServer
- Add-ADPermission
- Add-AvailabilityAddressSpace
- Set-AvailabilityConfig

Use the Exchange Management Shell to configure organization-wide free/busy information in an untrusted cross-forest topology

This example sets the organization-wide account on the availability configuration object to configure the access level for free/busy information in the target forest.

```
Set-AvailabilityConfig -OrgWideAccount "Contoso.com\User"
```

This example adds the Availability address space configuration object for the source forest, and you're prompted to enter the credentials for organization-wide user in Contoso.com domain.

```
Add-AvailabilityAddressspace -Forestname Contoso.com -Accessmethod OrgWideFB -Credential (Get-Credential)
```

Planning and deployment for Exchange Server

8/3/2020 • 5 minutes to read • Edit Online

This topic contains links to topics and information about planning for and then deploying Exchange Server 2016 or Exchange Server 2019.

> **IMPORTANT**
>
> Make sure you read the Release notes for Exchange Server topics before you begin your deployment. The release notes contain important information on issues you might encounter during and after your deployment.

Plan for Exchange Server

Use the information at following links to help plan your deployment of Exchange 2016 or Exchange 2019 into your organization.

> **IMPORTANT**
>
> See the Establish an Exchange 2016 or Exchange 2019 test environment section later in this topic for information about installing Exchange 2019 in a test environment.

Exchange architecture

> Learn about the Mailbox and Edge Transport server roles and more in Exchange.

Exchange Server system requirements

> Understand the system requirements that need to be satisfied in your organization before you can install Exchange.

Exchange Server prerequisites

> Learn about the Windows Server features and the other software that needs to be installed for a successful installation of Exchange.

Microsoft Exchange Server Deployment Assistant

> Use this tool to generate a customized checklist for planning, installing, or upgrading Exchange. Guidance is available for multiple scenarios, including an on-premises, hybrid, or cloud deployment.

Active Directory

> Learn about how Exchange uses Active Directory and how your Active Directory deployment affects your Exchange deployment.

Antispam and antimalware protection in Exchange Server

> Learn about the built-in antispam and antimalware protection options in Exchange.

Exchange Server Hybrid Deployments

> Learn about planning a hybrid deployment with Microsoft 365 or Office 365 and your on-premises Exchange organization.

Exchange Server virtualization

> Learn how you can deploy Exchange in a virtualized environment.

Exchange Online and Exchange development

> Learn about the application programming interfaces (APIs) that are available for applications that use Exchange 2019.

Establish an Exchange 2016 or Exchange 2019 test environment

Before you install your first Exchange server, we recommend that you install Exchange in an isolated test environment. This approach reduces the risk of end-user downtime and negative ramifications to the production environment.

The test environment will act as your "proof of concept" for your new Exchange design and make it possible to move forward or roll back any implementations before deploying into your production environments. Having an exclusive test environment for validation and testing allows you to do pre-installation checks for your future production environments. By installing in a test environment first, we believe that your organization will have a better likelihood of success in a full production implementation.

For many organizations, the costs of building a test lab may be high because of the need to duplicate the production environment. To reduce the hardware costs associated with a prototype lab, we recommend the use of virtualization by using Hyper-V technologies in Windows Server. Hyper-V enables server virtualization, allowing multiple virtual operating systems to run on a single physical machine.

For more detailed information about Hyper-V, see Server Virtualization. For information about the Microsoft support of production Exchange servers on hardware virtualization software, see Exchange Server virtualization.

Deploy Exchange 2016 or Exchange 2019

During the deployment phase, you install Exchange into your organization. Before you begin the deployment phase, you should plan your Exchange organization. For more information, see the Plan for Exchange Server section earlier in this topic.

Use the information at the following links to help you deploy Exchange.

Prepare Active Directory and domains for Exchange

> Learn about the steps you need to take to prepare your Active Directory forest for Exchange 2019 and the changes Exchange makes to your forest.

Install Exchange Mailbox servers using the Setup wizard

> Learn about using the Setup wizard to install Mailbox servers.

Use unattended mode in Exchange Setup

> Learn about using the unattended setup at the command line to install, remove, update, and recover Exchange servers.

Install Exchange Edge Transport servers using the Setup wizard

> Learn about using the Setup wizard to install Edge Transport servers in a perimeter network.

Upgrade Exchange to the latest Cumulative Update

> Learn about finding and installing the latest Cumulative Update (CU) for the Exchange servers in your organization.

Exchange Server Hybrid Deployments

> Read this topic for information that will help you deploy Exchange in an existing hybrid deployment.

Exchange Server post-installation tasks

> Learn about post-installation tasks to complete your Exchange installation.

Exchange Setup

You can use different types and modes of Exchange Setup to install and maintain the various editions and versions of Exchange.

Exchange editions and versions

Exchange is available in two server editions: Standard Edition and Enterprise Edition. The edition you install is defined by your product key (the only available download can install both versions). For more information, see Exchange licensing FAQs.

Types of Exchange Setup

You have the following options for Exchange Setup:

- **Exchange Setup wizard**: Running Setup.exe without any command line switches provides an interactive experience where you're guided by the Exchange 2019 Setup wizard.

- **Exchange unattended setup**: Running Setup.exe with command line switches enables you to install Exchange from an interactive command line or through a script.

Modes of Exchange Setup

Exchange setup includes the following modes:

- **Install**: Install a new server role (Mailbox server, Edge Transport server, or Management tools). This mode is available in the Exchange Setup wizard and unattended setup.

- **Uninstall**: Remove the Exchange installation from a computer. You can use this mode from both the Exchange Setup wizard and unattended setup.

- **Upgrade**: Install a CU on an existing Exchange server. You can use this mode from both the Exchange Setup wizard and unattended setup.

  > **NOTE**
  >
  > Exchange doesn't support in-place upgrades from previous versions. This mode is used only to install CUs.

- **RecoverServer**: You need to recover data from the Exchange server after a catastrophic failure. To do this, you install a new Windows server with the same FQDN as the failed server (for example, mailbox01.contoso.com), and then run Exchange Setup with the */Mode:RecoverServer* switch without specifying the Exchange server 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. To run in `RecoverServer` mode:

- Exchange can't be already installed on the server.
- The Exchange server object must exist in Active Directory.
- You can only use unattended setup.

> **NOTE**
>
> You must complete one mode of Setup before you can use another mode.

Exchange Server system requirements

8/3/2020 • 13 minutes to read • Edit Online

Before you install Exchange Server 2019, we recommend that you review this topic to ensure your network, hardware, software, clients, and other elements meet the requirements for Exchange 2019. Also, make sure you understand the coexistence scenarios that are supported for Exchange 2019 and earlier versions of Exchange.

To actually install Exchange 2019, see Deploy new installations of Exchange.

Supported coexistence scenarios for Exchange 2019

The supported coexistence scenarios between Exchange 2019 and earlier versions of Exchange are described in the following table:

EXCHANGE VERSION	EXCHANGE 2019 ORGANIZATION COEXISTENCE
Exchange 2010 and earlier versions	Not supported
Exchange 2013	Supported with Exchange 2013 Cumulative Update 21 (CU21) or later on all Exchange 2013 servers in the organization, including Edge Transport servers.
Exchange 2016	Supported with Exchange 2016 CU11 or later on all Exchange 2016 servers in the organization, including Edge Transport servers.
Mixed Exchange 2013 and Exchange 2016 organization	Supported if all Exchange 2013 and Exchange 2016 servers in the organization meet the requirements as previously described in this table.

Supported hybrid deployment scenarios for Exchange 2019

Exchange 2019 supports hybrid deployments with Microsoft 365 or Office 365 organizations that have been upgraded to the latest version of Microsoft 365 or Office 365. For more information about specific hybrid deployments, see Hybrid deployment prerequisites.

Network and directory server requirements for Exchange 2019

The requirements for the network and the directory servers in your Exchange 2019 organization are described in the following table:

COMPONENT	REQUIREMENT
Domain controllers	All domain controllers in the forest need to be running one of the following versions of Windows Server: • Windows Server 2019 Standard or Datacenter • Windows Server 2016 Standard or Datacenter • Windows Server 2012 R2 Standard or Datacenter
Active Directory forest	The Active Directory forest functional level is **Windows Server 2012 R2** or higher.

COMPONENT	REQUIREMENT
Active Directory site	The Active Directory site where you install the Exchange Server must contain at least one writeable domain controller that's also a global catalog server, or the installation will fail. Furthermore, you can't install the Exchange server and then remove the domain controller from the Active Directory site.
DNS namespace	Exchange 2019 supports the following DNS namespaces: • Contiguous • Noncontiguous • Single label domains • Disjoint For more information about DNS namespaces that are supported by Exchange, see KB2269838.
IPv6	Exchange 2013 and later support IPv6 only when IPv4 is also installed and enabled on the Exchange server. If you deploy Exchange in this configuration, and your network supports IPv4 and IPv6, all Exchange servers can send data to and receive data from devices, servers, and clients that use IPv6 addresses. For more information, see IPv6 Support in Exchange 2013.

Directory server architecture for Exchange 2019

Active Directory domain controllers on 64-bit hardware with a 64-bit version of Windows Server will increase directory service performance for Exchange 2019.

Installing Exchange 2019 on directory servers

For security and performance reasons, we don't recommend installing Exchange 2019 on Active Directory directory servers. Only install Exchange 2019 on member servers.

To learn more about the issues that you'll encounter when you install Exchange on a directory server, see Installing Exchange on a domain controller is not recommended [WarningInstallExchangeRolesOnDomainController]. After Exchange is installed, changing the server role from a member server to a directory server or vice-versa isn't supported.

Hardware requirements for Exchange 2019

For information about deploying Exchange in a virtualized environment, see Exchange Server virtualization.

COMPONENT	REQUIREMENT	NOTES
Processor	Either of the following types of 64-bit processors: • Intel processor that supports Intel 64 architecture (formerly known as Intel EM64T). • AMD processor that supports the AMD64 platform. **Notes**: • Intel Itanium IA64 processors are not supported. • Recommended supported processor sockets is up to 2 on physical machines.	See the Supported operating systems for Exchange 2019 section later in this topic for supported operating systems.

COMPONENT	REQUIREMENT	NOTES
Memory	Varies by Exchange server role: • **Mailbox**: 128 GB minimum recommended • **Edge Transport**: 64 GB minimum recommended.	Exchange 2019 has large memory support (up to 256 GB).
Paging file size	Set the paging file minimum and maximum value to the same size: 25% of installed memory.	None
Disk space	• At least 30 GB of free space on the drive where you're installing Exchange. • At least 200 MB of free space on the system drive. • At least 500 MB of free space on the drive that contains the message queue database.	None
Screen resolution	1024 x 768 pixels (XGA) or higher	None
File system	**NTFS**: Required on partitions that contain the following types of files: • The System partition. • Exchange binaries. • Files generated by Exchange diagnostic logging. • Transport database files (for example, the mail queue database). **ReFS**: Supported on partitions that contain the following types of Exchange files: • Mailbox databases. • Transaction logs.	None

Supported operating systems for Exchange 2019

EXCHANGE COMPONENT	REQUIREMENT
Mailbox and Edge Transport server roles	Windows Server 2019 Standard or Datacenter
Management tools	One of the following versions of Windows: • Windows Server 2019 Standard or Datacenter • 64-bit edition of Windows 10

Notes:

- Installing Exchange 2019 on a computer that's running Windows Server Core is fully supported and recommended. The Desktop Experience feature is no longer required.
- Installing Exchange 2019 on a computer that's running Nano Server isn't supported.

Supported Powershell versions for Exchange 2019 servers

Exchange 2019 servers support the version of PowerShell that's included in the release of Windows Server where Exchange is installed. Don't install stand-alone downloads of WMF or PowerShell on Exchange servers.

Installing other software on Exchange 2019 servers

We don't support installing Office client or Office server software on Exchange servers (for example, SharePoint Server, Skype for Business Server, Office Online Server, or Project Server). Other software that you want to install on an Exchange 2019 server needs to be designed to run on the same computer as Exchange Server.

Supported .NET Framework versions for Exchange 2019

We strongly recommend that you use the latest version of the .NET Framework that's supported by the release of Exchange you're installing.

> **IMPORTANT**
>
> • **Releases of .NET Framework that aren't listed in the table below aren't supported on any release of Exchange 2019**. This includes minor and patch-level releases of .NET Framework.
>
> • The complete prerequisite list for Exchange 2019 is available here.

EXCHANGE 2019 VERSION	.NET FRAMEWORK 4.8	.NET FRAMEWORK 4.7.2
CU4, CU5, CU6	Supported	
CU2, CU3	Supported	Supported
RTM, CU1		Supported

Supported clients (with latest updates) in Exchange 2019

- Microsoft 365 Apps for enterprise
- Outlook 2019
- Outlook 2016
- Outlook 2013
- Outlook for Mac for Office 365
- Outlook 2016 for Mac

> **IMPORTANT**
>
> You need KB3140245 to apply registry keys to enable TLS 1.1 & 1.2 support for Windows 7. Otherwise, Outlook 2013 and 2016 will not work on Windows 7.

Lync/Skype For Business Server integration with Exchange 2019

If you're integrating Lync presence and instant messaging with Exchange Server, Lync Server 2013 Cumulative Update 10 or later is required. If you're integrating Skype for Business presence and instant messaging with Exchange Server, Skype for Business Server Cumulative Update 7 or later is required.

Before you install Exchange Server 2016, we recommend that you review this topic to ensure your network, hardware, software, clients, and other elements meet the requirements for Exchange 2016. Also, make sure you understand the coexistence scenarios that are supported for Exchange 2016 and earlier versions of Exchange.

To actually install Exchange 2016, see Deploy new installations of Exchange.

Supported coexistence scenarios for Exchange 2016

The following table lists the scenarios in which coexistence between Exchange 2016 and earlier versions of Exchange is supported.

EXCHANGE VERSION	EXCHANGE ORGANIZATION COEXISTENCE
Exchange 2007 and earlier versions	Not supported
Exchange 2010	Supported with Update Rollup 11 for Exchange 2010 SP3 or later on all Exchange 2010 servers in the organization, including Edge Transport servers.
Exchange 2013	Supported with Exchange 2013 Cumulative Update 10 or later on all Exchange 2013 servers in the organization, including Edge Transport servers.
Mixed Exchange 2010 and Exchange 2013 organization	Supported with the following minimum versions of Exchange: • Update Rollup 11 Exchange 2010 SP3 or later on all Exchange 2010 servers in the organization, including Edge Transport servers. • Exchange 2013 Cumulative Update 10 or later on all Exchange 2013 servers in the organization, including Edge Transport servers.

Supported hybrid deployment scenarios for Exchange 2016

Exchange 2016 supports hybrid deployments with Microsoft 365 or Office 365 organizations that have been upgraded to the latest version of Microsoft 365 or Office 365. For more information about specific hybrid deployments, see Hybrid Deployment Prerequisites.

Network and directory server requirements for Exchange 2016

The following table lists the requirements for the network and the directory servers in your Exchange 2016 organization.

COMPONENT	REQUIREMENT
Domain controllers	All domain controllers in the forest need to be running one of the following versions of Windows Server: • Windows Server 2019 Standard or Datacenter • Windows Server 2016 Standard or Datacenter • Windows Server 2012 R2 Standard or Datacenter • Windows Server 2012 Standard or Datacenter • Windows Server 2008 R2 Standard or Enterprise • Windows Server 2008 R2 Datacenter RTM or later
Active Directory forest	The Active Directory forest functional level is Windows Server 2008 R2 or higher.

COMPONENT	REQUIREMENT
Active Directory site	• The Active Directory site where you install the Exchange Server must contain at least one writeable domain controller that's also a global catalog server, or the installation will fail. • Furthermore, you can't install the Exchange server and then remove the domain controller from the Active Directory site.
DNS namespace support	Exchange 2016 supports the following domain name system (DNS) namespaces: • Contiguous • Noncontiguous • Single label domains • Disjoint For more information about DNS namespaces supported by Exchange, see Microsoft Knowledge Base article 2269838, Microsoft Exchange compatibility with Single Label Domains, Disjoined Namespaces, and Discontiguous Namespaces.
IPv6 support	In Exchange 2016, IPv6 is supported only when IPv4 is also installed and enabled. If Exchange 2016 is deployed in this configuration, and the network supports IPv4 and IPv6, all Exchange servers can send data to and receive data from devices, servers, and clients that use IPv6 addresses. For more information, see IPv6 Support in Exchange 2013.

Directory server architecture for Exchange 2016

The use of 64-bit Active Directory domain controllers increases directory service performance for Exchange 2016.

> **NOTE**
>
> In multi-domain environments, on Windows Server 2008 domain controllers that have the Active Directory language locale set to Japanese (ja-jp), your servers may not receive some attributes that are stored on an object during inbound replication. For more information, see KB949189.

Installing Exchange 2016 on directory servers

For security and performance reasons, we recommend that you install Exchange 2016 only on member servers and not on Active Directory directory servers. To learn about the issues you can face when installing Exchange 2016 on a directory server, see Installing Exchange on a domain controller is not recommended [WarningInstallExchangeRolesOnDomainController]. After Exchange 2016 is installed, changing its role from a member server to a directory server, or vice versa, isn't supported.

Hardware requirements for Exchange 2016

For information about deploying Exchange in a virtualized environment, see Exchange Server virtualization.

COMPONENT	REQUIREMENT	NOTES

COMPONENT	REQUIREMENT	NOTES
Processor	Either of the following types of 64-bit processors: • Intel processor that supports Intel 64 architecture (formerly known as Intel EM64T). • AMD processor that supports the AMD64 platform. **Note**: Intel Itanium IA64 processors are not supported.	For more information, see Sizing Exchange 2016 Deployments. See the Supported operating systems for Exchange 2016 section later in this topic for supported operating systems.
Memory	Varies by Exchange server role: • **Mailbox**: 8 GB minimum. • **Edge Transport**: 4 GB minimum.	For more information, see Sizing Exchange 2016 Deployments.
Paging file size	Set the paging file minimum and maximum value to the same size: • **Less than 32 GB of RAM installed**: Physical RAM plus 10MB, up to a maximum value of 32GB (32,778MB). • **32 GB of RAM or more installed**: 32GB	None
Disk space	• At least 30 GB of free space on the drive where you're installing Exchange, plus an additional 500 MB for each Unified Messaging (UM) language pack that you plan to install. • At least 200 MB of free space on the System drive. • At least 500 MB of free space on the drive that contains the message queue database.	For more information, see Sizing Exchange 2016 Deployments.
Drive	DVD-ROM drive, local or network accessible.	None
Screen resolution	1024 x 768 pixels (XGA) or higher	None
File format	**NTFS**: Required on partitions that contain the following types of files: • The System partition. • Exchange binaries. • Files generated by Exchange diagnostic logging. • Transport database files (for example, the mail queue database). **ReFS**: Supported on partitions that contain the following types of Exchange files: • Mailbox databases. • Transaction logs. • Content indexing files.	None

Supported operating systems for Exchange 2016

Important: We don't support the installation of Exchange 2016 on a computer that's running Windows Server Core or Nano Server. The Windows Server Desktop Experience feature needs to be installed. To install Exchange 2016, you need to do one of the following to install the Desktop Experience on Windows Server prior to starting Exchange 2016 Setup:

- **Windows Server 2012 and Windows Server 2012 R2**: Run the following command in Windows PowerShell

  ```
  Install-WindowsFeature Server-Gui-Mgmt-Infra,Server-Gui-Shell -Restart
  ```

- **Windows Server 2016**: Install Windows Server 2016 and choose the **Desktop Experience** installation option. If a computer is running Windows Server 2016 Core mode and you want to install Exchange 2016 on it, you'll need to reinstall the operating system and choose the **Desktop Experience** installation option.

COMPONENT	REQUIREMENT
Mailbox and Edge Transport server roles	• Windows Server 2016 Standard or Datacenter* • Windows Server 2012 R2 Standard or Datacenter • Windows Server 2012 Standard or Datacenter
Management tools	One of the following versions of Windows: • Windows Server 2016 Standard or Datacenter* • Windows Server 2012 R2 Standard or Datacenter • Windows Server 2012 Standard or Datacenter • 64-bit edition of Windows 10 • 64-bit edition of Windows 8.1

* Requires Exchange Server 2016 Cumulative Update 3 or later.

Supported Windows Management Framework versions for Exchange 2016

Exchange 2016 only supports the version of Windows Management Framework that's built in to the release of Windows that you're installing Exchange on. Don't install versions of Windows Management Framework that are made available as stand-alone downloads on servers running Exchange.

Installing other software on Exchange 2016 servers

We don't support installing Office clients or other Office server products (for example, SharePoint Server, Skype for Business Server, Office Online Server, or Project Server) on Exchange 2016 servers. Software that you want to install on an Exchange 2016 server needs to be designed to run on the same computer as Exchange Server.

Supported .NET Framework versions for Exchange 2016

We strongly recommend that you use the latest version of .NET Framework that's supported by the release of Exchange you're installing.

> **IMPORTANT**
>
> • **Releases of .NET Framework that aren't listed in the table below are not supported on any release of Exchange 2016**. This includes minor and patch-level releases of .NET Framework.
>
> • The complete prerequisite list for Exchange 2016 is available here.

EXCHANGE 2016 VERSION	.NET FRAMEWORK 4.8	.NET FRAMEWORK 4.7.2	.NET FRAMEWORK 4.7.1	.NET FRAMEWORK 4.6.2
CU15, CU16, CU17	Supported			
CU13, CU14	Supported	Supported		
CU11, CU12		Supported	Supported	
CU10			Supported	
CU8, CU9			Supported	Supported
CU5, CU6, CU7				Supported

Note: For older versions, see Exchange Server supportability matrix.

Supported clients (with latest updates) in Exchange 2016

- Microsoft 365 Apps for enterprise
- Outlook 2019
- Outlook 2016
- Outlook 2013
- Outlook 2010 SP2
- Outlook 2016 for Mac
- Outlook for Mac for Office 365

Exchange third-party clients

Exchange Server offers several well-known protocols, and publishes APIs that third-party vendors often write clients for.

Microsoft makes no warranties, expressed or implied, as to the overall suitability, fitness, compatibility, or security of clients that are created by third-party developers.

If you want to use a third-party client that uses our protocols or APIs, we recommend that you thoroughly review and test all considerations (functionality, security, maintenance, management, and so on) before you deploy the client in the enterprise workspace. We also recommend that you make sure that the third-party vendor offers an appropriate Enterprise Support Agreement (ESA).

Exchange Server prerequisites

8/3/2020 • 10 minutes to read • Edit Online

This topic provides the steps for installing the necessary Windows Server operating system prerequisites for Exchange Server 2016 and Exchange Server 2019 Mailbox servers and Edge Transport servers, and also the Windows prerequisites for installing the Exchange Management Tools on Windows client computers.

After you've prepared your environment for Exchange Server, use the Exchange Deployment Assistant for the next steps in your actual deployment. For information on hybrid deployments, see Exchange Server Hybrid Deployments.

To actually install Exchange 2016 and Exchange 2019, see Deploy new installations of Exchange.

TIP

- Looking for Exchange 2013 prerequisites? See Exchange 2013 prerequisites.
- Remote Registry Service must be set to Automatic and cannot be Disabled. For recommended Security Guidelines, See Security Guidelines regarding Remote Registry.
- Have you heard about the Exchange Server Deployment Assistant? It's a free online tool that helps you quickly deploy Exchange Server in your organization by asking you a few questions and creating a customized deployment checklist just for you. If you want to learn more about it, go to Microsoft Exchange Server Deployment Assistant.

What do you need to know before you begin?

- Verify that your Active Directory meets the requirements for Exchange 2019: Exchange 2019 Network and directory servers.
- Verify that your Active Directory meets the requirements for Exchange 2016: Exchange 2016 Network and directory servers.
- The full installation option of Windows Server 2012 and Windows Server 2012 R2 must be used for all servers running Exchange 2016 server roles or management tools.
- Some prerequisites require you to reboot the server to complete installation.

NOTE

You can't upgrade Windows from one version to another, or from Standard to Datacenter, when Exchange is installed on the server.

- Verify the Supported operating systems for Exchange 2019 or Supported operating systems for Exchange 2016.

NOTE

New to Exchange 2019 is the ability to upgrade your operating system to a newer version while Exchange is installed on Windows Server 2019 or later.

- Verify the computer is joined to the appropriate internal Active Directory domain.

- Install the latest Windows updates on your computer.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Exchange 2019 prerequisites for preparing Active Directory

You can use any member of the Active Directory domain to prepare Active Directory for Exchange 2019.

1. The computer requires the following software:

   a. .NET Framework 4.8

   b. Visual C++ Redistributable Package for Visual Studio 2012

   > **NOTE**
   >
   > - The system requirements for the Visual C++ redistributable package do not mention support for Windows Server 2016 or Windows Server 2019, but the redistributable package is safe to install on these versions of Windows.
   > - An overview of the latest supported versions is available at: Visual C++ Redistributable versions.
   > - The Visual C++ Redistributable package is required if you're using the Exchange Setup Wizard to prepare Active Directory. If you're using unattended Setup from the command line to prepare Active Directory, this package isn't required. For more information, see Prepare Active Directory and domains.

2. Install the Remote Tools Administration Pack by running the following command in Windows PowerShell:

   ```
   Install-WindowsFeature RSAT-ADDS
   ```

> **NOTE**
>
> Using the Exchange Setup Wizard to prepare Active Directory requires the installation of the Management Tools Exchange role.

Windows Server 2019 prerequisites for Exchange 2019

The requirements to install Exchange 2019 on Windows Server 2019 computers are described in the following sections. We recommend either of the following methods to install the Windows prerequisites for Exchange 2019:

- Use the /InstallWindowsComponents switch in unattended Setup mode.
- Select the check box in the Exchange Setup Wizard to install Windows prerequisites.

When you use one of these options, you don't need to restart the computer after the Windows components have been added.

Exchange 2019 Mailbox servers on Windows Server 2019

1. Install the following software:

   a. .NET Framework 4.8

b. Visual C++ Redistributable Package for Visual Studio 2012

c. Visual C++ Redistributable Package for Visual Studio 2013

> **NOTE**
>
> - The system requirements for the Visual C++ redistributable package do not mention support for Windows Server 2016 or Windows Server 2019, but the redistributable package is safe to install on these versions of Windows.
> - An overview of the latest supported versions is available at: Visual C++ Redistributable versions

2. Add the required Lync Server or Skype for Business Server components:

   a. Install the Server Media Foundation windows feature by executing the following command in Windows PowerShell:

   ```
   Install-WindowsFeature Server-Media-Foundation
   ```

   b. Install Unified Communications Managed API 4.0. This package is available for download and in the \UCMARedist folder on the Exchange Server media.

   > **NOTE**
   >
   > When installing on Windows Server Core, you must use the installation package located in \UCMARedist on distributed media.

3. If you aren't going to use Exchange Setup to install the required Windows components (in the wizard or from the command line), run the one of the following commands in Windows PowerShell:

   - Desktop Experience:

   ```
   Install-WindowsFeature Server-Media-Foundation, NET-Framework-45-Features, RPC-over-HTTP-proxy, RSAT-Clustering, RSAT-Clustering-CmdInterface, RSAT-Clustering-Mgmt, RSAT-Clustering-PowerShell, 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, RSAT-ADDS
   ```

   - Server Core:

   ```
   Install-WindowsFeature Server-Media-Foundation, NET-Framework-45-Features, RPC-over-HTTP-proxy, RSAT-Clustering, RSAT-Clustering-CmdInterface, RSAT-Clustering-PowerShell, 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-Metabase, Web-Mgmt-Service, Web-Net-Ext45, Web-Request-Monitor, Web-Server, Web-Stat-Compression, Web-Static-Content, Web-Windows-Auth, Web-WMI, RSAT-ADDS
   ```

Exchange 2019 Edge Transport servers on Windows Server 2019

1. Install the Visual C++ Redistributable Package for Visual Studio 2012

> **NOTE**
>
> - The system requirements for the Visual C++ redistributable package do not mention support for Windows Server 2016 or Windows Server 2019, but the redistributable package is safe to install on these versions of Windows.
> - An overview of the latest supported versions is available at: Visual C++ Redistributable versions.

2. If you aren't going to use Exchange Setup to install the required Windows components (in the wizard or from the command line), run the following command in Windows PowerShell:

```
Install-WindowsFeature ADLDS
```

Windows 10 client prerequisites for the Exchange 2019 management tools

1. Install the Visual C++ Redistributable Package for Visual Studio 2012

> **NOTE**
>
> - The system requirements for the Visual C++ redistributable package do not mention support for Windows Server 2016 or Windows Server 2019, but the redistributable package is safe to install on these versions of Windows.
> - An overview of the latest supported versions is available at: Visual C++ Redistributable versions.

2. If you aren't going to use Exchange Setup to install the required Windows components (in the wizard or from the command line), run the following command in Windows PowerShell:

```
Enable-WindowsOptionalFeature -Online -FeatureName IIS-ManagementScriptingTools,IIS-ManagementScriptingTools,IIS-IIS6ManagementCompatibility,IIS-LegacySnapIn,IIS-ManagementConsole,IIS-Metabase,IIS-WebServerManagementTools,IIS-WebServerRole
```

Exchange 2016 prerequisites for preparing Active Directory

You can use any member of the Active Directory domain to prepare Active Directory for Exchange 2016.

1. The computer requires the following software:

   a. .NET Framework 4.8

   b. Visual C++ Redistributable Package for Visual Studio 2012

> **NOTE**
>
> - The system requirements for the Visual C++ redistributable package do not mention support for Windows Server 2016 or Windows Server 2019, but the redistributable package is safe to install on these versions of Windows.
> - An overview of the latest supported versions is available at: Visual C++ Redistributable versions.

2. Install the Remote Tools Administration Pack by running the following command in Windows PowerShell:

```
Install-WindowsFeature RSAT-ADDS
```

After you've installed the Remote Tools Administration Pack you can use the computer to prepare Active Directory. For more information, see Prepare Active Directory and domains.

Windows Server 2016 prerequisites for Exchange 2016

The prerequisites that are needed to install Exchange 2016 on computers running Windows Server 2016 depends on which Exchange role you want to install. Read the section below that matches the role you want to install.

> **IMPORTANT**
>
> Windows Server 2016 requires Exchange 2016 Cumulative Update 3 or later.

Exchange 2016 Mailbox servers on Windows Server 2016

1. Run the following command in Windows PowerShell to install the required Windows components:

```
Install-WindowsFeature NET-Framework-45-Features, Server-Media-Foundation, RPC-over-HTTP-proxy, RSAT-Clustering, RSAT-Clustering-CmdInterface, RSAT-Clustering-Mgmt, RSAT-Clustering-PowerShell, 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, RSAT-ADDS
```

2. Install the following software in order:

a. .NET Framework 4.8

b. December 13, 2016 (KB3206632) security update

> **NOTE**
>
> You can only install this update if your Windows Server 2016 version is 14393.576 or earlier (circa December, 2016). You can check your Windows Server version by running the **winver** command. If your Windows Server 2016 version is greater than 14393.576, you don't need this update or its replacement KB3213522, which was released one week later. Exchange 2016 Setup looks for the installation of this update, won't allow you to continue if this update is missing, and will clearly inform you if you need it.

c. Visual C++ Redistributable Package for Visual Studio 2012

d. Visual C++ Redistributable Package for Visual Studio 2013

NOTE

- The system requirements for the Visual C++ redistributable package do not mention support for Windows Server 2016 or Windows Server 2019, but the redistributable package is safe to install on these versions of Windows.
- An overview of the latest supported versions is available at: Visual C++ Redistributable versions.
- Only the Mailbox role requires the Visual C++ Redistributable Packages for Visual Studio 2013. Other Exchange installations (management tools and Edge Transport) only require the Visual C++ Redistributable Packages for Visual Studio 2012.

e. Microsoft Unified Communications Managed API 4.0, Core Runtime 64-bit

Exchange 2016 Edge Transport servers on Windows Server 2016

1. Run the following command in Windows PowerShell to install the required Windows components:

```
Install-WindowsFeature ADLDS
```

2. Install the following software in order:

a. .NET Framework 4.8

b. Visual C++ Redistributable Package for Visual Studio 2012

NOTE

- The system requirements for the Visual C++ redistributable package do not mention support for Windows Server 2016 or Windows Server 2019, but the redistributable package is safe to install on these versions of Windows.
- An overview of the latest supported versions is available at: Visual C++ Redistributable versions.

Windows Server 2012 and Windows Server 2012 R2 prerequisites for Exchange 2016

The prerequisites for Exchange 2016 on Windows Server 2012 or Windows Server 2012 R2 computers depend on the Exchange role that you're installing. Read the following section that matches the role you want to install.

Exchange 2016 Mailbox servers on Windows Server 2012 or Windows Server 2012 R2

1. Run the following command in Windows Powershell to install the required Windows components:

```
Install-WindowsFeature AS-HTTP-Activation, Server-Media-Foundation, NET-Framework-45-Features, RPC-over-HTTP-proxy, RSAT-Clustering, RSAT-Clustering-CmdInterface, RSAT-Clustering-Mgmt, RSAT-Clustering-PowerShell, 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, RSAT-ADDS
```

2. Install the following software in order:

a. .NET Framework 4.8

b. Visual C++ Redistributable Package for Visual Studio 2012

c. Visual C++ Redistributable Package for Visual Studio 2013

> **NOTE**
>
> - The system requirements for the Visual C++ redistributable package do not mention support for Windows Server 2016 or Windows Server 2019, but the redistributable package is safe to install on these versions of Windows.
> - An overview of the latest supported versions is available at: Visual C++ Redistributable versions.
> - Only the Mailbox role requires the Visual C++ Redistributable Packages for Visual Studio 2013. Installations of the Exchange management tools and Edge Transport servers only require the Visual C++ Redistributable Packages for Visual Studio 2012.

d. Microsoft Unified Communications Managed API 4.0, Core Runtime 64-bit

Exchange 2016 Edge Transport servers on Windows Server 2012 or Windows Server 2012 R2

1. Run the following command in Windows PowerShell to install the required Windows components:

```
Install-WindowsFeature ADLDS
```

2. Install the following software in order:

a. .NET Framework 4.8

b. Visual C++ Redistributable Package for Visual Studio 2012

> **NOTE**
>
> - The system requirements for the Visual C++ redistributable package do not mention support for Windows Server 2016 or Windows Server 2019, but the redistributable package is safe to install on these versions of Windows.
> - An overview of the latest supported versions is available at: Visual C++ Redistributable versions.

Windows client prerequisites for the Exchange 2016 management tools

Exchange 2016 management tools on Windows 10

1. Install Visual C++ Redistributable Package for Visual Studio 2012

> **NOTE**
>
> - The system requirements for the Visual C++ redistributable package do not mention support for Windows Server 2016 or Windows Server 2019, but the redistributable package is safe to install on these versions of Windows.
> - An overview of the latest supported versions is available at: Visual C++ Redistributable versions.

2. Run the following command in an elevated Windows PowerShell window (a Windows PowerShell window you open by selecting `Run as administrator`):

```
Enable-WindowsOptionalFeature -Online -FeatureName IIS-ManagementScriptingTools,IIS-
ManagementScriptingTools,IIS-IIS6ManagementCompatibility,IIS-LegacySnapIn,IIS-ManagementConsole,IIS-
Metabase,IIS-WebServerManagementTools,IIS-WebServerRole
```

Exchange 2016 management tools on Windows 8.1

1. Install .NET Framework 4.8

2. Install Visual C++ Redistributable Package for Visual Studio 2012

> **NOTE**
>
> - The system requirements for the Visual C++ redistributable package do not mention support for Windows Server 2016 or Windows Server 2019, but the redistributable package is safe to install on these versions of Windows.
>
> - An overview of the latest supported versions is available at: Visual C++ Redistributable versions.

3. Run the following command in an elevated Windows PowerShell window (a Windows PowerShell window you open by selecting **Run as administrator**):

```
Enable-WindowsOptionalFeature -Online -FeatureName IIS-ManagementScriptingTools,IIS-
ManagementScriptingTools,IIS-IIS6ManagementCompatibility,IIS-LegacySnapIn,IIS-ManagementConsole,IIS-
Metabase,IIS-WebServerManagementTools,IIS-WebServerRole
```

Install Office Online Server in an Exchange organization

8/3/2020 • 5 minutes to read • Edit Online

An optional prerequisite for Exchange 2016 Cumulative Update 1 (CU1) or later, as well as for Exchange 2019, is the installation of Office Online Server on one or more servers in your organization. Office Online Server enables users to view supported file attachments within Outlook on the web (formerly known as Outlook Web App) without downloading them first and without having a local installation of the program. Without Office Online Server installed, Outlook users need to download attachments to their local computer and then open them in a local application.

> **NOTE**
> Office Online Server is available for download as part of a volume licensing agreement. If you don't have a volume license agreement, you can skip the instructions in this step. However, without Office Online Server installed, Outlook users will need to download attachments to their local computer to view them; they won't be able to view them in Outlook.

You can configure an Office Online Server endpoint in two places in Exchange 2016 and later: at the organization level, and at the Mailbox server level. Where you configure the endpoint depends on the size of your organization and the location of your servers and users.

- **Organization**: There are a couple of reasons why you might configure the Office Online Server endpoint at the organization level:

  - **Single-server or single-location deployment**: You can configure the endpoint at the organization level if all of your Exchange 2016 Mailbox servers are in the same location and you don't plan on having geographically distributed Office Online Server servers.

  - **Fallback for large deployments**: You can configure endpoint at the organization level as a fallback if the endpoint configured on a Mailbox server isn't available. If an Office Web Apps server isn't available, the client will try to connect to the endpoint configured at the organization level.

    Notes:

    - If you have Exchange 2013 servers in your organization, don't configure an endpoint at the organization level. Doing so will direct Exchange 2013 servers to use the Office Online Server server. This isn't supported.

    - Previewing attachments in S/MIME messages in Outlook on the web isn't supported by Office Online Server.

- **Mailbox server**: If you want to distribute client requests between two or more Office Online Server servers, if you want to geographically distribute Office Online Server servers, or if you have Exchange 2013 in your organization, configure the endpoints at the Exchange Mailbox server level. When you configure an endpoint at the server level, mailboxes located on that server will send requests to the configured Office Online Server server.

If you want users outside of your network to view supported file attachments in Outlook, Office Online Server needs to be accessible from the Internet. TCP port 443 needs to be opened on your firewall and forwarded to the Office Online Server server. If you deploy more than one Office Online Server server, each server needs its own fully qualified domain name (FQDN). Each server also needs to be accessible from the Internet via TCP port 443.

Office Online Server system requirements

To set up Office Online Server, you will need the following:

- Windows Server 2012 R2 or Windows Server 2016
- Exchange 2016 Cumulative Update 1 (CU1) or later, or Exchange 2019

  > **NOTE**
  >
  > If you're running Windows Server 2016, you will need Exchange 2016 CU3 or later, as detailed in Exchange Server prerequisites.

- Microsoft .NET Framework 4.5.2
- Visual C++ Redistributable for Visual Studio 2015
- Visual C++ Redistributable Packages for Visual Studio 2013
- Microsoft.IdentityModel.Extention.dll
- All available Windows updates installed

> **NOTE**
>
> Office Online Server can't be installed on an Exchange server, SharePoint server, Active Directory domain controller, or any other computer with existing applications installed.

Install Office Online Server

1. To install Office Online Server, follow Steps 1 through 3 in the section **Prepare servers to run Office Online Server** of the article Deploy Office Online Server before proceeding.

2. Obtain and import an SSL certificate with the fully qualified domain name(s) (FQDN) of the Office Online Server server. If your organization is configured for split DNS, you only need to configure one FQDN on the certificate. For example, oos.contoso.com. If you have different internal and external FQDNs, you'll need to configure both FQDNs on the certificate. For example, oos.internal.contoso.com and oos.contoso.com.

3. Configure DNS records to point the FQDN(s) on the certificate to your Office Online Serverserver. If you have different DNS servers for internal and external users, you'll need to configure the appropriate FQDN on each server.

4. Open Windows PowerShell and run the following commands. When you run the commands, replace the example FQDNs and certificate friendly name with your own.

   ```
   New-OfficeWebAppsFarm -InternalURL "https://oos.contoso.com" -ExternalURL "https://oos.contoso.com" -
   CertificateName "Office Online Server Preview Certificate"`
   ```

   > **NOTE**
   >
   > You can configure different internal and external URLs, but in the next step you'll see that you can only configure one URL for Exchange. In this case, if you use the internal URL in the next step, this function will only work internally and external users will get an unexpected error. If you use the external URL, this function will only work for external users and internal users will get an unexpected error.

Configure the Office Online Server endpoint at the Mailbox server level

After you've configured the Office Online Server server, do the following on your Exchange 2016 server. This will allow Outlook to send requests to the Office Online Server server.

1. Open the Exchange Management Shell and run the following command. Replace the example server name and URL with your own.

   ```
   Set-MailboxServer MBX -WacDiscoveryEndpoint "https://oos.contoso.com/hosting/discovery"
   ```

2. Restart the MsExchangeOwaAppPool by running the following command.

   ```
   Restart-WebAppPool MsExchangeOwaAppPool
   ```

Configure the Office Online Server endpoint at the organization level

After you've configured the Office Online Server server, do the following on your Exchange 2016 server. This will allow Outlook to send requests to the Office Online Server server.

1. Open the Exchange Management Shell and run the following command. Replace the example URL with your own.

   ```
   Set-OrganizationConfig -WacDiscoveryEndpoint "https://oos.internal.contoso.com/hosting/discovery"
   ```

   > **IMPORTANT**
   >
   > If you have Exchange 2013 servers in your organization, don't configure an endpoint at the organization level. Doing so will direct Exchange 2013 servers to use the Office Online Server server. This isn't supported.

2. Restart the MsExchangeOwaAppPool by running the following command.

   ```
   Restart-WebAppPool MsExchangeOwaAppPool
   ```

Active Directory in Exchange Server organizations

8/3/2020 • 2 minutes to read • Edit Online

Exchange Server 2016 and Exchange Server 2019 use Active Directory to store and share directory information with Windows. Starting with Exchange 2013, we've made some changes to how Exchange works with Active Directory. These changes are described in this topic.

Active Directory driver

The Active Directory driver is the core Microsoft Exchange component that allows Exchange services to create, modify, delete, and query for Active Directory Domain Services (AD DS) data. In Exchange 2013 and later, all access to Active Directory is done using the Active Directory driver itself. In previous versions of Exchange, DSAccess provided directory lookup services for components such as SMTP, message transfer agent (MTA), and the Exchange store.

The Active Directory driver also uses the Microsoft Exchange Active Directory Topology (MSExchangeADTopology) server, which allows the Active Directory driver to use Directory Service Access (DSAccess) topology data. This data includes the list of available domain controllers and global catalog servers that are available to handle Exchange requests. For more information about the Active Directory Driver, see Active Directory Domain Services.

Active Directory schema changes

Exchange add new attributes to the Active Directory domain service schema and also make other modifications to existing classes and attributes. For more information about Active Directory changes when you install Exchange, see Active Directory schema changes in Exchange Server.

For more information

To learn more about how Exchange stores and retrieves information in Active Directory so that you can plan for access to it, see Access to Active Directory in Exchange Server.

For more information about Active Directory forest design, see AD DS Design Guide.

To learn more about computers running Windows in an Active Directory domain and deploying Exchange 2013 or later in a domain that has a disjoint namespace, see Disjoint Namespace Scenarios.

Access to Active Directory by Exchange servers

8/3/2020 • 5 minutes to read • Edit Online

Exchange Server 2016 and Exchange Server 2019 store all configuration and recipient information in the Active Directory directory service database. When an Exchange server requires information about recipients the configuration of the Exchange organization, it queries Active Directory. Active Directory servers must be available for Exchange to function correctly.

This topic explains how Exchange stores and retrieves information in Active Directory so that you can plan access to Active Directory. This topic also discusses issues you should be aware of if you try to recover deleted Exchange Active Directory objects.

Exchange information stored in Active Directory

The Active Directory database stores information in three types of logical partitions that are described in the following sections:

- Schema partition
- Configuration partition
- Domain partition

Schema partition

The schema partition stores the following two types of information:

- **Schema classes** define all the types of objects that can be created and stored in Active Directory.
- **Schema attributes** define all the properties that can be used to describe the objects that are stored in Active Directory.

When you install the first Exchange server in the forest (or run the Active Directory preparation process), the Active Directory preparation process adds many classes and attributes to the Active Directory schema. The classes that are added to the schema are used to create Exchange-specific objects (for example, agents and connectors). These attributes are used to configure the Exchange-specific objects and the mail-enabled users and groups. These attributes include properties, such as Outlook on the web (formerly known as Outlook Web App) settings.

Every domain controller and global catalog server in the forest contains a complete replica of the schema partition.

For more information about schema modifications in Exchange, see Active Directory schema changes in Exchange Server.

Configuration partition

The configuration partition stores information about the forest-wide configuration. This configuration information includes the configuration of Active Directory sites, Exchange global settings, transport settings, and mailbox policies. Each type of configuration information is stored in a container in the configuration partition. Exchange configuration information is stored in a subfolder under the configuration partition's Services container. The type of information that's stored in this container includes:

- Address lists
- Address book mailbox policies
- Administrative groups

- Client Access settings
- Connections
- Mobile mailbox Settings
- Global settings
- Monitoring Settings
- System policies
- Retention policies container
- Transport settings

Every domain controller and global catalog server in the forest contains a complete replica of the configuration partition.

Domain partition

The domain partition stores information in default containers and in organizational units that are created by the Active Directory administrator. These containers hold the domain-specific objects. This data includes Exchange system objects and information about the computers, users, and groups in that domain. When Exchange is installed, Exchange updates the objects in this partition to support Exchange functionality. This functionality affects how recipient information is stored and accessed.

Each domain controller contains a complete replica of the domain partition for the domain for which it is authoritative. Every global catalog server in the forest contains a subset of the information in every domain partition in the forest.

How Exchange accesses information in Active Directory

Exchange uses an Active Directory API to access information that's stored in Active Directory. This service reads information from all Active Directory partitions. The data that is retrieved is cached and is used by Exchange servers to discover the Active Directory site location of all Exchange services in the organization.

For more information about topology and service discovery in Exchange 2013 or later, see Planning to use Active Directory sites for routing Mail.

Exchange is an Active Directory site-aware application that prefers to communicate with the directory servers that are located in the same site as the Exchange server to optimize network traffic. Each Exchange server must communicate with Active Directory to retrieve information about recipients and information about the other Exchange servers. Mailbox servers store configuration information about mailbox users and mailbox stores in Active Directory. Additionally, the Mailbox server stores information in Active Directory for the Client Access protocols, Transport service, Mailbox databases, and so on. The Mailbox server handles all activity for the active mailboxes on that server.

By default, whenever an Exchange server starts, it binds to a randomly selected domain controller and global catalog server in its own site. You can view the selected directory servers by using the **Get-ExchangeServer** cmdlet in the Exchange Management Shell. You can also use the **Set-ExchangeServer** cmdlet to configure a static list of domain controllers that an Exchange 2016 server should bind to or a list of domain controllers that should be excluded.

> **IMPORTANT**
>
> You can't deploy an Exchange server in any site that contains only read-only directory servers.

Recovery of deleted Exchange objects

Active Directory Recycle Bin helps minimize directory service downtime by enhancing your ability to preserve and recover accidentally deleted Active Directory objects without restoring Active Directory data from backups, restarting Active Directory Domain Services (AD DS), or rebooting domain controllers.

The most important thing to understand about recovering deleted Exchange-related Active Directory objects is that Exchange objects don't exist in isolation. For example, when you mail-enable a user, several different policies and links are calculated for the user based on your current Exchange configuration. Two problems that may arise when you restore a deleted Exchange configuration or recipient object are:

- **Collisions**: Some Exchange attributes must be unique across a forest. For example, all email addresses on a mail-enabled object (also known as proxy addresses) must be unique. Two different mail-enabled objects can't have the same email address. Active Directory doesn't enforce proxy address uniqueness (Exchange administrative tools check for uniqueness). Exchange email address policies also automatically resolve possible conflicts in proxy address assignment based on deterministic rules. Therefore, restoring an Exchange user object might create a collision with proxy addresses or other attributes that should be unique.

- **Misconfigurations**: Exchange has automated rules that assign various policies or settings. If you delete a recipient, and then change the rules or policies, restoring an Exchange user object may result in a user being assigned to the wrong policy (or even to a policy that no longer exists).

The following guidelines will help you minimize problems or issues when you recover deleted Exchange-related objects:

- If you deleted an Exchange configuration object using Exchange management tools, don't restore the object. Instead, create the object again using the Exchange management tools (the Exchange admin center or the Exchange Management Shell).

- If you deleted an Exchange configuration object without using the Exchange management tools, recover the object as soon as possible. The more administrative and configuration changes that are mae after the deletion, the more likely that restoring the objects will result in misconfiguration.

- If you recover deleted Exchange recipients (contacts, users, or distribution groups), monitor closely for collisions and errors relating to the recovered objects. If Exchange policies or other recipient configuration settings were modified after the deletion, re-apply the current policies to the restored recipients to ensure that they're configured correctly.

For more information

Active Directory Recycle Bin Step-by-Step Guide

Introduction to Active Directory Administrative Center Enhancements (Level 100)

Advanced AD DS Management Using Active Directory Administrative Center (Level 200)

Active Directory schema changes in Exchange Server

8/3/2020 • 25 minutes to read • Edit Online

This reference topic provides a summary of the Active Directory schema changes that are made when you install Exchange Server 2016 or Exchange Server 2019 in your organization. Refer to the .ldf files for more information about changes to the Active Directory schema. The .ldf files are located in the \Setup\Data\ directory in the Exchange installation files.

Exchange schema updates are cumulative. Each Cumulative Update (CU) includes all of the changes that were included in previous releases. This means that if you skip a CU, you might still need to apply schema updates even if the CU that you're installing doesn't include its own changes.

> **NOTE**
>
> The Active Directory schema changes that are described in this topic might not apply to all editions of an Exchange 2019 version. To verify that Active Directory has been successfully prepared, see the Exchange Active Directory versions section in Prepare Active Directory and domains for Exchange 2019.

Exchange 2019 CU6 Active Directory schema changes

No changes are made to the Active Directory schema in Exchange 2019 CU6.

Exchange 2019 CU5 Active Directory schema changes

No changes are made to the Active Directory schema in Exchange 2019 CU5.

Exchange 2019 CU4 Active Directory schema changes

No changes are made to the Active Directory schema in Exchange 2019 CU4.

Exchange 2019 CU3 Active Directory schema changes

No changes are made to the Active Directory schema in Exchange 2019 CU3.

Exchange 2019 CU2 Active Directory schema changes

No changes are made to the Active Directory schema in Exchange 2019 CU2.

Exchange 2019 CU1 Active Directory schema changes

No changes are made to the Active Directory schema in Exchange 2019 CU1.

However, a reduction in Active Directory permissions is made: The AdminSDHolder object on the domain is updated to remove the "Allow" ACE that grants the "Exchange Trusted Subsystem" group the "Write DACL" right on the "Group" inherited object types. For more information, see KB 4490059.

Exchange 2019 RTM Active Directory schema changes

This section summarizes the changes that are made to the Active Directory schema when you install Exchange 2019 RTM. This section includes the following subsections:

- Classes added by Exchange 2019 RTM
- Classes modified by Exchange 2019 RTM
- Attributes added by Exchange 2019 RTM
- Global catalog attributes added by Exchange 2019 RTM
- Attributes modified by Exchange 2019 RTM
- Object IDs added by Exchange 2019 RTM
- Indexed attributes added by Exchange 2019 RTM
- Property sets modified by Exchange 2019 RTM
- MAPI IDs added by Exchange 2019 RTM
- Extended rights added by Exchange 2019 RTM

> **NOTE**
>
> No changes to the Active Directory schema were made between Exchange 2019 Preview and Exchange 2019 RTM.

Classes added by Exchange 2019 RTM

CLASS	CHANGE
Exch-Mapi-Virtual-Directory	ntdsSchemaAdd
Exch-Push-Notifications-App	ntdsSchemaAdd
ms-Exch-Account-Forest	ntdsSchemaAdd
ms-Exch-ActiveSync-Device-Autoblock-Threshold	ntdsSchemaAdd
ms-Exch-Auth-Auth-Config	ntdsSchemaAdd
ms-Exch-Auth-Auth-Server	ntdsSchemaAdd
ms-Exch-Auth-Partner-Application	ntdsSchemaAdd
ms-Exch-Auth-Policy	ntdsSchemaAdd
ms-Exch-Client-Access-Rule	ntdsSchemaModify
ms-Exch-Config-Settings	ntdsSchemaAdd
ms-Exch-Encryption-Virtual-Directory	ntdsSchemaAdd
ms-Exch-Exchange-Transport-Server	ntdsSchemaAdd
ms-Exch-Hosted-Content-Filter-Config	ntdsSchemaAdd
ms-Exch-Http-Delivery-Connector	ntdsSchemaAdd

CLASS	CHANGE
ms-Exch-Hygiene-Configuration	ntdsSchemaAdd
ms-Exch-Intra-Organization-Connector	ntdsSchemaModify
ms-Exch-Mailbox-Policy	ntdsSchemaAdd
ms-Exch-Mailflow-Policy	ntdsSchemaAdd
ms-Exch-Mailflow-Policy-Collection	ntdsSchemaAdd
ms-Exch-Malware-Filter-Config	ntdsSchemaAdd
ms-Exch-MSO-Forward-Sync-Divergence	ntdsSchemaAdd
ms-Exch-MSO-Sync-Service-Instance	ntdsSchemaAdd
ms-Exch-Organization-Upgrade-Policy	ntdsSchemaAdd
ms-Exch-Protocol-Cfg-SIP-Container	ntdsSchemaAdd
ms-Exch-Protocol-Cfg-SIP-FE-Server	ntdsSchemaAdd
ms-Exch-Resource-Policy	ntdsSchemaAdd
ms-Exch-Safe-Attachment-Protection-Config	ntdsSchemaAdd
ms-Exch-Smart-Links-Protection-Config	ntdsSchemaAdd
ms-Exch-Team-Mailbox-Provisioning-Policy	ntdsSchemaAdd
ms-Exch-Throttling-Policy	ntdsSchemaModify
ms-Exch-Unified-Policy	ntdsSchemaAdd
ms-Exch-Unified-Rule	ntdsSchemaAdd
ms-Exch-Workload-Policy	ntdsSchemaAdd

Classes modified by Exchange 2019 RTM

CLASS	CHANGE	ATTRIBUTE/CLASS
Mail-Recipient	add: mayContain	msExchAdministrativeUnitLink
Mail-Recipient	add: mayContain	msExchAuthPolicyLink
Mail-Recipient	add: mayContain	msExchImmutableSid
Mail-Recipient	add: mayContain	msExchUGEventSubscriptionLink

CLASS	CHANGE	ATTRIBUTE/CLASS
ms-Exch-Base-Class	add: mayContain	msExchUserHoldPolicies
ms-Exch-Configuration-Unit-Container	add: mayContain	msExchAuthPolicyLink
ms-Exch-Configuration-Unit-Container	add: mayContain	msExchMSOForwardSyncReplayList
ms-Exch-Container	add: mayContain	msExchScopeFlags
ms-Exch-Mail-Storage	add: mayContain	msExchDataEncryptionPolicyLink
ms-Exch-Organization-Container	add: mayContain	msExchDataEncryptionPolicyLink
Exch-Accepted-Domain	add:mayContain	msExchOfflineOrgIdHomeRealmRecord
Exch-Base-Class	add:mayContain	msExchCapabilityIdentifiers
Exch-Base-Class	add:mayContain	msExchObjectID
Exch-Base-Class	add:mayContain	msExchProvisioningTags
Exch-Configuration-Unit-Container	add:mayContain	msExchArchiveRelease
Exch-Configuration-Unit-Container	add:mayContain	msExchMailboxRelease
Exch-Exchange-Server	add:mayContain	msExchArchiveRelease
Exch-Exchange-Server	add:mayContain	msExchMailboxRelease
Exch-MDB-Availability-Group	add:mayContain	msExchEvictedMembersLink
Exch-OAB	add:mayContain	msExchLastUpdateTime
Exch-OWA-Mailbox-Policy	add:mayContain	msExchConfigurationXML
Exch-OWA-Virtual-Directory	add:mayContain	msExchConfigurationXML
Exch-On-Premises-Organization	add:mayContain	msExchTrustedDomainLink
Exch-Organization-Container	add:mayContain	msExchMaxABP
Exch-Organization-Container	add:mayContain	msExchMaxOAB
Exch-Organization-Container	add:mayContain	pFContacts
Exch-Team-Mailbox-Provisioning-Policy	add:mayContain	msExchConfigurationXML
Group	add: auxiliaryClass	msExchMailStorage
Mail-Recipient	add:mayContain	msExchLocalizationFlags

CLASS	CHANGE	ATTRIBUTE/CLASS
Mail-Recipient	add:mayContain	msExchRoleGroupType
Mail-Recipient	add:mayContain	ms-DS-GeoCoordinates-Altitude
Mail-Recipient	add:mayContain	ms-DS-GeoCoordinates-Latitude
Mail-Recipient	add:mayContain	ms-DS-GeoCoordinates-Longitude
Mail-Recipient	add:mayContain	msExchRecipientSoftDeletedStatus
Mail-Recipient	add:mayContain	msExchWhenSoftDeletedTime
Mail-Recipient	add:mayContain	msExchHomeMTASL
Mail-Recipient	add:mayContain	msExchMailboxMoveSourceArchiveMDBLinkSL
Mail-Recipient	add:mayContain	msExchMailboxMoveSourceMDBLinkSL
Mail-Recipient	add:mayContain	msExchMailboxMoveTargetArchiveMDBLinkSL
Mail-Recipient	add:mayContain	msExchMailboxMoveTargetMDBLinkSL
Mail-Recipient	add:mayContain	ms-exch-group-external-member-count
Mail-Recipient	add:mayContain	ms-exch-group-member-count
Mail-Recipient	add:mayContain	msExchGroupExternalMemberCount
Mail-Recipient	add:mayContain	msExchGroupMemberCount
Mail-Recipient	add:mayContain	msExchShadowWhenSoftDeletedTime
Mail-Recipient	add:mayContain	msExchPublicFolderMailbox
Mail-Recipient	add:mayContain	msExchPublicFolderSmtpAddress
Mail-Recipient	add: mayContain	msExchAuxMailboxParentObjectIdLink
Mail-Recipient	add: mayContain	msExchStsRefreshTokensValidFrom
Mail-Recipient	add:mayContain	msDS-ExternalDirectoryObjectId
Mail-Recipient	add:mayContain	msExchGroupSecurityFlags
Mail-Recipient	add:mayContain	msExchMultiMailboxDatabasesLink
Ms-Exch-Organization-Container	add:mayContain	ms-exch-organization-flags-2

CLASS	CHANGE	ATTRIBUTE/CLASS
Top	add:mayContain	msExchMultiMailboxDatabasesBL
Top	add:mayContain	msExchMultiMailboxLocationsBL
Top	add:mayContain	msExchAccountForestBL
Top	add:mayContain	msExchTrustedDomainBL
Top	add:mayContain	msExchAcceptedDomainBL
Top	add:mayContain	msExchHygieneConfigurationMalwareBL
Top	add:mayContain	msExchHygieneConfigurationSpamBL
Top	add:mayContain	msExchEvictedMembersBL
Top	add: mayContain	msExchOABGeneratingMailboxBL
Top	add: mayContain	msExchAuxMailboxParentObjectIdBL
Top	add: mayContain	msExchAdministrativeUnitBL
Top	add: mayContain	msExchAuthPolicyBL
Top	add: mayContain	msExchDataEncryptionPolicyBL
Top	add: mayContain	msExchUGEventSubscriptionBL
ms-Exch-Accepted-Domain	add:mayContain	msExchHygieneConfigurationLink
ms-Exch-Accepted-Domain	add:mayContain	msExchTransportResellerSettingsLinkSL
ms-Exch-Account-Forest	possSuperiors	msExchContainer
ms-Exch-Account-Forest	Add:mayContain	msExchPartnerId
ms-Exch-Active-Sync-Device	add:mayContain	msExchDeviceClientType
ms-Exch-Availability-Address-Space	add:mayContain	msExchFedTargetAutodiscoverEPR
ms-Exch-Base-Class	add:mayContain	msExchDirsyncAuthorityMetadata
ms-Exch-Base-Class	add:mayContain	msExchDirsyncStatusAck
ms-Exch-Base-Class	add:mayContain	msExchEdgeSyncConfigFlags
ms-Exch-Base-Class	add:mayContain	msExchHABRootDepaPreviewentLink
ms-Exch-Base-Class	add:mayContain	msExchDefaultPublicFolderMailbox

CLASS	CHANGE	ATTRIBUTE/CLASS
ms-Exch-Base-Class	add:mayContain	msExchForestModeFlag
ms-Exch-Base-Class	add:mayContain	msExchELCMailboxFlags
ms-Exch-Base-Class	add:mayContain	msExchCanaryData0
ms-Exch-Base-Class	add:mayContain	msExchCanaryData1
ms-Exch-Base-Class	add:mayContain	msExchCanaryData2
ms-Exch-Base-Class	add:mayContain	msExchCorrelationId
ms-Exch-Base-Class	Add:mayContain	msExchTenantCountry
ms-Exch-Base-Class	Add:mayContain	msExchConfigurationXML
ms-Exch-Base-Class	add: mayContain	msExchMultiMailboxGUIDs
ms-Exch-Base-Class	add: mayContain	msExchMultiMailboxLocationsLink
ms-Exch-Coexistence-Relationship	add:mayContain	msExchCoexistenceOnPremisesSmartHost
ms-Exch-Coexistence-Relationship	add:mayContain	msExchCoexistenceSecureMailCertificateThumbprint
ms-Exch-Coexistence-Relationship	add:mayContain	msExchCoexistenceTransportServers
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchDirsyncStatus
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchIsDirsyncStatusPending
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchDirSyncServiceInstance
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchOrganizationUpgradePolicyLink
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchManagementSiteLinkSL
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchOrganizationUpgradePolicyLinkSL
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchRelocateTenantCompletionTargetVector
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchRelocateTenantFlags
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchRelocateTenantSafeLockdownSchedule
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchRelocateTenantSourceForest

CLASS	CHANGE	ATTRIBUTE/CLASS
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchRelocateTenantStartLockdown
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchRelocateTenantStartRetired
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchRelocateTenantStartSync
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchRelocateTenantStatus
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchRelocateTenantTargetForest
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchRelocateTenantTransitionCounter
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchSyncCookie
ms-Exch-Control-Point-Config	add:mayContain	msExchRMSOnlineCertificationLocationUrl
ms-Exch-Control-Point-Config	add:mayContain	msExchRMSOnlineKeySharingLocationUrl
ms-Exch-Control-Point-Config	add:mayContain	msExchRMSOnlineLicensingLocationUrl
ms-Exch-Custom-Attributes	add:mayContain	msExchExtensionCustomAttribute1
ms-Exch-Custom-Attributes	add:mayContain	msExchExtensionCustomAttribute2
ms-Exch-Custom-Attributes	add:mayContain	msExchExtensionCustomAttribute3
ms-Exch-Custom-Attributes	add:mayContain	msExchExtensionCustomAttribute4
ms-Exch-Custom-Attributes	add:mayContain	msExchExtensionCustomAttribute5
ms-Exch-Domain-Content-Config	add:mayContain	msExchContentByteEncoderTypeFor7BitCharsets
ms-Exch-Domain-Content-Config	add:mayContain	msExchContentPreferredInternetCodePageForShiftJis
ms-Exch-Domain-Content-Config	add:mayContain	msExchContentRequiredCharSetCoverage
ms-Exch-Exchange-Server	add:mayContain	msExchWorkloadManagementPolicyLink
ms-Exch-Exchange-Server	add:mayContain	msExchMalwareFilteringDeferAttempts
ms-Exch-Exchange-Server	add:mayContain	msExchMalwareFilteringDeferWaitTime
ms-Exch-Exchange-Server	add:mayContain	msExchMalwareFilteringFlags

CLASS	CHANGE	ATTRIBUTE/CLASS
ms-Exch-Exchange-Server	add:mayContain	msExchMalwareFilteringPrimaryUpdatePath
ms-Exch-Exchange-Server	add:mayContain	msExchMalwareFilteringSecondaryUpdatePath
ms-Exch-Exchange-Server	add:mayContain	msExchMalwareFilteringUpdateFrequency
ms-Exch-Exchange-Server	add:mayContain	msExchMalwareFilteringUpdateTimeout
ms-Exch-Exchange-Server	add:mayContain	msExchMalwareFilteringScanTimeout
ms-Exch-Fed-OrgId	add:mayContain	msExchFedDelegationTrustSL
ms-Exch-Hosted-Content-Filter-Config	add:mayContain	msExchSpamCountryBlockList
ms-Exch-Hosted-Content-Filter-Config	add:mayContain	msExchSpamLanguageBlockList
ms-Exch-Hosted-Content-Filter-Config	add:mayContain	msExchSpamNotifyOutboundRecipients
ms-Exch-Hosted-Content-Filter-Config	add:mayContain	msExchSpamDigestFrequency
ms-Exch-Hosted-Content-Filter-Config	add:mayContain	msExchSpamQuarantineRetention
ms-Exch-MDB	add:mayContain	msExchCalendarLoggingQuota
ms-Exch-MRS-Request	add:mayContain	msExchMailboxMoveSourceMDBLinkSL
ms-Exch-MRS-Request	add:mayContain	msExchMailboxMoveStorageMDBLinkSL
ms-Exch-MRS-Request	add:mayContain	msExchMailboxMoveTargetMDBLinkSL
ms-Exch-MSO-Sync-Service-Instance	add:mayContain	msExchMSOForwardSyncNonRecipientCookie
ms-Exch-MSO-Sync-Service-Instance	add:mayContain	msExchMSOForwardSyncRecipientCookie
ms-Exch-MSO-Sync-Service-Instance	add:mayContain	msExchMSOForwardSyncReplayList
ms-Exch-MSO-Sync-Service-Instance	add:mayContain	msExchAccountForestLink
ms-Exch-MSO-Sync-Service-Instance	add:mayContain	msExchActiveInstanceSleepInterval
ms-Exch-MSO-Sync-Service-Instance	add:mayContain	msExchPassiveInstanceSleepInterval
ms-Exch-MSO-Sync-Service-Instance	add:mayContain	msExchSyncDaemonMaxVersion
ms-Exch-MSO-Sync-Service-Instance	add:mayContain	msExchSyncDaemonMinVersion

CLASS	CHANGE	ATTRIBUTE/CLASS
ms-Exch-MSO-Sync-Service-Instance	add:mayContain	msExchSyncServiceInstanceNewTenantMaxVersion
ms-Exch-MSO-Sync-Service-Instance	add:mayContain	msExchSyncServiceInstanceNewTenantMinVersion
ms-Exch-Mail-Gateway	add:mayContain	msExchHomeMDBSL
ms-Exch-Mail-Gateway	add:mayContain	msExchHomeMTASL
ms-Exch-Mail-Storage	add:mayContain	msExchPreviousArchiveDatabase
ms-Exch-Mail-Storage	add:mayContain	msExchTeamMailboxExpiration
ms-Exch-Mail-Storage	add:mayContain	msExchTeamMailboxOwners
ms-Exch-Mail-Storage	add:mayContain	msExchTeamMailboxSharePointLinkedBy
ms-Exch-Mail-Storage	add:mayContain	msExchTeamMailboxSharePointUrl
ms-Exch-Mail-Storage	add:mayContain	msExchTeamMailboxShowInClientList
ms-Exch-Mail-Storage	add:mayContain	msExchCalendarLoggingQuota
ms-Exch-Mail-Storage	add:mayContain	msExchArchiveDatabaseLinkSL
ms-Exch-Mail-Storage	add:mayContain	msExchDisabledArchiveDatabaseLinkSL
ms-Exch-Mail-Storage	add:mayContain	msExchHomeMDBSL
ms-Exch-Mail-Storage	add:mayContain	msExchMailboxMoveTargetMDBLinkSL
ms-Exch-Mail-Storage	add:mayContain	msExchPreviousArchiveDatabaseSL
ms-Exch-Mail-Storage	add:mayContain	msExchPreviousHomeMDBSL
ms-Exch-Mail-Storage	add: mayContain	msExchMailboxContainerGuid
ms-Exch-Mail-Storage	add: mayContain	msExchUnifiedMailbox
ms-Exch-Mail-Storage	add:mayContain	msExchUserCulture
ms-Exch-Mailflow-Policy	add:mayContain	msExchImmutableId
ms-Exch-Malware-Filter-Config	add:mayContain	msExchMalwareFilterConfigExternalSenderAdminAddress
ms-Exch-Malware-Filter-Config	add:mayContain	msExchMalwareFilterConfigInternalSenderAdminAddress
ms-Exch-OAB	add:mayContain	msExchOffLineABServerSL

CLASS	CHANGE	ATTRIBUTE/CLASS
ms-Exch-OAB	add: mayContain	msExchOABGeneratingMailboxLink
ms-Exch-OWA-Mailbox-Policy	add:mayContain	msExchOWASetPhotoURL
ms-Exch-OWA-Virtual-Directory	add:mayContain	msExchOWASetPhotoURL
ms-Exch-Organization-Container	add:mayContain	msExchOrganizationFlags2
ms-Exch-Organization-Container	add:mayContain	msExchUMAvailableLanguages
ms-Exch-Organization-Container	add:mayContain	msExchWACDiscoveryEndpoint
ms-Exch-Organization-Container	add:mayContain	msExchAdfsAuthenticationRawConfiguration
ms-Exch-Organization-Container	add:mayContain	msExchServiceEndPointURL
ms-Exch-Private-MDB	add:mayContain	msExchMailboxDatabaseTransportFlags
ms-Exch-Public-Folder	add:mayContain	msExchPublicFolderEntryId
ms-Exch-Resource-Policy	add:mayContain	msExchCustomerExpectationCritical
ms-Exch-Resource-Policy	add:mayContain	msExchDiscretionaryCritical
ms-Exch-Resource-Policy	add:mayContain	msExchInternalMaintenanceCritical
ms-Exch-Resource-Policy	add:mayContain	msExchUrgentCritical
ms-Exch-Routing-Group-Connector	add:mayContain	msExchHomeMTASL
ms-Exch-Safe-Attachment-Protection-Config	add:mayContain	msExchMalwareFilterConfigFlags
ms-Exch-Safe-Attachment-Protection-Config	add:mayContain	msExchMalwareFilterConfigFromAddress
ms-Exch-Safe-Attachment-Protection-Config	add:mayContain	msExchMalwareFilterConfigInternalBody
ms-Exch-Safe-Attachment-Protection-Config	add:mayContain	msExchMalwareFilterConfigInternalSenderAdminAddress
ms-Exch-Safe-Attachment-Protection-Config	add:mayContain	msExchMalwareFilterConfigInternalSubject
ms-Exch-Safe-Attachment-Protection-Config	add:mayContain	msExchMalwareFilteringScanTimeout
ms-Exch-Safe-Attachment-Protection-Config	add:mayContain	msExchMalwareFilteringUpdateFrequency

CLASS	CHANGE	ATTRIBUTE/CLASS
ms-Exch-Site-Connector	add:mayContain	msExchHomeMTASL
ms-Exch-Smart-Links-Protection-Config	add:mayContain	msExchAddressRewriteExceptionList
ms-Exch-Smart-Links-Protection-Config	add:mayContain	msExchSpamFlags
ms-Exch-Tenant-Perimeter-Settings	add:mayContain	msExchTransportResellerSettingsLinkSL
ms-Exch-Throttling-Policy	add:mayContain	msExchThrottlingPolicyFlags
ms-Exch-Throttling-Policy	add:mayContain	msExchAnonymousThrottlingPolicyStateEx
ms-Exch-Throttling-Policy	add:mayContain	msExchEASThrottlingPolicyStateEx
ms-Exch-Throttling-Policy	add:mayContain	msExchEWSThrottlingPolicyStateEx
ms-Exch-Throttling-Policy	add:mayContain	msExchGeneralThrottlingPolicyStateEx
ms-Exch-Throttling-Policy	add:mayContain	msExchIMAPThrottlingPolicyStateEx
ms-Exch-Throttling-Policy	add:mayContain	msExchOWAThrottlingPolicyStateEx
ms-Exch-Throttling-Policy	add:mayContain	msExchPOPThrottlingPolicyStateEx
ms-Exch-Throttling-Policy	add:mayContain	msExchPowershellThrottlingPolicyStateEx
ms-Exch-Throttling-Policy	add:mayContain	msExchRCAThrottlingPolicyStateEx
ms-Exch-Transport-Rule	add:mayContain	msExchTransportRuleImmutableId
ms-Exch-Transport-Rule	add:mayContain	msExchImmutableId
ms-Exch-Transport-Settings	add:mayContain	msExchTranspoPreviewaxRetriesForLocalSiteShadow
ms-Exch-Transport-Settings	add:mayContain	msExchTranspoPreviewaxRetriesForRemoteSiteShadow
ms-Exch-Transport-Settings	add:mayContain	msExchConfigurationXML
ms-Exch-Virtual-Directory	add:mayContain	msExchMRSProxyFlags
ms-Exch-Virtual-Directory	add:mayContain	msExchMRSProxyMaxConnections

Attributes added by Exchange 2019 RTM

- ms-DS-External-Directory-Object-Id
- ms-DS-GeoCoordinates-Altitude

- ms-DS-GeoCoordinates-Latitude
- ms-DS-GeoCoordinates-Longitude
- ms-Exch-Accepted-Domain-BL
- ms-Exch-Account-Forest-BL
- ms-Exch-Account-Forest-Link
- ms-Exch-ActiveSync-Device-AutoBlock-Duration
- ms-Exch-ActiveSync-Device-Autoblock-Threshold-Incidence-Duration
- ms-Exch-ActiveSync-Device-Autoblock-Threshold-Incidence-Limit
- ms-Exch-ActiveSync-Device-Autoblock-Threshold-Type
- ms-Exch-Adfs-Authentication-Raw-Configuration
- ms-Exch-Administrative-Unit-BL
- ms-Exch-Administrative-Unit-Link
- ms-Exch-Anonymous-Throttling-Policy-State-Ex
- ms-Exch-Archive-Database-Link-SL
- ms-Exch-Auth-Application-Identifier
- ms-Exch-Auth-App-Secret
- ms-Exch-Auth-Authorization-Url
- ms-Exch-Auth-Auth-Server-Type
- ms-Exch-Auth-Certificate-Data
- ms-Exch-Auth-Certificate-Thumbprint
- ms-Exch-Auth-Flags
- ms-Exch-Auth-Issuer-Name
- ms-Exch-Auth-Issuing-Url
- ms-Exch-Auth-Linked-Account
- ms-Exch-Auth-Metadata-Url
- ms-Exch-Auth-Policy-BL
- ms-Exch-Auth-Policy-Link
- ms-Exch-Auth-Realm
- ms-Exch-Aux-Mailbox-Parent-Object-Id-BL
- ms-Exch-Aux-Mailbox-Parent-Object-Id-Link
- ms-Exch-Canary-Data-0
- ms-Exch-Canary-Data-1
- ms-Exch-Canary-Data-2

- ms-Exch-Content-Byte-Encoder-Type-For-7-Bit-Charsets
- ms-Exch-Content-Preferred-Internet-Code-Page-For-Shift-Jis
- ms-Exch-Content-Required-Char-Set-Coverage
- ms-Exch-Correlation-Id
- ms-Exch-Customer-Expectation-Critical
- ms-Exch-Customer-Expectation-Overloaded
- ms-Exch-Customer-Expectation-Underloaded
- ms-Exch-Data-Encryption-Policy-BL
- ms-Exch-Data-Encryption-Policy-Link
- ms-Exch-Default-Public-Folder-Mailbox
- ms-Exch-Device-Client-Type
- ms-Exch-Dirsync-Authority-Metadata
- ms-Exch-Dir-Sync-Service-Instance
- ms-Exch-Dirsync-Status
- ms-Exch-Dirsync-Status-Ack
- ms-Exch-Disabled-Archive-Database-Link-SL
- ms-Exch-Discretionary-Critical
- ms-Exch-Discretionary-Overloaded
- ms-Exch-Discretionary-Underloaded
- ms-Exch-EAS-Throttling-Policy-State-Ex
- ms-Exch-Edge-Sync-Config-Flags
- ms-Exch-Encryption-Throttling-Policy-State-Ex
- ms-Exch-EWS-Throttling-Policy-State-Ex
- ms-Exch-Extension-Custom-Attribute-1
- ms-Exch-Extension-Custom-Attribute-2
- ms-Exch-Extension-Custom-Attribute-3
- ms-Exch-Extension-Custom-Attribute-4
- ms-Exch-Extension-Custom-Attribute-5
- ms-Exch-External-Directory-Object-Class
- ms-Exch-Fed-Delegation-Trust-SL
- ms-Exch-Forest-Mode-Flag
- ms-Exch-General-Throttling-Policy-State-Ex
- ms-Exch-Group-External-Member-Count

- ms-Exch-Group-Member-Count
- ms-Exch-Group-Security-Flags
- ms-Exch-Home-MDB-SL
- ms-Exch-Home-MTA-SL
- ms-Exch-Hosted-Content-Filter-Config-Link
- ms-Exch-Hygiene-Configuration-Link
- ms-Exch-Hygiene-Configuration-Malware-BL
- ms-Exch-Hygiene-Configuration-Spam-BL
- ms-Exch-IMAP-Throttling-Policy-State-Ex
- ms-Exch-Immutable-Sid
- ms-Exch-Internal-Maintenance-Critical
- ms-Exch-Internal-Maintenance-Overloaded
- ms-Exch-Internal-Maintenance-Underloaded
- ms-Exch-Is-Dirsync-Status-Pending,
- ms-Exch-Localization-Flags
- ms-Exch-Mailbox-Database-Transport-Flags
- ms-Exch-Mailbox-Move-Source-Archive-MDB-Link-SL
- ms-Exch-Mailbox-Move-Source-MDB-Link-SL
- ms-Exch-Mailbox-Move-Storage-MDB-Link-SL
- ms-Exch-Mailbox-Move-Target-Archive-MDB-Link-SL
- ms-Exch-Mailbox-Move-Target-MDB-Link-SL
- ms-Exch-Mailflow-Policy-Countries
- ms-Exch-Mailflow-Policy-Keywords
- ms-Exch-Mailflow-Policy-Publisher-Name
- ms-Exch-Mailflow-Policy-Transport-Rules-Template-Xml
- ms-Exch-Mailflow-Policy-Version
- ms-Exch-Malware-Filter-Config-Alert-Text
- ms-Exch-Malware-Filter-Config-External-Body
- ms-Exch-Malware-Filter-Config-External-Sender-Admin-Address
- ms-Exch-Malware-Filter-Config-External-Subject
- ms-Exch-Malware-Filter-Config-Flags
- ms-Exch-Malware-Filter-Config-From-Address
- ms-Exch-Malware-Filter-Config-From-Name

- ms-Exch-Malware-Filter-Config-Internal-Body
- ms-Exch-Malware-Filter-Config-Internal-Sender-Admin-Address
- ms-Exch-Malware-Filter-Config-Internal-Subject
- ms-Exch-Malware-Filter-Config-Link
- ms-Exch-Malware-Filtering-Defer-Attempts
- ms-Exch-Malware-Filtering-Defer-Wait-Time
- ms-Exch-Malware-Filtering-Flags
- ms-Exch-Malware-Filtering-Primary-Update-Path
- ms-Exch-Malware-Filtering-Scan-Timeout
- ms-Exch-Malware-Filtering-Secondary-Update-Path
- ms-Exch-Malware-Filtering-Update-Frequency
- ms-Exch-Malware-Filtering-Update-Timeout
- ms-Exch-Management-Site-Link-SL
- ms-Exch-MRS-Proxy-Flags
- ms-Exch-MRS-Proxy-Max-Connections
- ms-Exch-MSO-Forward-Sync-Divergence-Count
- ms-Exch-MSO-Forward-Sync-Divergence-Related-Object-Link
- ms-Exch-MSO-Forward-Sync-Divergence-Timestamp
- ms-Exch-Multi-Mailbox-Databases-BL
- ms-Exch-Multi-Mailbox-Databases-Link
- ms-Exch-Multi-Mailbox-GUID
- ms-Exch-Multi-Mailbox-Locations-BL
- ms-Exch-Multi-Mailbox-Locations-Link
- ms-Exch-OAB-Generating-Mailbox-BL
- ms-Exch-OAB-Generating-Mailbox-Link
- ms-Exch-Off-Line-AB-Server-SL
- ms-Exch-Organization-Flags-2
- ms-Exch-Organization-Upgrade-Policy-BL
- ms-Exch-Organization-Upgrade-Policy-Date
- ms-Exch-Organization-Upgrade-Policy-Enabled
- ms-Exch-Organization-Upgrade-Policy-Link
- ms-Exch-Organization-Upgrade-Policy-Link-SL
- ms-Exch-Organization-Upgrade-Policy-MaxMailboxes

- ms-Exch-Organization-Upgrade-Policy-Priority
- ms-Exch-Organization-Upgrade-Policy-Source-Version
- ms-Exch-Organization-Upgrade-Policy-Status
- ms-Exch-Organization-Upgrade-Policy-Target-Version
- ms-Exch-OWA-Set-Photo-URL
- ms-Exch-OWA-Throttling-Policy-State-Ex
- ms-Exch-POP-Throttling-Policy-State-Ex
- ms-Exch-Powershell-Throttling-Policy-State-Ex
- ms-Exch-Previous-Archive-Database
- ms-Exch-Previous-Archive-Database-SL
- ms-Exch-Previous-Home-MDB-SL
- ms-Exch-Public-Folder-EntryId
- ms-Exch-Public-Folder-Mailbox
- ms-Exch-Public-Folder-Smtp-Address
- ms-Exch-RCA-Throttling-Policy-State-Ex
- ms-Exch-Recipient-SoftDeleted-Status
- ms-Exch-Relocate-Tenant-Completion-Target-Vector
- ms-Exch-Relocate-Tenant-Flags
- ms-Exch-Relocate-Tenant-Safe-Lockdown-Schedule
- ms-Exch-Relocate-Tenant-Source-Forest
- ms-Exch-Relocate-Tenant-Start-Lockdown
- ms-Exch-Relocate-Tenant-Start-Retired
- ms-Exch-Relocate-Tenant-Start-Sync
- ms-Exch-Relocate-Tenant-Status
- ms-Exch-Relocate-Tenant-Target-Forest
- ms-Exch-Relocate-Tenant-Transition-Counter
- ms-Exch-Resource-Type
- ms-Exch-RMS-Computer-Accounts-Link-SL
- ms-Exch-RMSOnline-Certification-Location-Url
- ms-Exch-RMSOnline-Key-Sharing-Location-Url
- ms-Exch-RMSOnline-Licensing-Location-Url
- ms-Exch-RoleGroup-Type
- ms-Exch-Service-End-Point-URL

- ms-Exch-Shadow-When-Soft-Deleted-Time
- ms-Exch-Spam-Add-Header
- ms-Exch-Spam-Asf-Settings
- ms-Exch-Spam-Asf-Test-Bcc-Address
- ms-Exch-Spam-Country-Block-List
- ms-Exch-Spam-Digest-Frequency
- ms-Exch-Spam-False-Positive-Cc
- ms-Exch-Spam-Flags
- ms-Exch-Spam-Language-Block-List
- ms-Exch-Spam-Modify-Subject
- ms-Exch-Spam-Notify-Outbound-Recipients
- ms-Exch-Spam-Outbound-Spam-Cc
- ms-Exch-Spam-Quarantine-Retention
- ms-Exch-Spam-Redirect-Address
- ms-Exch-Sts-Refresh-Tokens-Valid-From
- ms-Exch-Sync-Cookie
- ms-Exch-Sync-Service-Instance-New-Tenant-Max-Version
- ms-Exch-Sync-Service-Instance-New-Tenant-Min-Version
- ms-Exch-Team-Mailbox-Expiration
- ms-Exch-Team-Mailbox-Expiry-Days
- ms-Exch-Team-Mailbox-Owners
- ms-Exch-Team-Mailbox-SharePoint-Linked-By
- ms-Exch-Team-Mailbox-SharePoint-Url
- ms-Exch-Team-Mailbox-Show-In-Client-List
- ms-Exch-Tenant-Country
- ms-Exch-Throttling-Policy-Flags
- ms-Exch-Transport-MaxRetriesForLocalSiteShadow
- ms-Exch-Transport-MaxRetriesForRemoteSiteShadow
- ms-Exch-Transport-Reseller-Settings-Link-SL
- ms-Exch-Transport-Rule-Immutable-Id
- ms-Exch-Trusted-Domain-BL
- ms-Exch-Trusted-Domain-Link
- ms-Exch-UG-Event-Subscription-BL

- ms-Exch-UG-Event-Subscription-Link
- ms-Exch-UG-Member-BL
- ms-Exch-UG-Member-Link
- ms-Exch-Urgent-Critical
- ms-Exch-Urgent-Overloaded
- ms-Exch-Urgent-Underloaded
- ms-Exch-WAC-Discovery-Endpoint
- ms-Exch-When-Soft-Deleted-Time
- ms-Exch-Workload-Classification
- ms-Exch-Workload-Management-Is-Enabled
- ms-Exch-Workload-Management-Policy
- ms-Exch-Workload-Management-Policy-BL
- ms-Exch-Workload-Management-Policy-Link

Global catalog attributes added by Exchange 2019 RTM

- ms-Exch-Administrative-Unit-BL
- ms-Exch-Administrative-Unit-Link
- ms-Exch-Archive-Database-Link-SL
- ms-Exch-Auth-Policy-Link
- ms-Exch-Correlation-Id
- ms-Exch-Data-Encryption-Policy-BL
- ms-Exch-Data-Encryption-Policy-Link
- ms-Exch-Default-Public-Folder-Mailbox
- ms-Exch-Device-Client-Type
- ms-Exch-Dirsync-Authority-Metadata
- ms-Exch-Dirsync-Status
- ms-Exch-Dirsync-Status-Ack
- ms-Exch-Disabled-Archive-Database-Link-SL
- ms-Exch-Edge-Sync-Config-Flags
- ms-Exch-EvictedMembers-Link
- ms-Exch-EvictedMembers-BL
- ms-Exch-Extension-Custom-Attribute-1
- ms-Exch-Extension-Custom-Attribute-2
- ms-Exch-Extension-Custom-Attribute-3
- ms-Exch-Extension-Custom-Attribute-4

- ms-Exch-Extension-Custom-Attribute-5
- ms-Exch-Group-External-Member-Count
- ms-Exch-Group-Member-Count
- ms-Exch-HAB-Root-DepaPreviewent-Link
- ms-Exch-Home-MDB-SL
- ms-Exch-Home-MTA-SL
- ms-Exch-Is-Dirsync-Status-Pending
- ms-Exch-Localization-Flags
- ms-Exch-Mailbox-Container-Guid
- ms-Exch-Mailbox-Move-Source-Archive-MDB-Link-SL
- ms-Exch-Mailbox-Move-Source-MDB-Link-SL
- ms-Exch-Mailbox-Move-Storage-MDB-Link-SL
- ms-Exch-Mailbox-Move-Target-Archive-MDB-Link-SL
- ms-Exch-Mailbox-Move-Target-MDB-Link-SL
- ms-Exch-Offline-OrgId-Home-Realm-Record
- ms-Exch-Previous-Archive-Database
- ms-Exch-Previous-Archive-Database-SL
- ms-Exch-Previous-Home-MDB-SL
- ms-Exch-Recipient-SoftDeleted-Status
- ms-Exch-Relocate-Tenant-Completion-Target-Vector,
- ms-Exch-Relocate-Tenant-Flags
- ms-Exch-Relocate-Tenant-Safe-Lockdown-Schedule
- ms-Exch-Relocate-Tenant-Source-Forest
- ms-Exch-Relocate-Tenant-Start-Lockdown
- ms-Exch-Relocate-Tenant-Start-Retired
- ms-Exch-Relocate-Tenant-Start-Sync
- ms-Exch-Relocate-Tenant-Status
- ms-Exch-Relocate-Tenant-Target-Forest
- ms-Exch-Relocate-Tenant-Transition-Counter
- ms-Exch-RMS-Computer-Accounts-Link-SL
- ms-Exch-RoleGroup-Type
- ms-Exch-Sync-Cookie
- ms-Exch-Team-Mailbox-Expiration

- ms-Exch-Team-Mailbox-Expiry-Days
- ms-Exch-Team-Mailbox-Owners
- ms-Exch-Team-Mailbox-SharePoint-Linked-By
- ms-Exch-Team-Mailbox-SharePoint-Url
- ms-Exch-Team-Mailbox-Show-In-Client-List
- ms-Exch-UG-Event-Subscription-BL
- ms-Exch-UG-Event-Subscription-Link
- ms-Exch-Unified-Mailbox
- ms-Exch-When-Soft-Deleted-Time

Attributes modified by Exchange 2019 RTM

CLASS	CHANGE	ATTRIBUTE/CLASS
Exch-Configuration-Unit-Container	rangeUpper	15254
Exch-Mailflow-Policy-Transport-Rules-Template-Xml	rangeUpper	256000
Mail-Recipient	Replace: mayContain	msExchUGMemberLink
ms-Exch-Accepted-Domain-Name	replace: searchFlags	9
ms-Exch-Archive-GUID	replace: searchFlags	9
ms-Exch-Bypass-Audit	replace: searchFlags	19
ms-Exch-Coexistence-On-Premises-Smart-Host	ntdsSchemaAdd	attributeID: 1.2.840.113556.1.4.7000.102.51992 isMemberOfPartialAttributeSet: FALSE (not in global catalogue) searchFlags: 0 (no index)
ms-Exch-Coexistence-Secure-Mail-Certificate-Thumbprint	ntdsSchemaAdd	attributeID: 1.2.840.113556.1.4.7000.102.51991 isMemberOfPartialAttributeSet: FALSE (not in global catalogue) searchFlags: 0 (no index)
ms-Exch-Coexistence-Secure-Mail-Certificate-Thumbprintms-Exch-Sync-Cookie	rangeUpper	1024
ms-Exch-Coexistence-Transport-Servers	ntdsSchemaAdd	attributeID: 1.2.840.113556.1.4.7000.102.51990 isMemberOfPartialAttributeSet: FALSE (not in global catalogue) searchFlags: 0 (no index)
ms-Exch-ELC-Mailbox-Flags	replace: attributeSecurityGuid	F6SzsVXskUGzJ7cuM+OK8g==

CLASS	CHANGE	ATTRIBUTE/CLASS
ms-Exch-Extension-Custom-Attribute-1	isMemberOfPartialAttributeSet:	TRUE
ms-Exch-Extension-Custom-Attribute-2	isMemberOfPartialAttributeSet:	TRUE
ms-Exch-Extension-Custom-Attribute-3	isMemberOfPartialAttributeSet:	TRUE
ms-Exch-Extension-Custom-Attribute-4	isMemberOfPartialAttributeSet:	TRUE
ms-Exch-Extension-Custom-Attribute-5	isMemberOfPartialAttributeSet	TRUE
ms-Exch-Group-External-Member-Count	ntdsSchemaModify	isMemberOfPartialAttributeSet: TRUE MAPIID:36066
ms-Exch-Group-Member-Count	ntdsSchemaModify	replace: isMemberOfPartialAttributeSetisMemberOfPartialAttributeSet: TRUE MAPIID: 36067
ms-Exch-Group-Security-Flags	ntdsSchemaModify	replace: mapiId: 36111
ms-Exch-HAB-Root-DepaPreviewent-Link	replace: isMemberOfPartialAttributeSet	TRUE
ms-Exch-Mailbox-Audit-Enable	replace: searchFlags	19
ms-Exch-Malware-Filtering-Update-Frequency	rangeUpper	38880
ms-Exch-MSO-Forward-Sync-Non-Recipient-Cookie	rangeUpper	20480
ms-Exch-MSO-Forward-Sync-Recipient-Cookie	rangeUpper	20480
ms-Exch-Role-Entries	rangeUpper	8192
ms-Exch-Schema-Version-Pt	rangeUpper	15137
ms-Exch-Schema-Version-Pt	rangeUpper	15281
Ms-exch-schema-version-pt	rangeUpper	15292
ms-Exch-Smtp-Receive-Tls-Certificate-Name	Replace: rangeUpper	1024
ms-Exch-Smtp-TLS-Certificate	replace: rangeUpper	1024
ms-Exch-Sync-Cookie	rangeUpper	262144
Top	Replace: mayContain	msExchUGMemberBL

Object IDs added by Exchange 2019 RTM

The following class object IDs are added when you install Exchange 2019 RTM:

- 1.2.840.113556.1.5.7000.62.50161
- 1.2.840.113556.1.5.7000.62.50162
- 1.2.840.113556.1.5.7000.62.50163
- 1.2.840.113556.1.5.7000.62.50164
- 1.2.840.113556.1.5.7000.62.50165
- 1.2.840.113556.1.5.7000.62.50166
- 1.2.840.113556.1.5.7000.62.50167
- 1.2.840.113556.1.5.7000.62.50170
- 1.2.840.113556.1.5.7000.62.50171
- 1.2.840.113556.1.5.7000.62.50172
- 1.2.840.113556.1.5.7000.62.50173
- 1.2.840.113556.1.5.7000.62.50174
- 1.2.840.113556.1.5.7000.62.50176
- 1.2.840.113556.1.5.7000.62.50177
- 1.2.840.113556.1.5.7000.62.50178
- 1.2.840.113556.1.5.7000.62.50187
- 1.2.840.113556.1.5.7000.62.50188
- 1.2.840.113556.1.5.7000.62.50189
- 1.2.840.113556.1.5.7000.62.50190
- 1.2.840.113556.1.5.7000.62.50191
- 1.2.840.113556.1.5.7000.62.50192
- 1.2.840.113556.1.5.7000.62.50202
- 1.2.840.113556.1.5.7000.62.50203
- 1.2.840.113556.1.5.7000.62.50204
- 1.2.840.113556.1.5.7000.62.50205
- 1.2.840.113556.1.5.7000.62.50212
- 1.2.840.113556.1.5.7000.62.50213
- 1.2.840.113556.1.5.7000.62.50214

The following attribute object IDs are added when you install Exchange 2019 RTM:

- 1.2.840.113556.1.4.2183
- 1.2.840.113556.1.4.2184
- 1.2.840.113556.1.4.2185

- 1.2.840.113556.1.4.7000.102.51773
- 1.2.840.113556.1.4.7000.102.51774
- 1.2.840.113556.1.4.7000.102.51775
- 1.2.840.113556.1.4.7000.102.51786
- 1.2.840.113556.1.4.7000.102.51787
- 1.2.840.113556.1.4.7000.102.51788
- 1.2.840.113556.1.4.7000.102.51789
- 1.2.840.113556.1.4.7000.102.51790
- 1.2.840.113556.1.4.7000.102.51791
- 1.2.840.113556.1.4.7000.102.51792
- 1.2.840.113556.1.4.7000.102.51794
- 1.2.840.113556.1.4.7000.102.51795
- 1.2.840.113556.1.4.7000.102.51796
- 1.2.840.113556.1.4.7000.102.51797
- 1.2.840.113556.1.4.7000.102.51798
- 1.2.840.113556.1.4.7000.102.51799
- 1.2.840.113556.1.4.7000.102.51800
- 1.2.840.113556.1.4.7000.102.51801
- 1.2.840.113556.1.4.7000.102.51805
- 1.2.840.113556.1.4.7000.102.51806
- 1.2.840.113556.1.4.7000.102.51807
- 1.2.840.113556.1.4.7000.102.51808
- 1.2.840.113556.1.4.7000.102.51809
- 1.2.840.113556.1.4.7000.102.51810
- 1.2.840.113556.1.4.7000.102.51811
- 1.2.840.113556.1.4.7000.102.51812
- 1.2.840.113556.1.4.7000.102.51813
- 1.2.840.113556.1.4.7000.102.51814
- 1.2.840.113556.1.4.7000.102.51815
- 1.2.840.113556.1.4.7000.102.51816
- 1.2.840.113556.1.4.7000.102.51818
- 1.2.840.113556.1.4.7000.102.51819
- 1.2.840.113556.1.4.7000.102.51820

- 1.2.840.113556.1.4.7000.102.51821
- 1.2.840.113556.1.4.7000.102.51822
- 1.2.840.113556.1.4.7000.102.51823
- 1.2.840.113556.1.4.7000.102.51824
- 1.2.840.113556.1.4.7000.102.51826
- 1.2.840.113556.1.4.7000.102.51827
- 1.2.840.113556.1.4.7000.102.51829
- 1.2.840.113556.1.4.7000.102.51830
- 1.2.840.113556.1.4.7000.102.51832
- 1.2.840.113556.1.4.7000.102.51833
- 1.2.840.113556.1.4.7000.102.51836
- 1.2.840.113556.1.4.7000.102.51837
- 1.2.840.113556.1.4.7000.102.51838
- 1.2.840.113556.1.4.7000.102.51839
- 1.2.840.113556.1.4.7000.102.51840
- 1.2.840.113556.1.4.7000.102.51851
- 1.2.840.113556.1.4.7000.102.51852
- 1.2.840.113556.1.4.7000.102.51859
- 1.2.840.113556.1.4.7000.102.51860
- 1.2.840.113556.1.4.7000.102.51861
- 1.2.840.113556.1.4.7000.102.51862
- 1.2.840.113556.1.4.7000.102.51863
- 1.2.840.113556.1.4.7000.102.51864
- 1.2.840.113556.1.4.7000.102.51865
- 1.2.840.113556.1.4.7000.102.51866
- 1.2.840.113556.1.4.7000.102.51867
- 1.2.840.113556.1.4.7000.102.51868
- 1.2.840.113556.1.4.7000.102.51869
- 1.2.840.113556.1.4.7000.102.51870
- 1.2.840.113556.1.4.7000.102.51871
- 1.2.840.113556.1.4.7000.102.51872
- 1.2.840.113556.1.4.7000.102.51873
- 1.2.840.113556.1.4.7000.102.51874

- 1.2.840.113556.1.4.7000.102.51875
- 1.2.840.113556.1.4.7000.102.51876
- 1.2.840.113556.1.4.7000.102.51877
- 1.2.840.113556.1.4.7000.102.51878
- 1.2.840.113556.1.4.7000.102.51879
- 1.2.840.113556.1.4.7000.102.51880
- 1.2.840.113556.1.4.7000.102.51881
- 1.2.840.113556.1.4.7000.102.51882
- 1.2.840.113556.1.4.7000.102.51883
- 1.2.840.113556.1.4.7000.102.51914
- 1.2.840.113556.1.4.7000.102.51915
- 1.2.840.113556.1.4.7000.102.51916
- 1.2.840.113556.1.4.7000.102.51917
- 1.2.840.113556.1.4.7000.102.51918
- 1.2.840.113556.1.4.7000.102.51919
- 1.2.840.113556.1.4.7000.102.51920
- 1.2.840.113556.1.4.7000.102.51921
- 1.2.840.113556.1.4.7000.102.51922
- 1.2.840.113556.1.4.7000.102.51923
- 1.2.840.113556.1.4.7000.102.51924
- 1.2.840.113556.1.4.7000.102.51925
- 1.2.840.113556.1.4.7000.102.51926
- 1.2.840.113556.1.4.7000.102.51927
- 1.2.840.113556.1.4.7000.102.51928
- 1.2.840.113556.1.4.7000.102.51929
- 1.2.840.113556.1.4.7000.102.51930
- 1.2.840.113556.1.4.7000.102.51931
- 1.2.840.113556.1.4.7000.102.51932
- 1.2.840.113556.1.4.7000.102.51933
- 1.2.840.113556.1.4.7000.102.51934
- 1.2.840.113556.1.4.7000.102.51935
- 1.2.840.113556.1.4.7000.102.51936
- 1.2.840.113556.1.4.7000.102.51937

- 1.2.840.113556.1.4.7000.102.51938
- 1.2.840.113556.1.4.7000.102.51939
- 1.2.840.113556.1.4.7000.102.51940
- 1.2.840.113556.1.4.7000.102.51941
- 1.2.840.113556.1.4.7000.102.51942
- 1.2.840.113556.1.4.7000.102.51943
- 1.2.840.113556.1.4.7000.102.51944
- 1.2.840.113556.1.4.7000.102.51945
- 1.2.840.113556.1.4.7000.102.51946
- 1.2.840.113556.1.4.7000.102.51947
- 1.2.840.113556.1.4.7000.102.51948
- 1.2.840.113556.1.4.7000.102.51949
- 1.2.840.113556.1.4.7000.102.51950
- 1.2.840.113556.1.4.7000.102.51951
- 1.2.840.113556.1.4.7000.102.51952
- 1.2.840.113556.1.4.7000.102.51953
- 1.2.840.113556.1.4.7000.102.51954
- 1.2.840.113556.1.4.7000.102.51955
- 1.2.840.113556.1.4.7000.102.51993
- 1.2.840.113556.1.4.7000.102.51994
- 1.2.840.113556.1.4.7000.102.51995
- 1.2.840.113556.1.4.7000.102.51996
- 1.2.840.113556.1.4.7000.102.51997
- 1.2.840.113556.1.4.7000.102.51998
- 1.2.840.113556.1.4.7000.102.52001
- 1.2.840.113556.1.4.7000.102.52002
- 1.2.840.113556.1.4.7000.102.52003
- 1.2.840.113556.1.4.7000.102.52004
- 1.2.840.113556.1.4.7000.102.52005
- 1.2.840.113556.1.4.7000.102.52006
- 1.2.840.113556.1.4.7000.102.52007
- 1.2.840.113556.1.4.7000.102.52008
- 1.2.840.113556.1.4.7000.102.52011

- 1.2.840.113556.1.4.7000.102.52012
- 1.2.840.113556.1.4.7000.102.52013
- 1.2.840.113556.1.4.7000.102.52014
- 1.2.840.113556.1.4.7000.102.52015
- 1.2.840.113556.1.4.7000.102.52016
- 1.2.840.113556.1.4.7000.102.52017
- 1.2.840.113556.1.4.7000.102.52018
- 1.2.840.113556.1.4.7000.102.52019
- 1.2.840.113556.1.4.7000.102.52020
- 1.2.840.113556.1.4.7000.102.52021
- 1.2.840.113556.1.4.7000.102.52022
- 1.2.840.113556.1.4.7000.102.52023
- 1.2.840.113556.1.4.7000.102.52024
- 1.2.840.113556.1.4.7000.102.52029
- 1.2.840.113556.1.4.7000.102.52030
- 1.2.840.113556.1.4.7000.102.52031
- 1.2.840.113556.1.4.7000.102.52032
- 1.2.840.113556.1.4.7000.102.52033
- 1.2.840.113556.1.4.7000.102.52034
- 1.2.840.113556.1.4.7000.102.52035
- 1.2.840.113556.1.4.7000.102.52036
- 1.2.840.113556.1.4.7000.102.52037
- 1.2.840.113556.1.4.7000.102.52039
- 1.2.840.113556.1.4.7000.102.52040
- 1.2.840.113556.1.4.7000.102.52041
- 1.2.840.113556.1.4.7000.102.52042
- 1.2.840.113556.1.4.7000.102.52043
- 1.2.840.113556.1.4.7000.102.52044
- 1.2.840.113556.1.4.7000.102.52045
- 1.2.840.113556.1.4.7000.102.52046
- 1.2.840.113556.1.4.7000.102.52047
- 1.2.840.113556.1.4.7000.102.52048
- 1.2.840.113556.1.4.7000.102.52049

- 1.2.840.113556.1.4.7000.102.52050
- 1.2.840.113556.1.4.7000.102.52051
- 1.2.840.113556.1.4.7000.102.52052
- 1.2.840.113556.1.4.7000.102.52053
- 1.2.840.113556.1.4.7000.102.52054
- 1.2.840.113556.1.4.7000.102.52055
- 1.2.840.113556.1.4.7000.102.52056
- 1.2.840.113556.1.4.7000.102.52057
- 1.2.840.113556.1.4.7000.102.52058
- 1.2.840.113556.1.4.7000.102.52059
- 1.2.840.113556.1.4.7000.102.52060
- 1.2.840.113556.1.4.7000.102.52061
- 1.2.840.113556.1.4.7000.102.52062
- 1.2.840.113556.1.4.7000.102.52063
- 1.2.840.113556.1.4.7000.102.52064
- 1.2.840.113556.1.4.7000.102.52065
- 1.2.840.113556.1.4.7000.102.52109
- 1.2.840.113556.1.4.7000.102.52110
- 1.2.840.113556.1.4.7000.102.52126
- 1.2.840.113556.1.4.7000.102.52127
- 1.2.840.113556.1.4.7000.102.52128
- 1.2.840.113556.1.4.7000.102.52129
- 1.2.840.113556.1.4.7000.102.52130
- 1.2.840.113556.1.4.7000.102.52151
- 1.2.840.113556.1.4.7000.102.52152
- 1.2.840.113556.1.4.7000.102.52155
- 1.2.840.113556.1.4.7000.102.52156
- 1.2.840.113556.1.4.7000.102.52157
- 1.2.840.113556.1.4.7000.102.52158
- 1.2.840.113556.1.4.7000.102.52159
- 1.2.840.113556.1.4.7000.102.52160
- 1.2.840.113556.1.4.7000.102.52161

Indexed attributes added by Exchange 2019 RTM

ATTRIBUTE	SEARCH FLAG VALUE
ms-DS-GeoCoordinates-Altitude	1
ms-DS-GeoCoordinates-Latitude	1
ms-DS-GeoCoordinates-Longitude	1
ms-Exch-Accepted-Domain-Name	9
ms-Exch-Archive-GUID	9
ms-Exch-Auth-Application-Identifier	1
ms-Exch-Auth-Issuer-Name	1
ms-Exch-Bypass-Audit	9
ms-Exch-Data-Encryption-Policy-Link	1
ms-Exch-Default-Public-Folder-Mailbox	19
ms-Exch-Device-Client-Type	1
ms-Exch-Extension-Custom-Attribute-1	1
ms-Exch-Extension-Custom-Attribute-2	1
ms-Exch-Extension-Custom-Attribute-3	1
ms-Exch-Extension-Custom-Attribute-4	1
ms-Exch-Extension-Custom-Attribute-5	1
ms-Exch-Home-MDB-SL	1
ms-Exch-Home-MTA-SL	1
ms-Exch-Is-Dirsync-Status-Pending	1
ms-Exch-Mailbox-Audit-Enable	19
ms-Exch-Mailbox-Database-Transport-Flags	16
ms-Exch-Mailbox-Move-Source-Archive-MDB-Link-SL	1
ms-Exch-Mailbox-Move-Source-MDB-Link-SL	1
ms-Exch-Mailbox-Move-Target-Archive-MDB-Link-SL	1
ms-Exch-Organization-Upgrade-Policy-Link	1

ATTRIBUTE	SEARCH FLAG VALUE
ms-Exch-Organization-Upgrade-Policy-Link-SL	1
ms-Exch-OWA-Set-Photo-URL	16
ms-Exch-Previous-Archive-Database-SL	8
ms-Exch-Previous-Home-MDB-SL	8
ms-Exch-Provisioning-Tags	1
ms-Exch-Public-Folder-EntryId	24
ms-Exch-Public-Folder-Mailbox	24
ms-Exch-Public-Folder-Smtp-Address	24
ms-Exch-Recipient-SoftDeleted-Status	27
ms-Exch-Relocate-Tenant-Completion-Target-Vector	8
ms-Exch-Relocate-Tenant-Flags	8
ms-Exch-Relocate-Tenant-Safe-Lockdown-Schedule	8
ms-Exch-Relocate-Tenant-Source-Forest	9
ms-Exch-Relocate-Tenant-Start-Lockdown	8
ms-Exch-Relocate-Tenant-Start-Retired	8
ms-Exch-Relocate-Tenant-Start-Sync	8
ms-Exch-Relocate-Tenant-Status,	9
ms-Exch-Relocate-Tenant-Target-Forest	9
ms-Exch-Relocate-Tenant-Transition-Counter	8
ms-Exch-Sync-Cookie	8
ms-Exch-Team-Mailbox-Expiration	16
ms-Exch-Team-Mailbox-Expiry-Days	16
ms-Exch-Team-Mailbox-Owners	16
ms-Exch-Team-Mailbox-SharePoint-Linked-By	16
ms-Exch-Team-Mailbox-SharePoint-Url	16

ATTRIBUTE	SEARCH FLAG VALUE
ms-Exch-Team-Mailbox-Show-In-Client-List	16
ms-Exch-Transport-Rule-Immutable-Id	1
ms-Exch-When-Soft-Deleted-Time	17

Property sets modified by Exchange 2019 RTM

The following property sets are modified when you install Exchange 2019 RTM:

- Exchange-Information

MAPI IDs added by Exchange 2019 RTM

The following MAPI IDs are added when you install Exchange 2019 RTM:

- 36066
- 36067

Extended rights added by Exchange 2019 RTM

The following table lists the extended rights that are added when you install Exchange 2019 RTM. Installing Exchange 2019 RTM doesn't modify any existing extended rights.

IDENTIFIER	VALUES
CN=ms-Exch-SMTP-Accept-XProxyFrom,CN=Extended-Rights,<ConfigurationContainerDN>	changetype: ntdsSchemaAdddisplayName: Accept XProxyFromobjectClass: controlAccessRightrightsGuid: 5bee2b72-50d7-49c7-ba66-39a25daa1e92validAccesses: 256

CURRENT EXCHANGE 2016 RELEASE INSTALLED	NEW EXCHANGE 2016 RELEASE BEING INSTALLED	ARE SCHEMA UPDATES REQUIRED?
Release to Manufacturing	Cumulative Update 4 through Cumulative Update 6	**Yes**, schema updates are required. You need to apply the CU1, CU2 and CU3 schema updates.
Cumulative Update 2	Cumulative Update 4 through Cumulative Update 6	**Yes**, schema updates are required. You need to apply the CU3 schema updates.
Cumulative Update 3	Cumulative Update 4 through Cumulative Update 6	**No**, no schema updates are required. No schema changes are made in CU4 through CU6.
Cumulative Update 6	Cumulative Update 7 through Cumulative Update 14	**Yes**, schema updates are required. You need to apply the CU14 schema updates.
Cumulative Update 7 or higher	Cumulative Update 8 through Cumulative Update 17	**No**, no schema updates are required. No schema changes are made in CU8 through CU17.

NOTE

The Active Directory schema changes that are described in this topic might not apply to all editions of an Exchange 2016 version. To verify that Active Directory has been successfully prepared, see the Exchange Active Directory versions section in Prepare Active Directory and domains for Exchange Server.

Exchange 2016 CU17 Active Directory schema changes

No changes are made to the Active Directory schema in Exchange 2016 in CU17.

Exchange 2016 CU16 Active Directory schema changes

No changes are made to the Active Directory schema in Exchange 2016 in CU16.

Exchange 2016 CU15 Active Directory schema changes

No changes are made to the Active Directory schema in Exchange 2016 in CU15.

Exchange 2016 CU14 Active Directory schema changes

No changes are made to the Active Directory schema in Exchange 2016 in CU14.

Exchange 2016 CU13 Active Directory schema changes

No changes are made to the Active Directory schema in Exchange 2016 in CU13.

Exchange 2016 CU12 Active Directory schema changes

No changes are made to the Active Directory schema in Exchange 2016 in CU12.

However, a reduction in Active Directory permissions is made: The AdminSDHolder object on the domain is updated to remove the "Allow" ACE that grants the "Exchange Trusted Subsystem" group the "Write DACL" right on the "Group" inherited object types. For more information, see KB 4490059.

Exchange 2016 CU11 Active Directory schema changes

No changes are made to the Active Directory schema in Exchange 2016 in CU11.

Exchange 2016 CU10 Active Directory schema changes

No changes are made to the Active Directory schema in Exchange 2016 in CU10.

Exchange 2016 CU9 Active Directory schema changes

No changes are made to the Active Directory schema in Exchange 2016 in CU9.

Exchange 2016 CU8 Active Directory schema changes

No changes are made to the Active Directory schema in Exchange 2016 in CU8.

Exchange 2016 CU7 Active Directory schema changes

This section summarizes the changes that are made to the Active Directory schema when you install Exchange 2016 CU7. This section includes the following subsections:

- Classes added by Exchange 2016 CU7
- Classes modified by Exchange 2016 CU7
- Attributes added by Exchange 2016 CU7
- Attributes modified by Exchange 2016 CU7
- Object IDs added by Exchange 2016 CU7

Classes added by Exchange 2016 CU7

This section contains the classes added in Exchange 2016 CU7.

CLASS	CHANGE
ms-Exch-Http-Delivery-Connector	ntdsSchemaAdd

Classes modified by Exchange 2016 CU7

This section contains the classes modified in Exchange 2016 CU7.

CLASS	CHANGE	ATTRIBUTE/CLASS
Mail-Recipient	add: mayContain	msExchImmutableSid

Attributes added by Exchange 2016 CU7

This section contains the attributes added in Exchange 2016 CU7.

- ms-Exch-Immutable-Sid

Attributes modified by Exchange 2016 CU7

This section contains the classes modified in Exchange 2016 CU7.

CLASS	CHANGE	ATTRIBUTE/CLASS
ms-Exch-Group-Security-Flags	ntdsSchemaModify	replace: mapild: 36111

Object IDs added by Exchange 2016 CU7

The following attribute object IDs are added when you install Exchange 2016 CU7:

- 1.2.840.113556.1.5.7000.62.50214
- 1.2.840.113556.1.4.7000.102.52161

Exchange 2016 CU6 Active Directory schema changes

No changes are made to the Active Directory schema in Exchange 2016 CU6.

Exchange 2016 CU5 Active Directory schema changes

No changes are made to the Active Directory schema in Exchange 2016 CU5.

Exchange 2016 CU4 Active Directory schema changes

No changes are made to the Active Directory schema in Exchange 2016 CU4.

Exchange 2016 CU3 Active Directory schema changes

This section summarizes the changes that are made to the Active Directory schema when you install Exchange 2016 CU3. This section includes the following subsections:

- Classes modified by Exchange 2016 CU3
- Attributes added by Exchange 2016 CU3
- Global catalog attributes added by Exchange 2016 CU3
- Object IDs added by Exchange 2016 CU3

Classes modified by Exchange 2016 CU3

This section contains the classes modified in Exchange 2016 CU3.

CLASS	CHANGE	ATTRIBUTE/CLASS
Mail-Recipient	add: mayContain	msExchUGEventSubscriptionLink
Top	add: mayContain	msExchUGEventSubscriptionBL

Attributes added by Exchange 2016 CU3

This section contains the attributes added in Exchange 2016 CU3.

- ms-Exch-UG-Event-Subscription-Link
- ms-Exch-UG-Event-Subscription-BL

Global catalog attributes added by Exchange 2016 CU3

This section contains the global catalog attributes added in Exchange 2016 CU3.

- ms-Exch-UG-Event-Subscription-Link
- ms-Exch-UG-Event-Subscription-BL

Object IDs added by Exchange 2016 CU3

The following attribute object IDs are added when you install Exchange 2016 CU3:

- 1.2.840.113556.1.4.7000.102.52159
- 1.2.840.113556.1.4.7000.102.52160

Exchange 2016 CU2 Active Directory schema changes

This section summarizes the changes that are made to the Active Directory schema when you install Exchange 2016 CU2. This section includes the following subsections:

- Classes modified by Exchange 2016 CU2
- Attributes added by Exchange 2016 CU2
- Global catalog attributes added by Exchange 2016 CU2
- Object IDs added by Exchange 2016 CU2

Classes modified by Exchange 2016 CU2

This section contains the classes modified in Exchange 2016 CU2.

CLASS	CHANGE	ATTRIBUTE/CLASS
Mail-Recipient	add: mayContain	msExchAdministrativeUnitLink
ms-Exch-Container	add: mayContain	msExchScopeFlags
Top	add: mayContain	msExchAdministrativeUnitBL
ms-Exch-Base-Class	add: mayContain	msExchUserHoldPolicies

Attributes added by Exchange 2016 CU2

This section contains the attributes added in Exchange 2016 CU2.

- ms-Exch-Administrative-Unit-Link
- ms-Exch-Administrative-Unit-BL

Global catalog attributes added by Exchange 2016 CU2

This section contains the global catalog attributes added in Exchange 2016 CU2.

- ms-Exch-Administrative-Unit-Link
- ms-Exch-Administrative-Unit-BL

Object IDs added by Exchange 2016 CU2

The following attribute object IDs are added when you install Exchange 2016 CU2:

- 1.2.840.113556.1.4.7000.102.52157
- 1.2.840.113556.1.4.7000.102.52158

Exchange 2016 CU1 Active Directory schema changes

This section summarizes the changes that are made to the Active Directory schema when you install Exchange 2016 CU1. This section includes the following subsections:

- Classes added by Exchange 2016 CU1
- Classes modified by Exchange 2016 CU1
- Attributes added by Exchange 2016 CU1
- Indexed attributes added by Exchange 2016 CU1
- Global catalog attributes added by Exchange 2016 CU1
- Object IDs added by Exchange 2016 CU1

Classes added by Exchange 2016 CU1

This section contains the classes added in Exchange 2016 CU1.

CLASS	CHANGE
ms-Exch-Mailbox-Policy	ntdsSchemaAdd
ms-Exch-Auth-Policy	ntdsSchemaAdd

Classes modified by Exchange 2016 CU1

This section contains the classes modified in Exchange 2016 CU1.

CLASS	CHANGE	ATTRIBUTE/CLASS
ms-Exch-Mail-Storage	add: mayContain	msExchDataEncryptionPolicyLink
ms-Exch-Organization-Container	add: mayContain	msExchDataEncryptionPolicyLink
Top	add: mayContain	msExchDataEncryptionPolicyBL
Top	add: mayContain	msExchAuthPolicyBL
Mail-Recipient	add: mayContain	msExchAuthPolicyLink
ms-Exch-Configuration-Unit-Container	add: mayContain	msExchAuthPolicyLink
ms-Exch-Configuration-Unit-Container	add: mayContain	msExchMSOForwardSyncReplayList

Attributes added by Exchange 2016 CU1

This section contains the attributes added in Exchange 2016 CU1.

- ms-Exch-Data-Encryption-Policy-Link
- ms-Exch-Data-Encryption-Policy-BL
- ms-Exch-Auth-Policy-Link
- ms-Exch-Auth-Policy-BL

Indexed attributes added by Exchange 2016 CU1

This section contains the indexed attributes added in Exchange 2016 CU1.

ATTRIBUTE	SEARCH FLAG VALUE
ms-Exch-Data-Encryption-Policy-Link	1

Global catalog attributes added by Exchange 2016 CU1

This section contains the global catalog attributes added in Exchange 2016 CU1.

- ms-Exch-Data-Encryption-Policy-Link
- ms-Exch-Data-Encryption-Policy-BL
- ms-Exch-Auth-Policy-Link

Object IDs added by Exchange 2016 CU1

The following class object IDs are added when you install Exchange 2016 CU1:

- 1.2.840.113556.1.5.7000.62.50212
- 1.2.840.113556.1.5.7000.62.50213

The following attribute object IDs are added when you install Exchange 2016 CU1:

- 1.2.840.113556.1.4.7000.102.52151
- 1.2.840.113556.1.4.7000.102.52152

- 1.2.840.113556.1.4.7000.102.52155
- 1.2.840.113556.1.4.7000.102.52156

Exchange 2016 RTM Active Directory schema changes

This section summarizes the changes that are made to the Active Directory schema when you install Release to Manufacturing (RTM) version of Exchange 2016. This section includes the following subsections:

- Classes added by Exchange 2016 RTM
- Classes modified by Exchange 2016 RTM
- Attributes added by Exchange 2016 RTM
- Global catalog attributes added by Exchange 2016 RTM
- Attributes modified by Exchange 2016 RTM
- Object IDs added by Exchange 2016 RTM
- Indexed attributes added by Exchange 2016 RTM
- Property sets modified by Exchange 2016 RTM
- MAPI IDs added by Exchange 2016 RTM
- Extended rights added by Exchange 2016 RTM

> **NOTE**
> No changes to the Active Directory schema were made between Exchange 2016 Preview and Exchange 2016 RTM.

Classes added by Exchange 2016 RTM

This section contains the classes added in Exchange 2016 RTM.

CLASS	CHANGE
Exch-Mapi-Virtual-Directory	ntdsSchemaAdd
Exch-Push-Notifications-App	ntdsSchemaAdd
ms-Exch-Account-Forest	ntdsSchemaAdd
ms-Exch-ActiveSync-Device-Autoblock-Threshold	ntdsSchemaAdd
ms-Exch-Auth-Auth-Config	ntdsSchemaAdd
ms-Exch-Auth-Auth-Server	ntdsSchemaAdd
ms-Exch-Auth-Partner-Application	ntdsSchemaAdd
ms-Exch-Client-Access-Rule	ntdsSchemaModify
ms-Exch-Config-Settings	ntdsSchemaAdd
ms-Exch-Encryption-Virtual-Directory	ntdsSchemaAdd

CLASS	CHANGE
ms-Exch-Exchange-Transport-Server	ntdsSchemaAdd
ms-Exch-Hosted-Content-Filter-Config	ntdsSchemaAdd
ms-Exch-Hygiene-Configuration	ntdsSchemaAdd
ms-Exch-Intra-Organization-Connector	ntdsSchemaModify
ms-Exch-MSO-Forward-Sync-Divergence	ntdsSchemaAdd
ms-Exch-MSO-Sync-Service-Instance	ntdsSchemaAdd
ms-Exch-Mailflow-Policy	ntdsSchemaAdd
ms-Exch-Mailflow-Policy-Collection	ntdsSchemaAdd
ms-Exch-Malware-Filter-Config	ntdsSchemaAdd
ms-Exch-Organization-Upgrade-Policy	ntdsSchemaAdd
ms-Exch-Protocol-Cfg-SIP-Container	ntdsSchemaAdd
ms-Exch-Protocol-Cfg-SIP-FE-Server	ntdsSchemaAdd
ms-Exch-Resource-Policy	ntdsSchemaAdd
ms-Exch-Safe-Attachment-Protection-Config	ntdsSchemaAdd
ms-Exch-Smart-Links-Protection-Config	ntdsSchemaAdd
ms-Exch-Team-Mailbox-Provisioning-Policy	ntdsSchemaAdd
ms-Exch-Throttling-Policy	ntdsSchemaModify
ms-Exch-Unified-Policy	ntdsSchemaAdd
ms-Exch-Unified-Rule	ntdsSchemaAdd
ms-Exch-Workload-Policy	ntdsSchemaAdd

Classes modified by Exchange 2016 RTM

This section contains the classes modified in Exchange 2016 RTM.

CLASS	CHANGE	ATTRIBUTE/CLASS
Exch-Accepted-Domain	add:mayContain	msExchOfflineOrgIdHomeRealmRecord
Exch-Base-Class	add:mayContain	msExchCapabilityIdentifiers
Exch-Base-Class	add:mayContain	msExchObjectID

CLASS	CHANGE	ATTRIBUTE/CLASS
Exch-Base-Class	add:mayContain	msExchProvisioningTags
Exch-Configuration-Unit-Container	add:mayContain	msExchArchiveRelease
Exch-Configuration-Unit-Container	add:mayContain	msExchMailboxRelease
Exch-Exchange-Server	add:mayContain	msExchArchiveRelease
Exch-Exchange-Server	add:mayContain	msExchMailboxRelease
Exch-MDB-Availability-Group	add:mayContain	msExchEvictedMembersLink
Exch-OAB	add:mayContain	msExchLastUpdateTime
Exch-OWA-Mailbox-Policy	add:mayContain	msExchConfigurationXML
Exch-OWA-Virtual-Directory	add:mayContain	msExchConfigurationXML
Exch-On-Premises-Organization	add:mayContain	msExchTrustedDomainLink
Exch-Organization-Container	add:mayContain	msExchMaxABP
Exch-Organization-Container	add:mayContain	msExchMaxOAB
Exch-Organization-Container	add:mayContain	pFContacts
Exch-Team-Mailbox-Provisioning-Policy	add:mayContain	msExchConfigurationXML
Group	add: auxiliaryClass	msExchMailStorage
Mail-Recipient	add:mayContain	msExchLocalizationFlags
Mail-Recipient	add:mayContain	msExchRoleGroupType
Mail-Recipient	add:mayContain	ms-DS-GeoCoordinates-Altitude
Mail-Recipient	add:mayContain	ms-DS-GeoCoordinates-Latitude
Mail-Recipient	add:mayContain	ms-DS-GeoCoordinates-Longitude
Mail-Recipient	add:mayContain	msExchRecipientSoftDeletedStatus
Mail-Recipient	add:mayContain	msExchWhenSoftDeletedTime
Mail-Recipient	add:mayContain	msExchHomeMTASL
Mail-Recipient	add:mayContain	msExchMailboxMoveSourceArchiveMDBLinkSL
Mail-Recipient	add:mayContain	msExchMailboxMoveSourceMDBLinkSL

CLASS	CHANGE	ATTRIBUTE/CLASS
Mail-Recipient	add:mayContain	msExchMailboxMoveTargetArchiveMDBLinkSL
Mail-Recipient	add:mayContain	msExchMailboxMoveTargetMDBLinkSL
Mail-Recipient	add:mayContain	ms-exch-group-external-member-count
Mail-Recipient	add:mayContain	ms-exch-group-member-count
Mail-Recipient	add:mayContain	msExchGroupExternalMemberCount
Mail-Recipient	add:mayContain	msExchGroupMemberCount
Mail-Recipient	add:mayContain	msExchShadowWhenSoftDeletedTime
Mail-Recipient	add:mayContain	msExchPublicFolderMailbox
Mail-Recipient	add:mayContain	msExchPublicFolderSmtpAddress
Mail-Recipient	add: mayContain	msExchAuxMailboxParentObjectIdLink
Mail-Recipient	add: mayContain	msExchStsRefreshTokensValidFrom
Mail-Recipient	add:mayContain	msDS-ExternalDirectoryObjectId
Mail-Recipient	add:mayContain	msExchGroupSecurityFlags
Mail-Recipient	add:mayContain	msExchMultiMailboxDatabasesLink
Ms-Exch-Organization-Container	add:mayContain	ms-exch-organization-flags-2
Top	add:mayContain	msExchMultiMailboxDatabasesBL
Top	add:mayContain	msExchMultiMailboxLocationsBL
Top	add:mayContain	msExchAccountForestBL
Top	add:mayContain	msExchTrustedDomainBL
Top	add:mayContain	msExchAcceptedDomainBL
Top	add:mayContain	msExchHygieneConfigurationMalwareBL
Top	add:mayContain	msExchHygieneConfigurationSpamBL
Top	add:mayContain	msExchEvictedMembersBL
Top	add: mayContain	msExchOABGeneratingMailboxBL

CLASS	CHANGE	ATTRIBUTE/CLASS
Top	add: mayContain	msExchAuxMailboxParentObjectIdBL
ms-Exch-Accepted-Domain	add:mayContain	msExchHygieneConfigurationLink
ms-Exch-Accepted-Domain	add:mayContain	msExchTransportResellerSettingsLinkSL
ms-Exch-Account-Forest	possSuperiors	msExchContainer
ms-Exch-Account-Forest	Add:mayContain	msExchPartnerId
ms-Exch-Active-Sync-Device	add:mayContain	msExchDeviceClientType
ms-Exch-Availability-Address-Space	add:mayContain	msExchFedTargetAutodiscoverEPR
ms-Exch-Base-Class	add:mayContain	msExchDirsyncAuthorityMetadata
ms-Exch-Base-Class	add:mayContain	msExchDirsyncStatusAck
ms-Exch-Base-Class	add:mayContain	msExchEdgeSyncConfigFlags
ms-Exch-Base-Class	add:mayContain	msExchHABRootDepaPreviewentLink
ms-Exch-Base-Class	add:mayContain	msExchDefaultPublicFolderMailbox
ms-Exch-Base-Class	add:mayContain	msExchForestModeFlag
ms-Exch-Base-Class	add:mayContain	msExchELCMailboxFlags
ms-Exch-Base-Class	add:mayContain	msExchCanaryData0
ms-Exch-Base-Class	add:mayContain	msExchCanaryData1
ms-Exch-Base-Class	add:mayContain	msExchCanaryData2
ms-Exch-Base-Class	add:mayContain	msExchCorrelationId
ms-Exch-Base-Class	Add:mayContain	msExchTenantCountry
ms-Exch-Base-Class	Add:mayContain	msExchConfigurationXML
ms-Exch-Base-Class	add: mayContain	msExchMultiMailboxGUIDs
ms-Exch-Base-Class	add: mayContain	msExchMultiMailboxLocationsLink
ms-Exch-Coexistence-Relationship	add:mayContain	msExchCoexistenceOnPremisesSmartHost
ms-Exch-Coexistence-Relationship	add:mayContain	msExchCoexistenceSecureMailCertificateThumbprint

CLASS	CHANGE	ATTRIBUTE/CLASS
ms-Exch-Coexistence-Relationship	add:mayContain	msExchCoexistenceTransportServers
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchDirsyncStatus
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchIsDirsyncStatusPending
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchDirSyncServiceInstance
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchOrganizationUpgradePolicyLink
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchManagementSiteLinkSL
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchOrganizationUpgradePolicyLinkSL
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchRelocateTenantCompletionTargetVector
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchRelocateTenantFlags
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchRelocateTenantSafeLockdownSchedule
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchRelocateTenantSourceForest
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchRelocateTenantStartLockdown
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchRelocateTenantStartRetired
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchRelocateTenantStartSync
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchRelocateTenantStatus
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchRelocateTenantTargetForest
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchRelocateTenantTransitionCounter
ms-Exch-Configuration-Unit-Container	add:mayContain	msExchSyncCookie
ms-Exch-Control-Point-Config	add:mayContain	msExchRMSOnlineCertificationLocationUrl
ms-Exch-Control-Point-Config	add:mayContain	msExchRMSOnlineKeySharingLocationUrl
ms-Exch-Control-Point-Config	add:mayContain	msExchRMSOnlineLicensingLocationUrl
ms-Exch-Custom-Attributes	add:mayContain	msExchExtensionCustomAttribute1
ms-Exch-Custom-Attributes	add:mayContain	msExchExtensionCustomAttribute2

CLASS	CHANGE	ATTRIBUTE/CLASS
ms-Exch-Custom-Attributes	add:mayContain	msExchExtensionCustomAttribute3
ms-Exch-Custom-Attributes	add:mayContain	msExchExtensionCustomAttribute4
ms-Exch-Custom-Attributes	add:mayContain	msExchExtensionCustomAttribute5
ms-Exch-Domain-Content-Config	add:mayContain	msExchContentByteEncoderTypeFor7BitCharsets
ms-Exch-Domain-Content-Config	add:mayContain	msExchContentPreferredInternetCodePageForShiftJis
ms-Exch-Domain-Content-Config	add:mayContain	msExchContentRequiredCharSetCoverage
ms-Exch-Exchange-Server	add:mayContain	msExchWorkloadManagementPolicyLink
ms-Exch-Exchange-Server	add:mayContain	msExchMalwareFilteringDeferAttempts
ms-Exch-Exchange-Server	add:mayContain	msExchMalwareFilteringDeferWaitTime
ms-Exch-Exchange-Server	add:mayContain	msExchMalwareFilteringFlags
ms-Exch-Exchange-Server	add:mayContain	msExchMalwareFilteringPrimaryUpdatePath
ms-Exch-Exchange-Server	add:mayContain	msExchMalwareFilteringSecondaryUpdatePath
ms-Exch-Exchange-Server	add:mayContain	msExchMalwareFilteringUpdateFrequency
ms-Exch-Exchange-Server	add:mayContain	msExchMalwareFilteringUpdateTimeout
ms-Exch-Exchange-Server	add:mayContain	msExchMalwareFilteringScanTimeout
ms-Exch-Fed-OrgId	add:mayContain	msExchFedDelegationTrustSL
ms-Exch-Hosted-Content-Filter-Config	add:mayContain	msExchSpamCountryBlockList
ms-Exch-Hosted-Content-Filter-Config	add:mayContain	msExchSpamLanguageBlockList
ms-Exch-Hosted-Content-Filter-Config	add:mayContain	msExchSpamNotifyOutboundRecipients
ms-Exch-Hosted-Content-Filter-Config	add:mayContain	msExchSpamDigestFrequency
ms-Exch-Hosted-Content-Filter-Config	add:mayContain	msExchSpamQuarantineRetention
ms-Exch-MDB	add:mayContain	msExchCalendarLoggingQuota

CLASS	CHANGE	ATTRIBUTE/CLASS
ms-Exch-MRS-Request	add:mayContain	msExchMailboxMoveSourceMDBLinkSL
ms-Exch-MRS-Request	add:mayContain	msExchMailboxMoveStorageMDBLinkSL
ms-Exch-MRS-Request	add:mayContain	msExchMailboxMoveTargetMDBLinkSL
ms-Exch-MSO-Sync-Service-Instance	add:mayContain	msExchMSOForwardSyncNonRecipientCookie
ms-Exch-MSO-Sync-Service-Instance	add:mayContain	msExchMSOForwardSyncRecipientCookie
ms-Exch-MSO-Sync-Service-Instance	add:mayContain	msExchMSOForwardSyncReplayList
ms-Exch-MSO-Sync-Service-Instance	add:mayContain	msExchAccountForestLink
ms-Exch-MSO-Sync-Service-Instance	add:mayContain	msExchActiveInstanceSleepInterval
ms-Exch-MSO-Sync-Service-Instance	add:mayContain	msExchPassiveInstanceSleepInterval
ms-Exch-MSO-Sync-Service-Instance	add:mayContain	msExchSyncDaemonMaxVersion
ms-Exch-MSO-Sync-Service-Instance	add:mayContain	msExchSyncDaemonMinVersion
ms-Exch-MSO-Sync-Service-Instance	add:mayContain	msExchSyncServiceInstanceNewTenantMaxVersion
ms-Exch-MSO-Sync-Service-Instance	add:mayContain	msExchSyncServiceInstanceNewTenantMinVersion
ms-Exch-Mail-Gateway	add:mayContain	msExchHomeMDBSL
ms-Exch-Mail-Gateway	add:mayContain	msExchHomeMTASL
ms-Exch-Mail-Storage	add:mayContain	msExchPreviousArchiveDatabase
ms-Exch-Mail-Storage	add:mayContain	msExchTeamMailboxExpiration
ms-Exch-Mail-Storage	add:mayContain	msExchTeamMailboxOwners
ms-Exch-Mail-Storage	add:mayContain	msExchTeamMailboxSharePointLinkedBy
ms-Exch-Mail-Storage	add:mayContain	msExchTeamMailboxSharePointUrl
ms-Exch-Mail-Storage	add:mayContain	msExchTeamMailboxShowInClientList
ms-Exch-Mail-Storage	add:mayContain	msExchCalendarLoggingQuota
ms-Exch-Mail-Storage	add:mayContain	msExchArchiveDatabaseLinkSL
ms-Exch-Mail-Storage	add:mayContain	msExchDisabledArchiveDatabaseLinkSL

CLASS	CHANGE	ATTRIBUTE/CLASS
ms-Exch-Mail-Storage	add:mayContain	msExchHomeMDBSL
ms-Exch-Mail-Storage	add:mayContain	msExchMailboxMoveTargetMDBLinkSL
ms-Exch-Mail-Storage	add:mayContain	msExchPreviousArchiveDatabaseSL
ms-Exch-Mail-Storage	add:mayContain	msExchPreviousHomeMDBSL
ms-Exch-Mail-Storage	add: mayContain	msExchMailboxContainerGuid
ms-Exch-Mail-Storage	add: mayContain	msExchUnifiedMailbox
ms-Exch-Mail-Storage	add:mayContain	msExchUserCulture
ms-Exch-Mailflow-Policy	add:mayContain	msExchImmutableId
ms-Exch-Malware-Filter-Config	add:mayContain	msExchMalwareFilterConfigExternalSenderAdminAddress
ms-Exch-Malware-Filter-Config	add:mayContain	msExchMalwareFilterConfigInternalSenderAdminAddress
ms-Exch-OAB	add:mayContain	msExchOffLineABServerSL
ms-Exch-OAB	add: mayContain	msExchOABGeneratingMailboxLink
ms-Exch-OWA-Mailbox-Policy	add:mayContain	msExchOWASetPhotoURL
ms-Exch-OWA-Virtual-Directory	add:mayContain	msExchOWASetPhotoURL
ms-Exch-Organization-Container	add:mayContain	msExchOrganizationFlags2
ms-Exch-Organization-Container	add:mayContain	msExchUMAvailableLanguages
ms-Exch-Organization-Container	add:mayContain	msExchWACDiscoveryEndpoint
ms-Exch-Organization-Container	add:mayContain	msExchAdfsAuthenticationRawConfiguration
ms-Exch-Organization-Container	add:mayContain	msExchServiceEndPointURL
ms-Exch-Private-MDB	add:mayContain	msExchMailboxDatabaseTransportFlags
ms-Exch-Public-Folder	add:mayContain	msExchPublicFolderEntryId
ms-Exch-Resource-Policy	add:mayContain	msExchCustomerExpectationCritical
ms-Exch-Resource-Policy	add:mayContain	msExchDiscretionaryCritical
ms-Exch-Resource-Policy	add:mayContain	msExchInternalMaintenanceCritical

CLASS	CHANGE	ATTRIBUTE/CLASS
ms-Exch-Resource-Policy	add:mayContain	msExchUrgentCritical
ms-Exch-Routing-Group-Connector	add:mayContain	msExchHomeMTASL
ms-Exch-Safe-Attachment-Protection-Config	add:mayContain	msExchMalwareFilterConfigFlags
ms-Exch-Safe-Attachment-Protection-Config	add:mayContain	msExchMalwareFilterConfigFromAddress
ms-Exch-Safe-Attachment-Protection-Config	add:mayContain	msExchMalwareFilterConfigInternalBody
ms-Exch-Safe-Attachment-Protection-Config	add:mayContain	msExchMalwareFilterConfigInternalSenderAdminAddress
ms-Exch-Safe-Attachment-Protection-Config	add:mayContain	msExchMalwareFilterConfigInternalSubject
ms-Exch-Safe-Attachment-Protection-Config	add:mayContain	msExchMalwareFilteringScanTimeout
ms-Exch-Safe-Attachment-Protection-Config	add:mayContain	msExchMalwareFilteringUpdateFrequency
ms-Exch-Site-Connector	add:mayContain	msExchHomeMTASL
ms-Exch-Smart-Links-Protection-Config	add:mayContain	msExchAddressRewriteExceptionList
ms-Exch-Smart-Links-Protection-Config	add:mayContain	msExchSpamFlags
ms-Exch-Tenant-Perimeter-Settings	add:mayContain	msExchTransportResellerSettingsLinkSL
ms-Exch-Throttling-Policy	add:mayContain	msExchThrottlingPolicyFlags
ms-Exch-Throttling-Policy	add:mayContain	msExchAnonymousThrottlingPolicyStateEx
ms-Exch-Throttling-Policy	add:mayContain	msExchEASThrottlingPolicyStateEx
ms-Exch-Throttling-Policy	add:mayContain	msExchEWSThrottlingPolicyStateEx
ms-Exch-Throttling-Policy	add:mayContain	msExchGeneralThrottlingPolicyStateEx
ms-Exch-Throttling-Policy	add:mayContain	msExchIMAPThrottlingPolicyStateEx
ms-Exch-Throttling-Policy	add:mayContain	msExchOWAThrottlingPolicyStateEx
ms-Exch-Throttling-Policy	add:mayContain	msExchPOPThrottlingPolicyStateEx

CLASS	CHANGE	ATTRIBUTE/CLASS
ms-Exch-Throttling-Policy	add:mayContain	msExchPowershellThrottlingPolicyStateEx
ms-Exch-Throttling-Policy	add:mayContain	msExchRCAThrottlingPolicyStateEx
ms-Exch-Transport-Rule	add:mayContain	msExchTransportRuleImmutableId
ms-Exch-Transport-Rule	add:mayContain	msExchImmutableId
ms-Exch-Transport-Settings	add:mayContain	msExchTranspoPreviewaxRetriesForLocalSiteShadow
ms-Exch-Transport-Settings	add:mayContain	msExchTranspoPreviewaxRetriesForRemoteSiteShadow
ms-Exch-Transport-Settings	add:mayContain	msExchConfigurationXML
ms-Exch-Virtual-Directory	add:mayContain	msExchMRSProxyFlags
ms-Exch-Virtual-Directory	add:mayContain	msExchMRSProxyMaxConnections

Attributes added by Exchange 2016 RTM

This section contains the attributes added in Exchange 2016 RTM.

- ms-DS-GeoCoordinates-Altitude
- ms-DS-GeoCoordinates-Latitude
- ms-DS-GeoCoordinates-Longitude
- ms-Exch-Accepted-Domain-BL
- ms-Exch-Account-Forest-BL
- ms-Exch-Account-Forest-Link
- ms-Exch-ActiveSync-Device-AutoBlock-Duration
- ms-Exch-ActiveSync-Device-Autoblock-Threshold-Incidence-Duration
- ms-Exch-ActiveSync-Device-Autoblock-Threshold-Incidence-Limit
- ms-Exch-ActiveSync-Device-Autoblock-Threshold-Type
- ms-Exch-Adfs-Authentication-Raw-Configuration
- ms-Exch-Anonymous-Throttling-Policy-State-Ex
- ms-Exch-Archive-Database-Link-SL
- ms-Exch-Auth-App-Secret
- ms-Exch-Auth-Application-Identifier
- ms-Exch-Auth-Auth-Server-Type
- ms-Exch-Auth-Authorization-Url

- ms-Exch-Auth-Certificate-Data
- ms-Exch-Auth-Certificate-Thumbprint
- ms-Exch-Auth-Flags
- ms-Exch-Auth-Issuer-Name
- ms-Exch-Auth-Issuing-Url
- ms-Exch-Auth-Linked-Account
- ms-Exch-Auth-Metadata-Url
- ms-Exch-Auth-Realm
- ms-Exch-Aux-Mailbox-Parent-Object-Id-BL
- ms-Exch-Aux-Mailbox-Parent-Object-Id-Link
- ms-Exch-Canary-Data-0
- ms-Exch-Canary-Data-1
- ms-Exch-Canary-Data-2
- ms-Exch-Content-Byte-Encoder-Type-For-7-Bit-Charsets
- ms-Exch-Content-Preferred-Internet-Code-Page-For-Shift-Jis
- ms-Exch-Content-Required-Char-Set-Coverage
- ms-Exch-Correlation-Id
- ms-Exch-Customer-Expectation-Critical
- ms-Exch-Customer-Expectation-Overloaded
- ms-Exch-Customer-Expectation-Underloaded
- ms-Exch-Default-Public-Folder-Mailbox
- ms-Exch-Device-Client-Type
- ms-Exch-Dir-Sync-Service-Instance
- ms-Exch-Dirsync-Authority-Metadata
- ms-Exch-Dirsync-Status
- ms-Exch-Dirsync-Status-Ack
- ms-Exch-Disabled-Archive-Database-Link-SL
- ms-Exch-Discretionary-Critical
- ms-Exch-Discretionary-Overloaded
- ms-Exch-Discretionary-Underloaded
- ms-Exch-EAS-Throttling-Policy-State-Ex
- ms-Exch-EWS-Throttling-Policy-State-Ex
- ms-Exch-Edge-Sync-Config-Flags

- ms-Exch-Encryption-Throttling-Policy-State-Ex
- ms-Exch-Extension-Custom-Attribute-1
- ms-Exch-Extension-Custom-Attribute-2
- ms-Exch-Extension-Custom-Attribute-3
- ms-Exch-Extension-Custom-Attribute-4
- ms-Exch-Extension-Custom-Attribute-5
- ms-Exch-External-Directory-Object-Class
- ms-Exch-Fed-Delegation-Trust-SL
- ms-Exch-Forest-Mode-Flag
- ms-Exch-General-Throttling-Policy-State-Ex
- ms-Exch-Group-External-Member-Count
- ms-Exch-Group-Member-Count
- ms-Exch-Home-MDB-SL
- ms-Exch-Home-MTA-SL
- ms-Exch-Hosted-Content-Filter-Config-Link
- ms-Exch-Hygiene-Configuration-Link
- ms-Exch-Hygiene-Configuration-Malware-BL
- ms-Exch-Hygiene-Configuration-Spam-BL
- ms-Exch-IMAP-Throttling-Policy-State-Ex
- ms-Exch-Internal-Maintenance-Critical
- ms-Exch-Internal-Maintenance-Overloaded
- ms-Exch-Internal-Maintenance-Underloaded
- ms-Exch-Is-Dirsync-Status-Pending,
- ms-Exch-Localization-Flags
- ms-Exch-MRS-Proxy-Flags
- ms-Exch-MRS-Proxy-Max-Connections
- ms-Exch-MSO-Forward-Sync-Divergence-Count
- ms-Exch-MSO-Forward-Sync-Divergence-Related-Object-Link
- ms-Exch-MSO-Forward-Sync-Divergence-Timestamp
- ms-Exch-Mailbox-Database-Transport-Flags
- ms-Exch-Mailbox-Move-Source-Archive-MDB-Link-SL
- ms-Exch-Mailbox-Move-Source-MDB-Link-SL
- ms-Exch-Mailbox-Move-Storage-MDB-Link-SL

- ms-Exch-Mailbox-Move-Target-Archive-MDB-Link-SL
- ms-Exch-Mailbox-Move-Target-MDB-Link-SL
- ms-Exch-Mailflow-Policy-Countries
- ms-Exch-Mailflow-Policy-Keywords
- ms-Exch-Mailflow-Policy-Publisher-Name
- ms-Exch-Mailflow-Policy-Transport-Rules-Template-Xml
- ms-Exch-Mailflow-Policy-Version
- ms-Exch-Malware-Filter-Config-Alert-Text
- ms-Exch-Malware-Filter-Config-External-Body
- ms-Exch-Malware-Filter-Config-External-Sender-Admin-Address
- ms-Exch-Malware-Filter-Config-External-Subject
- ms-Exch-Malware-Filter-Config-Flags
- ms-Exch-Malware-Filter-Config-From-Address
- ms-Exch-Malware-Filter-Config-From-Name
- ms-Exch-Malware-Filter-Config-Internal-Body
- ms-Exch-Malware-Filter-Config-Internal-Sender-Admin-Address
- ms-Exch-Malware-Filter-Config-Internal-Subject
- ms-Exch-Malware-Filter-Config-Link
- ms-Exch-Malware-Filtering-Defer-Attempts
- ms-Exch-Malware-Filtering-Defer-Wait-Time
- ms-Exch-Malware-Filtering-Flags
- ms-Exch-Malware-Filtering-Primary-Update-Path
- ms-Exch-Malware-Filtering-Scan-Timeout
- ms-Exch-Malware-Filtering-Secondary-Update-Path
- ms-Exch-Malware-Filtering-Update-Frequency
- ms-Exch-Malware-Filtering-Update-Timeout
- ms-Exch-Management-Site-Link-SL
- ms-Exch-Multi-Mailbox-GUID
- ms-Exch-Multi-Mailbox-Locations-Link
- ms-Exch-OAB-Generating-Mailbox-BL
- ms-Exch-OAB-Generating-Mailbox-Link
- ms-Exch-OWA-Set-Photo-URL
- ms-Exch-OWA-Throttling-Policy-State-Ex

- ms-Exch-Off-Line-AB-Server-SL
- ms-Exch-Organization-Flags-2
- ms-Exch-Organization-Upgrade-Policy-BL
- ms-Exch-Organization-Upgrade-Policy-Date
- ms-Exch-Organization-Upgrade-Policy-Enabled
- ms-Exch-Organization-Upgrade-Policy-Link
- ms-Exch-Organization-Upgrade-Policy-Link-SL
- ms-Exch-Organization-Upgrade-Policy-MaxMailboxes
- ms-Exch-Organization-Upgrade-Policy-Priority
- ms-Exch-Organization-Upgrade-Policy-Source-Version
- ms-Exch-Organization-Upgrade-Policy-Status
- ms-Exch-Organization-Upgrade-Policy-Target-Version
- ms-Exch-POP-Throttling-Policy-State-Ex
- ms-Exch-Powershell-Throttling-Policy-State-Ex
- ms-Exch-Previous-Archive-Database
- ms-Exch-Previous-Archive-Database-SL
- ms-Exch-Previous-Home-MDB-SL
- ms-Exch-Public-Folder-EntryId
- ms-Exch-Public-Folder-Mailbox
- ms-Exch-Public-Folder-Smtp-Address
- ms-Exch-RCA-Throttling-Policy-State-Ex
- ms-Exch-RMS-Computer-Accounts-Link-SL
- ms-Exch-RMSOnline-Certification-Location-Url
- ms-Exch-RMSOnline-Key-Sharing-Location-Url
- ms-Exch-RMSOnline-Licensing-Location-Url
- ms-Exch-Recipient-SoftDeleted-Status
- ms-Exch-Relocate-Tenant-Completion-Target-Vector
- ms-Exch-Relocate-Tenant-Flags
- ms-Exch-Relocate-Tenant-Safe-Lockdown-Schedule
- ms-Exch-Relocate-Tenant-Source-Forest
- ms-Exch-Relocate-Tenant-Start-Lockdown
- ms-Exch-Relocate-Tenant-Start-Retired
- ms-Exch-Relocate-Tenant-Start-Sync

- ms-Exch-Relocate-Tenant-Status
- ms-Exch-Relocate-Tenant-Target-Forest
- ms-Exch-Relocate-Tenant-Transition-Counter
- ms-Exch-Resource-Type
- ms-Exch-RoleGroup-Type
- ms-Exch-Service-End-Point-URL
- ms-Exch-Shadow-When-Soft-Deleted-Time
- ms-Exch-Spam-Add-Header
- ms-Exch-Spam-Asf-Settings
- ms-Exch-Spam-Asf-Test-Bcc-Address
- ms-Exch-Spam-Country-Block-List
- ms-Exch-Spam-Digest-Frequency
- ms-Exch-Spam-False-Positive-Cc
- ms-Exch-Spam-Flags
- ms-Exch-Spam-Language-Block-List
- ms-Exch-Spam-Modify-Subject
- ms-Exch-Spam-Notify-Outbound-Recipients
- ms-Exch-Spam-Outbound-Spam-Cc
- ms-Exch-Spam-Quarantine-Retention
- ms-Exch-Spam-Redirect-Address
- ms-Exch-Sts-Refresh-Tokens-Valid-From
- ms-Exch-Sync-Cookie
- ms-Exch-Sync-Service-Instance-New-Tenant-Max-Version
- ms-Exch-Sync-Service-Instance-New-Tenant-Min-Version
- ms-Exch-Team-Mailbox-Expiration
- ms-Exch-Team-Mailbox-Expiry-Days
- ms-Exch-Team-Mailbox-Owners
- ms-Exch-Team-Mailbox-SharePoint-Linked-By
- ms-Exch-Team-Mailbox-SharePoint-Url
- ms-Exch-Team-Mailbox-Show-In-Client-List
- ms-Exch-Tenant-Country
- ms-Exch-Throttling-Policy-Flags
- ms-Exch-Transport-MaxRetriesForLocalSiteShadow

- ms-Exch-Transport-MaxRetriesForRemoteSiteShadow
- ms-Exch-Transport-Reseller-Settings-Link-SL
- ms-Exch-Transport-Rule-Immutable-Id
- ms-Exch-Trusted-Domain-BL
- ms-Exch-Trusted-Domain-Link
- ms-Exch-UG-Member-BL
- ms-Exch-UG-Member-Link
- ms-Exch-Urgent-Critical
- ms-Exch-Urgent-Overloaded
- ms-Exch-Urgent-Underloaded
- ms-Exch-WAC-Discovery-Endpoint
- ms-Exch-When-Soft-Deleted-Time
- ms-Exch-Workload-Classification
- ms-Exch-Workload-Management-Is-Enabled
- ms-Exch-Workload-Management-Policy
- ms-Exch-Workload-Management-Policy-BL
- ms-Exch-Workload-Management-Policy-Link
- ms-DS-External-Directory-Object-Id
- ms-Exch-Group-Security-Flags
- ms-Exch-Multi-Mailbox-Locations-BL
- ms-Exch-Multi-Mailbox-Databases-Link
- ms-Exch-Multi-Mailbox-Databases-BL

Global catalog attributes added by Exchange 2016 RTM

The following global catalog attributes are added by Exchange 2016 RTM:

- ms-Exch-Archive-Database-Link-SL
- ms-Exch-Correlation-Id
- ms-Exch-Default-Public-Folder-Mailbox
- ms-Exch-Device-Client-Type
- ms-Exch-Dirsync-Authority-Metadata
- ms-Exch-Dirsync-Status
- ms-Exch-Dirsync-Status-Ack
- ms-Exch-Disabled-Archive-Database-Link-SL
- ms-Exch-Edge-Sync-Config-Flags
- ms-Exch-EvictedMembers-Link

- ms-Exch-EvictedMembers-BL
- ms-Exch-Extension-Custom-Attribute-1
- ms-Exch-Extension-Custom-Attribute-2
- ms-Exch-Extension-Custom-Attribute-3
- ms-Exch-Extension-Custom-Attribute-4
- ms-Exch-Extension-Custom-Attribute-5
- ms-Exch-Group-External-Member-Count
- ms-Exch-Group-Member-Count
- ms-Exch-HAB-Root-DepaPreviewent-Link
- ms-Exch-Home-MDB-SL
- ms-Exch-Home-MTA-SL
- ms-Exch-Is-Dirsync-Status-Pending
- ms-Exch-Localization-Flags
- ms-Exch-Mailbox-Container-Guid
- ms-Exch-Mailbox-Move-Source-Archive-MDB-Link-SL
- ms-Exch-Mailbox-Move-Source-MDB-Link-SL
- ms-Exch-Mailbox-Move-Storage-MDB-Link-SL
- ms-Exch-Mailbox-Move-Target-Archive-MDB-Link-SL
- ms-Exch-Mailbox-Move-Target-MDB-Link-SL
- ms-Exch-Offline-OrgId-Home-Realm-Record
- ms-Exch-Previous-Archive-Database
- ms-Exch-Previous-Archive-Database-SL
- ms-Exch-Previous-Home-MDB-SL
- ms-Exch-RMS-Computer-Accounts-Link-SL
- ms-Exch-Recipient-SoftDeleted-Status
- ms-Exch-Relocate-Tenant-Completion-Target-Vector,
- ms-Exch-Relocate-Tenant-Flags
- ms-Exch-Relocate-Tenant-Safe-Lockdown-Schedule
- ms-Exch-Relocate-Tenant-Source-Forest
- ms-Exch-Relocate-Tenant-Start-Lockdown
- ms-Exch-Relocate-Tenant-Start-Retired
- ms-Exch-Relocate-Tenant-Start-Sync
- ms-Exch-Relocate-Tenant-Status

- ms-Exch-Relocate-Tenant-Target-Forest
- ms-Exch-Relocate-Tenant-Transition-Counter
- ms-Exch-RoleGroup-Type
- ms-Exch-Sync-Cookie
- ms-Exch-Team-Mailbox-Expiration
- ms-Exch-Team-Mailbox-Expiry-Days
- ms-Exch-Team-Mailbox-Owners
- ms-Exch-Team-Mailbox-SharePoint-Linked-By
- ms-Exch-Team-Mailbox-SharePoint-Url
- ms-Exch-Team-Mailbox-Show-In-Client-List
- ms-Exch-Unified-Mailbox
- ms-Exch-When-Soft-Deleted-Time

Attributes modified by Exchange 2016 RTM

This section contains the attributes modified in Exchange 2016 RTM.

ATTRIBUTE	CHANGE	VALUE
Exch-Configuration-Unit-Container	rangeUpper	15254
Exch-Mailflow-Policy-Transport-Rules-Template-Xml	rangeUpper	256000
Mail-Recipient	Replace: mayContain	msExchUGMemberLink
Ms-exch-schema-version-pt	rangeUpper	15292
Top	Replace: mayContain	msExchUGMemberBL
ms-Exch-Accepted-Domain-Name	replace: searchFlags	9
ms-Exch-Archive-GUID	replace: searchFlags	9
ms-Exch-Bypass-Audit	replace: searchFlags	19
ms-Exch-Coexistence-On-Premises-Smart-Host	ntdsSchemaAdd	attributeID: 1.2.840.113556.1.4.7000.102.51992 isMemberOfPartialAttributeSet: FALSE (not in global catalogue) searchFlags: 0 (no index)
ms-Exch-Coexistence-Secure-Mail-Certificate-Thumbprint	ntdsSchemaAdd	attributeID: 1.2.840.113556.1.4.7000.102.51991 isMemberOfPartialAttributeSet: FALSE (not in global catalogue) searchFlags: 0 (no index)

ATTRIBUTE	CHANGE	VALUE
ms-Exch-Coexistence-Secure-Mail-Certificate-Thumbprintms-Exch-Sync-Cookie	rangeUpper	1024
ms-Exch-Coexistence-Transport-Servers	ntdsSchemaAdd	attributeID: 1.2.840.113556.1.4.7000.102.51990 isMemberOfPartialAttributeSet: FALSE (not in global catalogue) searchFlags: 0 (no index)
ms-Exch-ELC-Mailbox-Flags	replace: attributeSecurityGuid	F6SzsVXskUGzJ7cuM+OK8g==
ms-Exch-Extension-Custom-Attribute-1	isMemberOfPartialAttributeSet:	TRUE
ms-Exch-Extension-Custom-Attribute-2	isMemberOfPartialAttributeSet:	TRUE
ms-Exch-Extension-Custom-Attribute-3	isMemberOfPartialAttributeSet:	TRUE
ms-Exch-Extension-Custom-Attribute-4	isMemberOfPartialAttributeSet:	TRUE
ms-Exch-Extension-Custom-Attribute-5	isMemberOfPartialAttributeSet	TRUE
ms-Exch-Group-External-Member-Count	ntdsSchemaModify	isMemberOfPartialAttributeSet: TRUE MAPIID:36066
ms-Exch-Group-Member-Count	ntdsSchemaModify	replace: isMemberOfPartialAttributeSetisMemberOfPartialAttributeSet: TRUE MAPIID: 36067
ms-Exch-HAB-Root-DepaPreviewent-Link	replace: isMemberOfPartialAttributeSet	TRUE
ms-Exch-MSO-Forward-Sync-Non-Recipient-Cookie	rangeUpper	20480
ms-Exch-MSO-Forward-Sync-Recipient-Cookie	rangeUpper	20480
ms-Exch-Mailbox-Audit-Enable	replace: searchFlags	19
ms-Exch-Malware-Filtering-Update-Frequency	rangeUpper	38880
ms-Exch-Role-Entries	rangeUpper	8192
ms-Exch-Schema-Version-Pt	rangeUpper	15137
ms-Exch-Schema-Version-Pt	rangeUpper	15281
ms-Exch-Smtp-Receive-Tls-Certificate-Name	Replace: rangeUpper	1024

ATTRIBUTE	CHANGE	VALUE
ms-Exch-Smtp-TLS-Certificate	replace: rangeUpper	1024
ms-Exch-Sync-Cookie	rangeUpper	262144

Object IDs added by Exchange 2016 RTM

The following **class object IDs** are added when you install Exchange 2016 RTM:

- 1.2.840.113556.1.5.7000.62.50161
- 1.2.840.113556.1.5.7000.62.50162
- 1.2.840.113556.1.5.7000.62.50163
- 1.2.840.113556.1.5.7000.62.50164
- 1.2.840.113556.1.5.7000.62.50165
- 1.2.840.113556.1.5.7000.62.50166
- 1.2.840.113556.1.5.7000.62.50167
- 1.2.840.113556.1.5.7000.62.50170
- 1.2.840.113556.1.5.7000.62.50171
- 1.2.840.113556.1.5.7000.62.50172
- 1.2.840.113556.1.5.7000.62.50173
- 1.2.840.113556.1.5.7000.62.50174
- 1.2.840.113556.1.5.7000.62.50176
- 1.2.840.113556.1.5.7000.62.50177
- 1.2.840.113556.1.5.7000.62.50178
- 1.2.840.113556.1.5.7000.62.50187
- 1.2.840.113556.1.5.7000.62.50188
- 1.2.840.113556.1.5.7000.62.50189
- 1.2.840.113556.1.5.7000.62.50190
- 1.2.840.113556.1.5.7000.62.50191
- 1.2.840.113556.1.5.7000.62.50192
- 1.2.840.113556.1.5.7000.62.50202
- 1.2.840.113556.1.5.7000.62.50203
- 1.2.840.113556.1.5.7000.62.50204
- 1.2.840.113556.1.5.7000.62.50205

The following **attribute object IDs** are added when you install Exchange 2016 RTM:

- 1.2.840.113556.1.4.2183
- 1.2.840.113556.1.4.2184

- 1.2.840.113556.1.4.2185
- 1.2.840.113556.1.4.7000.102.51773
- 1.2.840.113556.1.4.7000.102.51774
- 1.2.840.113556.1.4.7000.102.51775
- 1.2.840.113556.1.4.7000.102.51786
- 1.2.840.113556.1.4.7000.102.51787
- 1.2.840.113556.1.4.7000.102.51788
- 1.2.840.113556.1.4.7000.102.51789
- 1.2.840.113556.1.4.7000.102.51790
- 1.2.840.113556.1.4.7000.102.51791
- 1.2.840.113556.1.4.7000.102.51792
- 1.2.840.113556.1.4.7000.102.51794
- 1.2.840.113556.1.4.7000.102.51795
- 1.2.840.113556.1.4.7000.102.51796
- 1.2.840.113556.1.4.7000.102.51797
- 1.2.840.113556.1.4.7000.102.51798
- 1.2.840.113556.1.4.7000.102.51799
- 1.2.840.113556.1.4.7000.102.51800
- 1.2.840.113556.1.4.7000.102.51801
- 1.2.840.113556.1.4.7000.102.51805
- 1.2.840.113556.1.4.7000.102.51806
- 1.2.840.113556.1.4.7000.102.51807
- 1.2.840.113556.1.4.7000.102.51808
- 1.2.840.113556.1.4.7000.102.51809
- 1.2.840.113556.1.4.7000.102.51810
- 1.2.840.113556.1.4.7000.102.51811
- 1.2.840.113556.1.4.7000.102.51812
- 1.2.840.113556.1.4.7000.102.51813
- 1.2.840.113556.1.4.7000.102.51814
- 1.2.840.113556.1.4.7000.102.51815
- 1.2.840.113556.1.4.7000.102.51816
- 1.2.840.113556.1.4.7000.102.51818
- 1.2.840.113556.1.4.7000.102.51819

- 1.2.840.113556.1.4.7000.102.51820
- 1.2.840.113556.1.4.7000.102.51821
- 1.2.840.113556.1.4.7000.102.51822
- 1.2.840.113556.1.4.7000.102.51823
- 1.2.840.113556.1.4.7000.102.51824
- 1.2.840.113556.1.4.7000.102.51826
- 1.2.840.113556.1.4.7000.102.51827
- 1.2.840.113556.1.4.7000.102.51829
- 1.2.840.113556.1.4.7000.102.51830
- 1.2.840.113556.1.4.7000.102.51832
- 1.2.840.113556.1.4.7000.102.51833
- 1.2.840.113556.1.4.7000.102.51836
- 1.2.840.113556.1.4.7000.102.51837
- 1.2.840.113556.1.4.7000.102.51838
- 1.2.840.113556.1.4.7000.102.51839
- 1.2.840.113556.1.4.7000.102.51840
- 1.2.840.113556.1.4.7000.102.51851
- 1.2.840.113556.1.4.7000.102.51852
- 1.2.840.113556.1.4.7000.102.51859
- 1.2.840.113556.1.4.7000.102.51860
- 1.2.840.113556.1.4.7000.102.51861
- 1.2.840.113556.1.4.7000.102.51862
- 1.2.840.113556.1.4.7000.102.51863
- 1.2.840.113556.1.4.7000.102.51864
- 1.2.840.113556.1.4.7000.102.51865
- 1.2.840.113556.1.4.7000.102.51866
- 1.2.840.113556.1.4.7000.102.51867
- 1.2.840.113556.1.4.7000.102.51868
- 1.2.840.113556.1.4.7000.102.51869
- 1.2.840.113556.1.4.7000.102.51870
- 1.2.840.113556.1.4.7000.102.51871
- 1.2.840.113556.1.4.7000.102.51872
- 1.2.840.113556.1.4.7000.102.51873

- 1.2.840.113556.1.4.7000.102.51874
- 1.2.840.113556.1.4.7000.102.51875
- 1.2.840.113556.1.4.7000.102.51876
- 1.2.840.113556.1.4.7000.102.51877
- 1.2.840.113556.1.4.7000.102.51878
- 1.2.840.113556.1.4.7000.102.51879
- 1.2.840.113556.1.4.7000.102.51880
- 1.2.840.113556.1.4.7000.102.51881
- 1.2.840.113556.1.4.7000.102.51882
- 1.2.840.113556.1.4.7000.102.51883
- 1.2.840.113556.1.4.7000.102.51914
- 1.2.840.113556.1.4.7000.102.51915
- 1.2.840.113556.1.4.7000.102.51916
- 1.2.840.113556.1.4.7000.102.51917
- 1.2.840.113556.1.4.7000.102.51918
- 1.2.840.113556.1.4.7000.102.51919
- 1.2.840.113556.1.4.7000.102.51920
- 1.2.840.113556.1.4.7000.102.51921
- 1.2.840.113556.1.4.7000.102.51922
- 1.2.840.113556.1.4.7000.102.51923
- 1.2.840.113556.1.4.7000.102.51924
- 1.2.840.113556.1.4.7000.102.51925
- 1.2.840.113556.1.4.7000.102.51926
- 1.2.840.113556.1.4.7000.102.51927
- 1.2.840.113556.1.4.7000.102.51928
- 1.2.840.113556.1.4.7000.102.51929
- 1.2.840.113556.1.4.7000.102.51930
- 1.2.840.113556.1.4.7000.102.51931
- 1.2.840.113556.1.4.7000.102.51932
- 1.2.840.113556.1.4.7000.102.51933
- 1.2.840.113556.1.4.7000.102.51934
- 1.2.840.113556.1.4.7000.102.51935
- 1.2.840.113556.1.4.7000.102.51936

- 1.2.840.113556.1.4.7000.102.51937
- 1.2.840.113556.1.4.7000.102.51938
- 1.2.840.113556.1.4.7000.102.51939
- 1.2.840.113556.1.4.7000.102.51940
- 1.2.840.113556.1.4.7000.102.51941
- 1.2.840.113556.1.4.7000.102.51942
- 1.2.840.113556.1.4.7000.102.51943
- 1.2.840.113556.1.4.7000.102.51944
- 1.2.840.113556.1.4.7000.102.51945
- 1.2.840.113556.1.4.7000.102.51946
- 1.2.840.113556.1.4.7000.102.51947
- 1.2.840.113556.1.4.7000.102.51948
- 1.2.840.113556.1.4.7000.102.51949
- 1.2.840.113556.1.4.7000.102.51950
- 1.2.840.113556.1.4.7000.102.51951
- 1.2.840.113556.1.4.7000.102.51952
- 1.2.840.113556.1.4.7000.102.51953
- 1.2.840.113556.1.4.7000.102.51954
- 1.2.840.113556.1.4.7000.102.51955
- 1.2.840.113556.1.4.7000.102.51993
- 1.2.840.113556.1.4.7000.102.51994
- 1.2.840.113556.1.4.7000.102.51995
- 1.2.840.113556.1.4.7000.102.51996
- 1.2.840.113556.1.4.7000.102.51997
- 1.2.840.113556.1.4.7000.102.51998
- 1.2.840.113556.1.4.7000.102.52001
- 1.2.840.113556.1.4.7000.102.52002
- 1.2.840.113556.1.4.7000.102.52003
- 1.2.840.113556.1.4.7000.102.52004
- 1.2.840.113556.1.4.7000.102.52005
- 1.2.840.113556.1.4.7000.102.52006
- 1.2.840.113556.1.4.7000.102.52007
- 1.2.840.113556.1.4.7000.102.52008

- 1.2.840.113556.1.4.7000.102.52011
- 1.2.840.113556.1.4.7000.102.52012
- 1.2.840.113556.1.4.7000.102.52013
- 1.2.840.113556.1.4.7000.102.52014
- 1.2.840.113556.1.4.7000.102.52015
- 1.2.840.113556.1.4.7000.102.52016
- 1.2.840.113556.1.4.7000.102.52017
- 1.2.840.113556.1.4.7000.102.52018
- 1.2.840.113556.1.4.7000.102.52019
- 1.2.840.113556.1.4.7000.102.52020
- 1.2.840.113556.1.4.7000.102.52021
- 1.2.840.113556.1.4.7000.102.52022
- 1.2.840.113556.1.4.7000.102.52023
- 1.2.840.113556.1.4.7000.102.52024
- 1.2.840.113556.1.4.7000.102.52029
- 1.2.840.113556.1.4.7000.102.52030
- 1.2.840.113556.1.4.7000.102.52031
- 1.2.840.113556.1.4.7000.102.52032
- 1.2.840.113556.1.4.7000.102.52033
- 1.2.840.113556.1.4.7000.102.52034
- 1.2.840.113556.1.4.7000.102.52035
- 1.2.840.113556.1.4.7000.102.52036
- 1.2.840.113556.1.4.7000.102.52037
- 1.2.840.113556.1.4.7000.102.52039
- 1.2.840.113556.1.4.7000.102.52040
- 1.2.840.113556.1.4.7000.102.52041
- 1.2.840.113556.1.4.7000.102.52042
- 1.2.840.113556.1.4.7000.102.52043
- 1.2.840.113556.1.4.7000.102.52044
- 1.2.840.113556.1.4.7000.102.52045
- 1.2.840.113556.1.4.7000.102.52046
- 1.2.840.113556.1.4.7000.102.52047
- 1.2.840.113556.1.4.7000.102.52048

- 1.2.840.113556.1.4.7000.102.52049
- 1.2.840.113556.1.4.7000.102.52050
- 1.2.840.113556.1.4.7000.102.52051
- 1.2.840.113556.1.4.7000.102.52052
- 1.2.840.113556.1.4.7000.102.52053
- 1.2.840.113556.1.4.7000.102.52054
- 1.2.840.113556.1.4.7000.102.52055
- 1.2.840.113556.1.4.7000.102.52056
- 1.2.840.113556.1.4.7000.102.52057
- 1.2.840.113556.1.4.7000.102.52058
- 1.2.840.113556.1.4.7000.102.52059
- 1.2.840.113556.1.4.7000.102.52060
- 1.2.840.113556.1.4.7000.102.52061
- 1.2.840.113556.1.4.7000.102.52062
- 1.2.840.113556.1.4.7000.102.52063
- 1.2.840.113556.1.4.7000.102.52064
- 1.2.840.113556.1.4.7000.102.52065
- 1.2.840.113556.1.4.7000.102.52109
- 1.2.840.113556.1.4.7000.102.52110
- 1.2.840.113556.1.4.7000.102.52126
- 1.2.840.113556.1.4.7000.102.52127
- 1.2.840.113556.1.4.7000.102.52128
- 1.2.840.113556.1.4.7000.102.52129
- 1.2.840.113556.1.4.7000.102.52130

Indexed attributes added by Exchange 2016 RTM

The following table lists the attributes that are added to the list of indexed attributes when you install Exchange 2016 RTM.

ATTRIBUTE	SEARCH FLAG VALUE
ms-DS-GeoCoordinates-Altitude	1
ms-DS-GeoCoordinates-Latitude	1
ms-DS-GeoCoordinates-Longitude	1
ms-Exch-Accepted-Domain-Name	9

ATTRIBUTE	SEARCH FLAG VALUE
ms-Exch-Archive-GUID	9
ms-Exch-Auth-Application-Identifier	1
ms-Exch-Auth-Issuer-Name	1
ms-Exch-Bypass-Audit	9
ms-Exch-Default-Public-Folder-Mailbox	19
ms-Exch-Device-Client-Type	1
ms-Exch-Extension-Custom-Attribute-1	1
ms-Exch-Extension-Custom-Attribute-2	1
ms-Exch-Extension-Custom-Attribute-3	1
ms-Exch-Extension-Custom-Attribute-4	1
ms-Exch-Extension-Custom-Attribute-5	1
ms-Exch-Home-MDB-SL	1
ms-Exch-Home-MTA-SL	1
ms-Exch-Is-Dirsync-Status-Pending	1
ms-Exch-Mailbox-Audit-Enable	19
ms-Exch-Mailbox-Database-Transport-Flags	16
ms-Exch-Mailbox-Move-Source-Archive-MDB-Link-SL	1
ms-Exch-Mailbox-Move-Source-MDB-Link-SL	1
ms-Exch-Mailbox-Move-Target-Archive-MDB-Link-SL	1
ms-Exch-OWA-Set-Photo-URL	16
ms-Exch-Organization-Upgrade-Policy-Link	1
ms-Exch-Organization-Upgrade-Policy-Link-SL	1
ms-Exch-Previous-Archive-Database-SL	8
ms-Exch-Previous-Home-MDB-SL	8
ms-Exch-Provisioning-Tags	1

ATTRIBUTE	SEARCH FLAG VALUE
ms-Exch-Public-Folder-EntryId	24
ms-Exch-Public-Folder-Mailbox	24
ms-Exch-Public-Folder-Smtp-Address	24
ms-Exch-Recipient-SoftDeleted-Status	27
ms-Exch-Relocate-Tenant-Completion-Target-Vector	8
ms-Exch-Relocate-Tenant-Flags	8
ms-Exch-Relocate-Tenant-Safe-Lockdown-Schedule	8
ms-Exch-Relocate-Tenant-Source-Forest	9
ms-Exch-Relocate-Tenant-Start-Lockdown	8
ms-Exch-Relocate-Tenant-Start-Retired	8
ms-Exch-Relocate-Tenant-Start-Sync	8
ms-Exch-Relocate-Tenant-Status,	9
ms-Exch-Relocate-Tenant-Target-Forest	9
ms-Exch-Relocate-Tenant-Transition-Counter	8
ms-Exch-Sync-Cookie	8
ms-Exch-Team-Mailbox-Expiration	16
ms-Exch-Team-Mailbox-Expiry-Days	16
ms-Exch-Team-Mailbox-Owners	16
ms-Exch-Team-Mailbox-SharePoint-Linked-By	16
ms-Exch-Team-Mailbox-SharePoint-Url	16
ms-Exch-Team-Mailbox-Show-In-Client-List	16
ms-Exch-Transport-Rule-Immutable-Id	1
ms-Exch-When-Soft-Deleted-Time	17

Property sets modified by Exchange 2016 RTM

The following property sets are modified when you install Exchange 2016 RTM:

- Exchange-Information

MAPI IDs added by Exchange 2016 RTM

The following MAPI IDs are added when you install Exchange 2016 RTM:

- 36066
- 36067

Extended rights added by Exchange 2016 RTM

The following table lists the extended rights that are added when you install Exchange 2016 RTM. Installing Exchange 2016 RTM doesn't modify any existing extended rights.

IDENTIFIER	VALUES
CN=ms-Exch-SMTP-Accept-XProxyFrom,CN=Extended-Rights,<ConfigurationContainerDN>	changetype: ntdsSchemaAdd displayName: Accept XProxyFrom objectClass: controlAccessRight rightsGuid: 5bee2b72-50d7-49c7-ba66-39a25daa1e92 validAccesses: 256

What changes in Active Directory when Exchange is installed?

8/3/2020 • 5 minutes to read • Edit Online

When you install Exchange Server 2016 or Exchange Server 2019, changes are made to your Active Directory forest and domains to store information about the Exchange servers, mailboxes, and other Exchange-related objects in your organization.

Three steps are required to prepare Active Directory for Exchange:

1. Extend the Active Directory schema

2. Prepare Active Directory containers, objects, and other items

3. Prepare Active Directory domains

After all three steps are done, your Active Directory forest is ready for Exchange. This topic explains what Exchange does at each step of Active Directory preparation.

You can make these changes before you install the first Exchange 2016 or Exchange 2019 server in the organization by running the */PrepareSchema*, */PrepareAD*, and */PrepareAllDomains* or */PrepareDomains* commands using Exchange command line Setup. For instructions, see Prepare Active Directory and domains for Exchange. Or, these changes are automatically made for you during the installation of the first Exchange server using the Exchange Setup wizard. For instructions, see Install Exchange Mailbox servers using the Setup wizard.

Extend the Active Directory schema

Extending the Active Directory schema adds and updates classes, attributes, and other items. These changes are needed so that Exchange can create containers and objects to store information about the Exchange organization. Because Exchange makes a lot of changes to the Active Directory schema, there's a topic dedicated to this step. To see all of the changes made to the schema, see Active Directory schema changes in Exchange Server.

After the schema has been extended by running the */PrepareSchema* command, the _/PrepareAD command, or installing the first Exchange server using the Exchange Setup wizard, the schema version is set in the `ms-Exch-Schema-Version-Pt` attribute. To verify that the Active Directory schema was extended successfully, you can check the value stored in this attribute. For more information, see Exchange Active Directory versions.

Prepare Active Directory containers, objects, and other items

With the schema extended, the next step is to add all of the containers, objects, attributes, and other items that Exchange uses to store information in Active Directory. Most of the changes made in this step are applied to the entire Active Directory forest. A smaller set of changes are made only to the local Active Directory domain where the */PrepareAD* command was run (or where the first Exchange server was installed using the Exchange Setup wizard).

Exchange makes the following changes to the Active Directory forest:

- The Microsoft Exchange container is created under CN=Services,CN=Configuration,DC=<*root domain*> if it doesn't already exist.

- The following containers and objects are created under CN=<*organization name*>,CN=Microsoft Exchange,CN=Services,CN=Configuration,DC=<*root domain*> if they don't already exist:

- CN=Address Lists Container
- CN=AddressBook Mailbox Policies
- CN=Addressing
- CN=Administrative Groups
- CN=Approval Applications
- CN=Auth Configuration
- CN=Availability Configuration
- CN=Client Access
- CN=Connections
- CN=ELC Folders Container
- CN=ELC Mailbox Policies
- CN=ExchangeAssistance
- CN=Federation
- CN=Federation Trusts
- CN=Global Settings
- CN=Hybrid Configuration
- CN=Mobile Mailbox Policies
- CN=Mobile Mailbox Settings
- CN=Monitoring Settings
- CN=OWA Mailbox Policies
- CN=Provisioning Policy Container
- CN=Push Notification Settings
- CN=RBAC
- CN=Recipient Policies
- CN=Remote Accounts Policies Container
- CN=Retention Policies Container
- CN=Retention Policy Tag Container
- CN=ServiceEndpoints
- CN=System Policies
- CN=Team Mailbox Provisioning Policies
- CN=Transport Settings
- CN=UM AutoAttendant Container (Exchange 2016 only)
- CN=UM DialPlan Container (Exchange 2016 only)

  - CN=UM IPGateway Container (Exchange 2016 only)
  - CN=UM Mailbox Policies (Exchange 2016 only)
  - CN=Workload Management Settings
- The following containers and objects are created under CN=Transport Settings,CN=<*Organization Name*>,CN=Microsoft Exchange,CN=Services,CN=Configuration,DC=<*root domain*> if they don't already exist:
  - CN=Accepted Domains
  - CN=ControlPoint Config
  - CN=DNS Customization
  - CN=Interceptor Rules
  - CN=Malware Filter
  - CN=Message Classifications
  - CN=Message Hygiene
  - CN=Rules
  - CN=MicrosoftExchange329e71ec88ae4615bbc36ab6ce41109e
- Permissions are set throughout the configuration partition in Active Directory.
- The Rights.ldf file is imported. This file adds permissions that are needed to install Exchange and configure Active Directory.
- The Microsoft Exchange Security Groups organizational unit (OU) is created in the root domain of the forest, and permissions are assigned to it.
- The following groups are created within the Microsoft Exchange Security Groups OU if they don't already exist:
  - Compliance Management
  - Delegated Setup
  - Discovery Management
  - Exchange Servers
  - Exchange Trusted Subsystem
  - Exchange Windows Permissions
  - ExchangeLegacyInterop
  - Help Desk
  - Hygiene Management
  - Managed Availability Servers
  - Organization Management
  - Public Folder Management
  - Recipient Management

  - Records Management
  - Server Management
  - View-Only Organization Management

- The new management role groups (which appear as universal security groups (USGs) in Active Directory) that were created in the Microsoft Exchange Security Groups OU are added to the **otherWellKnownObjects** attribute stored on the CN=Microsoft Exchange,CN=Services,CN=Configuration,DC=<*root domain*> container.
- In Exchange 2016 only, the Unified Messaging Voice Originator contact is created in the Microsoft Exchange System Objects container of the root domain.
- Only the domain where the */PrepareAD* command was run (or where the first Exchange server was installed using the Exchange Setup wizard) is prepared for Exchange. For information about what's done to prepare an Active Directory domain for Exchange, see the next section.

Prepare Active Directory domains

The final step of preparing Active Directory for Exchange is to prepare the Active Directory domains where Exchange servers will be installed or where mailbox-enabled users will be located (all domains in the forest using the */PrepareAllDomains* command, specific domains using the */PrepareDomains* command, or installing the first Exhange server using the Exchange Setup Wizard). This step is done automatically in the domain where the *PrepareAD* command was run (or where the first Exchange server was installed using the Exchange Setup wizard).

Exchange makes the follwing changes to the Active Directory domains:

- The Microsoft Exchange System Objects container is created in the root domain partition in Active Directory if it doesn't already exist.
- Permissions are set on the Microsoft Exchange System Objects container for the Exchange Servers, Organization Management, and Authenticated Users security groups.
- Modifying the **Default Domain Controllers GPO** to grant "Manage Auditing and Security Log policy" rights to **Exchange Enterprise Servers**.
- The Exchange Install Domain Servers domain global group is created in the current domain and placed in the Microsoft Exchange System Objects container.
- The Exchange Install Domain Servers group is added to the Exchange Servers USG in the root domain.
- Permissions are assigned at the domain level for the Exchange Servers USG and the Organization Management USG.
- The **objectVersion** property in the Microsoft Exchange System Objects container under DC=<*root domain*> is set. To verify that the Active Directory domains were successfully prepared, you can check the value stored in this attribute. For more information, see Exchange Active Directory versions.

Prepare Active Directory and domains for Exchange Server

8/3/2020 • 10 minutes to read • Edit Online

Exchange uses Active Directory to store information about mailboxes and the configuration of Exchange servers in the organization. Before you install Exchange Server 2016 or Exchange Server 2019 (even if you have earlier versions of Exchange installed in your organization), you need to prepare your Active Directory forest and its domains for the new version of Exchange. There are two ways to do this:

- **Let the Exchange Setup wizard do it for you**: If you don't have a large Active Directory deployment, and you don't have a separate team that manages Active Directory, we recommend using the Setup wizard. Your account needs to be a member of both the Schema Admins and Enterprise Admins security groups. For more information about how to use the Setup wizard, check out Install Exchange Mailbox servers using the Setup wizard.

- **Follow the steps in this topic**: If you have a large Active Directory deployment, or if a separate team manages Active Directory, this topic is for you. Following the steps in this topic gives you much more control over each stage of preparation, and who can do each step. For example, Exchange administrators might not have the required permissions to extend the Active Directory schema.

For details on new schema classes and attributes that Exchange adds to Active Directory, including those made by Cumulative Updates (CUs), see Active Directory schema changes in Exchange Server.

For details about what's happening when Active Directory is being prepared for Exchange, see What changes in Active Directory when Exchange is installed?.

If you aren't familiar with Active Directory forests or domains, check out Active Directory Domain Services Overview.

What do you need to know before you begin?

- Estimated time to complete: 10-15 minutes or more (not including Active Directory replication), depending on organization size and the number of child domains.

- The computer that you use for these procedures needs to meet the system requirements for Exchange.

- Verify that your Active Directory meets the requirements for Exchange:

  - **Exchange 2019**: Exchange 2019 Network and directory servers.

  - **Exchange 2016**: Exchange 2016 Network and directory servers.

- If your organization has multiple Active Directory domains, we recommend the following approach:

  - Do these procedures in an Active Directory site that contains an Active Directory server from every domain.

  - Install the first Exchange server in an Active Directory site that contains a writeable global catalog server from every domain.

- The computer that you use for all procedures in this topic requires access to Setup.exe in the Exchange installation files:

  1. Download the latest version of Exchange. For more information, see Updates for Exchange Server.

2. In File Explorer, right-click on the Exchange ISO image file that you downloaded, and then select **Mount**. Note the virtual DVD drive letter that's assigned.

3. Open a Windows Command Prompt window. For example:

   - Press the Windows key + 'R' to open the **Run** dialog, type cmd.exe, and then press **OK**.

   - Press **Start**. In the **Search** box, type **Command Prompt**, then in the list of results, select **Command Prompt**.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Step 1: Extend the Active Directory schema

> **TIP**
>
> If you don't have a separate team that manages your Active Directory schema, you can skip this step and go directly to Step 2: Prepare Active Directory. If you don't extend the schema in this step, the */PrepareAd* command in Step 2 will automatically extend the schema for you. If you skip this step, the requirements will also apply to Step 2.

When you extend the Active Directory schema for Exchange, the following requirements apply:

- Your account needs to be a member of the Schema Admins and Enterprise Admins security groups. If you have multiple Active Directory forests, make sure you're logged into the right one.

- The computer needs to be a member of the same Active Directory domain and site as the schema master.

- If you use the */DomainController:<DomainControllerFQDN>* switch, you need to specify the domain controller that's the schema master.

- The only supported way to extend the schema for Exchange is to use Setup.exe with */PrepareSchema*, */PrepareAD*, or the Exchange Setup wizard. Other ways of extending the schema aren't supported.

To extend the schema for Exchange, run the following command in a Windows Command Prompt window:

```
<Virtual DVD drive letter>:\Setup.exe /IAcceptExchangeServerLicenseTerms /PrepareSchema
```

For example, if the Exchange installation files are available on drive E:, run the following command:

```
E:\Setup.exe /IAcceptExchangeServerLicenseTerms /PrepareSchema
```

> **NOTE**
>
> When you run this command, a prerequisite check is performed that will tell you which requirements are missing.

After Setup finishes extending the schema, you'll need to wait while Active Directory replicates the changes to all of your domain controllers before you proceed. To check the progress of the replication, you can use the `repadmin` tool in Windows Server. For more information about how to use the `repadmin` tool, see Repadmin.

Step 2: Prepare Active Directory

After the Active Directory schema has been extended, you can prepare other parts of Active Directory for Exchange. During this step, Exchange will create containers, objects, and other items in Active Directory to store information. The collection of the Exchange containers, objects, attributes, and so on, is called the *Exchange organization*.

When you prepare Active Directory for Exchange, the following requirements apply:

- Your account needs to be a member of the Enterprise Admins security group. If you skipped Step 1 because you want the */PrepareAD* command to extend the schema, the account also needs to be a member of the Schema Admins security group.
- The computer needs to be a member of the same Active Directory domain and site as the schema master, and must be able to contact all of the domains in the forest on TCP port 389.
- Wait until Active Directory has finished replicating the schema changes from Step 1 to all domain controllers before you try to prepare Active Directory.
- You need to select a name for the Exchange organization. The organization name is used internally by Exchange, isn't typically seen by users, doesn't affect the functionality of Exchange, and doesn't determine what you can use for email addresses.
    - The organization name can't contain more than 64 characters, and can't be blank.
    - Valid characters are A to Z, a to z, 0 to 9, hyphen or dash (-), and space, but leading or trailing spaces aren't allowed.
    - You can't change the organization name after it's set.

To prepare Active Directory for Exchange, run the following command in a Windows Command Prompt window:

```
<Virtual DVD drive letter>:\Setup.exe /IAcceptExchangeServerLicenseTerms /PrepareAD  /OrganizationName:"
<Organization name>"
```

This example uses the Exchange installation files on drive E: and names the Exchange organization "Contoso Corporation".

```
E:\Setup.exe /IAcceptExchangeServerLicenseTerms /PrepareAD /OrganizationName:"Contoso Corporation"
```

IMPORTANT

If you have a hybrid deployment configured between your on-premises organization and Exchange Online, add the */TenantOrganizationConfig* switch to the command.

As in Step 1, you'll need to wait while Active Directory replicates the changes from this step to all of your domain controllers before you proceed, and you can use the `repadmin` tool to check the progress of the replication.

Step 3: Prepare Active Directory domains

TIP

If you have only one domain, you can skip this step because the */PrepareAD* command in Step 2 has already prepared the domain for you.

The final step is to prepare the Active Directory domain where Exchange servers will be installed or where mail-enabled users will be located. This step creates additional containers and security groups, and sets the permissions so Exchange can access them.

If you have multiple domains in your Active Directory forest, you have the following choices in how to prepare them:

- Prepare all domains in the Active Directory forest
- Choose the Active Directory domains to prepare

Regardless of the method you choose, wait until Active Directory has finished replicating the changes from Step 2 to all domain controllers before you proceed. Otherwise, you might get an error when you try to prepare the domains.

Prepare all domains in the Active Directory forest

When you prepare all domains in the Active Directory forest for Exchange, your account needs to be a member of the Enterprise Admins security group.

To prepare all domains in your Active Directory forest, run the following command in a Windows Command Prompt window:

```
<Virtual DVD drive letter>:\Setup.exe /IAcceptExchangeServerLicenseTerms /PrepareAllDomains
```

For example, if the Exchange installation files are available on drive E:, run the following command:

```
E:\Setup.exe /IAcceptExchangeServerLicenseTerms /PrepareAllDomains
```

Choose the Active Directory domains to prepare

> **TIP**
> You don't need to do this step in the domain where you ran the */PrepareAD* command in Step 2, because the */PrepareAD* command has automatically prepared that domain for you.

When you prepare specific domains in your Active Directory forest, the following requirements apply:

- You need to prepare every domain where an Exchange server will be installed.
- You need to prepare any domain that will contain mail-enabled users, even if the domain won't contain any Exchange servers.
- Your account needs to be a member of the Domain Admins group in the domain that you want to prepare.
- If the domain that you want to prepare was created **after** you ran /PrepareAD in Step 2, your account also needs to be a member of the Organization Management role group in Exchange.

To a prepare a specific domain in your Active Directory forest, run the following command in a Windows Command Prompt window:

```
<Virtual DVD drive letter>:\Setup.exe /IAcceptExchangeServerLicenseTerms /PrepareDomain[:<DomainFQDN>]
```

Notes:

- If the computer is a member of the domain that you want to prepare, you can use the */PrepareDomain* switch by itself. Otherwise, you need to specify the FQDN of the domain.

- You need to run this command for each Active Directory domain where you'll install an Exchange server or where mail-enabled users will be located.

This example uses the Exchange installation files on drive E: to prepare the engineering.corp.contoso.com domain:

```
E:\Setup.exe /IAcceptExchangeServerLicenseTerms /PrepareDomain:engineering.corp.contoso.com
```

This is the same example, but run on a computer that's a member of the engineering.corp.contoso.com domain:

```
E:\Setup.exe /IAcceptExchangeServerLicenseTerms /PrepareDomain
```

How do you know this worked?

To verify that you successfully prepared Active Directory and domains for Exchange, use any of the following steps:

- Use ADSI Edit and the information from the tables in the next section to verify that the specified objects have the correct values for the release of Exchange that you're installing. To learn more about ADSI Edit, see ADSI Edit (adsiedit.msc).

Caution

Never change values in ADSI Edit unless you're told to do so by Microsoft Customer Service and Support. Changing values in ADSI Edit can cause irreparable damage to your Exchange organization and Active Directory.

- Check the Exchange setup log to verify that Active Directory preparation has completed successfully. For more information, see Verify an Exchange installation. Note that you can't use the **Get-ExchangeServer** cmdlet as described in the topic until you've completed the installation of at least one Exchange Mailbox server in an Active Directory site.

Exchange Active Directory versions

The tables in the following sections contain the Exchange objects in Active Directory that are updated each time you install a new version of Exchange (a new installation or a CU). You can compare the object versions you see with the values in the tables to verify that Exchange successfully updated Active Directory during the installation.

- **rangeUpper** is located in the **Schema** naming context in the properties of the **ms-Exch-Schema-Version-Pt** container.
- **objectVersion (Default)** is the **objectVersion** attribute located in the **Default naming context** in the properties of the **Microsoft Exchange System Objects** container.
- **objectVersion (Configuration)** is the **objectVersion** attribute located in the **Configuration** naming context in **Services** > **Microsoft Exchange** in the properties of the **<Your Exchange Organization Name>** container.

Exchange 2019 Active Directory versions

EXCHANGE 2019 VERSION	RANGEUPPER	OBJECTVERSION (DEFAULT)	OBJECTVERSION (CONFIGURATION)
Exchange 2019 CU6	17001	13237	16754
Exchange 2019 CU5	17001	13237	16754

EXCHANGE 2019 VERSION	RANGEUPPER	OBJECTVERSION (DEFAULT)	OBJECTVERSION (CONFIGURATION)
Exchange 2019 CU4	17001	13237	16754
Exchange 2019 CU3	17001	13237	16754
Exchange 2019 CU2	17001	13237	16754
Exchange 2019 CU1	17000	13236	16752
Exchange 2019 RTM	17000	13236	16751
Exchange 2019 Preview	15332	13236	16213

Exchange 2016 Active Directory versions

EXCHANGE 2016 VERSION	RANGEUPPER	OBJECTVERSION (DEFAULT)	OBJECTVERSION (CONFIGURATION)
Exchange 2016 CU17	15332	13237	16217
Exchange 2016 CU16	15332	13237	16217
Exchange 2016 CU15	15332	13237	16217
Exchange 2016 CU14	15332	13237	16217
Exchange 2016 CU13	15332	13237	16217
Exchange 2016 CU12	15332	13236	16215
Exchange 2016 CU11	15332	13236	16214
Exchange 2016 CU10	15332	13236	16213
Exchange 2016 CU9	15332	13236	16213
Exchange 2016 CU8	15332	13236	16213
Exchange 2016 CU7	15332	13236	16213
Exchange 2016 CU6	15330	13236	16213
Exchange 2016 CU5	15326	13236	16213
Exchange 2016 CU4	15326	13236	16213
Exchange 2016 CU3	15326	13236	16212
Exchange 2016 CU2	15325	13236	16212
Exchange 2016 CU1	15323	13236	16211

EXCHANGE 2016 VERSION	RANGEUPPER	OBJECTVERSION (DEFAULT)	OBJECTVERSION (CONFIGURATION)
Exchange 2016 RTM	15317	13236	16210
Exchange 2016 Preview	15317	13236	16041

Deploy new installations of Exchange

8/3/2020 • 2 minutes to read • Edit Online

Before you begin your installation of Exchange Server, see Planning and deployment for important planning information, and information about system requirements and prerequisites.

The following topics provide information about deploying new installations of Exchange 2019 in your organization:

Install Exchange Mailbox servers using the Setup wizard

Install Exchange using unattended mode

Install Exchange Edge Transport servers using the Setup wizard

Delegate the installation of Exchange servers

Exchange dev/test environment in Azure

After you've completed your installation, see Exchange post-installation tasks.

Install Exchange Mailbox servers using the Setup wizard

8/3/2020 • 5 minutes to read • Edit Online

Before you install an Exchange Server 2016 or Exchange Server 2019 Mailbox server, verify the following prerequisites:

- Verify the network, computer hardware, operating system, and software requirements at: Exchange Server system requirements and Exchange Server prerequisites.
- The target server must be a member of an Active Directory domain.
- The account that you use to install Exchange requires the following permissions[*]:
  - **Enterprise Admins group membership**: Required if this is the first Exchange server in the organization.
  - **Schema Admins group membership**: Required if you haven't previously extended the Active Directory schema or prepared Active Directory for Exchange 2016 or Exchange 2019.
  - **Exchange Organization Management role group membership**: Required if you've already prepared the Active Directory domain that will contain the Exchange server, or if other Exchange servers already exist in the organization.

  [*] Members of the **Delegated Setup** role group can install Exchange on servers that have already been provisioned in Active Directory by an Exchange administrator. For more information, see Delegate the installation of Exchange servers.
- Verify that you've read the release notes at Release notes for Exchange Server.

For more information about planning and deploying Exchange, see Planning and deployment for Exchange Server.

To install the Edge Transport role on a computer, see Install Exchange Edge Transport servers using the Setup wizard. Note that you can't install the Edge Transport role on a Mailbox server.

What do you need to know before you begin?

- Estimated time to complete: 60 minutes
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

Caution

After you install Exchange on a server, you must not change the server name. Renaming a server after you've installed an Exchange server role is not supported.

Install the Exchange Mailbox server role

1. Download the latest version of version of Exchange. For more information, see Updates for Exchange Server.
2. In File Explorer, right-click on the Exchange ISO image file that you downloaded, and then select **Mount**. In the resulting virtual DVD drive that appears, start Exchange Setup by double-clicking `Setup.exe`.

3. The Exchange Server Setup wizard opens. On the **Check for Updates?** page, choose one of the following options, and then click **Next** to continue:

   - **Connect to the Internet and check for updates**: We recommend this option, which searches for updates to the version of Exchange *that you're currently installing* (it doesn't detect newer Cumulative Updates). This option takes you to the **Downloading Updates** page that searches for updates. Click **Next** to continue.

   - **Don't check for updates right now**

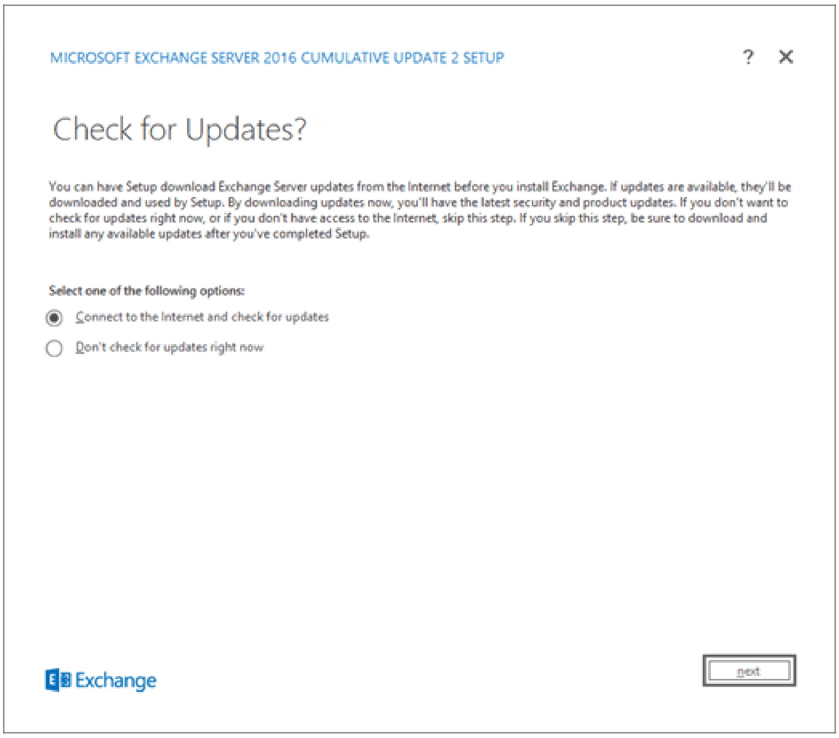

MICROSOFT EXCHANGE SERVER 2016 CUMULATIVE UPDATE 2 SETUP

Check for Updates?

You can have Setup download Exchange Server updates from the Internet before you install Exchange. If updates are available, they'll be downloaded and used by Setup. By downloading updates now, you'll have the latest security and product updates. If you don't want to check for updates right now, or if you don't have access to the Internet, skip this step. If you skip this step, be sure to download and install any available updates after you've completed Setup.

Select one of the following options:

● Connect to the Internet and check for updates

○ Don't check for updates right now

4. The **Copying Files** page shows the progress of copying files to the local hard drive. Typically, the files are copied to `%WinDir%\Temp\ExchangeSetup`, but you can confirm the location in the Exchange Setup log at `C:\ExchangeSetupLogs\ExchangeSetup.log`.

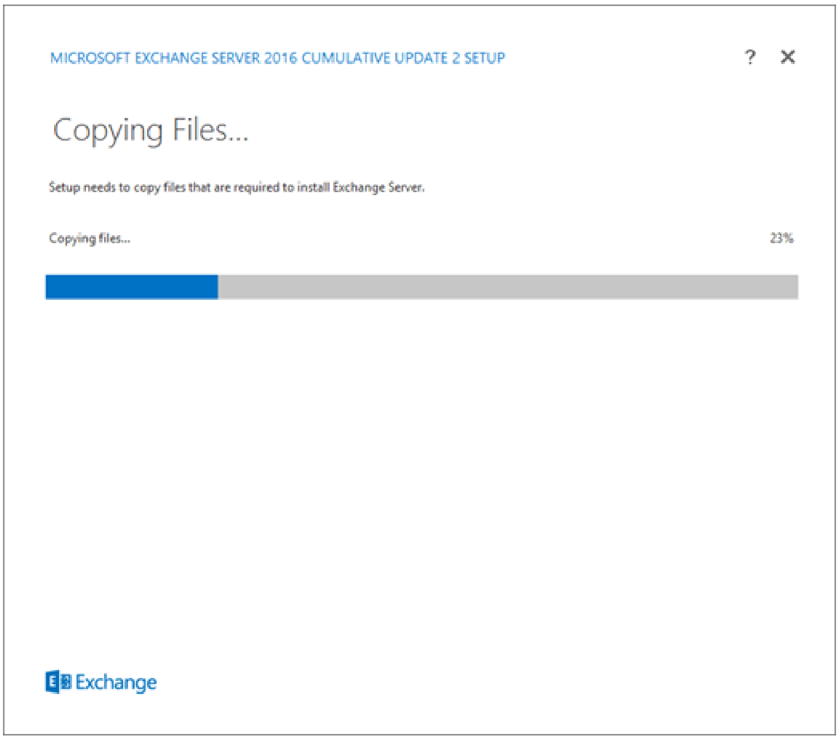

MICROSOFT EXCHANGE SERVER 2016 CUMULATIVE UPDATE 2 SETUP

Copying Files...

Setup needs to copy files that are required to install Exchange Server.

Copying files... 23%

5. On the **Introduction** page, we recommend that you visit the Exchange Server deployment planning links if you haven't already reviewed them. Click **Next** to continue.

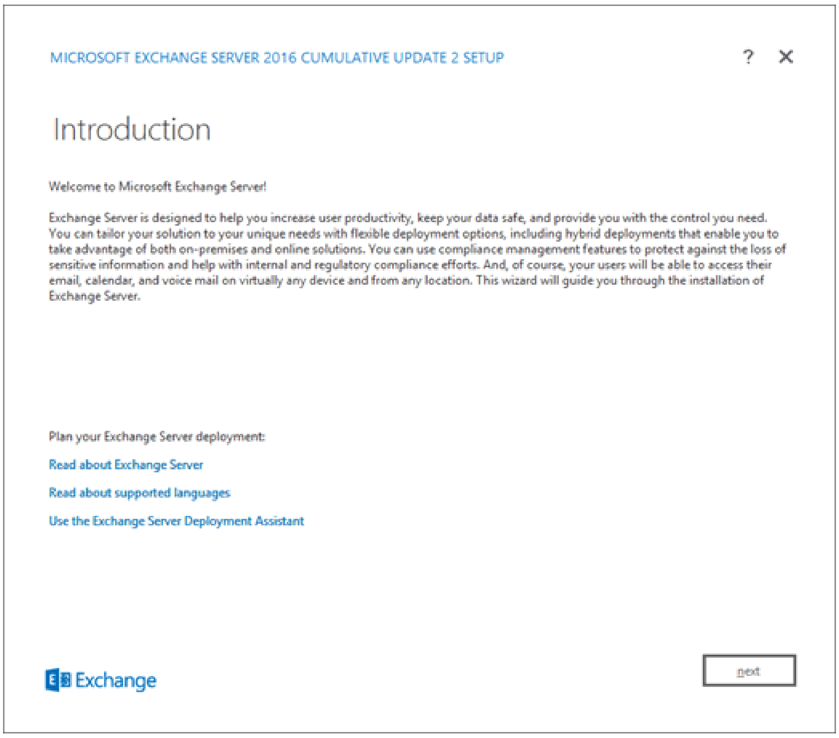

MICROSOFT EXCHANGE SERVER 2016 CUMULATIVE UPDATE 2 SETUP

Introduction

Welcome to Microsoft Exchange Server!

Exchange Server is designed to help you increase user productivity, keep your data safe, and provide you with the control you need. You can tailor your solution to your unique needs with flexible deployment options, including hybrid deployments that enable you to take advantage of both on-premises and online solutions. You can use compliance management features to protect against the loss of sensitive information and help with internal and regulatory compliance efforts. And, of course, your users will be able to access their email, calendar, and voice mail on virtually any device and from any location. This wizard will guide you through the installation of Exchange Server.

Plan your Exchange Server deployment:

Read about Exchange Server

Read about supported languages

Use the Exchange Server Deployment Assistant

6. On the **License Agreement** page, review the software license terms, select **I accept the terms in the license agreement**, and then click **Next** to continue.

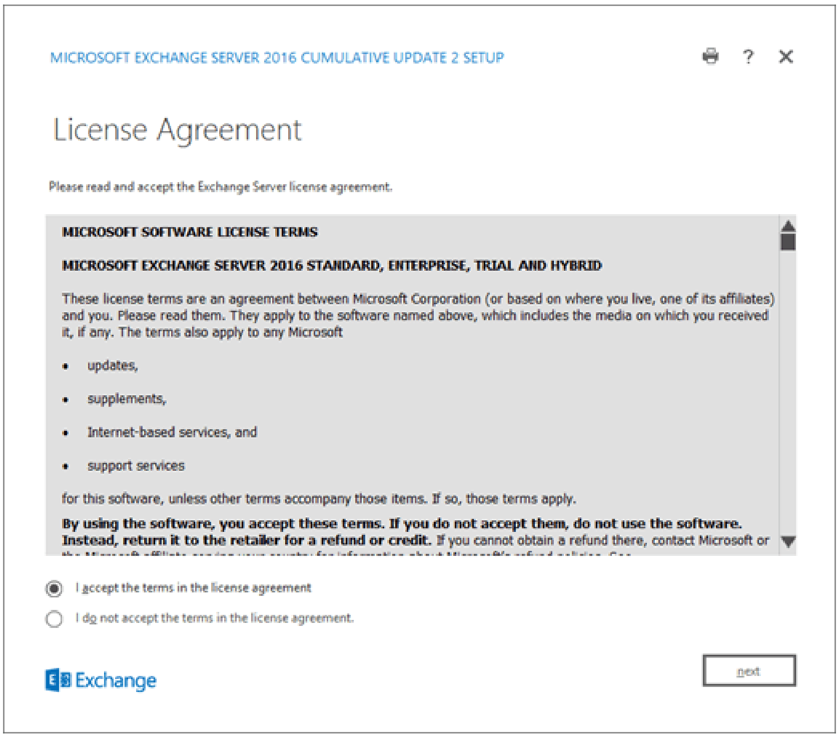

MICROSOFT EXCHANGE SERVER 2016 CUMULATIVE UPDATE 2 SETUP

License Agreement

Please read and accept the Exchange Server license agreement.

MICROSOFT SOFTWARE LICENSE TERMS

MICROSOFT EXCHANGE SERVER 2016 STANDARD, ENTERPRISE, TRIAL AND HYBRID

These license terms are an agreement between Microsoft Corporation (or based on where you live, one of its affiliates) and you. Please read them. They apply to the software named above, which includes the media on which you received it, if any. The terms also apply to any Microsoft

- updates,
- supplements,
- Internet-based services, and
- support services

for this software, unless other terms accompany those items. If so, those terms apply.

By using the software, you accept these terms. If you do not accept them, do not use the software. Instead, return it to the retailer for a refund or credit. If you cannot obtain a refund there, contact Microsoft or

● I accept the terms in the license agreement

○ I do not accept the terms in the license agreement.

7. On the **Recommended Settings** page, choose one of the following settings:

- **Use recommended settings**: Exchange automatically sends error reports and information about your computer hardware and how you use Exchange to Microsoft. For information about what's sent to Microsoft and how it's used, click **?** or the help links on the page.

- **Don't use recommended settings**: These settings are disabled, but you can enable them at any time after Setup completes.

Click **Next** to continue.

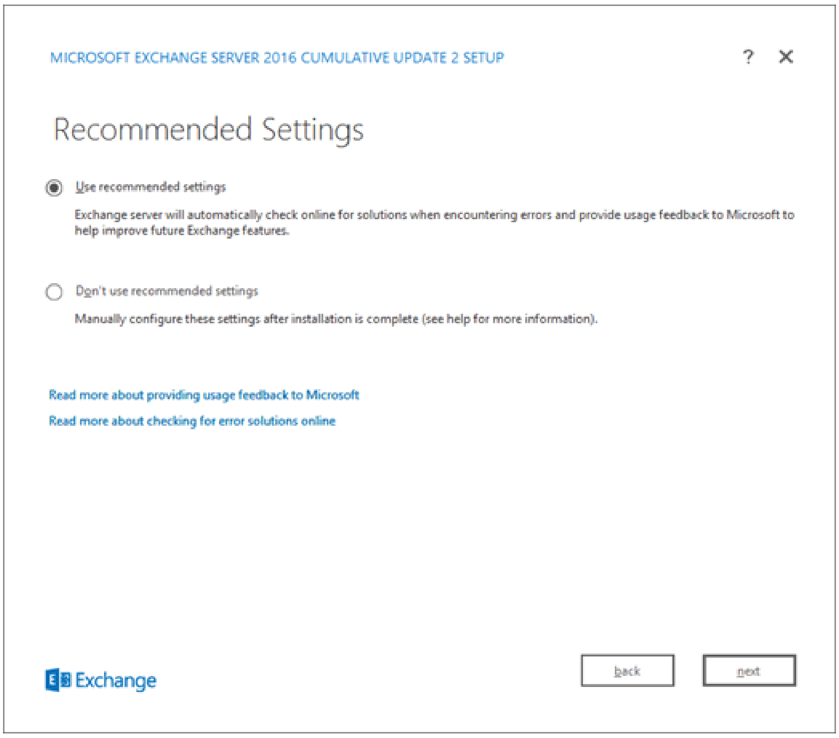

MICROSOFT EXCHANGE SERVER 2016 CUMULATIVE UPDATE 2 SETUP

Recommended Settings

● Use recommended settings

Exchange server will automatically check online for solutions when encountering errors and provide usage feedback to Microsoft to help improve future Exchange features.

○ Don't use recommended settings

Manually configure these settings after installation is complete (see help for more information).

Read more about providing usage feedback to Microsoft

Read more about checking for error solutions online

8. On the **Server Role Selection** page, configure the following options:

   - **Mailbox role**: Select this option, which also automatically installs the **Management Tools**.
   - **Automatically install Windows Server roles and features that are required to install Exchange**: Select this option to have the Setup wizard install the required Windows prerequisites. You might need to reboot the computer to complete the installation of some Windows features. If you don't select this option, you need to install the Windows features manually.

   Note: Selecting this option installs only the *Windows features* that are required by Exchange. You need to install other prerequisites manually. For more information, see Exchange Server prerequisites.

   Click **Next** to continue.

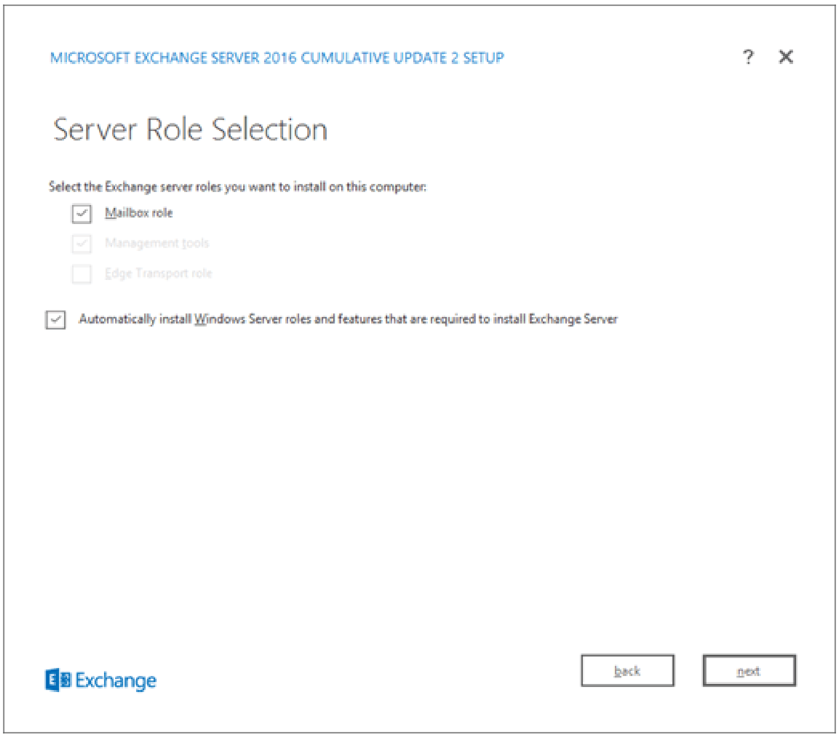

MICROSOFT EXCHANGE SERVER 2016 CUMULATIVE UPDATE 2 SETUP

Server Role Selection

Select the Exchange server roles you want to install on this computer:

☑ Mailbox role

☑ Management tools

□ Edge Transport role

☑ Automatically install Windows Server roles and features that are required to install Exchange Server

9. On the **Installation Space and Location** page, either accept the default installation location (`C:\Program Files\Microsoft\Exchange Server\V15`), or click **Browse** to choose a new location. Make sure

that you have enough disk space available in the location where you want to install Exchange. Click **Next** to continue.

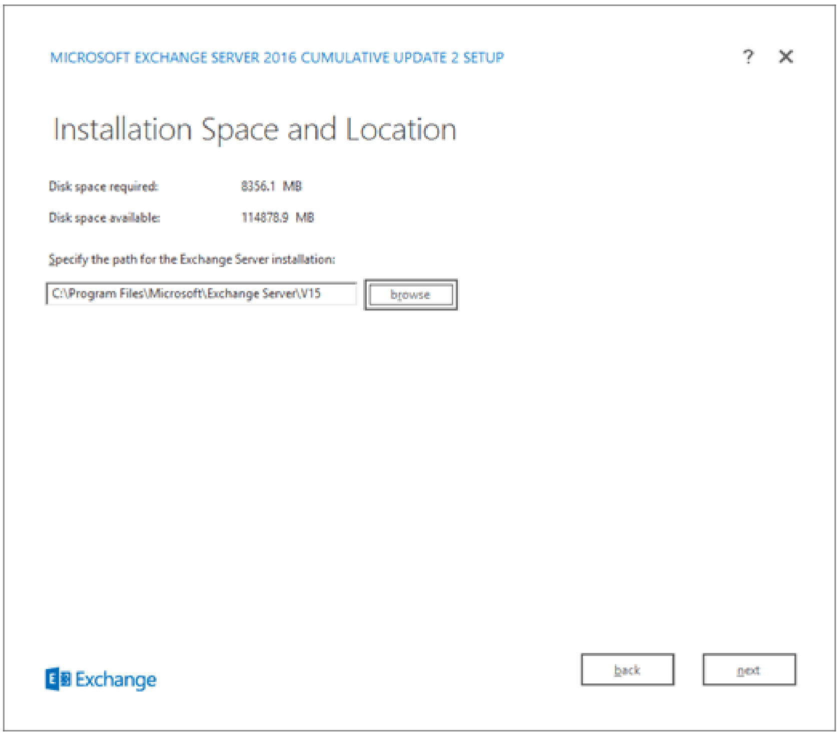

MICROSOFT EXCHANGE SERVER 2016 CUMULATIVE UPDATE 2 SETUP

Installation Space and Location

Disk space required: 8356.1 MB

Disk space available: 114878.9 MB

Specify the path for the Exchange Server installation:

C:\Program Files\Microsoft\Exchange Server\V15

10. If this is the first Exchange 2016 or Exchange 2019 server in your organization and you haven't already done the steps in Prepare Active Directory and domains for Exchange, you arrive on the **Exchange Organization** page. On this page, configure the following settings:

    - **Specify the name for this Exchange organization**: The default value is **First Organization**, but you typically use the company name for this value. The organization name is used internally by Exchange, isn't typically seen by users, doesn't affect the functionality of Exchange, and doesn't determine what you can use for email addresses.

        - The organization name can't contain more than 64 characters, and can't be blank.

        - Valid characters are A to Z, a to z, 0 to 9, hyphen or dash (-), and space, but leading or trailing spaces aren't allowed.

        - You can't change the organization name after it's set.

    - **Apply Active Directory split permission security model to the Exchange organization**: Most organizations don't need to select this option. If you need to separate management of Active Directory security principals and the Exchange configuration, split permissions might work for you. For more information, click **?**.

    Click **Next** to continue.

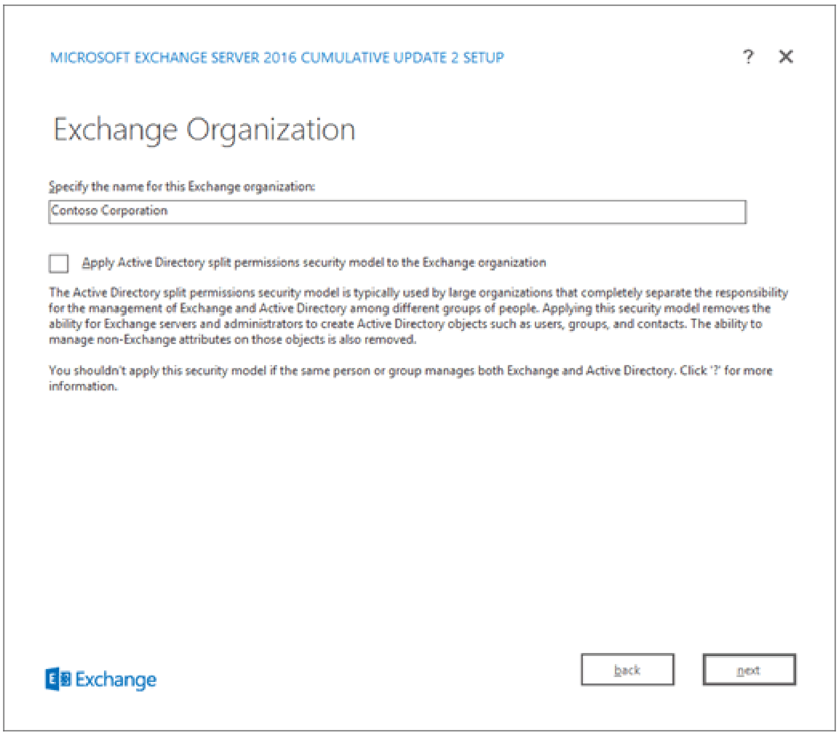
MICROSOFT EXCHANGE SERVER 2016 CUMULATIVE UPDATE 2 SETUP

Exchange Organization

Specify the name for this Exchange organization:

Contoso Corporation

□ Apply Active Directory split permissions security model to the Exchange organization

The Active Directory split permissions security model is typically used by large organizations that completely separate the responsibility for the management of Exchange and Active Directory among different groups of people. Applying this security model removes the ability for Exchange servers and administrators to create Active Directory objects such as users, groups, and contacts. The ability to manage non-Exchange attributes on those objects is also removed.

You shouldn't apply this security model if the same person or group manages both Exchange and Active Directory. Click '?' for more information.

11. On the **Malware Protection Settings** page, choose whether you want disable malware scanning. Malware scanning is enabled by default (the value **No** is selected). If you disable malware scanning, you can enable it in the future. Click **Next** to continue.

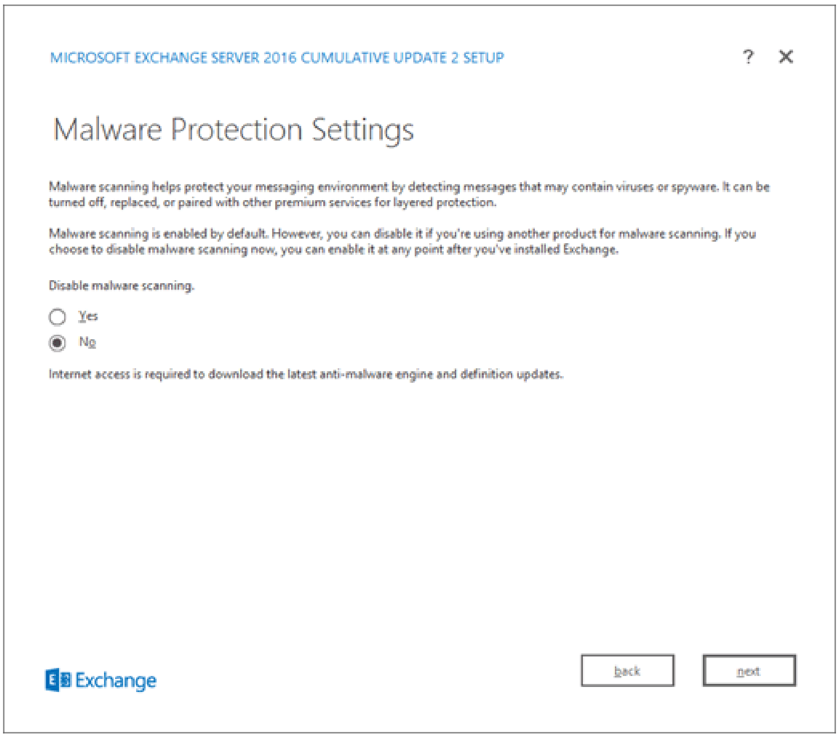
MICROSOFT EXCHANGE SERVER 2016 CUMULATIVE UPDATE 2 SETUP

Malware Protection Settings

Malware scanning helps protect your messaging environment by detecting messages that may contain viruses or spyware. It can be turned off, replaced, or paired with other premium services for layered protection.

Malware scanning is enabled by default. However, you can disable it if you're using another product for malware scanning. If you choose to disable malware scanning now, you can enable it at any point after you've installed Exchange.

Disable malware scanning.

○ Yes

● No

Internet access is required to download the latest anti-malware engine and definition updates.

12. On the **Readiness Checks** page, verify that the organization and server role prerequisite checks completed successfully. If they haven't, the only option on the page is **Retry**, so you need to resolve the errors before you can continue.

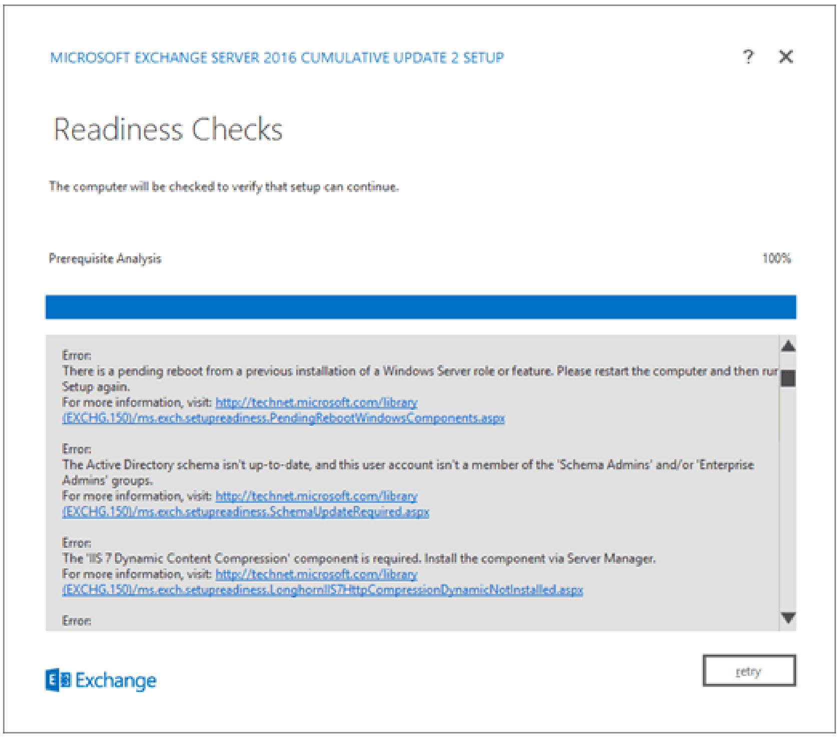

MICROSOFT EXCHANGE SERVER 2016 CUMULATIVE UPDATE 2 SETUP

Readiness Checks

The computer will be checked to verify that setup can continue.

Prerequisite Analysis 100%

Error:
There is a pending reboot from a previous installation of a Windows Server role or feature. Please restart the computer and then rur
Setup again.
For more information, visit: http://technet.microsoft.com/library
(EXCHG.150)/ms.exch.setupreadiness.PendingRebootWindowsComponents.aspx

Error:
The Active Directory schema isn't up-to-date, and this user account isn't a member of the 'Schema Admins' and/or 'Enterprise
Admins' groups.
For more information, visit: http://technet.microsoft.com/library
(EXCHG.150)/ms.exch.setupreadiness.SchemaUpdateRequired.aspx

Error:
The 'IIS 7 Dynamic Content Compression' component is required. Install the component via Server Manager.
For more information, visit: http://technet.microsoft.com/library
(EXCHG.150)/ms.exch.setupreadiness.LonghornIIS7HttpCompressionDynamicNotInstalled.aspx

Error:

After you resolve the errors, click **Retry** to run the prerequisite checks again. You can fix some errors without exiting Setup, while the fix for other errors requires you to restart the computer. If you restart the computer, you need to start over at Step 1.

When no more errors are detected on the **Readiness Checks** page, the **Retry** button changes to **Install** so you can continue. Be sure to review any warnings, and then click **Install** to install Exchange.

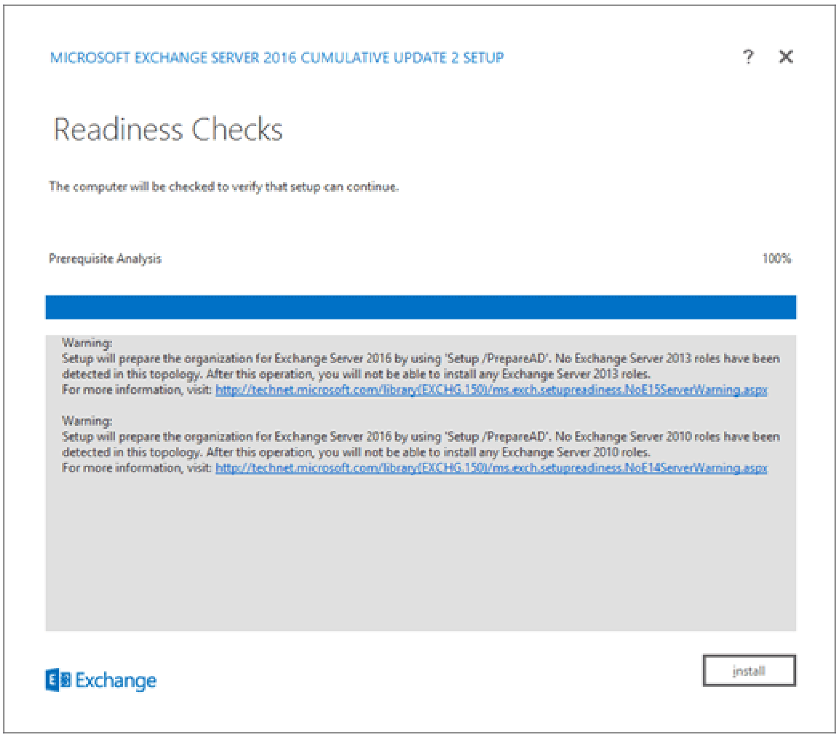

MICROSOFT EXCHANGE SERVER 2016 CUMULATIVE UPDATE 2 SETUP

Readiness Checks

The computer will be checked to verify that setup can continue.

Prerequisite Analysis 100%

Warning:
Setup will prepare the organization for Exchange Server 2016 by using 'Setup /PrepareAD'. No Exchange Server 2013 roles have been
detected in this topology. After this operation, you will not be able to install any Exchange Server 2013 roles.
For more information, visit: http://technet.microsoft.com/library(EXCHG.150)/ms.exch.setupreadiness.NoE15ServerWarning.aspx

Warning:
Setup will prepare the organization for Exchange Server 2016 by using 'Setup /PrepareAD'. No Exchange Server 2010 roles have been
detected in this topology. After this operation, you will not be able to install any Exchange Server 2010 roles.
For more information, visit: http://technet.microsoft.com/library(EXCHG.150)/ms.exch.setupreadiness.NoE14ServerWarning.aspx

13. On the **Setup Progress** page, a progress bar indicates how the installation is proceeding.

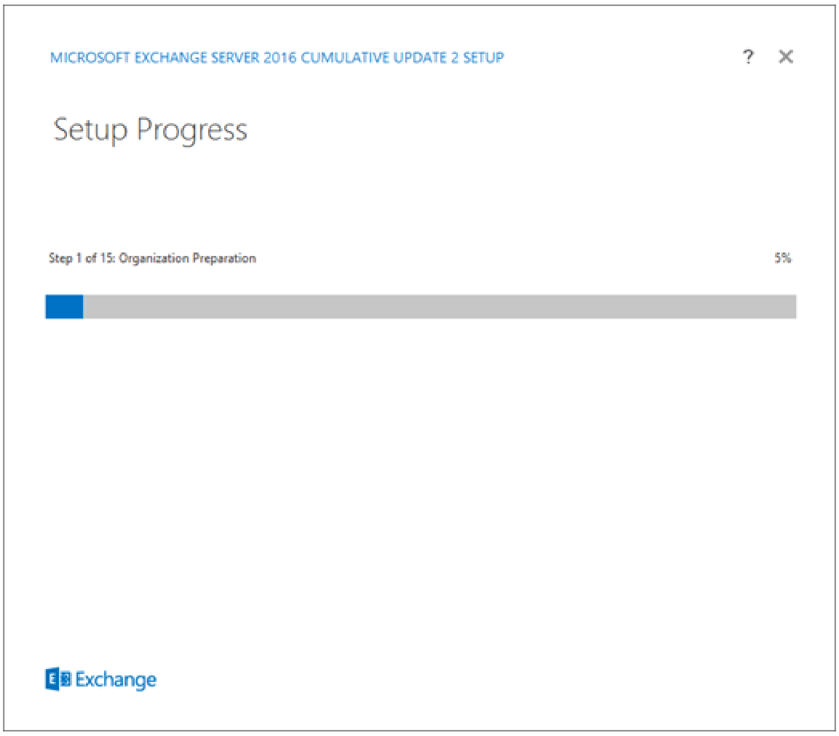

Setup Progress

Step 1 of 15: Organization Preparation

5%

14. On the **Setup Completed** page, click **Finish**, and then restart the computer.

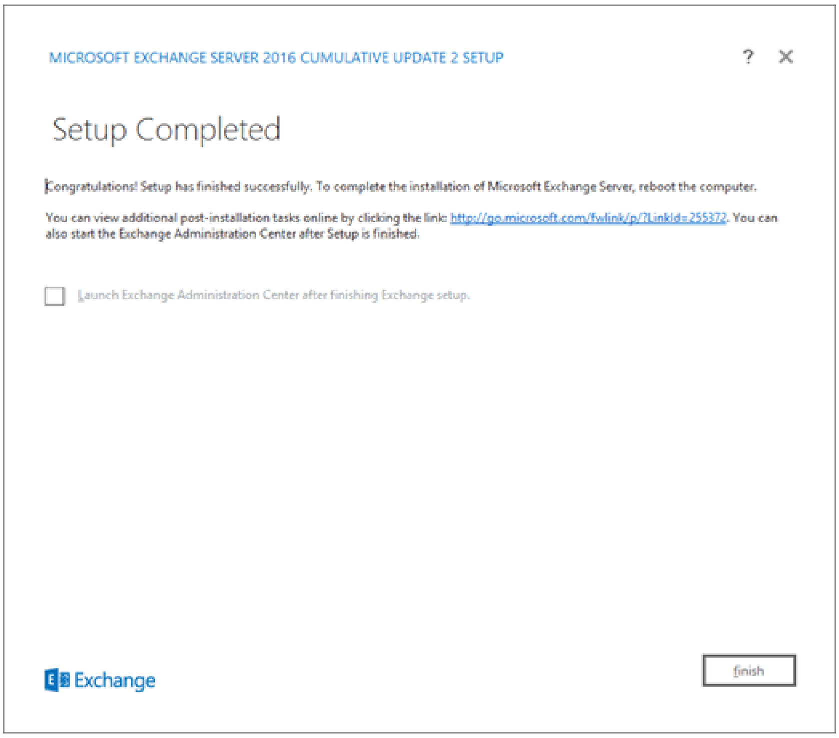

Setup Completed

Congratulations! Setup has finished successfully. To complete the installation of Microsoft Exchange Server, reboot the computer.

You can view additional post-installation tasks online by clicking the link: http://go.microsoft.com/fwlink/p/?LinkId=255372. You can also start the Exchange Administration Center after Setup is finished.

□ Launch Exchange Administration Center after finishing Exchange setup.

Next steps

- To verify that you've successfully installed Exchange, see Verify an Exchange installation.
- Complete your deployment by performing the tasks provided in Exchange post-installation tasks.
- Having problems? Ask for help in the Exchange forums. Visit the forums at Exchange Server.

Install Exchange Edge Transport servers using the Setup wizard

8/3/2020 • 4 minutes to read • Edit Online

Before you install an Exchange Server 2016 or Exchange Server 2019 Edge Transport server, verify the following prerequisites:

- We recommend that you install Edge Transport servers in a perimeter network that's outside of your organization's internal Active Directory forest. Installing the Edge Transport server role on domain-joined computers only enables domain management of Windows features and settings. Edge Transport servers don't directly access Active Directory. Instead, they use Active Directory Lightweight Directory Services (AD LDS) to store configuration and recipient information. For more information about the Edge Transport role, see Edge Transport servers.

> **NOTE**
> When Exchange Server 2019 Edge Transport server is domain-joined, a user from that domain must run the Exchange Management Shell. If a local user signs into the server, cmdlets in the Exchange Management Shell will result in "Access Denied".

- Verify the network, computer hardware, operating system, and software requirements at: Exchange Server system requirements and Exchange Server prerequisites.
- Verify the local account on the target computer is a member of the local Administrators group.
- Verify that you've read the release notes at Release notes for Exchange Server.

For more information about planning and deploying Exchange, see Planning and deployment for Exchange Server.

To install the Mailbox role on a computer, see Install Exchange Mailbox servers using the Setup wizard. Note that you can't install the Edge Transport role on a Mailbox server.

What do you need to know before you begin?

- Estimated time to complete: 40 minutes
- You need to configure the primary DNS suffix on the computer. For example, if the fully qualified domain name of your computer is edge.contoso.com, the DNS suffix for the computer is contoso.com. For more information, see Primary DNS Suffix is missing [ms.exch.setupreadiness.FqdnMissing].
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

Caution

After you install Exchange on a server, you must not change the server name. Renaming a server after you've installed an Exchange server role is not supported.

Install the Exchange Edge Transport server role

1. Download the latest version of Exchange. For more information, see Updates for Exchange Server.
2. In File Explorer, right-click on the Exchange ISO image file that you downloaded, and then select **Mount**. In

the resulting virtual DVD drive that appears, start Exchange Setup by double-clicking `Setup.exe`.

3. The Exchange Server Setup wizard opens. On the **Check for Updates?** page, choose one of the following options, and then click **Next** to continue:

   - **Connect to the Internet and check for updates**: We recommend this option, which searches for updates to the version of Exchange *that you're currently installing* (it doesn't detect newer Cumulative Updates). This option takes you to the **Downloading Updates** page that searches for updates. Click **Next** to continue.

   - **Don't check for updates right now**

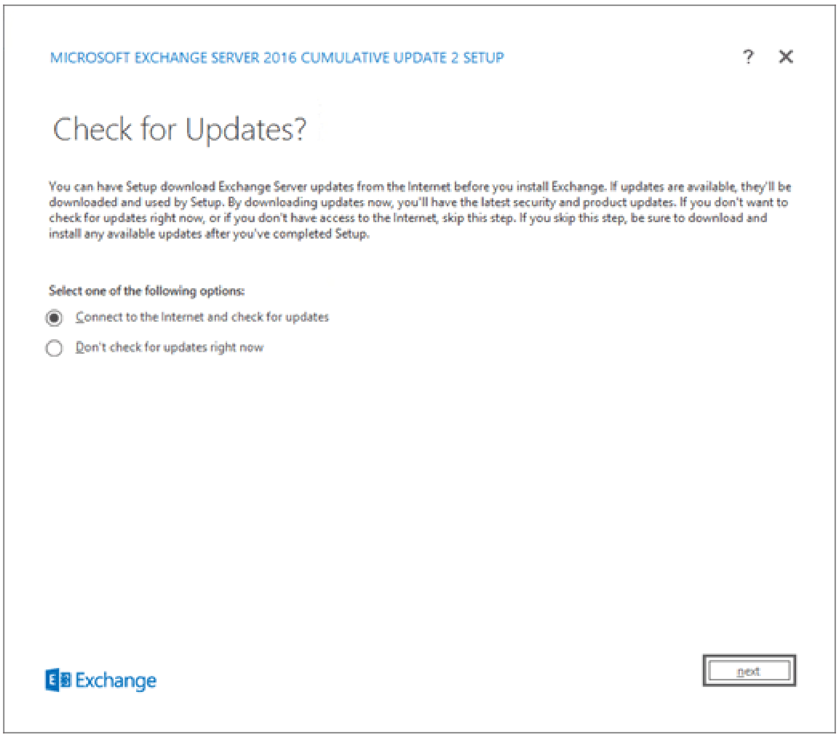

MICROSOFT EXCHANGE SERVER 2016 CUMULATIVE UPDATE 2 SETUP

Check for Updates?

You can have Setup download Exchange Server updates from the Internet before you install Exchange. If updates are available, they'll be downloaded and used by Setup. By downloading updates now, you'll have the latest security and product updates. If you don't want to check for updates right now, or if you don't have access to the Internet, skip this step. If you skip this step, be sure to download and install any available updates after you've completed Setup.

Select one of the following options:

● Connect to the Internet and check for updates

○ Don't check for updates right now

4. The **Copying Files** page shows the progress of copying files to the local hard drive. Typically, the files are copied to `%WinDir%\Temp\ExchangeSetup`, but you can confirm the location in the Exchange Setup log at `C:\ExchangeSetupLogs\ExchangeSetup.log`.

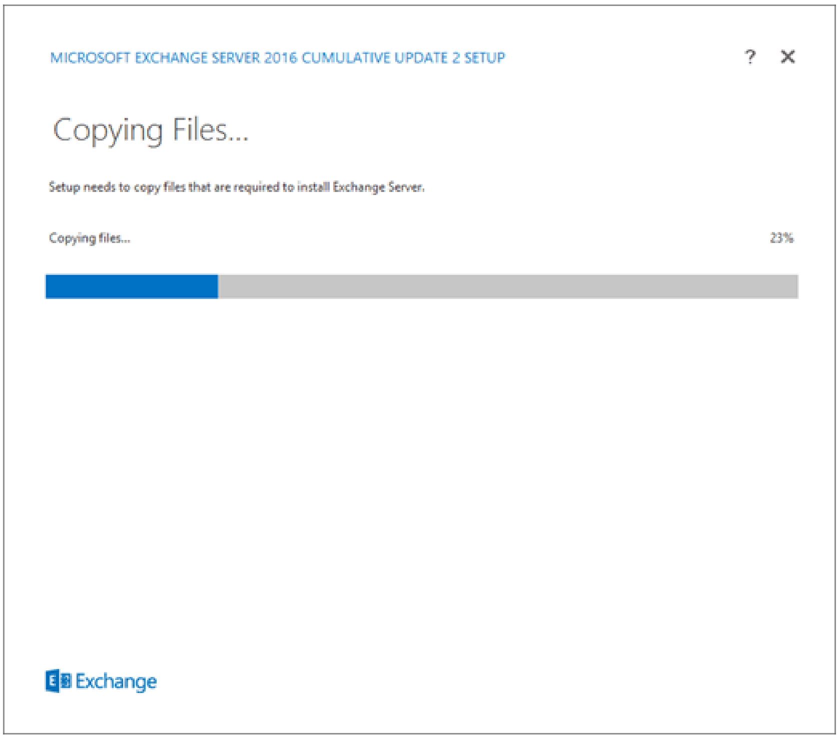

MICROSOFT EXCHANGE SERVER 2016 CUMULATIVE UPDATE 2 SETUP

Copying Files...

Setup needs to copy files that are required to install Exchange Server.

Copying files... 23%

5. On the **Introduction** page, we recommend that you visit the Exchange Server deployment planning links if you haven't already reviewed them. Click **Next** to continue.

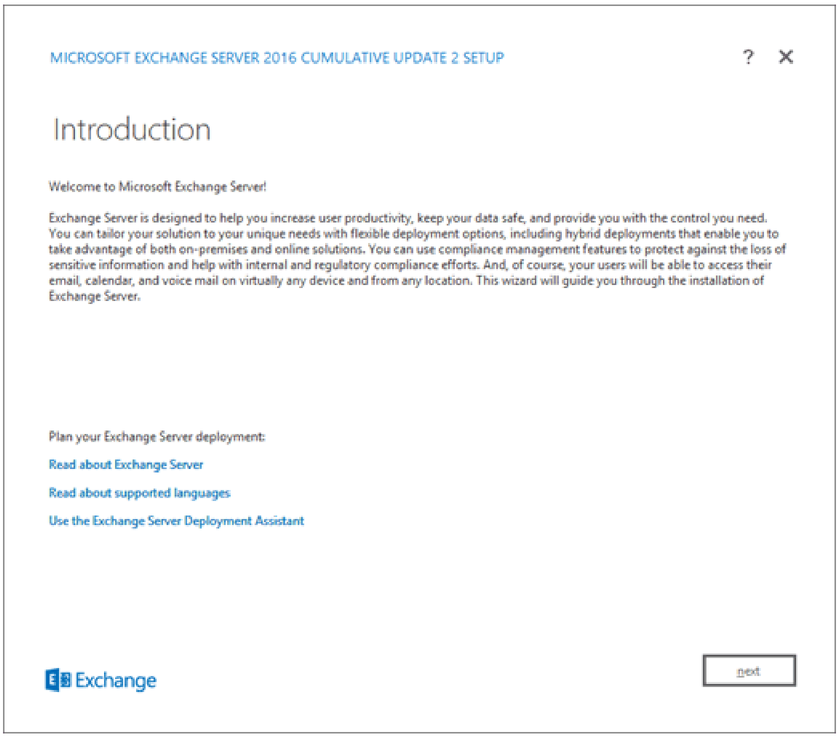

MICROSOFT EXCHANGE SERVER 2016 CUMULATIVE UPDATE 2 SETUP

Introduction

Welcome to Microsoft Exchange Server!

Exchange Server is designed to help you increase user productivity, keep your data safe, and provide you with the control you need. You can tailor your solution to your unique needs with flexible deployment options, including hybrid deployments that enable you to take advantage of both on-premises and online solutions. You can use compliance management features to protect against the loss of sensitive information and help with internal and regulatory compliance efforts. And, of course, your users will be able to access their email, calendar, and voice mail on virtually any device and from any location. This wizard will guide you through the installation of Exchange Server.

Plan your Exchange Server deployment:

Read about Exchange Server

Read about supported languages

Use the Exchange Server Deployment Assistant

6. On the **License Agreement** page, review the software license terms, select **I accept the terms in the license agreement**, and then click **Next** to continue.

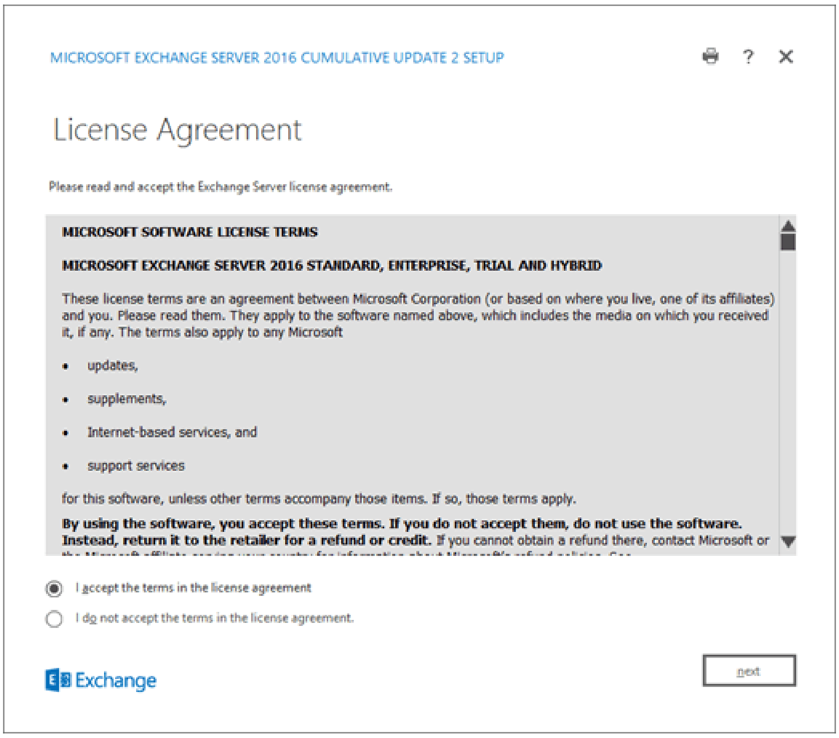

MICROSOFT EXCHANGE SERVER 2016 CUMULATIVE UPDATE 2 SETUP

License Agreement

Please read and accept the Exchange Server license agreement.

MICROSOFT SOFTWARE LICENSE TERMS

MICROSOFT EXCHANGE SERVER 2016 STANDARD, ENTERPRISE, TRIAL AND HYBRID

These license terms are an agreement between Microsoft Corporation (or based on where you live, one of its affiliates) and you. Please read them. They apply to the software named above, which includes the media on which you received it, if any. The terms also apply to any Microsoft

- updates,
- supplements,
- Internet-based services, and
- support services

for this software, unless other terms accompany those items. If so, those terms apply.

By using the software, you accept these terms. If you do not accept them, do not use the software. Instead, return it to the retailer for a refund or credit. If you cannot obtain a refund there, contact Microsoft or

● I accept the terms in the license agreement

○ I do not accept the terms in the license agreement.

7. On the **Recommended Settings** page, choose one of the following settings:

   - **Use recommended settings**: Exchange automatically sends error reports and information about your computer hardware and how you use Exchange to Microsoft. For information about what's sent to Microsoft and how it's used, click **?** or the help links on the page.

   - **Don't use recommended settings**: These settings are disabled, but you can enable them at any time after Setup completes.

   Click **Next** to continue.

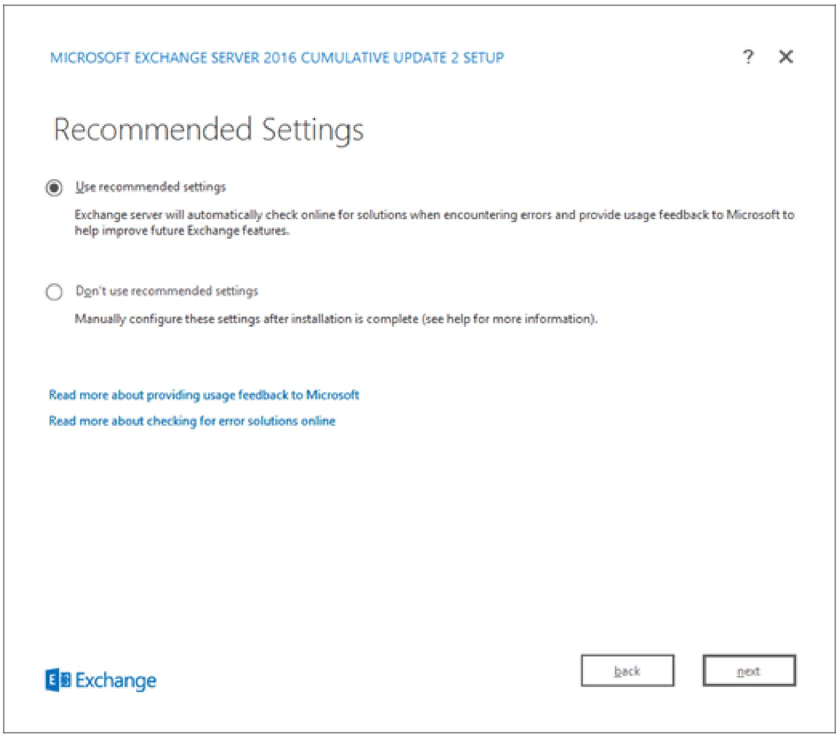

MICROSOFT EXCHANGE SERVER 2016 CUMULATIVE UPDATE 2 SETUP

Recommended Settings

● Use recommended settings

Exchange server will automatically check online for solutions when encountering errors and provide usage feedback to Microsoft to help improve future Exchange features.

○ Don't use recommended settings

Manually configure these settings after installation is complete (see help for more information).

Read more about providing usage feedback to Microsoft

Read more about checking for error solutions online

8. On the **Server Role Selection** page, configure the following options:

   - **Edge Transport role**: Select this option, which also automatically installs the **Management Tools**.

   - **Automatically install Windows Server roles and features that are required to install Exchange**: Select this option to have the Setup wizard install the required Windows prerequisites. You might need to reboot the computer to complete the installation of some Windows features. If you don't select this option, you need to install the Windows features manually.

   Note: Selecting this option installs only the *Windows features* that are required by Exchange. You need to install other prerequisites manually. For more information, see Exchange Server prerequisites.

   Click **Next** to continue.

9. On the **Installation Space and Location** page, either accept the default installation location (`C:\Program Files\Microsoft\Exchange Server\V15`), or click **Browse** to choose a new location. Make sure that you have enough disk space available in the location where you want to install Exchange. Click **Next** to continue.

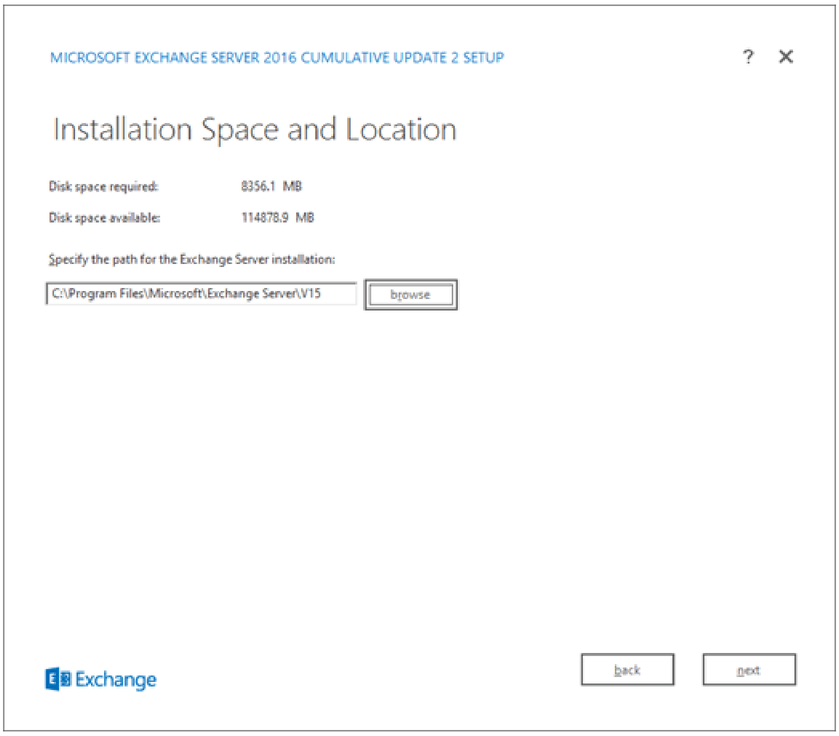

MICROSOFT EXCHANGE SERVER 2016 CUMULATIVE UPDATE 2 SETUP

Installation Space and Location

Disk space required: 8356.1 MB

Disk space available: 114878.9 MB

Specify the path for the Exchange Server installation:

C:\Program Files\Microsoft\Exchange Server\V15

10. On the **Readiness Checks** page, verify that the organization and server role prerequisite checks completed successfully. If they haven't, the only option on the page is **Retry**, so you need to resolve the errors before you can continue.

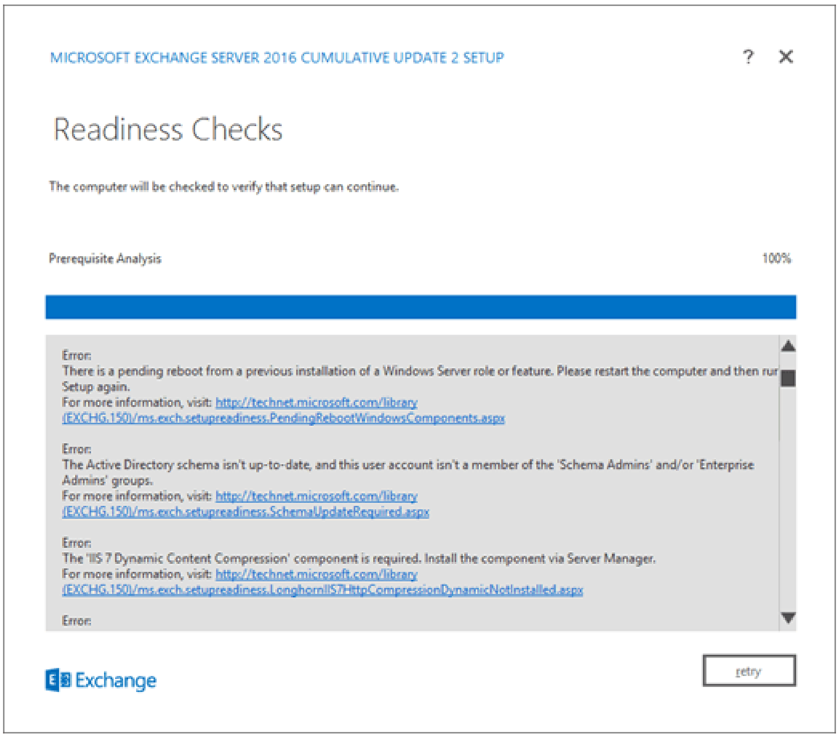

MICROSOFT EXCHANGE SERVER 2016 CUMULATIVE UPDATE 2 SETUP

Readiness Checks

The computer will be checked to verify that setup can continue.

Prerequisite Analysis 100%

```
Error:
There is a pending reboot from a previous installation of a Windows Server role or feature. Please restart the computer and then rur
Setup again.
For more information, visit: http://technet.microsoft.com/library
(EXCHG.150)/ms.exch.setupreadiness.PendingRebootWindowsComponents.aspx

Error:
The Active Directory schema isn't up-to-date, and this user account isn't a member of the 'Schema Admins' and/or 'Enterprise
Admins' groups.
For more information, visit: http://technet.microsoft.com/library
(EXCHG.150)/ms.exch.setupreadiness.SchemaUpdateRequired.aspx

Error:
The 'IIS 7 Dynamic Content Compression' component is required. Install the component via Server Manager.
For more information, visit: http://technet.microsoft.com/library
(EXCHG.150)/ms.exch.setupreadiness.LonghornIIS7HttpCompressionDynamicNotInstalled.aspx

Error:
```

After you resolve the errors, click **Retry** to run the prerequisite checks again. You can fix some errors without exiting Setup, while the fix for other errors requires you to restart the computer. If you restart the computer, you need to start over at Step 1.

When no more errors are detected on the **Readiness Checks** page, the **Retry** button changes to **Install** so you can continue. Be sure to review any warnings, and then click **Install** to install Exchange.

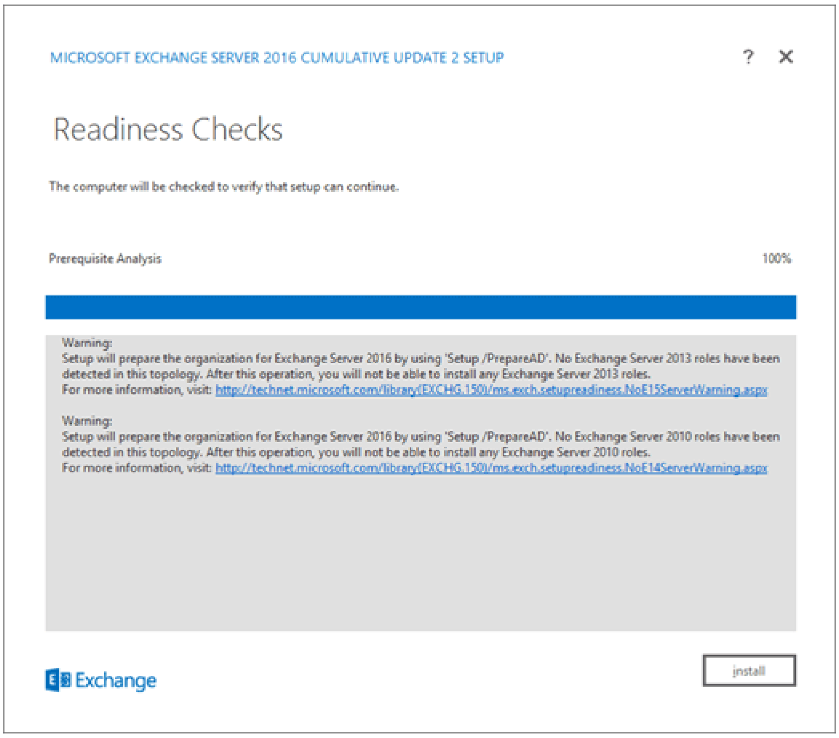
MICROSOFT EXCHANGE SERVER 2016 CUMULATIVE UPDATE 2 SETUP

Readiness Checks

The computer will be checked to verify that setup can continue.

Prerequisite Analysis 100%

Warning:
Setup will prepare the organization for Exchange Server 2016 by using 'Setup /PrepareAD'. No Exchange Server 2013 roles have been detected in this topology. After this operation, you will not be able to install any Exchange Server 2013 roles.
For more information, visit: http://technet.microsoft.com/library(EXCHG.150)/ms.exch.setupreadiness.NoE15ServerWarning.aspx

Warning:
Setup will prepare the organization for Exchange Server 2016 by using 'Setup /PrepareAD'. No Exchange Server 2010 roles have been detected in this topology. After this operation, you will not be able to install any Exchange Server 2010 roles.
For more information, visit: http://technet.microsoft.com/library(EXCHG.150)/ms.exch.setupreadiness.NoE14ServerWarning.aspx

11. On the **Setup Progress** page, a progress bar indicates how the installation is proceeding.

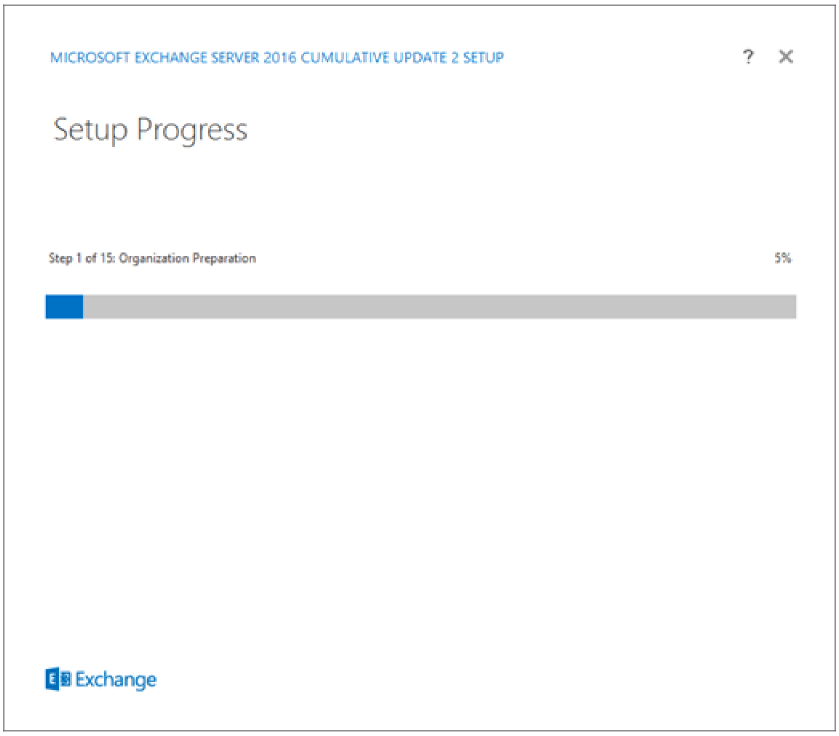
MICROSOFT EXCHANGE SERVER 2016 CUMULATIVE UPDATE 2 SETUP

Setup Progress

Step 1 of 15: Organization Preparation 5%

12. On the **Setup Completed** page, click **Finish**, and then restart the computer.

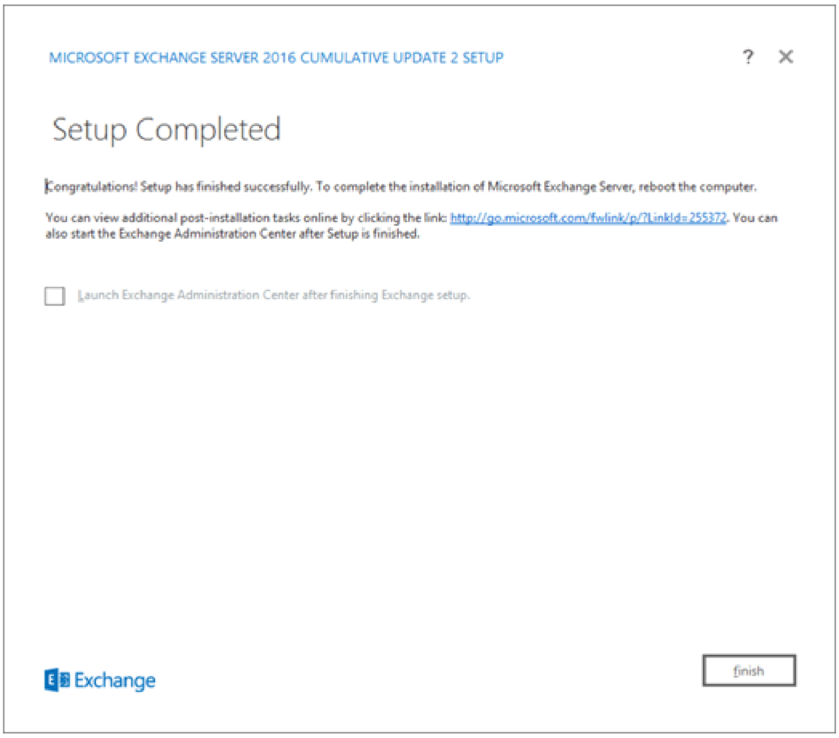

MICROSOFT EXCHANGE SERVER 2016 CUMULATIVE UPDATE 2 SETUP

Setup Completed

Congratulations! Setup has finished successfully. To complete the installation of Microsoft Exchange Server, reboot the computer.

You can view additional post-installation tasks online by clicking the link: http://go.microsoft.com/fwlink/p/?LinkId=255372. You can also start the Exchange Administration Center after Setup is finished.

□ Launch Exchange Administration Center after finishing Exchange setup.

Next steps

- To verify that you've successfully installed Exchange, see Verify an Exchange installation.
- Complete your deployment by performing the tasks provided in Exchange post-installation tasks.
- Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Use unattended mode in Exchange Setup

8/3/2020 • 14 minutes to read • Edit Online

Running Exchange Setup from the command line allows you to automate the installation of Exchange do and other related tasks on Exchange servers (for example, remove an existing Exchange server or recover a failed Exchange server).

This topic describes the available command line switches, and provides examples.

For more information about planning for Exchange 2016 or Exchange 2019, see Planning and deployment for Exchange Server.

For information about tasks to complete after installation, see Exchange Server post-installation tasks.

Primary command line switches for unattended mode

The primary (top-level, scenario-defining) command line switches that are available in unattended Setup mode in Exchange 2016 or Exchange 2019 are described in the following table:

SWITCH	DESCRIPTION
/IAcceptExchangeServerLicenseTerms	This switch is required in all unattended setup commands (whenever you run Setup.exe with any additional switches). If you don't use this switch, you'll get an error. To read the license terms, visit Microsoft License Terms.
/Mode:<InstallationMode> (or */m:<InstallationMode>*)	Valid values are: • **Install**: Installs Exchange on a new server using the Exchange server roles specified by the */Roles* switch. This is the default value if the command doesn't use the */Mode* switch. • **Uninstall**: Uninstalls Exchange from a working server. • **Upgrade**: Installs a Cumulative Update (CU) on an Exchange server. • **RecoverServer**: Recovers an Exchange server using the existing Exchange server object in Active Directory after a catastrophic hardware or software failure on the server. For instructions, see Recover Exchange servers.
/Roles:<ServerRole> (or */Role:<ServerRole>* or */r:<ServerRole>*)	This switch is required in `/Mode:Install` commands. Valid values are: • **Mailbox (or mb)**: Installs the Mailbox server role and the Exchange management tools on the local server. This is the default value. You can't use this value with **EdgeTransport**. • **EdgeTransport (or et)**: Installs the Edge Transport server role and the Exchange management tools on the local server. You can't use this value with **Mailbox**. • **ManagementTools (or mt or t)**: Installs the Exchange management tools on clients or other Windows servers that aren't running Exchange.
/PrepareAD (or */p*) */PrepareSchema* (or */ps*) */PrepareDomain:<DomainFQDN>* (or */pd:<DomainFQDN>*) */PrepareAllDomains* (or */pad*)	Use these switches to extend the Active Directory schema for Exchange, prepare Active Directory for Exchange, and prepare some or all Active Directory domains for Exchange. For more information, see Prepare Active Directory and domains for Exchange
/NewProvisionedServer[:<ServerName>] (or */nprs[:<ServerName>]* */RemoveProvisionedServer:<ServerName>* (or */rprs:<ServerName>*)	The */NewProvisionedServer* switch creates the Exchange server object in Active Directory. After that, a member of the Delegated Setup role group can install Exchange on the server. For more information, see Delegate the installation of Exchange servers. The */RemoveProvisionedServer* switch removes a provisioned Exchange server object from Active Directory *before* Exchange is installed on the server.
/AddUmLanguagePack:<Culture1>,<Culture2>...<CultureN> */RemoveUmLanguagePack:<Culture1>,<Culture2>...<CultureN>*	**Note**: These switches aren't available in Exchange 2019. They're only available in Exchange 2016. Adds or removes Unified Messaging (UM) language packs from existing Exchange 2016 Mailbox servers. UM language packs enable callers and Outlook Voice Access users to interact with the UM system in those languages. You can't add or remove the en-US language pack. You can install language packs on existing Mailbox servers by using the */AddUmLanguagePack* switch or by running the UMLanguagePack.<Culture>.exe file directly. You can only remove installed language packs by using the */RemoveUmLanguagePack* switch. For more information, see UM languages, prompts, and greetings.

Optional command line switches for unattended mode

The optional (supporting) command line switches that are available in unattended Setup mode in Exchange 2016 or Exchange 2019 are described in the following table:

SWITCH	VALID VALUES	DEFAULT VALUE	**AVAILABLE WITH	DESCRIPTION**
/ActiveDirectorySplitPermissions:<TrueOrFalse>	True or False	False	`/Mode:Install /Roles:Mailbox` or */PrepareAD* commands for the first Exchange server in the organization.	Specifies the Active Directory split permissions model when preparing Active Directory. For more information, see the "Active Directory split permissions" section in Understanding split permissions.
/AdamLdapPort:<TCPPortNumber>	A valid TCP port number	50389	`/Mode:Install /Roles:EdgeTransport` commands	Specifies a custom LDAP port to use for the Active Directory Lightweight Directory Services (AD LDS) instance on Edge Transport servers. The value is stored in the registry at `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Exchange` .
/AdamSslPort:<TCPPortNumber>	A valid TCP port number	50636	`/Mode:Install /Roles:EdgeTransport` commands	Specifies a custom SSL (TLS) port to use for the AD LDS instance on Edge Transport servers. The value is stored in the registry at `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Exchange` .

SWITCH	VALID VALUES	DEFAULT VALUE	**AVAILABLE WITH	DESCRIPTION**
/AnswerFile:"<PathAndFileName>" (or *af:"<PathAndFileName>"*)	The name and location of a text file (for example,"D:\Server data\answer.txt").	n/a	`/Mode:Install /Roles:Mailbox` or `/Mode:Install /Roles:EdgeTransport` commands	Use this switch to create a text file that you can use to install Exchange on multiple computers with the same settings. You can use the following switches in the answer file: *AdamLdapPort*, *AdamSslPort*, *CustomerFeedbackEnabled*, *DbFilePath*, *DisableAMFiltering*, *DoNotStartTransport*, *EnableErrorReporting*, *IAcceptExchangeServerLicenseTerms*, *LogFolderPath*, *Mdbname*, *OrganizationName*, *TenantOrganizationConfig*, and *UpdatesDir*. Don't use the forward slash character (/) with the switches in the answer file. Put each switch or switch/value pair on one line in the file.
/CustomerFeedbackEnabled:<TrueOrFalse>	True or False	False	`/Mode:Install` and */PrepareAD* commands	Specifies whether to allow or prevent Exchange from providing usage feedback to Microsoft to help improve future Exchange features. You can enable or disable error reporting on the server after setup is complete by using the *ErrorReportingEnabled* parameter on the **Set-ExchangeServer** cmdlet.
/DbFilePath:"<Path>\<FileName>.edb"	A folder path and an .edb filename (for example, "D:\Exchange Database Files\DB01\db01.edb").	**%ExchangeInstallPath%Mailbox\<DatabaseName>\<DatabaseName>.edb** where: • <DatabaseName> is **Mailbox Database <10DigitNumber>** that matches the default name of the database **or** the value you specified with the */MdbName* switch (without the .edb file name extension). • %ExchangeInstallPath% is **%ProgramFiles%\Microsoft\Exchange Server\V15\** or the location you specified with the */TargetDir* switch.	`/Mode:Install /Roles:Mailbox` commands	Specifies the location of the first mailbox database that's created on the new Mailbox server. You can specify the name of the database file with the */MdbName* switch and the location of the database transaction log files with the */LogFolderPath* switch.
/DisableAMFiltering	n/a	n/a	`/Mode:Install /Roles:Mailbox` commands	Disables the built-in Exchange antimalware filtering on Mailbox servers. For more information about antimalware filtering, see Antimalware protection in Exchange Server.
/DomainController:<ServerNameOrFQDN> (or */dc:<ServerNameOrFQDN>*)	The server name (for example, DC01) or FQDN (for example, dc01.contoso.com) of the domain controller.	A randomly-selected domain controller in the same Active Directory site as the target server where you're running Setup.	All */Mode* commands (except when you're installing an Edge Transport server) or */PrepareAD*, */PrepareSchema*, */PrepareDomain* and */PrepareAllDomains* commands	Specifies the domain controller that Exchange Setup uses to read from and write to Active Directory. The domain controller must meet the minimum requirements for Exchange 2016 or Exchange 2019. If you use this switch in */PrepareSchema* or */PrepareAD* commands that extend the Active Directory schema for Exchange, you must specify the schema master; otherwise, you'll get an error.
/DoNotStartTransport	n/a	n/a	`/Mode:Install /Roles:Mailbox` , `/Mode:Install /Roles:EdgeTransport` , and `/Mode:RecoverServer` commands.	Tells Setup to not start the Microsoft Exchange Transport service (mail flow) on Mailbox servers or Edge Transport servers after Setup is complete. You can use this switch to configure additional settings before the server accepts email messages (for example, configure antispam agents or move the queue database back onto a recovered Exchange server.)
/EnableErrorReporting	n/a	Disabled	`/Mode:Install` , `/Mode:Upgrade` , and `/Mode:RecoverServer` commands	Specifies whether to allow Exchange to automatically check online for solutions to errors that it encounters. You can enable or disable error reporting on the server after setup is complete by using the *ErrorReportingEnabled* parameter on the **Set-ExchangeServer** cmdlet.

SWITCH	VALID VALUES	DEFAULT VALUE	**AVAILABLE WITH	DESCRIPTION**
/InstallWindowsComponents	n/a	n/a	`/Mode:Install` commands	Installs the required Windows roles and features for the specified Exchange server role. If a reboot is required, Setup will resume where the installation ended.
/LogFolderPath:"<Path>"	A folder path (for example, "E:\Exchange Database Logs").	**%ExchangeInstallPath%Mailbox\\<DatabaseName>** where: • <DatabaseName> is **Mailbox Database <10DigitNumber>** that matches the default name of the database **or** the value you specified with the */MdbName* switch (without the .edb file name extension). • %ExchangeInstallPath% is **%ProgramFiles%\Microsoft\Exchange Server\V15\** or the location you specified with the */TargetDir* switch.	`/Mode:Install /Roles:Mailbox` commands	Specifies the location of the transaction log files for the first mailbox database that's created on the new Mailbox server. You can specify the location of the database files with the */DbFilePath* switch.
/MdbName:"<FileName>"	A database filename without the .edb extension (for example, "db01")	**Mailbox Database <10DigitNumber>** (for example, **Mailbox Database 0139595516**).	`/Mode:Install /Roles:Mailbox` commands	Specifies the name of the first mailbox database that's created on the new Mailbox server. You can specify the location of the database files with the */DbFilePath* switch.
/OrganizationName:"<Organization Name>" (or */on:"<Organization Name>"*)	A text string (for example, "Contoso Corporation").	Blank in command line setup; **First Organization** in the Exchange Setup wizard.	`/Mode:Install /Roles:Mailbox` or */PrepareAD* commands for the first Exchange server in the organization.	The organization name is used internally by Exchange, isn't typically seen by users, doesn't affect the functionality of Exchange, and doesn't determine what you can use for email addresses. • The organization name can't contain more than 64 characters, and can't be blank. • Valid characters are A to Z, a to z, 0 to 9, hyphen or dash (-), and space, but leading or trailing spaces aren't allowed. • You can't change the organization name after it's set.
/SourceDir:"<Path>" (or */s:"<Path>"*)	A folder path (for example, "Z:\Exchange).	The ServerRoles\UnifiedMessaging folder on the Exchange installation media.	*/AddUmLanguagePack* commands in Exchange 2016 (not available in Exchange 2019)	Specifies the location of the language packs (UMLanguagePack.<Culture>.exe files) to install on existing Exchange 2016 Mailbox servers.
/TargetDir:"<Path>" (or */t:"<Path>"*)	A folder path (for example, "D:\Program Files\Microsoft\Exchange").	**%ProgramFiles%\Microsoft\Exchange Server\V15\**	`/Mode:Install` and `/Mode:RecoverServer` commands	Specifies where to install Exchange on the server. You can't install Exchange in the root of a drive (for example, C:\), or on a ROM drive, RAM disk, network drive, removable disk, or unknown drive type. When you recover a failed Exchange server that was installed using a custom installation path, you need to use this switch to specify the custom path during the recovery.
/TenantOrganizationConfig:"<Path>"	A folder path (for example "C:\Data")	n/a	`/Mode:Install` or */PrepareAD* commands.	Required in hybrid deployments between on-premises organizations and Microsoft 365 or Office 365 to specify the location of the text file that contains the configuration information for your Microsoft 365 or Office 365 organization. You create this file by running the **Get-OrganizationConfig** cmdlet in Exchange Online PowerShell in your Microsoft 365 or Office 365 organization.
/UpdatesDir:"<Path>" (or */u:"<Path>"*)	A folder path (for example, "D:\Downloads\Exchange Updates").	The Updates folder at the root of the Exchange installation media.	`/Mode:Install`, `/Mode:Upgrade`, `/Mode:RecoverServer`, and */AddUmLanguagePack* commands.	Specifies the source location of updates for Setup to install. You can only specify one folder for updates. Any UM language packs located in this folder will be **automatically** installed on the target Exchange 2016 Mailbox server.

What do you need to know before you begin?

- Download the latest version of Exchange on the target computer. For more information, see Updates for Exchange Server.
- Verify the network, computer hardware, operating system, and software requirements at: Exchange Server system requirements and Exchange Server prerequisites.
- Verify that you've read the release notes at Release notes for Exchange Server.

  Caution

  After you install Exchange on a server, you must not change the server name. Renaming a server after you've installed an Exchange server role is not supported.

- For Mailbox servers:
  - Estimated time to complete: 60 minutes
  - The target server must be a member of an Active Directory domain.
  - The account that you use to install Exchange requires the following permissions:*:
    - **Enterprise Admins group membership**: Required if this is the first Exchange server in the organization.
    - **Schema Admins group membership**: Required if you haven't previously extended the Active Directory schema or prepared Active Directory for Exchange.
    - **Exchange Organization Management role group membership**: Required if you've already prepared the Active Directory domain that will contain the Exchange server, or if other Exchange servers already exist in the organization.

    * Members of the **Delegated Setup** role group can install Exchange on servers that have already been provisioned in Active Directory by an Exchange administrator. For more information, see Delegate the installation of Exchange servers.
- For Edge Transport servers:
  - Estimated time to complete: 40 minutes
  - We recommend that you install Edge Transport servers in a perimeter network that's outside of your organization's internal Active Directory forest. Installing the Edge Transport server role on domain-joined computers only enables domain management of Windows features and settings. Edge Transport servers don't directly access Active Directory. Instead, they use Active Directory Lightweight Directory Services (AD LDS) to store configuration and recipient information. For more information about the Edge Transport role, see Edge Transport servers.
  - Verify the local account on the target computer is a member of the local Administrators group on the target server.
  - You need to configure the primary DNS suffix on the computer. For example, if the fully qualified domain name of your computer is edge.contoso.com, the DNS suffix for the computer is contoso.com. For more information, see Primary DNS Suffix is missing [ms.exch.setupreadiness.FqdnMissing].
  - In coexistence scenarios, Exchange 2010 Hub Transport servers need an update before you can subscribe a Exchange 2016 Edge Transport server to an Active Directory site that contains Exchange 2010 Hub Transport servers. If you don't install this update, the EdgeSync Subscription won't work correctly for Exchange 2010 Hub Transport server that participate in EdgeSync synchronization. For more information, see Supported coexistence scenarios for Exchange 2016.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

Use Setup.exe to install Exchange in unattended mode

1. On the target server, open File Explorer, right-click on the Exchange ISO image file that you downloaded, and then select **Mount**. Note the virtual DVD drive letter that's assigned.
2. Open a Windows Command Prompt window. For example:
   - Press the Windows key + 'R' to open the **Run** dialog, type cmd.exe, and then press **OK**.
   - Press **Start**. In the **Search** box, type **Command Prompt**, then in the list of results, select **Command Prompt**.
3. In the Command Prompt window, use the following syntax:

   ```
   <Virtual DVD drive letter>:\Setup.exe /IAcceptExchangeServerLicenseTerms [Switches]
   ```

   Setup copies the setup files to the local computer.

   Setup checks the prerequisites, including all prerequisites specific to the server roles that you're installing. If you haven't met all the prerequisites, Setup fails and returns an error message that explains the reason for the failure. If you've met all the prerequisites, Setup installs Exchange.
4. Restart the server after the Exchange installation is complete.
5. Complete your deployment by performing the tasks provided in Exchange Server post-installation tasks.

Unattended mode examples

Prepare Active Directory for Exchange in unattended mode

This example configures "Fabrikam Ltd" as the Exchange organization name in Active Directory and prepares Active Directory for the version of Exchange that's being installed.

```
Setup.exe /IAcceptExchangeServerLicenseTerms /PrepareAD /OrganizationName:"Fabrikam Ltd"
```

For more information, see Prepare Active Directory and domains for Exchange.

Install Mailbox servers in unattended mode

- This example installs the first Exchange server (Mailbox server) in the organization, configures "Contoso Corporation" as the Exchange organization name in Active Directory, and installs the Exchange management tools on the local server.

  ```
  Setup.exe /IAcceptExchangeServerLicenseTerms /Mode:Install /Roles:Mailbox /on:"Contoso Corporation"
  ```

- This example installs the Mailbox server role and the management tools in the default folder on the local server in an organization where Active Directory has already been prepared for the version of Exchange that's being installed.

  ```
  Setup.exe /IAcceptExchangeServerLicenseTerms /mode:Install /r:MB
  ```

- This example installs the Mailbox server role and the management tools in the "C:\Exchange Server" folder on the local server.

  ```
  Setup.exe /IAcceptExchangeServerLicenseTerms /Mode:Install /Role:Mailbox /TargetDir:"C:\Exchange Server"
  ```

- This example installs the Mailbox server role on the local server by using the settings in the ExchangeConfig.txt file.

  ```
  Setup.exe /IAcceptExchangeServerLicenseTerms /mode:Install /role:Mailbox /AnswerFile:c:\ExchangeConfig.txt
  ```

- This example uses the domain controller named DC01 to read from and write to Active Directory while installing the Mailbox server role and the management tools on the local server.

  ```
  Setup.exe /IAcceptExchangeServerLicenseTerms /mode:Install /role:Mailbox /DomainController:DC01
  ```

- This example updates Exchange Setup with patches from the specified folder, and then installs the Mailbox server role and the management tools on the local server. In Exchange 2016 only, if any UM language packs are located in this folder, the language packs are automatically installed.

  ```
  Setup.exe /IAcceptExchangeServerLicenseTerms /role:Mailbox /UpdatesDir:"C:\ExchangeServer\New Patches"
  ```

Install Edge Transport servers in unattended mode

- This example installs the Edge Transport server role and the management tools in the default location on the local server.

  ```
  Setup.exe /IAcceptExchangeServerLicenseTerms /mode:Install /r:EdgeTransport
  ```

- This example installs the Edge Transport server role and the management tools in the specified folder on the local server.

  ```
  Setup.exe /IAcceptExchangeServerLicenseTerms /mode:Install /r:ET /TargetDir:"D:\Exchange Server"
  ```

Uninstall Exchange from servers in unattended mode

This example completely removes Exchange from the local server and removes the server's Exchange configuration from Active Directory.

```
Setup.exe /IAcceptExchangeServerLicenseTerms /mode:Uninstall
```

Remove provisioned Exchange server objects from Active Directory in unattended mode

This example removes the provisioned Exchange server object named Exchange03 from Active Directory *before* Exchange is installed on the server (if Exchange is already installed on the server, the command won't work).

```
Setup.exe /IAcceptExchangeServerLicenseTerms /rprs:Exchange03
```

For more information, see Delegate the installation of Exchange servers.

Add and remove UM language packs from existing Exchange 2016 Mailbox servers in unattended mode

> **NOTE**
>
> These procedures aren't available in Exchange 2019.

- This example installs the Russian and Spain Spanish language packs on the local Exchange 2016 Mailbox server from the specified folder.

  ```
  Setup.exe /IAcceptExchangeServerLicenseTerms /AddUmLanguagePack:ru-RU,es-ES /SourceDir:"D:\UM Language Packs"
  ```

- This example uninstalls the Korean UM language pack from the local Exchange 2016 Mailbox server.

  ```
  Setup.exe /IAcceptExchangeServerLicenseTerms /RemoveUmLanguagePack:ko-KR
  ```

Next steps

- To verify that you've successfully installed Exchange in unattended mode, see Verify Exchange Server installations.
- Complete your deployment by performing the tasks provided in Exchange post-installation tasks.
- Having problems? Ask for help in the Exchange forums. Visit the forums at Exchange Server.

Delegate the installation of Exchange servers

8/3/2020 • 4 minutes to read • Edit Online

In large companies, people who install and configure new Windows servers often aren't Exchange administrators. In Exchange 2016 and Exchange 2019, these users can still install Exchange on Windows servers, but only *after* an Exchange administrator *provisions* the Exchange server object in Active Directory. Provisioning an Exchange server object makes all of the required Active Directory changes independently of the actual installation of Exchange on a server. An Exchange administrator can provision a new Exchange server object hours or even days before Exchange is installed.

After an Exchange administrator provisions the Exchange server object, the only requirement for installing Exchange on the server is membership in the Delegated Setup role group, which allows members to install Exchange on provisioned servers. If this sounds like something you want to do, then this topic is for you.

What do you need to know before you begin?

- Estimated time to complete this procedure: Less than 10 minutes.
- You can only provision an Exchange server from the command line (Unattended Setup). You can't use the Exchange Setup wizard.
- You can't provision the first Exchange server object in your organization for the installation of Exchange by a delegate. An Exchange administrator needs to install the first Exchange server in the organization. After that, you can provision *additional* Exchange server objects so users who aren't Exchange administrators can install Exchange using delegated setup.
- A delegated user can't uninstall an Exchange server. To uninstall an Exchange server, you need to be an Exchange administrator.
- Download and use the latest available release of Updates for Exchange Server.
- To provision an Exchange server object, you need to be a member of the Organization Management role group.
- You can provision the Exchange server object in Active Directory from the target server itself, or from another computer.
- Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Command Prompt to provision Exchange 2019 servers

1. In File Explorer, right-click on the Exchange ISO image file that you downloaded, and then select **Mount**. Note the virtual DVD drive letter that's assigned.
2. Open a Windows Command Prompt window. For example:
   - Press the Windows key + 'R' to open the **Run** dialog, type cmd.exe, and then press **OK**.
   - Press **Start**. In the **Search** box, type **Command Prompt**, then in the list of results, select **Command Prompt**.
3. In the Command Prompt window, use the following syntax:

```
<Virtual DVD drive letter>:\Setup.exe /IAcceptExchangeServerLicenseTerms /NewProvisionedServer[:
<ServerName>]
```

If you run the command on the target server, you can use the */NewProvisionedServer* switch by itself. Otherwise, you need to specify the Name of the server to provision.

This example uses the Exchange installation files on drive E: to provision the server Mailbox01:

```
E:\Setup.exe /IAcceptExchangeServerLicenseTerms /NewProvisionedServer:Mailbox01
```

This example uses the Exchange installation files on drive E: to provision the local server where you're running the command:

```
E:\Setup.exe /IAcceptExchangeServerLicenseTerms /NewProvisionedServer
```

Note: To remove a provisioned Exchange server object from Active Directory *before* Exchange is installed on it, replace the */NewProvisionedServer* switch with */RemoveProvisionedServer*.

4. Add the appropriate users to the Delegated Setup role group so they can install Exchange on the provisioned server. To add users to a role group, see Add members to a role group. The delegates can use the procedures in Install Exchange Mailbox servers using the Setup wizard to install Exchange on the provisioned server.

How do you know this worked?

To verify that you've successfully provisioned an Exchange server for a delegate installation of Exchange, do the following steps:

1. In Active Directory Users & Computers, select **Microsoft Exchange Security Groups**, double-click **Exchange Servers**, and then select the **Members** tab.

2. On the **Members** tab, verify that the provisioned server is a member of the security group. A member of the Delegated Setup role group can now install Exchange on the server.

If your server is listed as a member of the Exchange Servers security group, it was properly provisioned. Someone who's a member of the Delegated Setup role group can now install Exchange on that server.

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

More information

An Exchange administrator might need to complete the deployment by performing the tasks provided in Exchange post-installation tasks.

The high-level Active Directory changes that are made when you provision an Exchange server object are described in the following list:

- A server object is created in the **CN=Servers,CN=Exchange Administrative Group (FYDIBOHF23SPDLT),CN=Administrative Groups,CN=<Organization Name>,CN=Microsoft Exchange,CN=Services,CN=Configuration,DC=<Root Domain>** configuration partition.

- The following access control entries (ACEs) are added to the server object within the configuration partition for the Delegated Setup role group:

  - Full Control on the server object and its child objects

  - Deny access control entry for the Send As extended right
  - Deny access control entry for the Receive As extended right
  - Deny CreateChild and DeleteChild permissions for Exchange Public Folder Store objects

    Note: Public folders are administered at an organizational level; therefore, the creation and deletion of public folder stores is restricted to Exchange administrators.

- The Active Directory computer account for the server is added to the Exchange Servers group.
- The server is added as a provisioned server in the Exchange admin center (EAC).

Only members of the Organization Management role group in Exchange have the permissions required to make these changes to Active Directory.

Exchange dev/test environments in Azure

8/3/2020 • 12 minutes to read • Edit Online

This topic steps you through creating an Exchange 2016 or Exchange 2019 dev/test deployment in Microsoft Azure. Here is the resulting configuration.

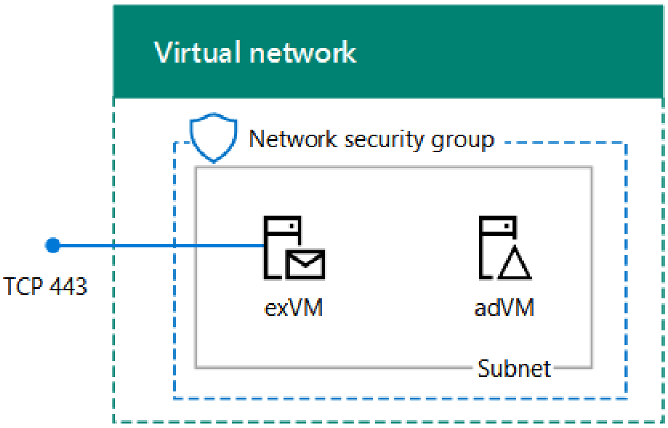

This configuration consists of a single Exchange server and a Windows Server Active Directory (AD) domain controller in a subnet of an Azure virtual network. This provides a basis and common starting point from which you can demonstrate Exchange and develop Exchange Server applications. This configuration is only for internal email and application testing on the Exchange server. No external email flow is configured.

There are three major phases to setting up this dev/test environment:

1. Set up the virtual network and domain controller (adVM).

2. Add the Exchange server (exVM).

3. Configure Exchange.

If you don't already have an Azure subscription, you can sign up for an Azure Free Trial. If you have an MSDN or Visual Studio subscription, see Monthly Azure credit for Visual Studio subscribers.

> **NOTE**
>
> Because Exchange makes changes to the schema in Windows Server AD, this configuration cannot use Azure Active Directory Domain Services.

Phase 1: Deploy the virtual network and a domain controller

You can create a new Azure virtual network with a domain controller with Azure PowerShell. You can run the following PowerShell commands from a Windows PowerShell command prompt or in the PowerShell Integrated Script Environment (ISE). If you have not installed Azure PowerShell, see Get started with Azure PowerShell cmdlets.

> **NOTE**
>
> These commands are for Azure PowerShell 1.0.0 and later.

1. Sign into your Azure account.

```
Connect-AzAccount
```

2. Get your subscription name using the following command.

   ```
   Get-AZSubscription | Sort SubscriptionName | Select SubscriptionName
   ```

3. Set your Azure subscription with the following commands. Set the `$subscr` variable by replacing everything within the quotes, including the < and > characters, with the correct name.

   ```
   $subscrName="<subscription name>"
   Select-AzSubscription -SubscriptionName $subscrName
   ```

4. Create a new resource group. To determine a unique resource group name, use this command to list your existing resource groups.

   ```
   Get-AZResourceGroup | Sort ResourceGroupName | Select ResourceGroupName
   ```

   Create your new resource group with these commands. Set the variables by replacing everything within the quotes, including the < and > characters, with the correct names.

   ```
   $rgName="<resource group name>"
   $locName="<location name, such as West US>"
   New-AZResourceGroup -Name $rgName -Location $locName
   ```

5. Resource Manager-based virtual machines require a Resource Manager-based storage account. You must pick a globally unique name for your storage account *that contains only lowercase letters and numbers*. You can use this command to list the existing storage accounts.

   ```
   Get-AZStorageAccount | Sort StorageAccountName | Select StorageAccountName
   ```

   Use this command to test whether a proposed storage account name is unique.

   ```
   Get-AZStorageAccountNameAvailability "<proposed name>"
   ```

   Create a new storage account for your new test environment with these commands.

   ```
   $rgName="<your new resource group name>"
   $saName="<storage account name>"
   $locName=(Get-AZResourceGroup -Name $rgName).Location
   New-AZStorageAccount -Name $saName -ResourceGroupName $rgName -Type Standard_LRS -Location $locName
   ```

6. Create the EXSrvrVnet Azure Virtual Network that will host the EXSrvrSubnet subnet and protect it with a network security group.

```
$rgName="<name of your new resource group>"
$locName=(Get-AZResourceGroup -Name $rgName).Location
$exSubnet=New-AZVirtualNetworkSubnetConfig -Name EXSrvrSubnet -AddressPrefix 10.0.0.0/24
New-AZVirtualNetwork -Name EXSrvrVnet -ResourceGroupName $rgName -Location $locName -AddressPrefix
10.0.0.0/16 -Subnet $exSubnet -DNSServer 10.0.0.4
$rule1 = New-AZNetworkSecurityRuleConfig -Name "RDPTraffic" -Description "Allow RDP to all VMs on the
subnet" -Access Allow -Protocol Tcp -Direction Inbound -Priority 100 -SourceAddressPrefix Internet -
SourcePortRange * -DestinationAddressPrefix * -DestinationPortRange 3389
$rule2 = New-AZNetworkSecurityRuleConfig -Name "ExchangeSecureWebTraffic" -Description "Allow HTTPS to
the Exchange server" -Access Allow -Protocol Tcp -Direction Inbound -Priority 101 -SourceAddressPrefix
Internet -SourcePortRange * -DestinationAddressPrefix "10.0.0.5/32" -DestinationPortRange 443
New-AZNetworkSecurityGroup -Name EXSrvrSubnet -ResourceGroupName $rgName -Location $locName -
SecurityRules $rule1, $rule2
$vnet=Get-AZVirtualNetwork -ResourceGroupName $rgName -Name EXSrvrVnet
$nsg=Get-AZNetworkSecurityGroup -Name EXSrvrSubnet -ResourceGroupName $rgName
Set-AZVirtualNetworkSubnetConfig -VirtualNetwork $vnet -Name EXSrvrSubnet -AddressPrefix "10.0.0.0/24" -
NetworkSecurityGroup $nsg
$vnet | Set-AzVirtualNetwork
```

7. Create the adVM virtual machine in Azure. adVM is a domain controller for the corp.contoso.com Windows Server AD domain and a DNS server for the virtual machines of the EXSrvrVnet virtual network.

   First, fill in the name of your resource group, Azure location, and storage account name and run these commands at the Azure PowerShell command prompt on your local computer to create an Azure virtual machine for adVM.

```
$rgName="<resource group name>"
# Create an availability set for domain controller virtual machines
New-AZAvailabilitySet -ResourceGroupName $rgName -Name dcAvailabilitySet -Location $locName -Sku Aligned
-PlatformUpdateDomainCount 5 -PlatformFaultDomainCount 2
# Create the domain controller virtual machine
$vnet=Get-AZVirtualNetwork -Name EXSrvrVnet -ResourceGroupName $rgName
$pip = New-AZPublicIpAddress -Name adVM-NIC -ResourceGroupName $rgName -Location $locName -
AllocationMethod Dynamic
$nic = New-AZNetworkInterface -Name adVM-NIC -ResourceGroupName $rgName -Location $locName -SubnetId
$vnet.Subnets[0].Id -PublicIpAddressId $pip.Id -PrivateIpAddress 10.0.0.4
$avSet=Get-AZAvailabilitySet -Name dcAvailabilitySet -ResourceGroupName $rgName
$vm=New-AZVMConfig -VMName adVM -VMSize Standard_D1_v2 -AvailabilitySetId $avSet.Id
$vm=Set-AZVMOSDisk -VM $vm -Name adVM-OS -DiskSizeInGB 128 -CreateOption FromImage -StorageAccountType
"Standard_LRS"
$diskConfig=New-AZDiskConfig -AccountType "Standard_LRS" -Location $locName -CreateOption Empty -
DiskSizeGB 20
$dataDisk1=New-AZDisk -DiskName adVM-DataDisk1 -Disk $diskConfig -ResourceGroupName $rgName
$vm=Add-AZVMDataDisk -VM $vm -Name adVM-DataDisk1 -CreateOption Attach -ManagedDiskId $dataDisk1.Id -Lun
1
$cred=Get-Credential -Message "Type the name and password of the local administrator account for adVM."
$vm=Set-AZVMOperatingSystem -VM $vm -Windows -ComputerName adVM -Credential $cred -ProvisionVMAgent -
EnableAutoUpdate
$vm=Set-AZVMSourceImage -VM $vm -PublisherName MicrosoftWindowsServer -Offer WindowsServer -Skus 2012-
R2-Datacenter -Version "latest"
$vm=Add-AZVMNetworkInterface -VM $vm -Id $nic.Id
New-AZVM -ResourceGroupName $rgName -Location $locName -VM $vm
```

You will be prompted for a username and password. This article will refer to this username as ADMIN_NAME. Use a strong password and record both in a secure location.

Note: The password that you specify cannot be "pass@word1". It must be between 8-123 characters long and must satisfy at least 3 of the following password complexity requirements:

- Contains an uppercase letter
- Contains an lowercase letter

- Contains a numeric digit
- Contains a special character

It can take a few minutes for Azure to build the virtual machine.

Connect to the domain controller virtual machine using local administrator account credentials

1. In the Azure portal, click **Resource Groups > <your resource group name>> adVM > Connect**.
2. Run the adVM.rdp file that is downloaded, and then click **Connect**.
3. In **Windows Security**, click **Use another account**. In **User name**, type **adVM**<ADMIN_NAME>.
4. In **Password**, type the password of the ADMIN_NAME account, and then click **OK**.
5. When prompted, click **Yes**.
6. Add an extra data disk as a new volume with the drive letter F: with these commands at an administrator-level Windows PowerShell command prompt on adVM.

```
$disk=Get-Disk | where {$_.PartitionStyle -eq "RAW"}
$diskNumber=$disk.Number
Initialize-Disk -Number $diskNumber
New-Partition -DiskNumber $diskNumber -UseMaximumSize -AssignDriveLetter
Format-Volume -DriveLetter F
```

7. Configure adVM as a domain controller and DNS server for the corp.contoso.com domain. Run these commands at an administrator-level Windows PowerShell command prompt on adVM.

```
Install-WindowsFeature AD-Domain-Services -IncludeManagementTools
Install-ADDSForest -DomainName corp.contoso.com -DatabasePath "F:\NTDS" -SysvolPath "F:\SYSVOL" -LogPath
"F:\Logs"
```

Note that these commands can take a few minutes to complete.

After adVM restarts, reconnect to the adVM virtual machine.

Connect to the domain controller virtual machine using domain credentials

1. In the Azure portal, click **Resource Groups >** <the name of your new resource group> **> adVM > Connect**.
2. Run the adVM.rdp file that is downloaded, and then click **Connect**.
3. In **Windows Security**, click **Use another account**. In **User name**, type **CORP**<ADMIN_NAME>.
4. In **Password**, type the password of the ADMIN_NAME account, and then click **OK**.
5. When prompted, click **Yes**.
6. From the desktop, open an administrator-level Windows PowerShell command prompt and run the following command:

```
Add-WindowsFeature RSAT-ADDS-Tools
```

Here is the result of Phase 1.

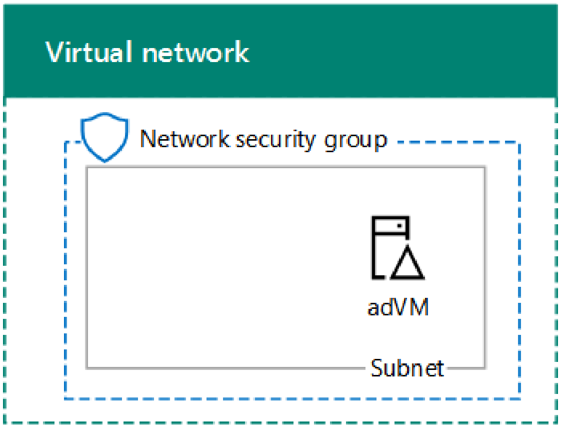

Phase 2: Create the Exchange virtual machine

In this phase, you create an Exchange virtual machine in the EXSrvrVNet virtual network and make it a member of the CORP domain.

To create the Exchange virtual machine with Azure PowerShell, first log in to Azure with your Azure account from the Windows PowerShell command prompt (if needed).

```
Connect-AzAccount
```

You must determine a globally unique DNS name for the exVM virtual machine. You must pick a globally unique DNS name *that contains only lowercase letters and numbers*. You can do this with the following PowerShell commands:

```
$vmDNSName="<DNS name to test>"
$rgName="<resource group name>"
$locName=(Get-AZResourceGroup -Name $rgName).Location
Test-AZDnsAvailability -DomainQualifiedName $vmDNSName -Location $locName
```

If you see "True", your proposed name is globally unique.

Next, fill in the variable values and run the resulting block at the PowerShell prompt.

```
# Set up key variables
$subscrName="<name of your Azure subscription>"
$rgName="<your resource group name>"
$vmDNSName="<unique, public DNS name for the Exchange server>"
# Set the Azure subscription
Select-AzSubscription -SubscriptionName $subscrName
# Get the Azure location and storage account names
$locName=(Get-AZResourceGroup -Name $rgName).Location
$saName=(Get-AZStorageaccount | Where {$_.ResourceGroupName -eq $rgName}).StorageAccountName
# Create an availability set for Exchange virtual machines
New-AZAvailabilitySet -ResourceGroupName $rgName -Name exAvailabilitySet -Location $locName -Sku Aligned  -
PlatformUpdateDomainCount 5 -PlatformFaultDomainCount 2
# Specify the virtual machine name and size
$vmName="exVM"
$vmSize="Standard_D3_v2"
$vnet=Get-AZVirtualNetwork -Name "EXSrvrVnet" -ResourceGroupName $rgName
$avSet=Get-AZAvailabilitySet -Name exAvailabilitySet -ResourceGroupName $rgName
$vm=New-AZVMConfig -VMName $vmName -VMSize $vmSize -AvailabilitySetId $avSet.Id
# Create the NIC for the virtual machine
$nicName=$vmName + "-NIC"
$pipName=$vmName + "-PublicIP"
$pip=New-AZPublicIpAddress -Name $pipName -ResourceGroupName $rgName -DomainNameLabel $vmDNSName -Location
$locName -AllocationMethod Dynamic
$nic=New-AZNetworkInterface -Name $nicName -ResourceGroupName $rgName -Location $locName -SubnetId
$vnet.Subnets[0].Id -PublicIpAddressId $pip.Id -PrivateIpAddress "10.0.0.5"
# Create and configure the virtual machine
$cred=Get-Credential -Message "Type the name and password of the local administrator account for exVM."
$vm=Set-AZVMOSDisk -VM $vm -Name ($vmName +"-OS") -DiskSizeInGB 128 -CreateOption FromImage -
StorageAccountType "Standard_LRS"
$vm=Set-AZVMOperatingSystem -VM $vm -Windows -ComputerName $vmName -Credential $cred -ProvisionVMAgent -
EnableAutoUpdate
$vm=Set-AZVMSourceImage -VM $vm -PublisherName MicrosoftWindowsServer -Offer WindowsServer -Skus 2019-
Datacenter -Version "latest"
$vm=Add-AZVMNetworkInterface -VM $vm -Id $nic.Id
New-AZVM -ResourceGroupName $rgName -Location $locName -VM $vm
```

NOTE

This command block uses a standard storage account created in phase 1 to reduce costs for this dev/test environment. For a production Exchange server, you must use a premium storage account.

From the Azure portal, connect to the exVM virtual machine using the credentials of the local administrator account.

Next, join exVM to the Windows AD domain with these commands at a Windows PowerShell prompt.

```
Add-Computer -DomainName "corp.contoso.com"
Restart-Computer
```

Note that you must supply domain account credentials after entering the **Add-Computer** command. Use the CORP\<ADMIN_NAME> account and password.

Here is the result of Phase 2.

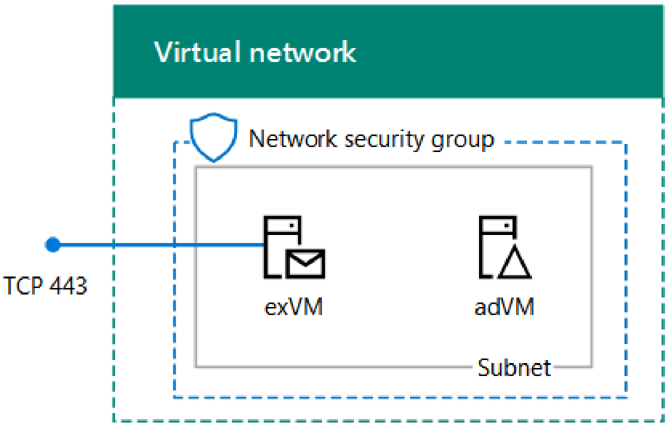

Phase 3: Configure Exchange

In this phase, you configure Exchange on exVM and test mail delivery between two mailboxes.

Prepare Windows Server AD

1. At the PowerShell command prompt on your local computer, run the following command:

```
Write-Host (Get-AZPublicIpaddress -Name "exVM-PublicIP" -ResourceGroup $rgName).DnsSettings.Fqdn
```

2. Note or copy the full DNS name from the display of the command. This is the Internet DNS name of the exVM virtual machine. You will need this value later.

3. If needed, connect to the adVM virtual machine with the Azure portal using the CORP\<ADMIN_NAME> account and password.

4. From the Start screen of adVM, type **Active Directory**, and then click **Active Directory Domains and Trusts**.

5. Right-click **Active Directory Domains and Trusts**, and then click **Properties**.

6. In **Alternative UPN suffixes**, type or copy the Internet DNS name of the exVM virtual machine from step 2, click **Add**, and then click **OK**.

7. Close the remote desktop session with adVM.

Install Exchange

1. Connect to the exVM virtual machine with the Azure portal using the CORP\<ADMIN_NAME> account and password.

2. From exVM, open an administrator-level Windows PowerShell command prompt and run the following commands.

```
Install-WindowsFeature 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, RSAT-ADDS-Tools
Restart-Computer
```

3. Connect to the exVM virtual machine with the Azure portal using the CORP\<ADMIN_NAME> account and password.

4. From Server Manager, click **Local Server**. In the **Properties** for exVM, click **On** for **IE Enhanced Security Configuration**. In **Internet Explorer Enhanced Security Configuration**, click **Off** for both Administrators and Users, and then click **OK**.

5. From the Start screen, click **Internet Explorer**, and then download the Unified Communications Managed API 4.0 Runtime from https://www.microsoft.com/download/details.aspx?id=34992. When prompted, click **Run**.

6. When prompted with the Microsoft Unified Communications Managed API 4.0, Runtime Setup, click **Next**.

7. Click **I have read and accept the license terms**, and then click **Install**. On the **Installation is Complete** page, click **Finish**.

8. From Internet Explorer, download the latest version of Exchange. For more information, see Updates for Exchange Server.

9. Click **Save** to store the ISO file in the Downloads folder.

10. Click **Open Folder**, right-click the Exchange ISO file, and then click **Mount**.

11. From an administrator-level Windows PowerShell command prompt on exVM, run the following:

```
e:
.\setup.exe /mode:Install /role:Mailbox /OrganizationName:Contoso /IacceptExchangeServerLicenseTerms
Restart-Computer
```

Wait until Exchange setup completes, which can take some time, and exVM restarts.

Add two mailboxes to the Exchange server

1. Connect to the exVM virtual machine with the Azure portal using the CORP\<ADMIN_NAME> account and password.

2. From the Start screen, type **Exchange**, and then click **Exchange Management Shell**.

3. Copy the following commands to Notepad, insert the Internet DNS name of the exVM virtual machine for the **$dnsName** variable, and then copy and paste the resulting commands into the Exchange Management Shell.

```
$dnsName="<Internet DNS name of the exVM virtual machine>"
$user1Name="chris@" + $dnsName
$user2Name="janet@" + $dnsName
$db=Get-MailboxDatabase
$dbName=$db.Name
$password = Read-Host "Enter password" -AsSecureString
```

4. Record the password specified in a safe place. Next, run these commands to create two mailboxes.

```
New-Mailbox -UserPrincipalName $user1Name -Alias chris -Database $dbName -Name ChrisAshton -
OrganizationalUnit Users -Password $password -FirstName Chris -LastName Ashton -DisplayName "Chris
Ashton"
New-Mailbox -UserPrincipalName $user2Name -Alias janet -Database $dbName -Name JanetSchorr -
OrganizationalUnit Users -Password $password -FirstName Janet -LastName Schorr -DisplayName "Janet
Schorr"
```

Test email delivery between mailboxes

1. From the browser on your local computer, access the web site **https://**<Internet DNS name of the exVM virtual machine> **/owa**. When prompted with an error page for the website's security certificate, click

Continue to this website. On the Outlook sign-in page, use the corp\chris account name with its password.

2. When prompted to specify the language and time zone, select the appropriate value for each, and then click **Save**.
3. From Chris Ashton's inbox, click **New**. In **To**, type **janet** and then click **Search Directory**. For **Subject**, type **Test message**, and then click **Send**.
4. Click the user icon in the upper right part of the Mail web page, and then click **Sign out**.
5. On the Outlook sign-in page, use the corp\janet account name with its password. When prompted to specify the language and time zone, select the appropriate value for each, and then click **Save**.
6. Verify that the inbox contains the test message from Chris Ashton. Click it, then click **Reply all**. In the body of the message, type **Replied**, and then click **Send**.
7. Click the user icon in the upper right part of the Mail web page, and then click **Sign out**.
8. On the Outlook sign-in page, use the corp\chris account name with its password. Verify that the reply email message sent from Janet is in the inbox.

You are now ready to test Exchange features or applications.

Stop and start the virtual machines

Azure virtual machines incur an ongoing cost when they are running. To help minimize the cost of your Exchange dev/test environment, use these commands to stop the virtual machines:

```
$rgName="<your resource group name>"
Stop-AZVM -Name exVM -ResourceGroupName $rgName -Force
Stop-AZVM -Name adVM -ResourceGroupName $rgName -Force
```

To start them again, use these commands:

```
$rgName="<your resource group name>"
Start-AZVM -Name adVM -ResourceGroupName $rgName
Start-AZVM -Name exVM -ResourceGroupName $rgName
```

See also

Troubleshoot outbound SMTP connectivity issues in Azure

Deploy new installations of Exchange

Exchange Server system requirements

Exchange Server

What's new in Exchange Server

Cloud adoption Test Lab Guides (TLGs)

Upgrade Exchange to the latest Cumulative Update

8/3/2020 • 4 minutes to read • Edit Online

If you have Exchange Server 2016 or Exchange Server 2019 installed, you can upgrade the Exchange servers to the latest Cumulative Update (CU). Because each CU is a full installation of Exchange that includes updates and changes from all previous CUs, you don't need to install any previous CUs or Exchange 2016 RTM or Exchange 2019 RTM first. For more information about the latest available Exchange CUs, see Updates for Exchange Server.

Caution

After you upgrade Exchange to a newer CU, you can't uninstall the new version to revert to the previous version. Uninstalling the new version completely removes Exchange from the server.

What do you need to know before you begin?

- Estimated time to complete: 60 minutes
- The account that you'll use to install the CU requires membership in the Exchange Organization Management role group. If the CU requires Active Directory schema updates or domain preparation, the account will likely require additional permissions. For more information, see Prepare Active Directory and domains for Exchange Server.
- Check the Release notes before you install the CU.
- Verify the target server meets the potentially new system requirements and prerequisites for the CU. For more information, see Exchange Server system requirements and Exchange Server prerequisites.

  Caution

  Any customized Exchange or Internet Information Server (IIS) settings that you made in Exchange XML application configuration files on the Exchange server (for example, web.config files or the EdgeTransport.exe.config file) **will be overwritten** when you install an Exchange CU. Be sure save this information so you can easily re-apply the settings after the install. After you install the Exchange CU, you need to re-configure these settings.
- After you install an Exchange CU, you need to restart the computer so that changes can be made to the registry and operating system.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Best Practices

- Always keep your servers as up to date as possible. This especially applies to the installation of a new server.
- Always install the latest Cumulative Update when creating a new server.
- There is no need to install the RTM build or previous builds and then upgrade to the latest Cumulative Update. This is because each Cumulative Update is a full build of the product.
- Reboot the server beforehand.
- Test the new update in a non-production environment first to avoid any problems in the new update affecting the running production environment.

- Have a tested and working backup of both the Active Directory and your Exchange Server.
- Backup any and all customizations. They will not survive the update.
- Use an elevated command prompt to run the Cumulative Update.
- Temporarily disable any anti-virus software during the update process.
- Reboot your server upon completion of the update.

Install an Exchange CU using the Setup wizard

1. Download the latest version of Exchange on the target computer. For more information, see Updates for Exchange Server.

2. In File Explorer, right-click on the Exchange CU ISO image file that you downloaded, and then select **Mount**. In the resulting virtual DVD drive that appears, start Exchange Setup by double-clicking `Setup.exe`.

3. The Exchange Server Setup wizard opens. On the **Check for Updates?** page, choose one of the following options, and then click **Next** to continue:

   - **Connect to the Internet and check for updates**: We recommend this option, which searches for updates to the version of Exchange *that you're currently installing* (it doesn't detect newer CUs). This option takes you to the **Downloading Updates** page that searches for updates. Click **Next** to continue.

   - **Don't check for updates right now**

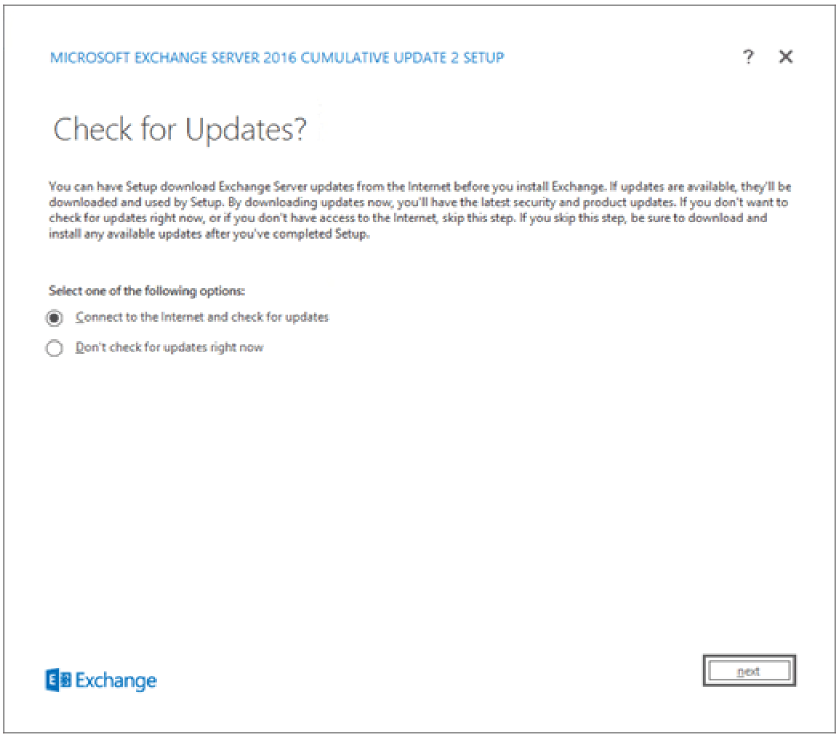

MICROSOFT EXCHANGE SERVER 2016 CUMULATIVE UPDATE 2 SETUP

Check for Updates?

You can have Setup download Exchange Server updates from the Internet before you install Exchange. If updates are available, they'll be downloaded and used by Setup. By downloading updates now, you'll have the latest security and product updates. If you don't want to check for updates right now, or if you don't have access to the Internet, skip this step. If you skip this step, be sure to download and install any available updates after you've completed Setup.

Select one of the following options:

● Connect to the Internet and check for updates

○ Don't check for updates right now

4. The **Copying Files** page shows the progress of copying files to the local hard drive. Typically, the files are copied to `%WinDir%\Temp\ExchangeSetup`, but you can confirm the location in the Exchange Setup log at `C:\ExchangeSetupLogs\ExchangeSetup.log`.

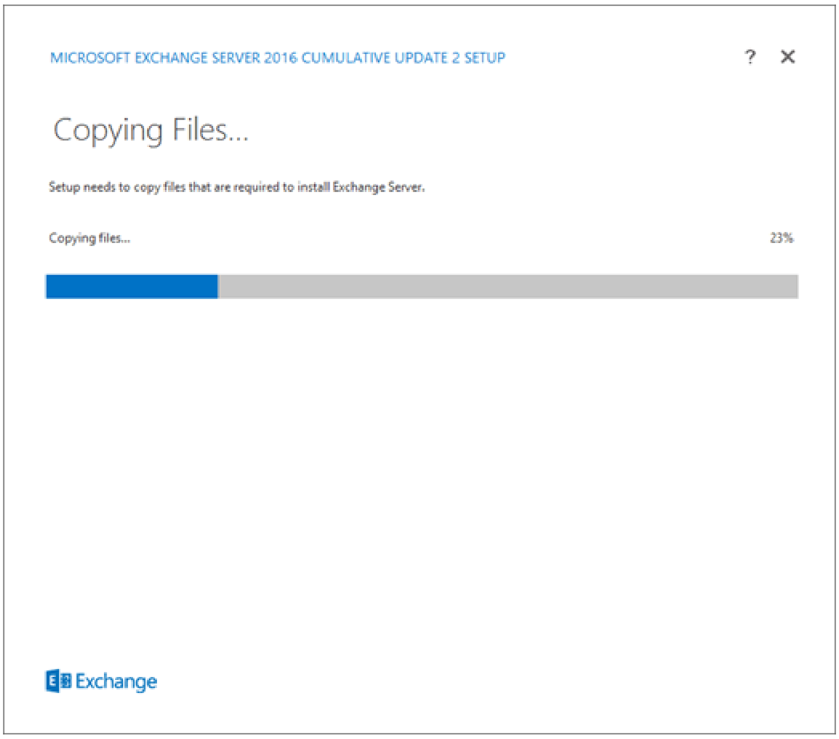

MICROSOFT EXCHANGE SERVER 2016 CUMULATIVE UPDATE 2 SETUP

Copying Files...

Setup needs to copy files that are required to install Exchange Server.

Copying files... 23%

5. The **Upgrade** page shows that Setup detected the existing installation of Exchange, so you're upgrading Exchange on the server (not installing a new Exchange server). Click **Next** to continue.

6. On the **License Agreement** page, review the software license terms, select **I accept the terms in the license agreement**, and then click **Next** to continue.

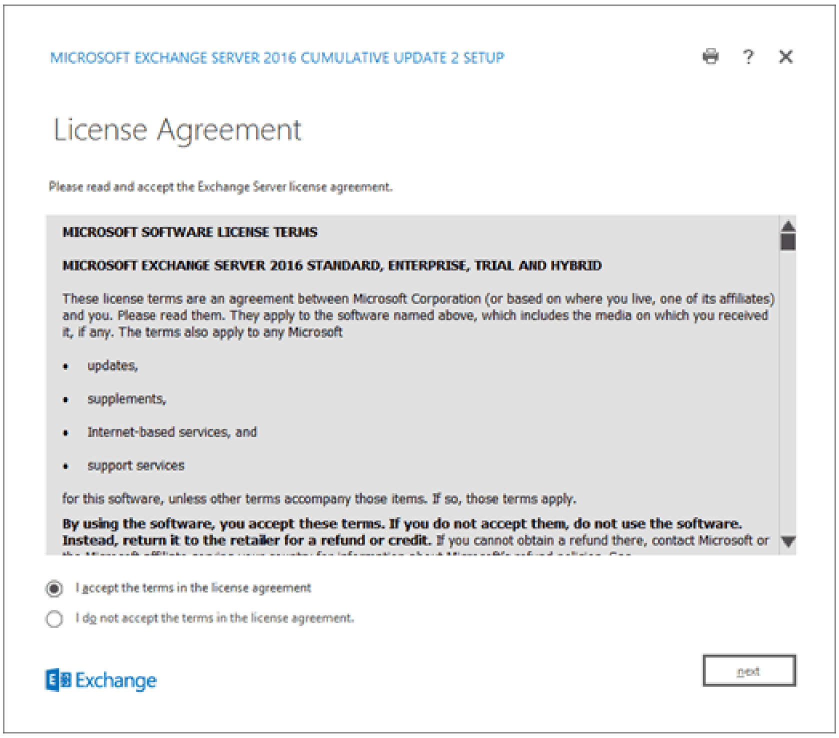

MICROSOFT EXCHANGE SERVER 2016 CUMULATIVE UPDATE 2 SETUP

License Agreement

Please read and accept the Exchange Server license agreement.

MICROSOFT SOFTWARE LICENSE TERMS

MICROSOFT EXCHANGE SERVER 2016 STANDARD, ENTERPRISE, TRIAL AND HYBRID

These license terms are an agreement between Microsoft Corporation (or based on where you live, one of its affiliates) and you. Please read them. They apply to the software named above, which includes the media on which you received it, if any. The terms also apply to any Microsoft

- updates,
- supplements,
- Internet-based services, and
- support services

for this software, unless other terms accompany those items. If so, those terms apply.

By using the software, you accept these terms. If you do not accept them, do not use the software. Instead, return it to the retailer for a refund or credit. If you cannot obtain a refund there, contact Microsoft or

● I accept the terms in the license agreement

○ I do not accept the terms in the license agreement.

7. On the **Readiness Checks** page, verify that the prerequisite checks completed successfully. If they haven't, the only option on the page is **Retry**, so you need to resolve the errors before you can continue.

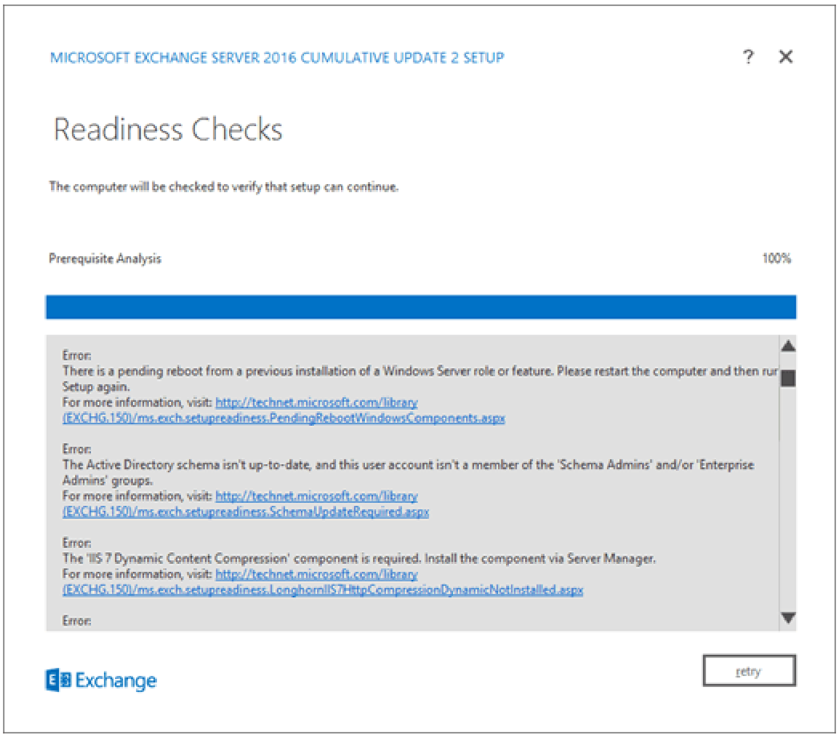

Readiness Checks

The computer will be checked to verify that setup can continue.

Prerequisite Analysis 100%

Error:
There is a pending reboot from a previous installation of a Windows Server role or feature. Please restart the computer and then rur
Setup again.
For more information, visit: http://technet.microsoft.com/library
(EXCHG.150)/ms.exch.setupreadiness.PendingRebootWindowsComponents.aspx

Error:
The Active Directory schema isn't up-to-date, and this user account isn't a member of the 'Schema Admins' and/or 'Enterprise
Admins' groups.
For more information, visit: http://technet.microsoft.com/library
(EXCHG.150)/ms.exch.setupreadiness.SchemaUpdateRequired.aspx

Error:
The 'IIS 7 Dynamic Content Compression' component is required. Install the component via Server Manager.
For more information, visit: http://technet.microsoft.com/library
(EXCHG.150)/ms.exch.setupreadiness.LonghornIIS7HttpCompressionDynamicNotInstalled.aspx

Error:

After you resolve the errors, click **Retry** to run the prerequisite checks again. You can fix some errors without exiting Setup, while the fix for other errors requires you to restart the computer. If you restart the computer, you need to start over at Step 1.

When no more errors are detected on the **Readiness Checks** page, the **Retry** button changes to **Install** so you can continue. Be sure to review any warnings, and then click **Install** to install Exchange.

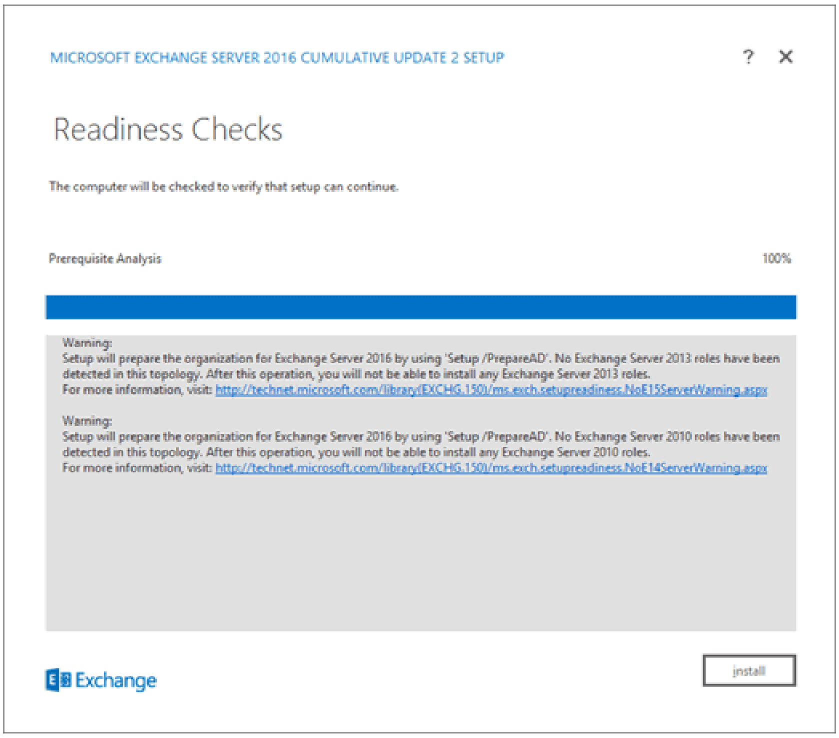

Readiness Checks

The computer will be checked to verify that setup can continue.

Prerequisite Analysis 100%

Warning:
Setup will prepare the organization for Exchange Server 2016 by using 'Setup /PrepareAD'. No Exchange Server 2013 roles have been detected in this topology. After this operation, you will not be able to install any Exchange Server 2013 roles.
For more information, visit: http://technet.microsoft.com/library(EXCHG.150)/ms.exch.setupreadiness.NoE15ServerWarning.aspx

Warning:
Setup will prepare the organization for Exchange Server 2016 by using 'Setup /PrepareAD'. No Exchange Server 2010 roles have been detected in this topology. After this operation, you will not be able to install any Exchange Server 2010 roles.
For more information, visit: http://technet.microsoft.com/library(EXCHG.150)/ms.exch.setupreadiness.NoE14ServerWarning.aspx

8. On the **Setup Progress** page, a progress bar indicates how the installation is proceeding.

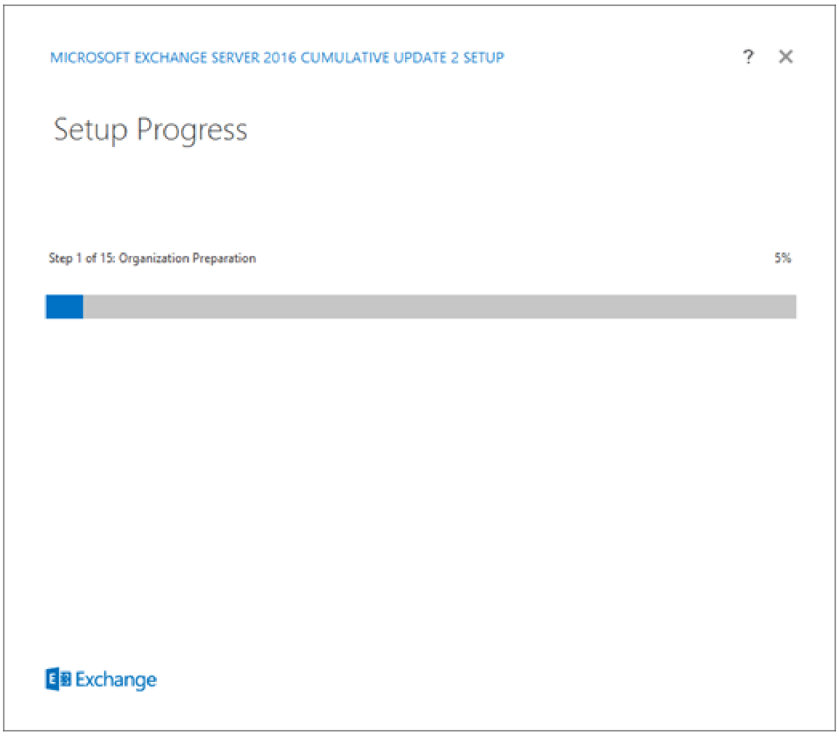

Setup Progress

Step 1 of 15: Organization Preparation 5%

9. On the **Setup Completed** page, click **Finish**, and then restart the computer.

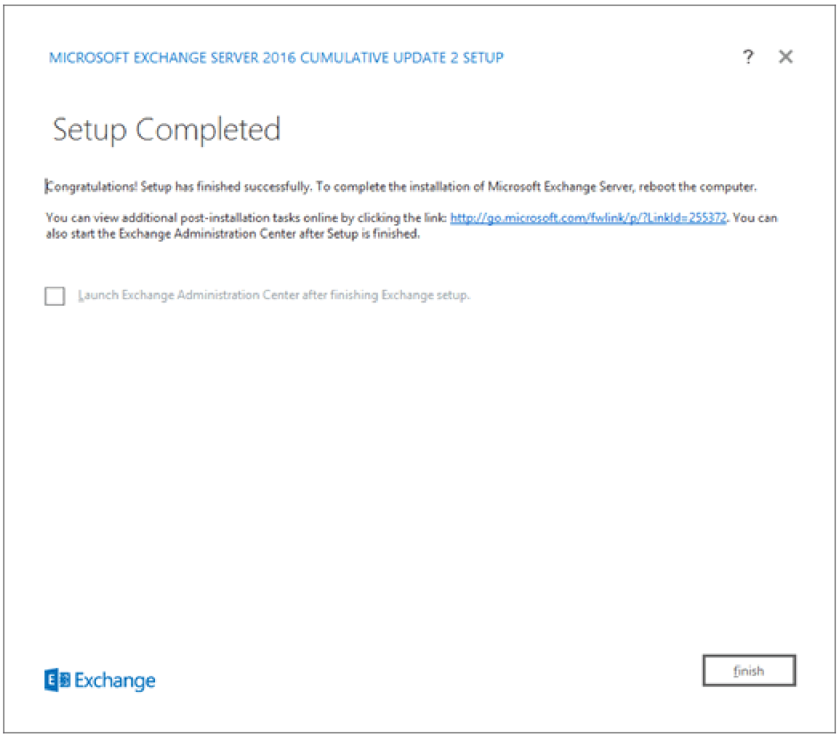

Setup Completed

Congratulations! Setup has finished successfully. To complete the installation of Microsoft Exchange Server, reboot the computer.

You can view additional post-installation tasks online by clicking the link: http://go.microsoft.com/fwlink/p/?LinkId=255372. You can also start the Exchange Administration Center after Setup is finished.

□ Launch Exchange Administration Center after finishing Exchange setup.

Install an Exchange CU using unattended Setup from the command line

To install an Exchange CU from the command line, use the following syntax:

```
<Virtual DVD drive letter>:\Setup.exe /IAcceptExchangeServerLicenseTerms /Mode:Upgrade [/DomainController:
<ServerFQDN>] [/EnableErrorReporting]
```

Notes:

- The optional */DomainController* switch specifies the domain controller that Setup uses to read from an write to Active Directory.

- The optional */EnableErrorReporting* switch enables Setup to automatically submit critical error reports to Microsoft. Microsoft uses this information to diagnose problems and provide solutions.

This example uses the Exchange CU files on drive E: to install the CU on the local server, and uses the domain controller dc01.contoso.com to read from and write to Active Directory.

```
E:\Setup.exe /IAcceptExchangeServerLicenseTerms /Mode:Upgrade /DomainController:dc01.contoso.com
```

For more information about unattended Setup from the command line, see Install Exchange using unattended mode.

How do you know this worked?

To verify that you've successfully installed an Exchange CU, see Verify Exchange Server installations.

Exchange Server supportability matrix

8/3/2020 • 12 minutes to read • Edit Online

The Exchange Server supportability matrix provides a central source for Exchange administrators to easily locate information about the level of support available for any configuration or required component for supported versions of Microsoft Exchange Server.

Release model

The following table identifies the release model for each supported version of Exchange.

In Exchange Server 2010 and earlier, each update rollup package (RU) is cumulative. An RU for Exchange Server 2010 includes all fixes for Exchange Server from all previous update rollup packages, so you only need to install the latest RU to apply all of the fixes that were released up to that point. However, individual updates or hotfixes for Exchange 2010 or earlier do not contain all previous fixes for Exchange Server. The updated files that are included in an individual update or hotfix include all updates that were applied only to those specific files by all previous updates, but any other files on Exchange Server will not be updated. For more information, see Exchange 2010 Servicing.

In Exchange Server 2013 or later, we changed the way we deliver hotfixes and service packs by using a scheduled delivery model. In this model, cumulative updates (CUs) are released quarterly (every three months). Each CU is a full installation of Exchange that includes updates and changes from all previous CUs, so you don't need to install any previous CUs or Exchange Server RTM first. For more information, see Updates for Exchange Server.

SERVICING RELEASE MODEL	EXCHANGE 2019	EXCHANGE 2016	EXCHANGE 2013	EXCHANGE 2010
Cumulative updates (CUs)	Yes	Yes	Yes	No
Update rollups (RUs)	No	No	No	Yes
Security hotfixes delivered separately	Yes	Yes	Yes	No

NOTE

At this time, no additional CUs are planned for Exchange Server 2013 and no additional RUs are planned for Exchange Server 2010.

Support lifecycle

For more information about the support lifecycle for specific versions of Exchange, Windows Server, or Windows client operating systems, see the Microsoft Support Lifecycle page. For more information about the Microsoft Support Lifecycle, see the Microsoft Support Lifecycle Policy FAQ.

Exchange Server 2007 End-of life

Exchange 2007 reached end of support on April 11, 2017, per the Microsoft Lifecycle Policy. There will be no new security updates, non-security updates, free or paid assisted support options, or online technical content updates.

Furthermore, as adoption of Microsoft 365 or Office 365 accelerates and cloud usage increases, custom support options for Office products will not be available. This includes Exchange Server, as well as Microsoft Office, SharePoint Server, Office Communications Server, Lync Server, Skype for Business Server, Project Server, and Visio. At this time, we encourage customers to complete their migration and upgrade plans. We recommend that customers leverage deployment benefits provided by Microsoft and Microsoft Certified Partners including Microsoft FastTrack for cloud migrations, and Software Assurance Planning Services for on-premises upgrades.

Supported operating system platforms

The following tables identify the operating system platforms on which each version of Exchange can run.

IMPORTANT

Releases of Windows Server and Windows client that aren't listed in the tables below are not supported for use with any version or release of Exchange.

SERVER OPERATING SYSTEM	EXCHANGE 2019	EXCHANGE 2016 CU3 AND LATER	EXCHANGE 2016 CU2 AND EARLIER	EXCHANGE 2013 SP1 AND LATER	EXCHANGE 2010 SP3
Windows Server 2019	Supported	Not supported	Not supported	Not supported	Not supported
Windows Server 2016	Not supported	Supported	Not supported	Not supported	Not supported
Windows Server 2012 R2	Not supported	Supported	Supported	Supported	Supported
Windows Server 2012	Not supported	Supported	Supported	Supported	Supported
Windows Server 2008 R2 SP1	Not supported	Not supported	Not supported	Supported	Supported
Windows Server 2008 SP2	Not supported	Not supported	Not supported	Not supported	Supported

NOTE

Client operating systems only support the Exchange management tools.

CLIENT OPERATING SYSTEM	EXCHANGE 2019	EXCHANGE 2016 CU3 AND LATER	EXCHANGE 2013 SP1 AND LATER	EXCHANGE 2010 SP3
Windows 10	Supported	Supported	Not supported	Not supported
Windows 8.1	Not supported	Supported	Supported	Not supported

CLIENT OPERATING SYSTEM	EXCHANGE 2019	EXCHANGE 2016 CU3 AND LATER	EXCHANGE 2013 SP1 AND LATER	EXCHANGE 2010 SP3
Windows 8	Not supported	Not supported	Supported	Supported

Supported Active Directory environments

The following table identifies the Active Directory environments that Exchange can communicate with. An Active Directory server refers to both writable global catalog servers and to writable domain controllers. Read-only global catalog servers and read-only domain controllers are not supported.

OPERATING SYSTEM ENVIRONMENT	EXCHANGE 2019	EXCHANGE 2016 CU12 AND LATER	EXCHANGE 2016 CU7 AND LATER	EXCHANGE 2016 CU3 TO CU6	EXCHANGE 2016 CU2 AND EARLIER	EXCHANGE 2013 SP1 AND LATER	EXCHANGE 2010 SP3 RU22 OR LATER	EXCHANGE 2010 SP3 RU5 - RU21
Windows Server 2019 Active Directory servers	Supported	Supported	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported
Windows Server 2016 Active Directory servers	Supported	Supported	Supported	Supported	Supported	Supported	Supported	Not supported
Windows Server 2012 R2 Active Directory servers	Supported	Supported	Supported	Supported	Supported	Supported	Supported	Supported
Windows Server 2012 Active Directory servers	Not supported	Supported	Supported	Supported	Supported	Supported	Supported	Supported
Windows Server 2008 R2 SP1 Active Directory servers	Not supported	Supported	Supported	Supported	Supported	Supported	Supported	Supported

OPERATING SYSTEM ENVIRONMENT	EXCHANGE 2019	EXCHANGE 2016 CU12 AND LATER	EXCHANGE 2016 CU7 AND LATER	EXCHANGE 2016 CU3 TO CU6	EXCHANGE 2016 CU2 AND EARLIER	EXCHANGE 2013 SP1 AND LATER	EXCHANGE 2010 SP3 RU22 OR LATER	EXCHANGE 2010 SP3 RU5 - RU21
Windows Server 2008 SP2 Active Directory servers	Not supported	Not supported	Supported	Supported	Supported	Supported	Supported	
Windows Server 2003 SP2 Active Directory servers	Not supported	Not supported	Not supported	Not supported	Supported	Supported	Supported	Supported

AD FOREST FUNCTIONAL LEVEL	EXCHANGE 2019	EXCHANGE 2016 CU7 AND LATER	EXCHANGE 2016 CU3 TO CU6	EXCHANGE 2016 CU2 AND EARLIER	EXCHANGE 2013 SP1 AND LATER	EXCHANGE 2010 SP3 RU22 OR LATER	EXCHANGE 2010 SP3 RU5 - RU21
Windows Server 2016	Supported	Supported	Supported	Supported	Supported	Supported	Not supported
Windows Server 2012 R2	Supported	Supported	Supported	Supported	Supported	Supported	Supported
Windows Server 2012	Not supported	Supported	Supported	Supported	Supported	Supported	Supported
Windows Server 2008 R2	Not supported	Supported	Supported	Supported	Supported	Supported	Supported
Windows Server 2008	Not supported	Not supported	Supported	Supported	Supported	Supported	Supported
Windows Server 2003	Not supported	Not supported	Not supported	Supported	Supported	Supported	Supported

Web browsers supported for use with the premium version of Outlook Web App or Outlook on the web

The following table identifies the web browsers supported for use together with the premium version of Outlook Web App or Outlook on the web.

BROWSER	EXCHANGE 2019	EXCHANGE 2016	EXCHANGE 2013 SP1 AND LATER	EXCHANGE 2010 SP3
Microsoft Edge	Supported	Supported	Not supported	Not supported
Internet Explorer 11	Supported	Supported	Supported	Supported
Internet Explorer 10	Not supported	Not supported	Supported	Supported
Internet Explorer 9	Not supported	Not supported	Supported	Supported
Internet Explorer 8	Not supported	Not supported	Supported	Supported
Internet Explorer 7	Not supported	Not supported	Not supported	Supported
Firefox	Current release of Firefox*	Current release of Firefox*	Not supported	Not supported
Firefox 3.0.1 or later	Not supported	Not supported	Not supported	Supported
Firefox 12 or later	Not supported	Not supported	Supported	Supported
Safari	Current release of Safari	Current release of Safari	Not supported	Not supported
Safari 3.1 or later	Not supported	Not supported	Not supported	Supported
Safari 5.0 or later	Not supported	Not supported	Supported	Supported
Chrome	Current release of Chrome*	Current release of Chrome*	Not supported	Not supported
Chrome 3.0.195 or later	Not supported	Not supported	Not supported	Supported
Chrome 18 or later	Not supported	Not supported	Supported	Supported

* Current release of Firefox or Chrome refers to the latest version or the immediately previous version.

Web browsers supported for use with the basic version of Outlook Web App or Outlook on the web

The following table identifies the web browsers supported for use together with the light (basic) version of Outlook Web App or Outlook on the web.

NOTE

Outlook Web App Basic (Outlook Web App Light) is supported for use in mobile browsers. However, if rendering or authentication issues occur in a mobile browser, determine whether the issue can be reproduced by using Outlook Web App Light in the full client of a supported browser. For example, test the use of Outlook Web App Light in Safari, Chrome, or Internet Explorer. If the issue can't be reproduced in the full client, we recommend that you contact the mobile device vendor for help. In these cases, we collaborate with the vendor as appropriate.

BROWSER	EXCHANGE 2019	EXCHANGE 2016	EXCHANGE 2013	EXCHANGE 2010 SP3
Microsoft Edge	Supported	Supported	Not supported	Not supported
Internet Explorer 11	Supported	Supported	Supported	Supported
Internet Explorer 10	Not supported	Not supported	Supported	Supported
Internet Explorer 9	Not supported	Not supported	Supported	Supported
Internet Explorer 8	Not supported	Not supported	Supported	Supported
Internet Explorer 7	Not supported	Not supported	Supported	Supported
Safari	Current release of Safari	Current release of Safari	Supported	Supported
Firefox	Not supported	Current release of Firefox*	Supported	Supported
Chrome	Not supported	Current release of Chrome*	Not supported	Not supported
Opera	Not supported	Not supported	Supported	Supported

* Current release of Firefox or Chrome refers to the latest version or the immediately previous version.

Web browsers supported for use of S/MIME with Outlook Web App or Outlook on the web

The following table identifies the web browsers supported for the use of S/MIME together with Outlook Web App or Outlook on the web.

BROWSER	EXCHANGE 2019	EXCHANGE 2016	EXCHANGE 2013 SP1 AND LATER	EXCHANGE 2010 SP3
Microsoft Edge	Supported	Not supported	Not supported	Not supported
Internet Explorer 11	Supported	Supported	Supported	Supported
Internet Explorer 10	Not supported	Not supported	Supported	Supported
Internet Explorer 9	Not supported	Not supported	Supported	Supported
Internet Explorer 8	Not supported	Not supported	Not supported	Supported
Internet Explorer 7	Not supported	Not supported	Not supported	Supported

Clients

The following tables identify the mail clients that are supported for use together with each version of Exchange.

CLIENT	EXCHANGE 2019	EXCHANGE 2016	EXCHANGE 2013 SP1 AND LATER	EXCHANGE 2010 SP3
Microsoft 365 Apps for enterprise	Supported	Supported	Supported	Not supported
Outlook 2019	Supported	Supported	Supported	Not supported
Outlook 2016	Supported[1]	Supported[1]	Supported	Supported
Outlook 2013	Supported[1]	Supported[1]	Supported	Supported
Outlook 2010	Not supported	Supported[1]	Supported[2]	Supported
Outlook 2007	Not supported	Not supported	Supported[3]	Supported
Outlook for Mac for Office 365	Not supported	Supported[1]	Supported	Supported
Entourage 2008 (EWS)	Not supported	Supported[4]	Supported[4]	Supported[4]

[1] Requires the latest Office service pack and the latest public update.

[2] Requires Outlook 2010 Service Pack 1 and the latest public update.

[3] Requires Outlook 2007 Service Pack 3 and the latest public update.

[4] EWS only. There is no DAV support for Exchange 2010.

Microsoft .NET Framework

The following tables identify the versions of the Microsoft .NET Framework that can be used with the specified versions of Exchange.

IMPORTANT

Versions of the .NET Framework that aren't listed in the tables below are not supported on any version of Exchange. This includes minor and patch-level releases of the .NET Framework.

If you are upgrading Exchange Server from an unsupported CU to the current CU and no intermediate CUs are available, you should first upgrade to the latest version of .NET that's supported by your version of Exchange Server and then immediately upgrade to the current CU. This method doesn't replace the need to keep your Exchange servers up to date and on the latest supported CU. Microsoft makes no claim that an upgrade failure will not occur using this method, which may result in the need to contact Microsoft Support Services.

Exchange 2019

.NET FRAMEWORK VERSION	CU6, CU5, CU4	CU3, CU2	CU1, RTM
4.8	Supported	Supported	Not supported
4.7.2	Not supported	Supported	Supported

Exchange 2016

.NET FRAMEWORK VERSION	CU17, CU16, CU15	CU14, CU13	CU12, CU11	CU10	CU9, CU8	CU7, CU6, CU5	CU4, CU3	CU2	CU1, RTM
4.8	Supported	Supported	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported
4.7.2	Not supported	Supported	Supported	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported
4.7.1	Not supported	Not supported	Supported	Supported	Supported	Not supported	Not supported	Not supported	Not supported
4.6.2	Not supported	Not supported	Not supported	Not supported	Supported	Supported	Supported	Not supported	Not supported
4.6.1*	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported	Supported	Supported	Not supported
4.5.2	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported	Supported	Supported	Supported

* .NET Framework 4.6.1 also requires a hotfix, and a different hotfix is required for different versions of Windows. For more information, see Released: June 2016 Quarterly Exchange Updates.

Exchange 2013

.NET FRAMEWORK VERSION	CU23	CU21, CU22	CU19, CU20	CU16, CU17, CU18	CU15	CU13, CU14	CU12 TO SP1	CU3 TO RTM
4.8	Supported	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported
4.7.2	Supported	Supported	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported
4.7.1	Not supported	Supported	Supported	Not supported	Not supported	Not supported	Not supported	Not supported
4.6.2	Not supported	Not supported	Supported	Supported	Supported	Not supported	Not supported	Not supported
4.6.1*	Not supported	Not supported	Not supported	Not supported	Supported	Supported	Not supported	Not supported
4.5.2	Not supported	Not supported	Not supported	Not supported	Supported	Supported	Supported	Not supported

.NET FRAMEWORK VERSION	CU23	CU21, CU22	CU19, CU20	CU16, CU17, CU18	CU15	CU13, CU14	CU12 TO SP1	CU3 TO RTM
4.5.1	Not supported	Not supported	Not supported	Not supported	Supported	Supported	Supported	Not supported
4.5	Not supported	Not supported	Not supported	Not supported	Not supported	Supported	Supported	Supported

* .NET Framework 4.6.1 also requires a hotfix, and a different hotfix is required for different versions of Windows. For more information, see Released: June 2016 Quarterly Exchange Updates.

Exchange 2010 SP3

.NET FRAMEWORK VERSION	EXCHANGE 2010 SP3
.NET Framework 4.5	Supported[1,2]
.NET Framework 4.0	Supported[1,2]
.NET Framework 3.5 SP1	Supported
.NET Framework 3.5	Supported[1]

[1] On Windows Server 2012, you need to install the .NET Framework 3.5 before you can use Exchange 2010 SP3.

[2] Exchange 2010 uses only the .NET Framework 3.5 and the .NET Framework 3.5 SP1 libraries. It doesn't use the .NET Framework 4.5 libraries if they're installed on the server. We support the installation of any version of the .NET Framework 4.5 (for example, .NET Framework 4.5.1, .NET Framework 4.5.2, etc.) as long as the .NET Framework 3.5 or the .NET Framework 3.5 SP1 is also installed on the server.

Windows PowerShell

- Exchange 2013 or later requires the version of Windows PowerShell that's included in Windows (unless otherwise specified by an Exchange Setup-enforced prerequisite rule).
- Exchange 2010 requires Windows PowerShell 2.0 on all supported versions of Windows.
- Exchange does not support the use of Windows Management Framework add-ons on any version of Windows PowerShell or Windows.
- If there are other installed versions of Windows PowerShell or PowerShell Core that support side-by-side operation, Exchange will use only the version that it requires.

Microsoft Management Console

The following table identifies the version of Microsoft Management Console (MMC) that can be used together with each version of Exchange.

MMC	EXCHANGE 2016	EXCHANGE 2013 SP1 AND LATER	EXCHANGE 2010 SP3
MMC 3.0	Supported	Supported	Supported

Windows Installer

The following table identifies the version of Windows Installer that is used together with each version of Exchange.

WINDOWS INSTALLER	EXCHANGE 2016	EXCHANGE 2013 SP1 AND LATER	EXCHANGE 2010 SP3
Windows Installer 4.5	Supported	Supported	Supported
Windows Installer 5.0	Supported	Supported	Not supported

Exchange Server virtualization

8/3/2020 • 10 minutes to read • Edit Online

You can deploy Exchange Server 2016 and Exchange Server 2019 in a virtualized environment. This topic provides an overview of the scenarios that are supported for deploying Exchange on hardware virtualization software.

The following terms are used in this discussion of Exchange virtualization:

- **Cold boot**: When bringing a system from a power-off state into a clean start of the operating system, the action is a *cold boot*. No operating system state has been persisted in this case.

- **Saved state**: When a virtual machine is powered off, hypervisors typically have the ability to save the state of the virtual machine, so when the machine is powered back on, it returns to that *saved state* rather than going through a cold boot startup.

- **Planned migration**: When a system administrator initiates the move of a virtual machine from one hypervisor host to another, the action is a *planned migration*. The action could be a single migration, or a system administrator could configure automation to move the virtual machine on a timed basis. A planned migration could also be the result of some other event that occurs in the system, other than hardware or software failure.

  The key point of a planned migration is the Exchange virtual machine is operating normally and needs to be relocated for some reason. This relocation can be done via technology (for example, Live Migration or vMotion). However, if the Exchange virtual machine or the hypervisor host where the virtual machine is located experiences some sort of failure condition, the outcome isn't characterized as a planned migration.

Requirements for hardware virtualization

Microsoft supports Exchange 2016 and Exchange 2019 in production on hardware virtualization software only when all the following conditions are true:

- The hardware virtualization software is running one of the following:

  - Any version of Windows Server with Hyper-V technology or Microsoft Hyper-V Server

  - Any third-party hypervisor that has been validated under the Windows Server Virtualization Validation Program.

  > **NOTE**
  > Deployment of Exchange 2016 or Exchange 2019 on Infrastructure-as-a-Service (IaaS) providers is supported if all supportability requirements are met. In the case of providers who are provisioning virtual machines, these requirements include ensuring that the hypervisor being used for Exchange virtual machines is fully supported, and that the infrastructure to be utilized by Exchange meets the performance requirements that were determined during the sizing process. Deployment on Microsoft Azure virtual machines is supported if all storage volumes used for Exchange databases and database transaction logs (including transport databases) are configured for Azure Premium Storage.

- The Exchange guest virtual machine has the following conditions:

  - It's running Exchange 2016 or Exchange 2019.

  - It's deployed on a supported version of Windows Server for Exchange.

For deployments of Exchange 2016 or Exchange 2019:

- All Exchange server roles are supported in a virtual machine.

- Exchange server virtual machines (including Exchange virtual machines that are part of a database availability group, or DAG), may be combined with host-based failover clustering and migration technology, as long as the virtual machines are configured such that they won't save and restore state on disk when moved or taken offline. All failover activity occurring at the hypervisor level must result in a cold boot when the virtual machine is activated on the target node. All planned migration must either result in shutdown and cold boot, or an online migration that makes use of a technology like Hyper-V Live Migration. Hypervisor migration of virtual machines is supported by the hypervisor vendor; therefore, you must ensure that your hypervisor vendor has tested and supports migration of Exchange virtual machines. Microsoft supports Hyper-V Live Migration of these virtual machines.

- Only management software (for example, antivirus software, backup software, or virtual machine management software) can be deployed on the physical host machine. No other server-based applications (for example, Exchange, SQL Server, Active Directory, or SAP) should be installed on the host machine. The host machine should be dedicated to running guest virtual machines.

- Some hypervisors include features for taking snapshots of virtual machines. Virtual machine snapshots capture the state of a virtual machine while it's running. This feature enables you to take multiple snapshots of a virtual machine and then revert the virtual machine to any of the previous states by applying a snapshot to the virtual machine. However, virtual machine snapshots aren't application aware, and using them can have unintended and unexpected consequences for a server application that maintains state data, such as Exchange. As a result, making virtual machine snapshots of an Exchange guest virtual machine isn't supported.

- Many hardware virtualization products allow you to specify the number of virtual processors that should be allocated to each guest virtual machine. The virtual processors located in the guest virtual machine share a fixed number of physical processor cores in the physical system. Exchange supports a virtual processor-to-physical processor core ratio no greater than 2:1, although we recommend a ratio of 1:1. For example, a dual processor system using quad core processors contains a total of 8 physical processor cores in the host system. On a system with this configuration, don't allocate more than a total of 16 virtual processors to all guest virtual machines combined.

- When calculating the total number of virtual processors required by the host machine, you must also account for both I/O and operating system requirements. In most cases, the equivalent number of virtual processors required in the host operating system for a system hosting Exchange virtual machines is 2. This value should be used as a baseline for the host operating system virtual processor when calculating the overall ratio of physical cores to virtual processors. If performance monitoring of the host operating system indicates you're consuming more processor utilization than the equivalent of 2 processors, you should reduce the count of virtual processors assigned to guest virtual machines accordingly and verify that the overall virtual processor-to-physical core ratio is no greater than 2:1.

- It's possible that guest virtual machines may be prevented from directly communicating with Fibre Channel or SCSI host bus adapters (HBAs) installed in the host machine. In this event, you must configure the adapters in the host machine's operating system and present the logical unit numbers (LUNs) to guest virtual machines as either a virtual disk or a pass-through disk.

- The only supported way to send emails to external domains from Azure compute resources is via an SMTP relay (also known as an SMTP smart host). The Azure compute resource sends the email to the SMTP relay and then the SMTP relay provider delivers the email to the external domain. Microsoft Exchange Online Protection is one provider of an SMTP relay, but there are a number of third-party providers as well. For more information, see Troubleshoot outbound SMTP connectivity issues in Azure.

Host machine storage requirements

The minimum disk space requirements for each host machine are described in the following list:

- Host machines in some hardware virtualization applications may require storage space for an operating system and its components. Additional storage space is also required to support the operating system's paging file, management software, and crash recovery (dump) files.
- Some hypervisors maintain files on the host machine that are unique to each guest virtual machine. For example, in a Hyper-V environment, a temporary memory storage file (BIN file) is created and maintained for each guest machine. The size of each BIN file is equal to the amount of memory allocated to the guest machine. In addition, other files may also be created and maintained on the host machine for each guest machine.
- If your host machine is running Windows Server 2012 Hyper-V or Hyper-V 2012, and you're configuring a host-based failover cluster that will host Exchange Mailbox servers in a DAG, we recommend following the guidance in KB2872325.

Exchange storage requirements

Requirements for storage connected to a virtualized Exchange server are as follows:

- Each Exchange guest machine must be allocated sufficient storage space on the host machine for the fixed disk that contains the guest's operating system, any temporary memory storage files in use, and related virtual machine files that are hosted on the host machine. In addition, for each Exchange guest machine, you must also allocate sufficient storage for the message queues and sufficient storage for the databases and log files on Mailbox servers.
- The storage used by the Exchange guest machine for storage of Exchange data (for example, mailbox databases and transport queues) can be virtual storage of a fixed size (for example, fixed virtual hard disks (VHD or VHDX) in a Hyper-V environment), dynamic virtual storage when using VHDX files with Hyper-V, SCSI pass-through storage, or Internet SCSI (iSCSI) storage. Pass-through storage is storage that's configured at the host level and dedicated to one guest machine. All storage used by an Exchange guest machine for storage of Exchange data must be block-level storage because Exchange doesn't support the use of network attached storage (NAS) volumes, other than in the SMB 3.0 scenario outlined later in this topic. Also, NAS storage that's presented to the guest as block-level storage via the hypervisor isn't supported.
- Fixed VHDs may be stored on SMB 3.0 files that are backed by block-level storage if the guest machine is running on Windows Server 2012 Hyper-V (or a later version of Hyper-V). The only supported usage of SMB 3.0 file shares is for storage of fixed VHDs. Such file shares can't be used for direct storage of Exchange data. When using SMB 3.0 file shares to store fixed VHDs, the storage backing the file share should be configured for high availability to ensure the best possible availability of the Exchange service.
- Storage used by Exchange should be hosted in disk spindles that are separate from the storage that's hosting the guest virtual machine's operating system.
- Configuring iSCSI storage to use an iSCSI initiator inside an Exchange guest virtual machine is supported. However, there is reduced performance in this configuration if the network stack inside a virtual machine isn't full-featured (for example, not all virtual network stacks support jumbo frames).

Exchange memory requirements and recommendations

Some hypervisors have the ability to oversubscribe/overcommit or dynamically adjust the amount of memory available to a specific guest machine based on the perceived usage of memory in the guest machine as compared to the needs of other guest machines managed by the same hypervisor. This technology makes sense for workloads in which memory is needed for brief periods of time and then can be surrendered for other uses. However, it doesn't make sense for workloads that are designed to use memory on an ongoing basis. Exchange (like many server applications with optimizations for performance that involve caching of data in memory) is susceptible to poor system performance and an unacceptable client experience if it doesn't have full control over the memory allocated to the physical or virtual machine on which it's running. As a result, using dynamic memory or memory overcommit features for Exchange isn't supported.

Host-based failover clustering and migration for Exchange

The following are answers to some frequently asked questions about host-based failover clustering and migration technology with Exchange DAGs:

- **Does Microsoft support third-party migration technology?**

  Microsoft can't make support statements for the integration of third-party hypervisor products using these technologies with Exchange, because these technologies aren't part of the Server Virtualization Validation Program (SVVP). The SVVP covers the other aspects of Microsoft support for third-party hypervisors. You need to ensure that your hypervisor vendor supports the combination of their migration and clustering technology with Exchange. If your hypervisor vendor supports their migration technology with Exchange, Microsoft supports Exchange with their migration technology.

- **How does Microsoft define host-based failover clustering?**

  Host-based failover clustering refers to any technology that provides the automatic ability to react to host-level failures and start affected virtual machines on alternate servers. Use of this technology is supported given that, in a failure scenario, the virtual machine is coming up from a cold boot on the alternate host. This technology helps to make sure that the virtual machine never comes up from a saved state that's persisted on disk because it will be stale relative to the rest of the DAG members.

- **What does Microsoft mean by migration support?**

  Migration technology refers to any technology that allows a planned move of a virtual machine from one host machine to another host machine. This move could also be an automated move that occurs as part of resource load balancing, but it isn't related to a failure in the system. Migrations are supported as long as the virtual machines never come up from a saved state that's persisted on disk. This means that technology that moves a virtual machine by transporting the state and virtual machine memory over the network with no perceived downtime is supported for use with Exchange. A third-party hypervisor vendor must provide support for the migration technology, while Microsoft provides support for Exchange when used in this configuration.

Plan Exchange 2016 integration with SharePoint and Skype for Business

8/3/2020 • 10 minutes to read • Edit Online

Exchange 2016 integration with SharePoint Server 2016 and Skype for Business allow for services that provide the ability to preserve, archive, and then quickly search email, documents, and other content. Together, these enterprise applications make possible scenarios such as eDiscovery and collaboration using site mailboxes to let your organization preserve important data. Critical in most organizations these days is the ability to archive and then locate email and documents as required to meet compliance and regulatory requirements. You can use Exchange 2016 along with SharePoint 2016 and Skype for Business to:

- Archive Exchange mailboxes
- Archive Skype for Business content
- Preserve SharePoint Server 2016 documents and websites
- Search across stores using eDiscovery
- Authenticate seamlessly across servers

The eDiscovery Center introduced in SharePoint 2013 provides content identification, preservation, collection, processing, and analysis. In an Exchange environment, eDiscovery lets you archive content discovered across SharePoint Server 2016, Skype for Business, andExchange. You can use the eDiscovery Center to create eDiscovery Case sites that are used to organize in-place holds, queries, and exports for a specific case.

Exchange 2016, SharePoint Server 2016, and Skype for Business Server use the standard protocol, Open Authorization (OAuth), for server-to-server authentication to provide the cross-product functionality described here. Using the same protocol allows these applications to seamlessly and securely authenticate to each other. The authorization method supports authentication as an application by means of a linked account and user impersonation where the access request is made in the user context. You can learn more about OAuth later in this article in the Server-to-server authentication using OAuth section.

> **NOTE**
>
> For enterprises that use Lync Server 2013, you can still make full use of the features described in this topic.

Archive Skype for Business content in Exchange 2016

With Exchange 2016 and Lync Server 2013 deployed in an organization, you can configure Skype for Business to archive instant message and on-line meeting content, including shared presentations or documents in the user's Exchange 2016 mailbox. Archiving Skype for Business data in Exchange 2016 allows you to apply retention policies to the data. Archived Skype for Business content also surfaces in any eDiscovery searches. For more details about Skype for Business archiving and how to deploy it, see the following topics:

- Planning for Archiving
- Deploying Archiving

Preserve documents in SharePoint Server 2016

You can create a query-based hold to preserve items that meet your specified criteria with an In-Place Hold.

For example, Litigation Hold preserves until the hold is removed any deleted items as well as original versions of modified items . You can optionally specify a hold duration that preserves a mailbox item for the named duration period. If you specify a hold duration period, it's calculated from the date a message is received or a mailbox item is created. For details, see Create or remove an In-Place Hold.

For more details on eDiscovery see the following topics:

- In-Place eDiscovery in Exchange Server
- In-Place Hold and Litigation Hold in Exchange Server
- Configure eDiscovery in SharePoint Server
- What's new in eDiscovery in SharePoint Server
- Configure Exchange for SharePoint eDiscovery Center

Search across applications by using eDiscovery

SharePoint Server 2016 provides the eDiscovery Center to help you locate and then transfer relevant content as needed to meet regulatory requirements. eDiscovery is the process of finding, preserving, analyzing, and producing content in digital format required by litigation or investigations. You can use eDiscovery across Exchange 2016, SharePoint Server 2016, and Skype for Business files. You can help protect content in-place that you've identified with eDiscovery queries and then export the results into an offline format to hand off for legal review. In-Place Hold in eDiscovery lets you:

- Protect content in-place and in real time at reduced storage costs, without affecting your users' daily work.
- Query to collect up-to-date, relevant content and statistics quickly answer questions.
- Export relevant content in an offline and portable format.

If your organization adheres to legal discovery requirements, that is, anything related to organizational policy, compliance, or lawsuits, In-Place eDiscovery in Exchange Server 2016 can help you perform discovery searches for relevant content within mailboxes. You can also use In-Place eDiscovery in an Exchange hybrid environment to search on-premises and cloud-based mailboxes in the same search.

When you configure server-to-server authentication betweenExchange 2016 and SharePoint Server 2016 in on-premises deployments, administrators and compliance officers can use the eDiscovery Center. For more information, see Configure Exchange for SharePoint eDiscovery Center. In hybrid deployments, for more information see Using Oauth Authentication to Support eDiscovery in an Exchange Hybrid Deployment

You can identify and reduce your data set by using keyword syntax, property restrictions, and refinements. The query experience focuses on statistics for individual sources and query fragments to help you make decisions about the content you're searching across. You can also preview SharePoint 2016 and Exchange 2016 content to confirm that you have identified the right set of results.

Server-to-server authentication using OAuth

The OAuth protocol is used by many web sites and web services to let clients access resources without having to provide a username and password. An authorization server trusted by the resource owner provides the client with an access token that grants access to a specific set of resources for a specified period. Exchange 2016 allows other applications to use OAuth to authenticate to Exchange. You'll need to configure the applications in Exchange as partner applications.

There are two configuration objects used for OAuth andExchange 2016 partner applications: AuthConfig and the partner application configuration.

- **AuthConfig**: Exchange 2016 Setup creates AuthConfig to publish the auth metadata. You only need to manage AuthConfig to provision a new certificate when the existing certificate is close to expiration. When this happens, you can renew the existing certificate and configure the new certificate as the next certificate in the AuthConfig along with its effective date.

  Exchange 2016 Setup creates a self-signed certificate with the friendly name Microsoft Exchange Server Auth Certificate and replicates the certificate to all front-end servers in the Exchange organization. The certificate's thumbprint is specified in the authorization configuration for Exchange 2016, along with its service name, which is a well-known GUID that represents on-premises Exchange 2016. Exchange uses the authorization configuration to publish its auth metadata document.

- **Partner applications**: You enable partner applications by creating a partner application configuration to request access tokens from Exchange. Exchange 2016 provides the `Configure-EnterprisePartnerApplication.ps1` script that lets you quickly and easily create partner application configurations and minimize configuration errors.

  When Exchange 2016 receives an access request from a partner application via Exchange Web Services (EWS), the following events take place.

  - EWS parses the `www-authenticate` header of the https request that contains the access token signed by the calling server using its private key.

  - The auth module validates the access token using the partner application configuration.

  - The module then grants access to resources based on the RBAC permissions granted to the application. If the access token is on behalf of a user, the RBAC permissions granted to the user are checked.

    For example, if a user performs an eDiscovery search using the eDiscovery Center in SharePoint 2016, Exchange checks whether the user is a member of the Discovery Management role group or has the Mailbox Search role assigned and the mailboxes being searched are within the scope of the RBAC role assignment. For more details, see Permissions.

In on-premises deployments, Exchange 2016, SharePoint Server 2016, and Skype for Business Server 2015 do not require an authorization server to issue tokens. Each application issues self-signed tokens to access the resources provided by other applications. The application that provides access to resources, for example Exchange 2016, trusts the self-signed tokens presented by the calling application. Trust is established by creating a *partner application* configuration for the calling application, which includes the calling application's ApplicationID, certificate, and AuthMetadataUrl. Exchange 2016, SharePoint 2016, and Skype for Business publish their auth metadata document in a well-known URL.

Auth metadata URLs

SERVER	AUTHMETADATAURL
Exchange 2016	`https://<serverfqdn>/autodiscover/metadata/json/1`
SharePoint Server 2016	`https://<serverfqdn>/_layouts/15/metadata/json/1`
Skype for Business	`https://<serverfqdn>/metadata/json/1`

In hybrid deployments, you need to configure OAuth authorization protocol between your on-premises Exchange 2016 and Exchange Online organizations. Hybrid deployments by default continue to use the federation trust process.

Certain Exchange 2016 features are only fully available across your organization by using the new OAuth protocol. For example, before you can use In-Place eDiscovery to search on-premises and cloud-based mailboxes in an Exchange hybrid organization, you need to configure OAuth authentication between your Exchange on-premises and Exchange Online organizations. The Hybrid Configuration Wizard doesn't manage the OAuth authorization connection. For more information, see Configure OAuth Authentication Between Exchange and Exchange Online Organizations.

In online deployments, Exchange Online, SharePoint Online and Skype for Business Online need to be configured for a modern authentication connection. Modern authentication brings Active Directory Authentication Library (ADAL)-based sign in to Office 2013 Windows clients. Office 2013 client applications sign in to the Microsoft 365 or Office 365 service to gain access to Exchange Online, SharePoint Online, and Skype for Business Online. We recommend that you enable Exchange Online for modern authentication when enabling modern authentication for Skype for Business. Modern authentication is enabled by default in SharePoint Online. For more information, see Enable or disable modern authentication for Outlook in Exchange Online.

The per service default state of modern authentication is:

- Skype for Business Online - OFF by default
- Skype for Business Online - OFF by default
- SharePoint Online - ON by default.

IMPORTANT

The default Server Auth Certificate created by Exchange 2016 is valid for five years. You need to make sure that the authorization configuration includes a current certificate.

Manage SharePoint site mailboxes

In many organizations, information resides in two different stores: email in Exchange and documents in SharePoint. There are two different interfaces to access these stores. This makes for a disjointed user experience that impedes effective collaboration. Site mailboxes in SharePoint let users collaborate effectively by bringing together Exchange emails and SharePoint documents. For users, a site mailbox serves as a central filing cabinet, providing a place to file project emails and documents that can only be accessed and edited by site members. Site mailboxes are visible in Outlook 2016 to give users easy access to the email and documents for the projects they care about. Additionally, the same set of content can be accessed directly from the SharePoint site itself.

In a site mailbox, content is kept where it belongs. Exchange stores the email, providing users with the same message view for email conversations that they use every day for their own mailboxes. SharePoint stores the documents, which allows for document coauthoring and versioning. Exchange synchronizes just enough metadata from SharePoint to create the document view in Outlook (that is, document title, last modified date, last modified author, and size).

You can provision and manage site mailboxes from SharePoint Server 2016. For more information, including how to configure site mailboxes, see the following topics.

- Site Mailboxes
- Configure email integration for a SharePoint Server farm

Manage access to unified contact store

The unified contact store (UCS) feature provides a consistent contact experience across Office products. This feature lets users store all contact information in their Exchange 2016 mailbox so that the same contact information is available globally across Skype for Business, SharePoint, Exchange, Outlook and Outlook on the web. When you deploy aSkype for Business Server and publish the topology, UCS is enabled for all users by default and no additional action is needed. For more information, see Configure Skype for Business Server to use the unified contact store.

A user's contacts are automatically migrated to the Exchange 2016 server when the user:

- Is assigned a user services policy that has UcsAllowed set to True.

- Was provisioned with an Exchange 2016 mailbox and has signed into the mailbox at least once.

- Logs in by using a Skype for Business rich client.

After you have installed SharePoint Server 2016 in an environment with Exchange 2016 and you have configured server-to-server authentication between the two, users can initiate the migration of existing contacts from SharePoint 2016 or Skype for Business Server 2015 to Exchange 2016. For details, see Planning and Deploying Unified Contact Store.

Manage access to high-resolution user photos

The user photos feature lets you store high resolution user photos in Exchange 2016 that can be accessed by client applications, including Outlook, Outlook on the web, SharePoint 2016, Skype for Business, and mobile email clients. A low-resolution photo is also stored in Active Directory. The cmdlet *Set-UserPhoto* stores a copy of a high resolution image in the user's Exchange mailbox, and stores a 64×64 pixel copy of the photo as an image in the Active Directory attribute thumbnailPhoto.

As with UCS, user photos allow your organization to maintain a consistent user profile photo that can be consumed by client applications without requiring each application to have its own user photos and different ways to add and manage them. Users can manage their own photos by using Outlook on the web, SharePoint 2016 or Skype for Business. For detail about managing photos on Outlook on the web, see Change your photo and account information in Outlook on the web.

Configure OAuth authentication with SharePoint 2013 and Lync 2013

8/3/2020 • 2 minutes to read • Edit Online

Exchange 2016 supports partner applications such as SharePoint Server 2016 and Skype for Business Server 2015 by using OAuth configuration with the script, `Configure-EnterpriseApplication.ps1`. You can automate the task using the script to more easily configure authentication with partner applications and reduce configuration errors. The script performs the following tasks:

1. Configures an Enterprise partner application that self-issues OAuth tokens to successfully authenticate to Exchange.

2. Assigns Role Based Access Control (RBAC) roles to the partner application to authorize it for calling specific Exchange Web Services APIs.

What do you need to know before you begin?

- Estimated time to complete: 5 minutes.

- The partner application needs to publish an authentication metadata document for Exchange 2016 to establish a direct trust to this application and accept authentication requests.

- Examples in this topic use the following default location of the `\Scripts` directory: `C:\Program Files\Microsoft\Exchange Server\V15\Scripts`.

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Partner applications - configure" entry in the Sharing and collaboration permissions topic.

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Configure OAuth authentication with a partner application

This procedure uses the `Configure-EnterpriseApplication.ps1` script to configure OAuth authentication with partner applications. Access to resources depends on the permissions assigned to the partner application and/or the user it impersonates by using RBAC. After configuring OAuth authentication from Exchange, the partner application can use Exchange 2016 resources.

1. This example configures OAuth authentication for SharePoint 2016.

```
Cd C:\Program Files\Microsoft\Exchange Server\V15\Scripts
Configure-EnterprisePartnerApplication.ps1 -AuthMetaDataUrl
https://sharepoint.contoso.com/_layouts/15/metadata/json/1 -ApplicationType SharePoint
```

2. This example configures OAuth authentication for Skype for Business or Lync Server 2013.

   ```
   Cd C:\Program Files\Microsoft\Exchange Server\V15\Scripts
   Configure-EnterprisePartnerApplication.ps1 -AuthMetaDataUrl https://lync.contoso.com/metadata/json/1 -
   ApplicationType Lync
   ```

If Exchange 2016 also needs to access resources offered by the partner application, you must also configure OAuth authentication in the partner application.

How do you know this worked?

To verify that you have successfully configured an enterprise partner application to authenticate to Exchange 2016 , run the Get-PartnerApplication cmdlet in the Exchange Management Shell to retrieve the configuration. You can also run the Test-OAuthConnectivity cmdlet to test OAuth connectivity with a partner application for a user.

Hybrid and on-premises deployments

- In hybrid deployments, you can use OAuth authentication between your on-premises Exchange 2016 organization and the Exchange Online organization. For more information, see Using Oauth Authentication to Support eDiscovery in an Exchange Hybrid Deployment.

- In on-premises deployments, you can configure server-to-server authentication between Exchange 2016 and SharePoint 2016 so administrators and compliance officers can search Exchange 2016 by using the SharePoint 2016 eDiscovery Center.. For more information, see Configure Exchange for SharePoint eDiscovery Center.

Exchange Server post-installation tasks

8/3/2020 • 2 minutes to read • Edit Online

Read the following topics to help you configure your new Exchange 2016 or Exchange 2016 organization.

TOPIC	DESCRIPTION
Enter your Exchange product key	Learn how to license your Exchange server.
Configure mail flow and client access on Exchange servers	Learn how to configure mail flow to and from the Internet and configure Exchange to accept client connections from the Internet.
Verify Exchange Server installations	Learn how to verify that Exchange 2016 was installed successfully in your organization.
Install the Exchange management tools	Learn how to install the Exchange Management Shell and Exchange Toolbox on client workstations or other non-Exchange servers in your organization.
Configure instant messaging integration with Outlook on the web in Exchange	Learn how to configure instant messaging (IM) integration between Skype for Business Server and Outlook on the web (formerly known as Outlook Web App)
Change the offline address book generation schedule in Exchange	Learn how to change the offline address book (OAB) generation schedule on specific Exchange servers or for the whole organization
Configure certificate based authentication in Exchange 2016	Learn how to configure CBA in Exchange 2016 CU1 or later
Edge Subscriptions	Learn how to configure an EdgeSync Subscription between a new Edge Transport server in the perimeter network and the Exchange Mailbox servers in an internal Active Directory site.

Note:

If you've enabled the Scripting Agent in your Exchange organization, and you keep a customized %ExchangeInstallPath%Bin\CmdletExtensionAgents\ScriptingAgentConfig.xml file on all of your Mailbox servers, you need to copy that file to every new Mailbox server that you deploy in your organization (the file isn't used on Edge Transport servers).

- The default value of %ExchangeInstallationPath% is %ProgramFiles%\Microsoft\Exchange Server\V15, but the actual value is wherever you installed Exchange on the server.
- The default name of the file on a new Exchange server is %ExchangeInstallPath%Bin\CmdletExtensionAgents\ScriptingAgentConfig.xml.sample. As part of enabling the Scripting Agent in your organization, you need to rename this file to ScriptingAgentConfig.xml and customize it or replace it with your existing ScriptingAgentConfig.xml file.

For more information about the Scripting Agent in Exchange 2013 (which still applies to Exchange 2016 and 2019), see Scripting Agent.

Enter your Exchange Server product key

8/3/2020 • 4 minutes to read • Edit Online

A product key tells Exchange Server 2016 or Exchange Server 2019 that you've purchased a Standard or Enterprise Edition license. If the product key you purchased is for an Enterprise Edition license, it lets you mount more than five databases per server in addition to everything that's available with a Standard Edition license. If you want to read more about Exchange licensing, see Exchange Server editions and versions.

If you don't enter a product key, your server is automatically licensed as a trial edition. The trial edition functions just like an Exchange Standard Edition server and is helpful if you want to try out Exchange before you buy it, or to run tests in a lab. The only difference is that you can only use an Exchange server licensed as a trial edition for up to 180 days. If you want to keep using the server beyond 180 days, you'll need to enter a product key or the Exchange admin center (EAC) will start to show reminders that you need to enter a product key to license the server.

Note: If you want to install or activate Office, check out:

- Install Office
- Need help with your Office product key?

If you want to enter a product key on an older version of Exchange, check out Enter an Exchange 2010 product key.

If you want to enter a product key on an Exchange 2016 or Exchange 2019 server, you're in the right place! Read on.

What do you need to know before you begin?

- Estimated time to complete this procedure: less than 5 minutes.
- To open the EAC, see Exchange admin center in Exchange Server. To open the Exchange Management Shell, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Product key" entry in the Exchange infrastructure and PowerShell permissions topic.
- After you license an Exchange Mailbox server, you need to restart the Microsoft Exchange Information Store service on the server after you enter the product key.
- You can upgrade from a Standard Edition license to an Enterprise Edition license. You can't downgrade from an Enterprise Edition license to a Standard Edition license without reinstalling Exchange.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to enter the product key

1. In the EAC. go to **Servers** > **Servers**, select the server you want to license, and then do either of the following steps:

   - Click **Edit** ✎.

   - In the details pane, click **Enter Product Key**. Note that this link is only available for unlicensed servers.

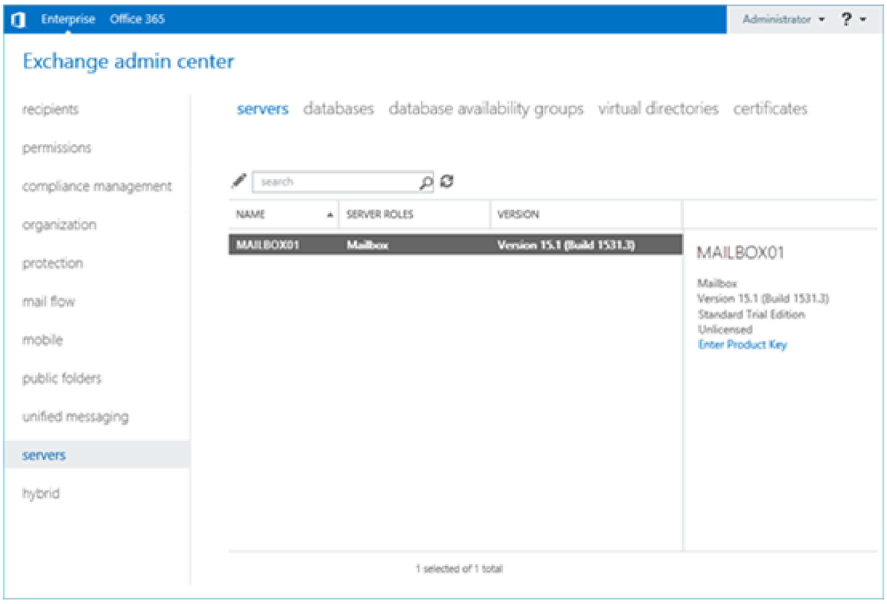

NAME	SERVER ROLES	VERSION
MAILBOX01	Mailbox	Version 15.1 (Build 1531.3)

MAILBOX01

Mailbox
Version 15.1 (Build 1531.3)
Standard Trial Edition
Unlicensed
Enter Product Key

2. The Exchange server properties window opens. On the **General** tab, do one of the following steps:

   - **License an unlicensed server**: Enter the product key in the **Enter a valid product key** text boxes.

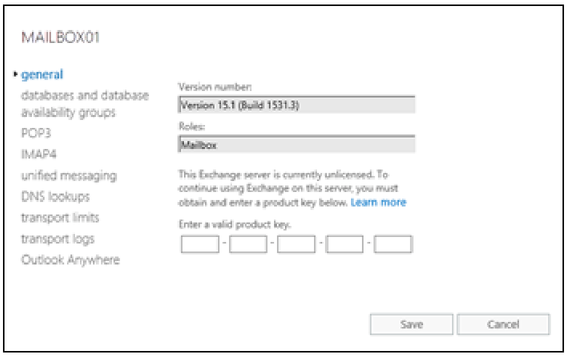

Version number:
Version 15.1 (Build 1531.3)
Roles:
Mailbox

This Exchange server is currently unlicensed. To continue using Exchange on this server, you must obtain and enter a product key below. Learn more

Enter a valid product key.
□ - □ - □ - □ - □

   - **Change or upgrade the product key on a licensed server**: Select **Change product key** and enter the product key in the **Enter a valid product key** text boxes. Note that you'll only see **Change product key** if the server is already licensed.

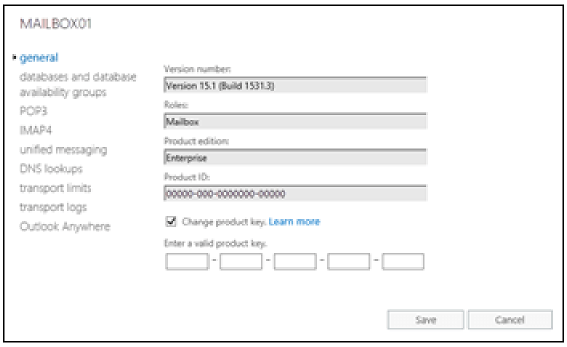

Version number:
Version 15.1 (Build 1531.3)
Roles:
Mailbox
Product edition:
Enterprise
Product ID:
00000-000-0000000-00000

☑ Change product key. Learn more

Enter a valid product key.
□ - □ - □ - □ - □

   When you're finished, click **Save**.

After you license a Mailbox server, do the following steps to restart the Microsoft Exchange Information Store service:

1. On the Exchange server, open the Windows Services console. For example:
   - Run the command `services.msc` from the **Run** dialog, a Command Prompt window, or the Exchange Management Shell.
   - Open Server Manager, and then click **Tools** > **Services**.
2. In the list of services, right-click on **Microsoft Exchange Information Store**, and then click **Restart**.

Use the Exchange Management Shell to enter the product key

To enter the product key in the Exchange Management Shell, use this syntax:

```
Set-ExchangeServer <ServerName> -ProductKey <ProductKey>
```

Note that this command works to license an unlicensed server or to upgrade a licensed server from a Standard Edition license to an Enterprise Edition license.

This example licenses the Exchange server named Mailbox01.

```
Set-ExchangeServer Mailbox01 -ProductKey 12345-12345-12345-12345-12345
```

For detailed syntax and parameter information, see Set-ExchangeServer.

After you license a Mailbox server, run the following command in the Exchange Management Shell to restart the Microsoft Exchange Information Store service:

```
Restart-Service MSExchangeIS
```

How do you know this worked?

To verify that you've successfully licensed the Exchange server, do any of the following steps:

- In the EAC, go to **Servers** > **Servers**, and select the server you licensed. In the details pane, verify the Exchange edition value (**Standard** or **Enterprise**) and whether the value **Licensed** is present.

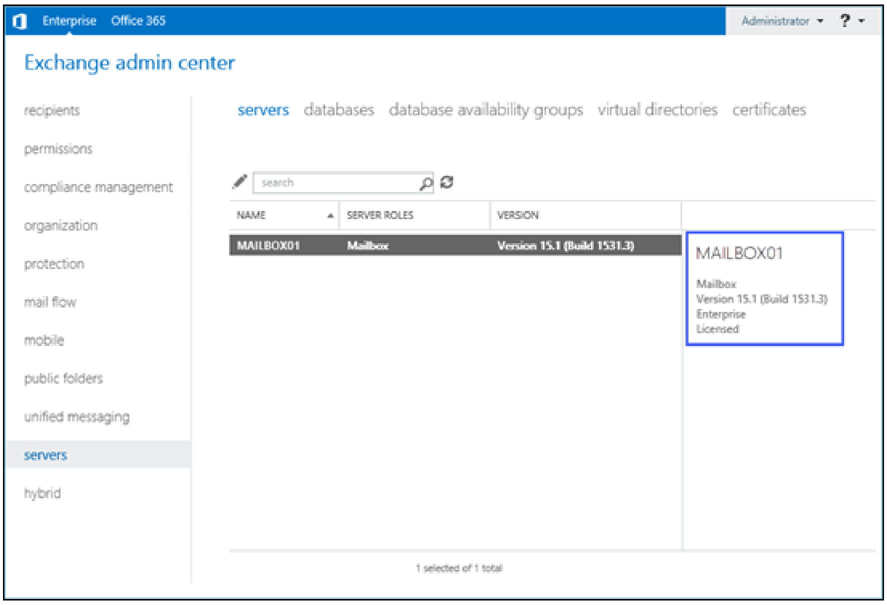

NAME	SERVER ROLES	VERSION
MAILBOX01	Mailbox	Version 15.1 (Build 1531.3)

MAILBOX01
Mailbox
Version 15.1 (Build 1531.3)
Enterprise
Licensed

1 selected of 1 total

- In the Exchange Management Shell, replace *<ServerName>* with the name of the Exchange server you licensed, and run the following command to verify the property values:

```
Get-ExchangeServer <ServerName> | Format-List Name,Edition,*Trial*
```

- In the Exchange Management Shell, run the following command to view the licensing status of all Exchange servers in your organization:

```
Get-ExchangeServer | Format-Table -Auto Name,Edition,*Trial*
```

Configure mail flow and client access on Exchange servers

8/3/2020 • 13 minutes to read • Edit Online

After you've installed Exchange Server 2016 or Exchange 2019 in your organization, you need to configure Exchange for mail flow and client access. Without these additional steps, you won't be able to send mail to the internet and external clients (for example, Microsoft Outlook, and Exchange ActiveSync devices) won't be able to connect to your Exchange organization.

The steps in this topic assume a basic Exchange deployment with a single Active Directory site and a single simple mail transport protocol (SMTP) namespace.

IMPORTANT

This topic uses example values such as Mailbox01, contoso.com, mail.contoso.com, and 172.16.10.11. Replace the example values with the server names, FQDNs, and IP addresses for your organization.

For additional management tasks related to mail flow and clients and devices, see Mail flow and the transport pipeline and Clients and mobile.

What do you need to know before you begin?

- Estimated time to complete this task: 50 minutes
- You might receive certificate warnings when you connect to the Exchange admin center (EAC) website until you configure a secure sockets layer (SSL) certificate on the Mailbox server. You'll be shown how to do this later in this topic.
- To open the EAC, see Exchange admin center in Exchange Server. To open the Exchange Management Shell, see Open the Exchange Management Shell.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Step 1: Create an internet Send connector

Before you can send mail to the internet, you need to create a Send connector on the Mailbox server. For instructions, see Create a Send connector in Exchange Server to send mail to the internet.

NOTE

By default, a Receive connector named "Default Frontend <ServerName>_" is created when Exchange is installed. This Receive connector accepts anonymous SMTP connections from external servers. You don't need to do any additional configuration if this is the functionality you want. If you want to restrict inbound connections from external servers, modify the **Default Frontend <Mailbox server>** Receive connector on the Mailbox server. For more information, see Default Receive connectors created during setup.

Step 2: Add additional accepted domains

By default, Exchange uses the Active Directory domain where Setup /PrepareAD was run for email addresses. If you want recipients to receive and send messages to and from another domain, you need to add the domain as an accepted domain. For instructions, see Create accepted domains and Configure Exchange to accept mail for multiple authoritative domains.

IMPORTANT

To receive email from the internet for a domain, you need an MX resource record in your public DNS for that domain. Each MX record should resolve to the internet-facing server that receives email for your organization.

Step 3: Configure the default email address policy

You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Email address policies" entry in the Email address and address book permissions topic.

If you added an accepted domain in the previous step and you want that domain to be added to every recipient in the organization, you need to update the default email address policy. For instructions, see Modify email address policies and Apply email address policies to recipients.

NOTE

We recommend that you configure a user principal name (UPN) that matches the primary email address of each user. If you don't provide a UPN that matches the email address of a user, the user will be required to manually provide their domain\username or UPN in addition to their email address. If their UPN matches their email address, Outlook on the web (formerly known as Outlook on the web), ActiveSync, and Outlook will automatically match their email address to their UPN.

Step 4: Configure external URLs

You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the " *<Service>* virtual directory settings" entry in the Clients and mobile devices permissions topic.

Before clients can connect to your new server from the internet, you need to configure the external domains (or URLs) on the virtual directories in the Client Access (frontend) services on the Mailbox server and then in your public DNS records. The steps below configure the same external domain on the external URL of each virtual directory. If you want to configure different external domains on one or more virtual directory external URLs, you need to configure the external URLs manually. For more information, see Default settings for Exchange virtual directories.

1. Open the EAC and go to **Servers** > **Servers**, select your internet-facing Mailbox server that your clients will connect to, and then click **Edit** ✎.

2. In the Exchange server properties window that opens, select the **Outlook Anywhere** tab, configure the following settings:

   Specify the external host name...: Enter the externally accessible FQDN that your external clients will use to connect to their mailboxes (for example, mail.contoso.com).

   Specify the internal host name...: Enter the internally accessible FQDN (for example, mail.contoso.com).

   When you're finished, click **Save**.

3. Go to **Servers** > **Virtual directories** and then select **Configure external access domain** 🔧.

4. In the **Configure external access domain** window opens, configure the following settings:

   a. **Select the Mailbox servers to use with the external URL**: Click **Add** +

   b. In the **Select a server** dialog that opens, select the Mailbox server you want to configure and then click **Add**. After you've added all of the Mailbox servers that you want to configure, click **OK**.

   c. **Enter the domain name you will use with your external Mailbox servers**: Enter the external domain that you want to apply (for example, mail.contoso.com). When you're finished, click **Save**.

Some organizations use a unique Outlook on the web FQDN to protect against future changes to the underlying server FQDN. Many organizations use owa.contoso.com for their Outlook on the web FQDN instead of mail.contoso.com. If you want to configure a unique Outlook on the web FQDN, do the following steps. This checklist assumes you have configured a unique Outlook on the web FQDN.

1. Back at **Servers** > **Virtual directories**, select **owa (Default Web Site)** on the server that you want to configure, and then click **Edit** ✏.

2. The **owa (Default web site)** window opens. On the **General** tab in the **External URL** field, enter the following information:

   - https://
   - The unique Outlook on the web FQDN you want to use (for example, owa.contoso.com), and then append /owa. For example, https://owa.contoso.com/owa.
   - /owa

   In this example, the final value would be https://owa.contoso.com/owa.

   When you're finished, click **Save**.

3. Back at **Servers** > **Virtual directories**, select **ecp (Default Web Site)** on the server that you want to configure, and click **Edit** ✏.

4. In the **ecp (Default web site)** window that opens, enter the same URL from the previous step, but append the value /ecp instead of /owa (for example, https://owa.contoso.com/ecp). When you're finished, click **Save**.

After you've configured the external URL in the Client Access services virtual directories on the Mailbox server, you need to configure your public DNS records for Autodiscover, Outlook on the web, and mail flow. The public DNS records should point to the external IP address or FQDN of your internet-facing Mailbox server and use the externally accessible FQDNs that you've configured on your Mailbox server. The recommended DNS records that you should create to enable mail flow and external client connectivity are described in the following table:

FQDN	DNS RECORD TYPE	VALUE
Contoso.com	MX	Mail.contoso.com

FQDN	DNS RECORD TYPE	VALUE
Mail.contoso.com	A	172.16.10.11
Owa.contoso.com	CNAME	Mail.contoso.com
Autodiscover.contoso.com	CNAME	Mail.contoso.com

How do you know this step worked?

To verify that you've successfully configured the external URLs in the Client Access services virtual directories on the Mailbox server, do the following steps:

1. In the EAC, go to **Servers** > **Virtual directories**.
2. In the **Select server** field, select the internet-facing Mailbox server.
3. Select a virtual directory and then, in the virtual directory details pane, verify that the **External URL** field is populated with the correct FQDN and service as shown in the following table:

VIRTUAL DIRECTORY	EXTERNAL URL VALUE
Autodiscover	No external URL displayed
ECP	https://owa.contoso.com/ecp
EWS	https://mail.contoso.com/EWS/Exchange.asmx
Microsoft-Server-ActiveSync	https://mail.contoso.com/Microsoft-Server-ActiveSync
OAB	https://mail.contoso.com/OAB
OWA	https://owa.contoso.com/owa
PowerShell	http://mail.contoso.com/PowerShell

To verify that you've successfully configured your public DNS records, do the following steps:

1. Open a command prompt and run `nslookup.exe`.
2. Change to a DNS server that can query your public DNS zone.
3. In `nslookup`, look up the record of each FQDN you created. Verify that the value that's returned for each FQDN is correct.
4. In `nslookup`, type `set type=mx` and then look up the accepted domain you added in Step 1. Verify that the value returned matches the FQDN of the Mailbox server.

Step 5: Configure internal URLs

You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the " *<Service>* virtual directory settings" entry in the Clients and mobile devices permissions topic.

Before clients can connect to your new server from your internal network, you need to configure the internal domains (or URLs) on the virtual directories in the Client Access (frontend) services on the Mailbox server and then in your internal DNS records.

The procedure below lets you choose whether you want users to use the same URL on your intranet and on the internet to access your Exchange server or whether they should use a different URL. What you choose depends on the addressing scheme you have in place already or that you want to implement. If you're implementing a new addressing scheme, we recommend that you use the same URL for both internal and external URLs. Using the same URL makes it easier for users to access your Exchange server because they only have to remember one address.

Regardless of your decision, you need to configure a private DNS zone for the address space you choose. For more information about administering DNS zones, see Administering DNS Server.

For more information about internal and external URLs on virtual directories, see Default settings for Exchange virtual directories Virtual Directory Management.

Configure internal and external URLs to be the same

1. Open the Exchange Management Shell on your Mailbox server.

2. Store the host name of your Mailbox server in a variable that will be used in the next step. For example, Mailbox01.

   ```
   $HostName = "Mailbox01"
   ```

3. Run each of the following commands in the Exchange Management Shell to configure each internal URL to match the virtual directory's external URL.

   ```
   Set-EcpVirtualDirectory "$HostName\ECP (Default Web Site)" -InternalUrl ((Get-EcpVirtualDirectory
   "$HostName\ECP (Default Web Site)").ExternalUrl)
   ```

   ```
   Set-WebServicesVirtualDirectory "$HostName\EWS (Default Web Site)" -InternalUrl ((Get-
   WebServicesVirtualDirectory "$HostName\EWS (Default Web Site)").ExternalUrl)
   ```

   ```
   Set-ActiveSyncVirtualDirectory "$HostName\Microsoft-Server-ActiveSync (Default Web Site)" -InternalUrl
   ((Get-ActiveSyncVirtualDirectory "$HostName\Microsoft-Server-ActiveSync (Default Web
   Site)").ExternalUrl)
   ```

   ```
   Set-OabVirtualDirectory "$HostName\OAB (Default Web Site)" -InternalUrl ((Get-OabVirtualDirectory
   "$HostName\OAB (Default Web Site)").ExternalUrl)
   ```

   ```
   Set-OwaVirtualDirectory "$HostName\OWA (Default Web Site)" -InternalUrl ((Get-OwaVirtualDirectory
   "$HostName\OWA (Default Web Site)").ExternalUrl)
   ```

   ```
   Set-PowerShellVirtualDirectory "$HostName\PowerShell (Default Web Site)" -InternalUrl ((Get-
   PowerShellVirtualDirectory "$HostName\PowerShell (Default Web Site)").ExternalUrl)
   ```

After you've configured the internal URL on the Mailbox server virtual directories, you need to configure your private DNS records for Outlook on the web and other connectivity. Depending on your configuration, you'll need to configure your private DNS records to point to the internal or external IP address or FQDN of your Mailbox server. Examples of recommended DNS records that you should create are described in the following table:

FQDN	DNS RECORD TYPE	VALUE
Mail.contoso.com	CNAME	Mailbox01.corp.contoso.com
Owa.contoso.com	CNAME	Mailbox01.corp.contoso.com

How do you know this step worked?

To verify that you've successfully configured the internal URL on the Mailbox server virtual directories, do the following:

1. In the EAC, go to **Servers** > **Virtual directories**.
2. In the **Select server** field, select the internet-facing Mailbox server.
3. Select a virtual directory and then click **Edit** ✎.
4. Verify that the **Internal URL** field is populated with the correct FQDN and service as shown in the following table:

VIRTUAL DIRECTORY	INTERNAL URL VALUE
Autodiscover	No internal URL displayed
ECP	https://owa.contoso.com/ecp
EWS	https://mail.contoso.com/EWS/Exchange.asmx
Microsoft-Server-ActiveSync	https://mail.contoso.com/Microsoft-Server-ActiveSync
OAB	https://mail.contoso.com/OAB
OWA	https://owa.contoso.com/owa
PowerShell	http://mail.contoso.com/PowerShell

To verify that you have successfully configured your private DNS records, do the following:

1. Open a command prompt and run `nslookup.exe`.
2. Change to a DNS server that can query your private DNS zone.
3. In `nslookup`, look up the record of each FQDN you created. Verify that the value that's returned for each FQDN is correct.

Configure different internal and external URLs

1. Open the EAC, and go to **Servers** > **Virtual directories**,
2. On the internet-facing Mailbox server, select the virtual directory that you want to configure, and then click **Edit** ✎.
3. The virtual directory properties window opens. In the **Internal URL** field, replace the existing host name value in the URL (likely, the FQDN of the Mailbox server) with the new value that you want to use (for example, internal.contoso.com).

   For example, in the properties of the Exchange Web Services (EWS) virtual directory, change the existing value from https://**Mailbox01.corp.contoso.com**/ews/exchange.asmx to https://**internal.contoso.com**/ews/exchange.asmx.

When you're finished, click **Save**.

4. Repeat the previous steps for each virtual directory you want to change.

> **NOTE**
>
> The ECP and OWA virtual directory internal URLs must be the same. You can't set an internal URL on the Autodiscover virtual directory.

After you've configured the internal URL on the Mailbox server virtual directories, you need to configure your private DNS records for Outlook on the web, and other connectivity. Depending on your configuration, you'll need to configure your private DNS records to point to the internal or external IP address or FQDN of your Mailbox server. An example of the recommended DNS record that you should create is described in the following table:

FQDN	DNS RECORD TYPE	VALUE
internal.contoso.com	CNAME	Mailbox01.corp.contoso.com

How do you know this step worked?

To verify that you've successfully configured the internal URLs in the Client Access services virtual directories on the Mailbox server, do the following steps:

1. In the EAC, go to **Servers** > **Virtual directories**.

2. In the **Select server** field, select the internet-facing Mailbox server.

3. Select a virtual directory and then click **Edit** ✎.

4. Verify that the **Internal URL** field is populated with the correct FQDN. For example, you may have set the internal URLs to use internal.contoso.com.

VIRTUAL DIRECTORY	INTERNAL URL VALUE
Autodiscover	No internal URL displayed
ECP	https://internal.contoso.com/ecp
EWS	https://internal.contoso.com/EWS/Exchange.asmx
Microsoft-Server-ActiveSync	https://internal.contoso.com/Microsoft-Server-ActiveSync
OAB	https://internal.contoso.com/OAB
OWA	https://internal.contoso.com/owa
PowerShell	http://internal.contoso.com/PowerShell

To verify that you've successfully configured your private DNS records, do the following:

1. Open a command prompt and run `nslookup.exe`.

2. Change to a DNS server that can query your private DNS zone.

3. In `nslookup`, look up the record of each FQDN you created. Verify that the value that's returned for each FQDN is correct.

Step 6: Configure an SSL certificate

Some services, such as Outlook Anywhere and Exchange ActiveSync, require certificates to be configured on your Exchange server. The following steps show you how to configure an SSL certificate from a third-party certificate authority (CA):

1. Create an Exchange Server certificate request for a certification authority.
    - You should request a certificate from a third-party CA so your clients automatically trust the certificate. For more information, see Best practices for Exchange certificates.
    - If you configured your internal and external URLs to be the same, **Outlook on the web (when accessed from the internet)** and **Outlook on the web (when accessed from the Intranet)** should both show owa.contoso.com. **OAB (when accessed from the internet)** and **OAB (when accessed from the Intranet)** should show mail.contoso.com.
    - If you configured the internal URLs to be internal.contoso.com, **Outlook on the web (when accessed from the internet)** should show owa.contoso.com and **Outlook on the web (when accessed from the Intranet)** should show internal.contoso.com.
2. Complete a pending Exchange Server certificate request.
3. Assign certificates to Exchange Server services
    - At minimum, you should select **SMTP** and **IIS**.
    - If you receive the warning **Overwrite the existing default SMTP certificate?**, click **Yes**.

How do you know this step worked?

To verify that you've successfully added a new certificate, do the following steps:

1. In the EAC, go to **Servers** > **Certificates**.
2. Select the new certificate and then, in the certificate details pane, verify that the following are true:
    - **Status** shows **Valid**
    - **Assigned to services** shows, at minimum, **IIS** and **SMTP**.

How do you know this task worked?

To verify that you've configured mail flow and external client access, do the following steps:

1. In Outlook, on an Exchange ActiveSync device, or on both, create a new profile. Verify that Outlook or the mobile device successfully creates the new profile.
2. In Outlook, or on the mobile device, send a new message to an external recipient. Verify the external recipient receives the message.
3. In the external recipient's mailbox, reply to the message you just sent from the Exchange mailbox. Verify the Exchange mailbox receives the message.
4. Go to https://owa.contoso.com/owa and verify that there are no certificate warnings.

Verify Exchange Server installations

8/3/2020 • 2 minutes to read • Edit Online

After you install Exchange Server 2016 or Exchange Server 2019, we recommend that you verify the installation by running the **Get-ExchangeServer** cmdlet and by reviewing the Exchange Setup log. If the setup process fails or errors occur during installation, you can use the Setup log to find the source of the problem.

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Run Get-ExchangeServer

To verify that Exchange installed successfully, run the following commands in the Exchange Management Shell. To open the Exchange Management Shell, see Open the Exchange Management Shell.

This command returns a summary list of the names, Active Directory sites, Exchange server roles, Exchange editions, and Exchange versions of all Exchange servers in the organization.

```
Get-ExchangeServer
```

This example returns additional details about the Exchange server named Mailbox01.

```
Get-ExchangeServer -Identity Mailbox01 | Format-List
```

For detailed syntax and parameter information, see Get-ExchangeServer.

Review the Windows Application log and the Exchange Setup log

- Exchange Setup logs events in the **Application** log of the Windows Server. This log contains a history of each action that the system takes during Exchange setup and any errors that occurred (By default, the logging method is set to Verbose). You can use the Windows **Event Viewer** to find the messages related to Exchange setup.

- The Exchange Setup log is available at *<system drive>*:\ExchangeSetupLogs\ExchangeSetup.log (*<system drive>* is the drive where Windows is installed). The Setup log tracks the progress of every task during the Exchange installation and configuration. The file contains information about the status of the prerequisite and system readiness checks before installation starts, the application installation progress, and the configuration changes that are made to the system. Check this log file to verify that Exchange was installed as expected.

We recommend that you start your review of the Windows Application log and/or the Exchange Setup log by searching for errors. If you find an error entry, read the associated text to determine the cause of the error.

Install the Exchange management tools

8/3/2020 • 5 minutes to read • Edit Online

The management tools in Exchange Server 2016 and Exchange Server 2019 include the Exchange Management Shell and the Exchange Toolbox. You can install the management tools on other client computers or servers in the Active Directory domain to help you manage your Exchange organization. The management tools have similar operating system, .NET Framework, and Windows Management Framework (Windows PowerShell) requirements as an Exchange server. The notable exception is: you can install the management tools on client versions of Windows. For more information, see Exchange Server system requirements and Exchange Server prerequisites.

> **NOTE**
>
> The management tools don't include the Exchange admin center (EAC). The EAC is a web-based console that's hosted on Exchange 2016 Mailbox servers, and like any web site, you can access the EAC from other computers. For more information about the EAC, see Exchange admin center in Exchange Server.

For more information about the Exchange Management Shell, see Exchange Server PowerShell (Exchange Management Shell).

What do you need to know before you begin?

- Estimated time to complete: 20 minutes
- The computer where you want to install the Exchange management tools requires access to Setup.exe in the Exchange installation files. To download the latest version of Exchange, see Updates for Exchange Server.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange 2016 Setup wizard to install the Exchange management tools

1. In File Explorer on the computer where you want to install the management tools, right-click on the Exchange ISO image file that you downloaded, and then select **Mount**. In the resulting virtual DVD drive that appears, start Exchange Setup by double-clicking `Setup.exe`.

2. The Exchange Server Setup wizard opens. On the **Check for Updates?** page, choose one of the following options, and then click **Next** to continue:

   Connect to the Internet and check for updates: We recommend this option, which searches for updates to the version of Exchange *that you're currently installing* (it doesn't detect newer Cumulative Updates). This option takes you to the **Downloading Updates** page that searches for updates. Click **Next** to continue.

   - **Don't check for updates right now**

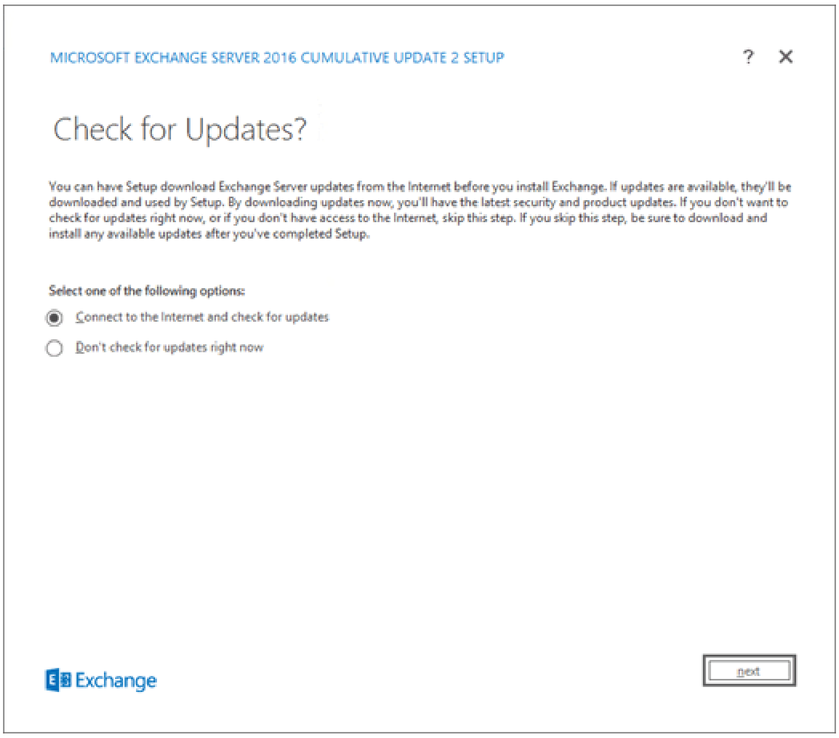

MICROSOFT EXCHANGE SERVER 2016 CUMULATIVE UPDATE 2 SETUP

Check for Updates?

You can have Setup download Exchange Server updates from the Internet before you install Exchange. If updates are available, they'll be downloaded and used by Setup. By downloading updates now, you'll have the latest security and product updates. If you don't want to check for updates right now, or if you don't have access to the Internet, skip this step. If you skip this step, be sure to download and install any available updates after you've completed Setup.

Select one of the following options:

● Connect to the Internet and check for updates

○ Don't check for updates right now

3. The **Copying Files** page shows the progress of copying files to the local hard drive. Typically, the files are copied to `%WinDir%\Temp\ExchangeSetup`, but you can confirm the location in the Exchange Setup log at `C:\ExchangeSetupLogs\ExchangeSetup.log`.

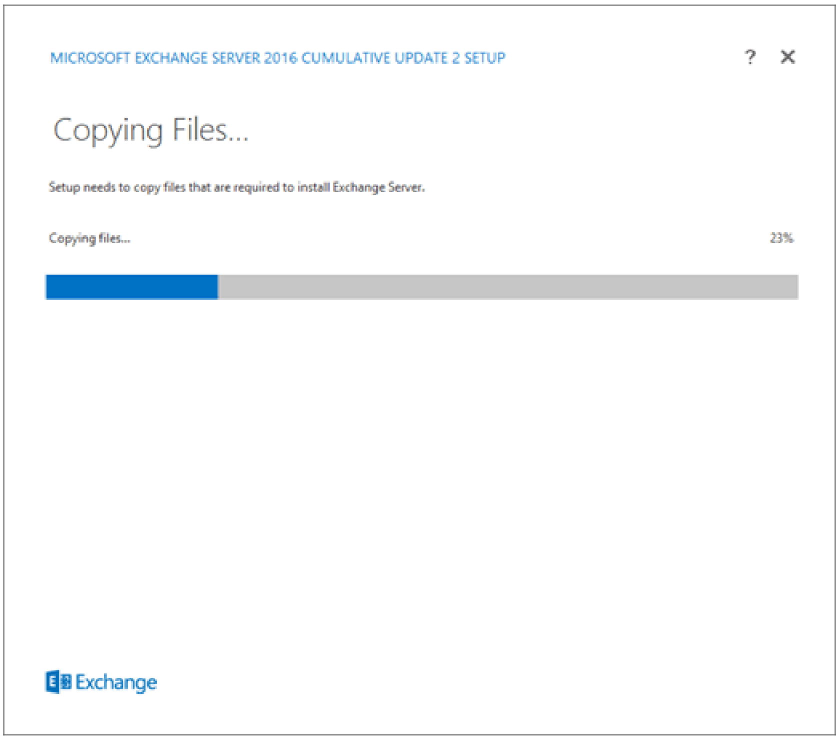

MICROSOFT EXCHANGE SERVER 2016 CUMULATIVE UPDATE 2 SETUP

Copying Files...

Setup needs to copy files that are required to install Exchange Server.

Copying files... 23%

4. On the **Introduction** page, click **Next** to continue.

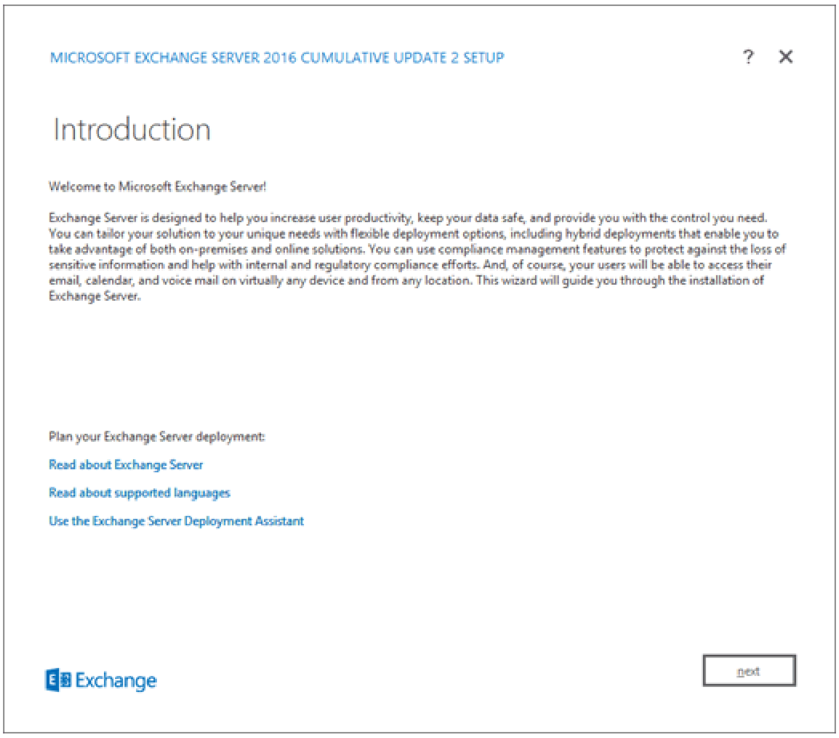

Introduction

Welcome to Microsoft Exchange Server!

Exchange Server is designed to help you increase user productivity, keep your data safe, and provide you with the control you need. You can tailor your solution to your unique needs with flexible deployment options, including hybrid deployments that enable you to take advantage of both on-premises and online solutions. You can use compliance management features to protect against the loss of sensitive information and help with internal and regulatory compliance efforts. And, of course, your users will be able to access their email, calendar, and voice mail on virtually any device and from any location. This wizard will guide you through the installation of Exchange Server.

Plan your Exchange Server deployment:

Read about Exchange Server

Read about supported languages

Use the Exchange Server Deployment Assistant

5. On the **License Agreement** page, review the software license terms, select **I accept the terms in the license agreement**, and then click **Next** to continue.

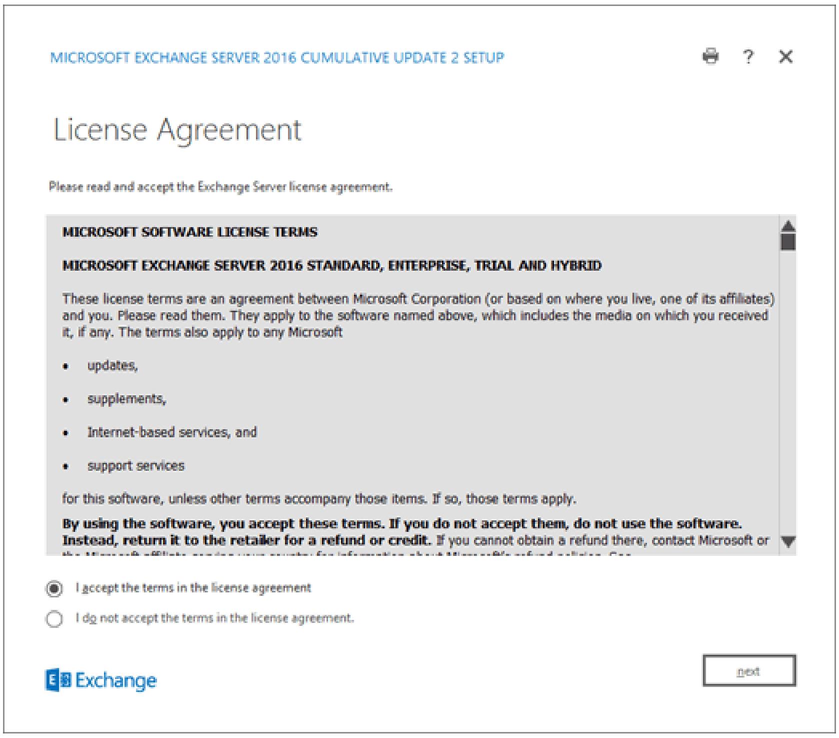

License Agreement

Please read and accept the Exchange Server license agreement.

MICROSOFT SOFTWARE LICENSE TERMS

MICROSOFT EXCHANGE SERVER 2016 STANDARD, ENTERPRISE, TRIAL AND HYBRID

These license terms are an agreement between Microsoft Corporation (or based on where you live, one of its affiliates) and you. Please read them. They apply to the software named above, which includes the media on which you received it, if any. The terms also apply to any Microsoft

- updates,
- supplements,
- Internet-based services, and
- support services

for this software, unless other terms accompany those items. If so, those terms apply.

By using the software, you accept these terms. If you do not accept them, do not use the software. Instead, return it to the retailer for a refund or credit. If you cannot obtain a refund there, contact Microsoft or

● I accept the terms in the license agreement

○ I do not accept the terms in the license agreement.

6. On the **Recommended Settings** page, choose one of the following settings:

   - **Use recommended settings**: Exchange automatically sends error reports and information about your computer hardware and how you use Exchange to Microsoft. For information about what's sent to Microsoft and how it's used, click **?** or the help links on the page.

   - **Don't use recommended settings**: These settings are disabled, but you can enable them at any time after Setup completes.

   Click **Next** to continue.

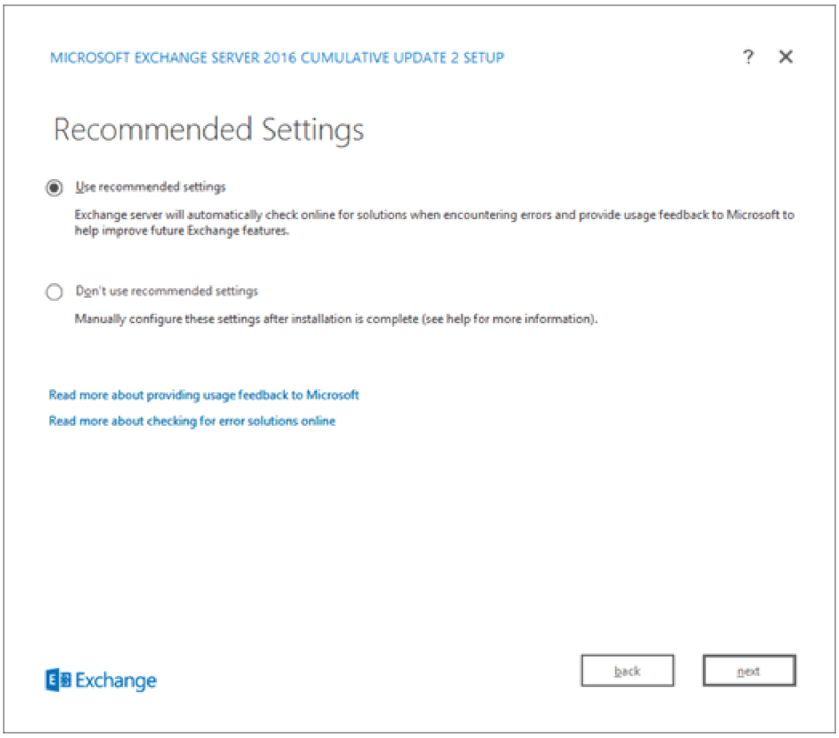

Recommended Settings

● Use recommended settings

Exchange server will automatically check online for solutions when encountering errors and provide usage feedback to Microsoft to help improve future Exchange features.

○ Don't use recommended settings

Manually configure these settings after installation is complete (see help for more information).

Read more about providing usage feedback to Microsoft

Read more about checking for error solutions online

7. On the **Server Role Selection** page, configure the following settings:

   - Select **Management tools**.

   - **Automatically install Windows Server roles and features that are required to install Exchange**: Select this option to have the Setup wizard install the required Windows prerequisites. You might need to reboot the computer to complete the installation of some Windows features. If you don't select this option, you need to install the Windows features manually.

     Note: Selecting this option installs only the *Windows features* that are required by Exchange. You need to install other prerequisites manually. For more information, see Exchange Server prerequisites.

   Click **Next** to continue.

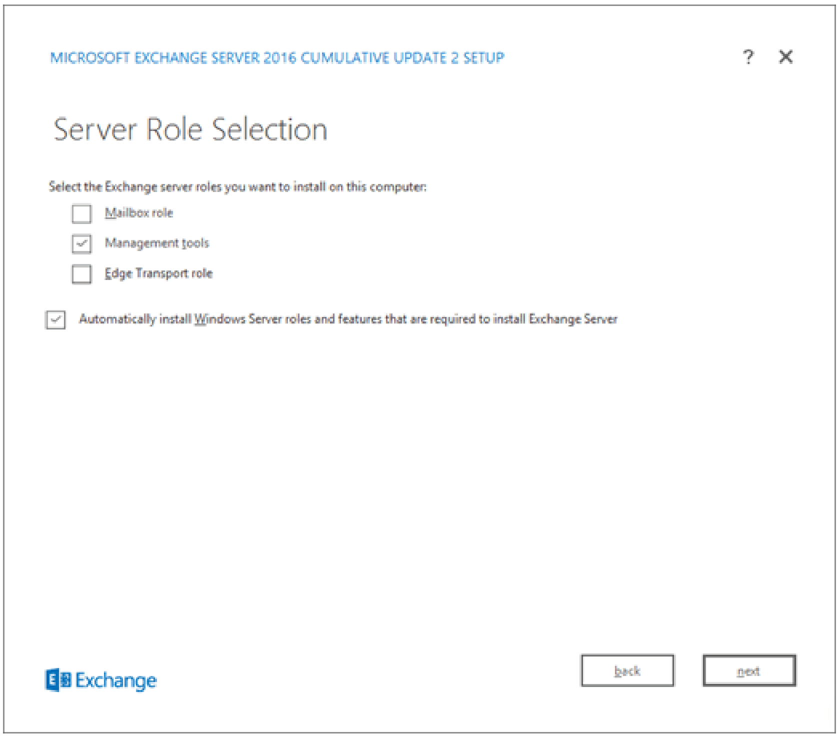

Server Role Selection

Select the Exchange server roles you want to install on this computer:

□ Mailbox role

☑ Management tools

□ Edge Transport role

☑ Automatically install Windows Server roles and features that are required to install Exchange Server

8. On the **Installation Space and Location** page, either accept the default installation location (

`C:\Program Files\Microsoft\Exchange Server\V15`), or click **Browse** to choose a new location. Make sure that you have enough disk space available in the location where you want to install the management tools. Click **Next** to continue.

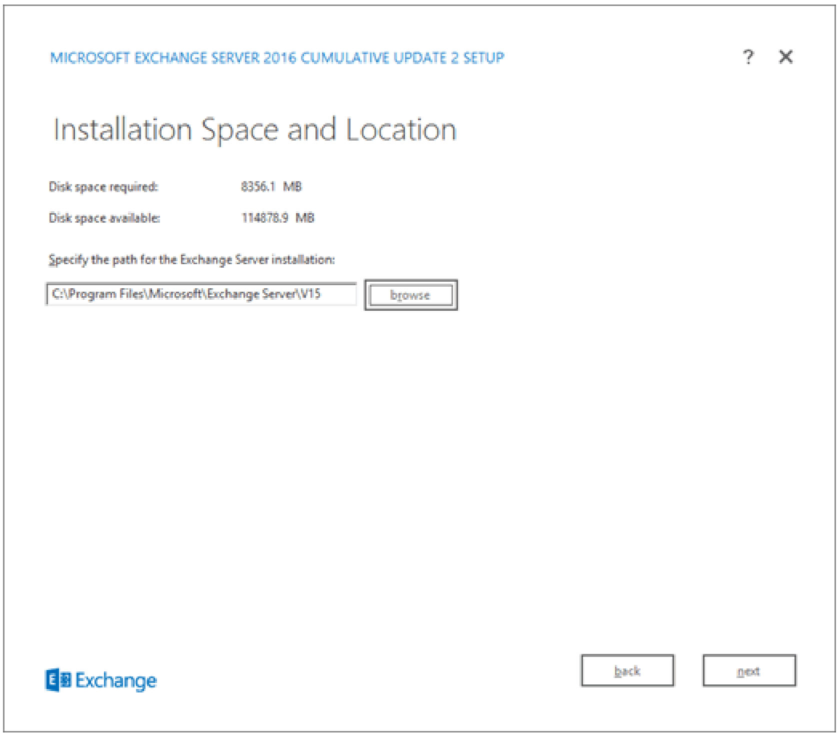

MICROSOFT EXCHANGE SERVER 2016 CUMULATIVE UPDATE 2 SETUP

Installation Space and Location

Disk space required: 8356.1 MB

Disk space available: 114878.9 MB

Specify the path for the Exchange Server installation:

C:\Program Files\Microsoft\Exchange Server\V15

9. If this is the first installation of Exchange in your organization (Exchange server or the management tools), you arrive on the **Exchange Organization** page. On this page, configure the following settings:

   - **Specify the name for this Exchange organization**: The default value is **First Organization**, but you typically use the company name for this value. The organization name is used internally by Exchange, isn't typically seen by users, doesn't affect the functionality of Exchange, and doesn't determine what you can use for email addresses.

     - The organization name can't contain more than 64 characters, and can't be blank.

     - Valid characters are A to Z, a to z, 0 to 9, hyphen or dash (-), and space, but leading or trailing spaces aren't allowed.

     - You can't change the organization name after it's set.

   - **Apply Active Directory split permission security model to the Exchange organization**: Most organizations don't need to select this option. If you need to separate management of Active Directory security principals and the Exchange configuration, split permissions might work for you. For more information, click **?**.

   Click **Next** to continue.

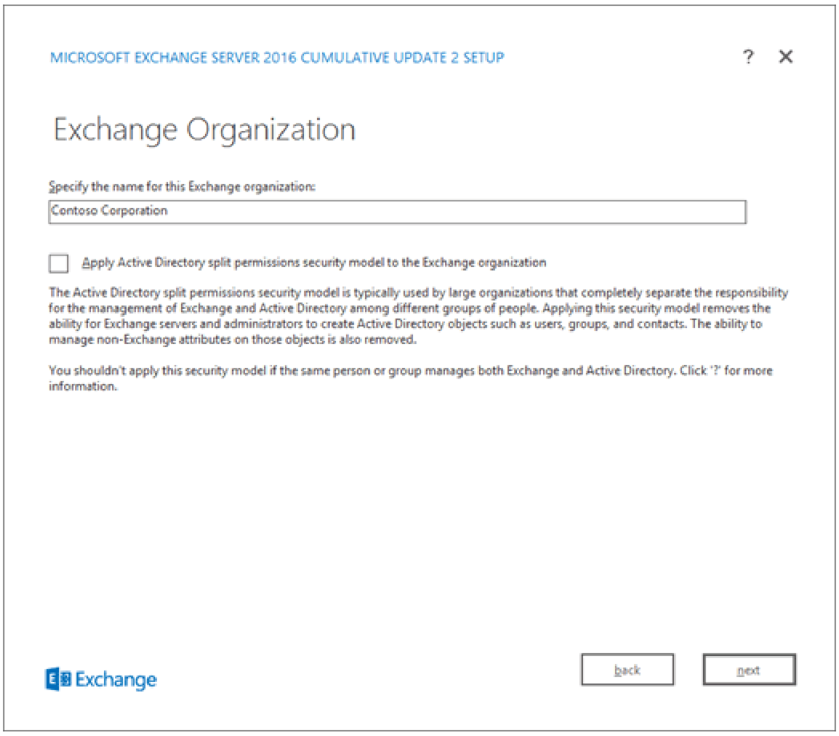

Exchange Organization

Specify the name for this Exchange organization:

Contoso Corporation

□ Apply Active Directory split permissions security model to the Exchange organization

The Active Directory split permissions security model is typically used by large organizations that completely separate the responsibility for the management of Exchange and Active Directory among different groups of people. Applying this security model removes the ability for Exchange servers and administrators to create Active Directory objects such as users, groups, and contacts. The ability to manage non-Exchange attributes on those objects is also removed.

You shouldn't apply this security model if the same person or group manages both Exchange and Active Directory. Click '?' for more information.

10. On the **Readiness Checks** page, verify that the organization and server role prerequisite checks completed successfully. If they haven't, the only option on the page is **Retry**, so you need to resolve the errors before you can continue.

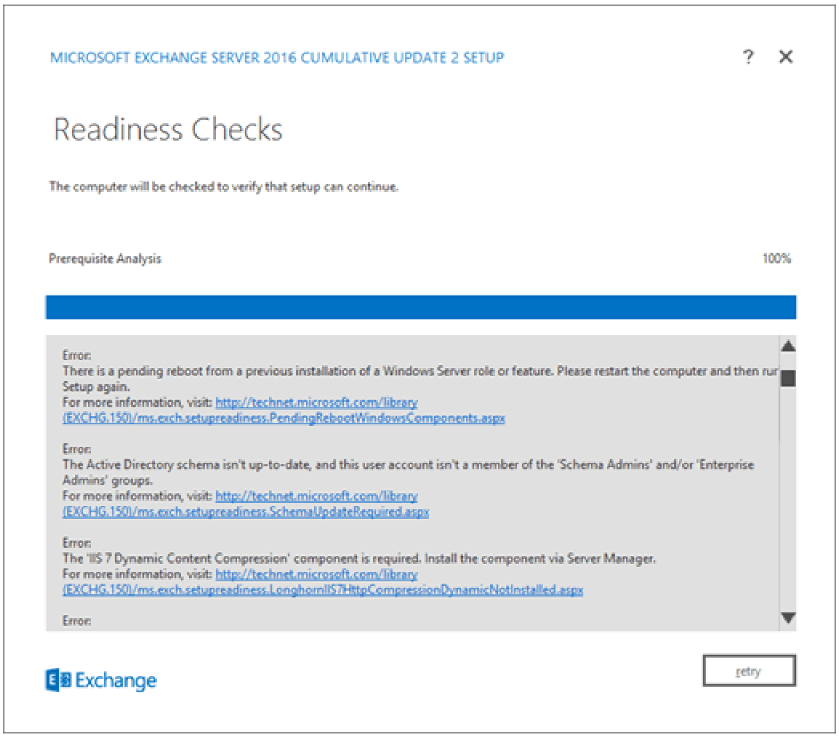

Readiness Checks

The computer will be checked to verify that setup can continue.

Prerequisite Analysis 100%

Error:
There is a pending reboot from a previous installation of a Windows Server role or feature. Please restart the computer and then rur
Setup again.
For more information, visit: http://technet.microsoft.com/library
(EXCHG.150)/ms.exch.setupreadiness.PendingRebootWindowsComponents.aspx

Error:
The Active Directory schema isn't up-to-date, and this user account isn't a member of the 'Schema Admins' and/or 'Enterprise Admins' groups.
For more information, visit: http://technet.microsoft.com/library
(EXCHG.150)/ms.exch.setupreadiness.SchemaUpdateRequired.aspx

Error:
The 'IIS 7 Dynamic Content Compression' component is required. Install the component via Server Manager.
For more information, visit: http://technet.microsoft.com/library
(EXCHG.150)/ms.exch.setupreadiness.LonghornIIS7HttpCompressionDynamicNotInstalled.aspx

Error:

After you resolve the errors, click **Retry** to run the prerequisite checks again. You can fix some errors without exiting Setup, while the fix for other errors requires you to restart the computer. If you restart the computer, you need to start over at Step 2.

When no more errors are detected on the **Readiness Checks** page, the **Retry** button changes to **Install** so you can continue. Be sure to review any warnings, and then click **Install** to install the management tools.

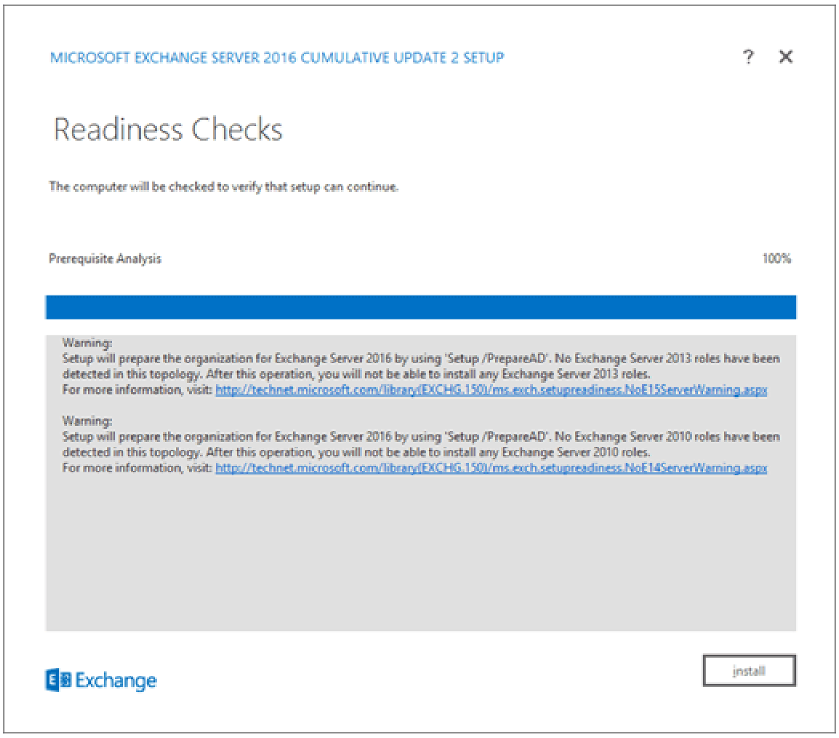

Readiness Checks

The computer will be checked to verify that setup can continue.

Prerequisite Analysis 100%

Warning:
Setup will prepare the organization for Exchange Server 2016 by using 'Setup /PrepareAD'. No Exchange Server 2013 roles have been detected in this topology. After this operation, you will not be able to install any Exchange Server 2013 roles.
For more information, visit: http://technet.microsoft.com/library(EXCHG.150)/ms.exch.setupreadiness.NoE15ServerWarning.aspx

Warning:
Setup will prepare the organization for Exchange Server 2016 by using 'Setup /PrepareAD'. No Exchange Server 2010 roles have been detected in this topology. After this operation, you will not be able to install any Exchange Server 2010 roles.
For more information, visit: http://technet.microsoft.com/library(EXCHG.150)/ms.exch.setupreadiness.NoE14ServerWarning.aspx

11. On the **Setup Completed** page, click **Finish**, and then restart the computer.

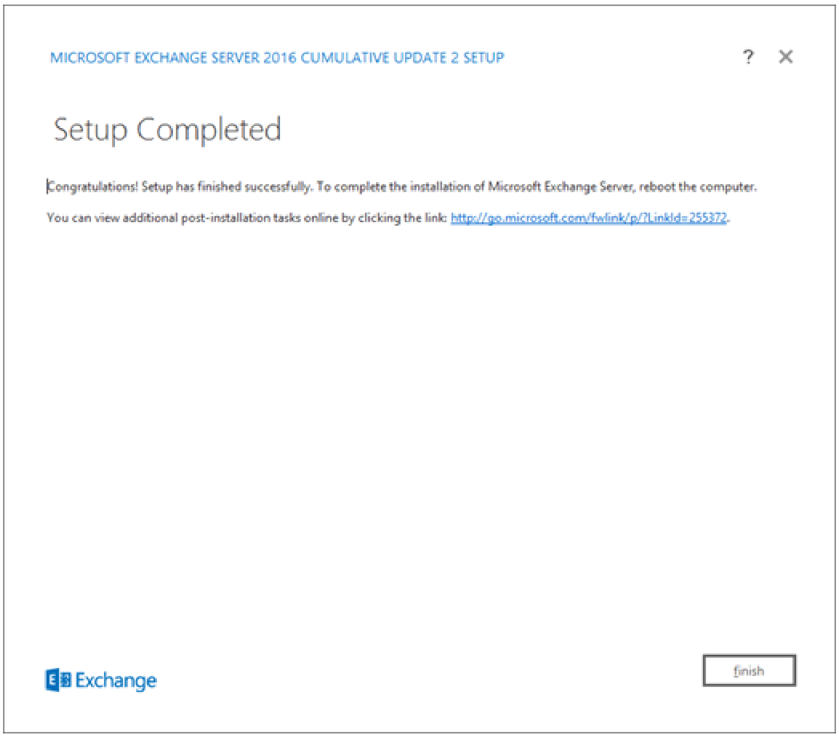

Setup Completed

Congratulations! Setup has finished successfully. To complete the installation of Microsoft Exchange Server, reboot the computer.

You can view additional post-installation tasks online by clicking the link: http://go.microsoft.com/fwlink/p/?LinkId=255372.

Use Exchange unattended Setup mode to install the Exchange management tools

1. In File Explorer on the computer where you want to install the Exchange management tools, right-click on the Exchange ISO image file that you downloaded, and then select **Mount**.

2. To install the Exchange management tools from the command line, use the following syntax in elevated command prompt (a Command Prompt window you opened by selecting **Run as administrator**):

```
<Virtual DVD drive letter>:\Setup.exe /IAcceptExchangeServerLicenseTerms /Role:ManagementTools
[/EnableErrorReporting] [/CustomerFeedbackEnabled:<True | False>] [/InstallWindowsComponents]
[/TargetDir:<Target folder>] [/OrganizationName:<Name>]
```

This example uses the Exchange Setup files on drive E: to install the management tools on the local server

```
E:\Setup.exe /IAcceptExchangeServerLicenseTerms /Role:ManagementTools
```

For more information, see Install Exchange using unattended mode.

Configure instant messaging integration with Outlook on the web in Exchange

8/3/2020 • 4 minutes to read • Edit Online

To configure instant messaging (IM) integration between Skype for Business Server and Outlook on the web (formerly known as Outlook Web App) in Exchange 2016 or Exchange 2019, you need to use the Exchange Management Shell. This is different than previous versions of Exchange where you needed to edit the web.config file. If you edit the web.config file instead of using the steps in this topic, the settings are ignored and Outlook on the web users receive the following error message:

```
There's a problem with instant messaging. Please try again later.
```

Also, the following health set errors are generated on the Exchange server:

- **HealthSet**: `OWA.Protocol.Dep`
- **Subject**:
  ```
  OWA.Protocol.Dep health set unhealthy (OwaIMInitializationFailedMonitor/OWA.Protocol.Dep) - Owa
  InstantMessaging provider failed to intialize
  ```
- **Message**:
  ```
  Owa InstantMessaging provider failed to initialize due to incorrect IM configuration on the server.
  Signin attempts to OWA IM will fail. Error Message: {Instant Messaging Certificate Thumbprint is null or
  empty on web.config).
  ```

Use the procedures in this topic to fix these errors and configure IM integration between Skype for Business Server and Exchange 2016 or Exchange 2019. IM integration between Lync Server 2013 and Exchange 2016 or later isn't supported. For details on setting up Skype for Business Server with Outlook on the web (formerly known as Outlook Web App), see Configure integration between on-premises Skype for Business Server and Outlook Web App

What do you need to know before you begin?

- Estimated time to complete: 5 minutes
- Exchange and Skype for Business integration requires server certificates that are trusted by all of the servers involved. The procedures in this topic assume that you already have the required certificates. For more information, see Plan to integrate Skype for Business Server 2015 and Exchange. The required IM certificate thumbprint refers to the Exchange Server certificate assigned to the IIS service.
- You can only use PowerShell to perform this procedure. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Client Access virtual directory settings" entry in the Clients and mobile devices permissions topic.
- Depending on your Skype for Business Server topology, you may have multiple FrontEnd pools, you should pick the regional endpoint (closest pool to the exchange AD site): `IMServerName=<Skype Server\pool Name>`.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to configure IM integration with Outlook on the web

Step 1: Specify the IM server and IM certificate thumbprint

Use the following syntax in the Exchange Management Shell to specify the IM server and IM certificate thumbprint:

```
New-SettingOverride -Name "<UniqueOverrideName>" -Component OwaServer -Section IMSettings -Parameters
@("IMServerName=<Skype server/pool  name>","IMCertificateThumbprint=<Certificate Thumbprint>") -Reason "
<DescriptiveReason>" [-Server <ServerName>]
```

Notes:

- To configure the same settings on all Exchange 2016 and Exchange 2019 servers in the Active Directory forest, don't use the *Server* parameter.
- To configure the settings on a specific Exchange 2016 or Exchange 2019 server, use the *Server* parameter and the name of the server (don't use the fully qualified domain name or FQDN). This method is useful when you need to specify different settings on different Exchange servers.

This example specifies the IM server and IM certificate thumbprint on all Exchange 2016 and Exchange 2019 servers in the organization.

- **Setting override name**: "IM Override" (must be unique)
- **Skype for Business server name**: skype01.contoso.com
- **Certificate thumbprint**: CDF34A740E9D225A1A06193A9D44B2CE22775308
- **Override reason**: Configure IM

```
New-SettingOverride -Name "IM Override"  -Component OwaServer -Section IMSettings -Parameters
@("IMServerName=skype01.contoso.com","IMCertificateThumbprint=CDF34A740E9D225A1A06193A9D44B2CE22775308") -
Reason "Configure IM"
```

This example specifies the IM server and IM certificate thumbprint, but only on the server named Mailbox01.

```
New-SettingOverride -Name "Mailbox01 IM Override"  -Component OwaServer -Section IMSettings -Parameters
@("IMServerName=skype01.contoso.com","IMCertificateThumbprint=CDF34A740E9D225A1A06193A9D44B2CE22775308") -
Reason "Configure IM" -Server Mailbox01
```

Step 2: Refresh the IM settings on the Exchange server

Use the following syntax in the Exchange Management Shell to refresh the IM settings on the server. You need to do this on every Exchange 2016 or Exchange 2019 server that's used for Outlook on the web.

```
Get-ExchangeDiagnosticInfo -Server <ServerName> -Process Microsoft.Exchange.Directory.TopologyService -
Component VariantConfiguration -Argument Refresh
```

This example refreshes the IM settings on the server named Mailbox01.

```
Get-ExchangeDiagnosticInfo -Server Mailbox01 -Process Microsoft.Exchange.Directory.TopologyService -Component
VariantConfiguration -Argument Refresh
```

Step 3: Restart the Outlook on the web pool on the Exchange server

Run the following command in the Exchange Management Shell or in Windows PowerShell on the server. You need to do this on every Exchange 2016 or Exchange 2019 server that's used for Outlook on the web.

```
Restart-WebAppPool MSExchangeOWAAppPool
```

How do you know this worked?

You'll know that you've successfully configured IM integration with Outlook on the web when the error message goes away, and clients are able to sign in to IM.

To verify the values of the **IMServerName** and **IMCertificateThumbprint** properties on an Exchange server, replace *<ServerName>* with the name of the server (not the FQDN), and run the following command:

```
[xml]$diag=Get-ExchangeDiagnosticInfo -Server <ServerName> -Process MSExchangeMailboxAssistants -Component
VariantConfiguration -Argument "Config,Component=OwaServer";
$diag.Diagnostics.Components.VariantConfiguration.Configuration.OwaServer.IMSettings
```

Note: In Exchange 2016 CU3 or earlier, you need to use different values for some of the parameters:

- *Process*: `Microsoft.Exchange.Directory.TopologyService` (instead of `MSExchangeMailboxAssistants`).
- *Argument*: `Config` (instead of `"Config,Component=OwaServer"`).

Change the offline address book generation schedule in Exchange

8/3/2020 • 3 minutes to read • Edit Online

An offline address book (OAB) is a copy of an address book that's been downloaded so that an Outlook user can access the information it contains while disconnected from the server. By default, a new OAB is generated every 8 hours in Exchange Server 2016 and Exchange Server 2019, but you can change the interval by using the Exchange Management Shell.

For additional management tasks related to OABs, see Procedures for offline address books in Exchange Server.

What do you need to know before you begin?

- Estimated time to complete this procedure: 5 minutes.
- You can only use PowerShell to perform this procedure. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Offline address books" entry in the Email address and address book permissions topic.

> **TIP**
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Change the offline address book generation schedule

Changing the OAB generation schedule is a two-step process:

1. Change the OAB generation schedule.
2. Apply the new OAB generation schedule.

Step 1: Use the Exchange Management Shell to change the OAB generation schedule

To change the OAB generation schedule, use this syntax:

```
New-SettingOverride -Name "<UniqueOverrideName>" -Component TimeBasedAssistants -Section OABGeneratorAssistant
-Parameters @("WorkCycle=<Timespan>") -Reason "<DescriptiveReason>" [-Server <ServerName>]
```

Notes:

- To specify a *<TimeSpan>* value, use the syntax `d.hh:mm:ss`, where *d* = days, *hh* = hours, *mm* = minutes, and *ss* = seconds.
- To configure the OAB generation schedule on all Exchange 2016 and Exchange 2019 Mailbox servers in the Active Directory forest, don't use the *Server* parameter.
- To configure the OAB generation schedule on a specific Exchange 2016 or Exchange 2019 Mailbox server, use the *Server* parameter and the name (not the fully qualified domain name or FQDN) of the server. This

method is useful when you need to specify different OAB generation schedules on different Exchange servers.

- In Exchange 2016 Cumulative Update 3 (CU3) or earlier, the *Component* parameter value is `MailboxAssistants`.

This example specifies that the OAB is generated every two hours on all Exchange 2016 and Exchange 2019 servers in the organization that are responsible for generating OABs.

- **Setting override name**: "OAB Generation Override" (must be unique)
- **WorkCycle**: `02:00:00` (2 hours)
- **Override reason**: Generate OAB every 2 hours

```
New-SettingOverride -Name "OAB Generation Override" -Component TimeBasedAssistants -Section
OABGeneratorAssistant -Parameters @("WorkCycle=02:00:00") -Reason "Generate OAB every 2 hours"
```

This example specifies the same OAB generation schedule, but only on the server named Mailbox01.

```
New-SettingOverride -Name "Mailbox01 OAB Generation Override" -Component TimeBasedAssistants -Section
OABGeneratorAssistant -Parameters @("WorkCycle=02:00:00") -Reason "Generate OAB every 2 hours" -Server
Mailbox01
```

Step 2: Use the Exchange Management Shell to apply the new OAB generation schedule

To apply the new OAB generation schedule, use this syntax:

```
Get-ExchangeDiagnosticInfo -Process Microsoft.Exchange.Directory.TopologyService -Component
VariantConfiguration -Argument Refresh [-Server <ServerName>]
```

Notes:

- If you didn't use the *Server* parameter in Step 1, don't use it here. If you used the *Server* parameter in Step 1, use the same server name here.
- If you delete the custom OAB generation schedule by using the **Remove-SettingOverride** cmdlet, you still need to run this command to change the generation schedule back to the default value of 8 hours.

This example applies the new OAB generation schedule on all Exchange 2016 and Exchange 2019 Mailbox servers in the organization.

```
Get-ExchangeDiagnosticInfo -Process Microsoft.Exchange.Directory.TopologyService -Component
VariantConfiguration -Argument Refresh
```

This example applies the new OAB generation schedule on the server named Mailbox01.

```
Get-ExchangeDiagnosticInfo -Process Microsoft.Exchange.Directory.TopologyService -Component
VariantConfiguration -Argument Refresh -Server Mailbox01
```

How do you know this worked?

To verify that you've configured the OAB generation schedule on one or more Exchange servers, replace *<ServerName>* with the name of the server (not the FQDN), and run the following command to verify the value of the **WorkCycle** property:

```
[xml]$diag=Get-ExchangeDiagnosticInfo -Server <ServerName> -Process MSExchangeMailboxAssistants -Component
VariantConfiguration -Argument "Config,Component=TimeBasedAssistants";
$diag.Diagnostics.Components.VariantConfiguration.Configuration.TimeBasedAssistants.OABGeneratorAssistant
```

Note: In Exchange 2016 CU3 or earlier, you need to run this command instead:

```
[xml]$diag=Get-ExchangeDiagnosticInfo -Server <ServerName> -Process
Microsoft.Exchange.Directory.TopologyService -Component VariantConfiguration -Argument Config;
$diag.Diagnostics.Components.VariantConfiguration.Configuration.MailboxAssistants.OABGeneratorAssistant
```

.

See also

Procedures for offline address books in Exchange Server

Configure certificate based authentication in Exchange 2016

8/3/2020 • 8 minutes to read • Edit Online

Certificate based authentication (CBA) in Exchange allows Outlook on the web (formerly known as Outlook Web App) and Exchange ActiveSync clients to be authenticated by client certificates instead of entering a username and password.

Before you configure Exchange, you need to issue a client certificate to each user. Because of the sheer number of certificates involved, you should use an automated internal public key infrastructure (PKI) to issue and manage the client certificates. An example of an automated internal PKI is Active Directory Certificate Services (AD CS). For more information about AD CS, see Active Directory Certificate Services Overview. Here's more information about the certificate requirements:

- The client certificate must be issued for client authentication (for example, the default **User** certificate template in AD CS).
- The client certificate must contain the user principal name (UPN) of the user (in the certificate's **Subject** or **Subject Alternative Name** fields).
- The client certificate must be associated with the user account in Active Directory.
- All servers and devices that are involved in access to Outlook on the web and ActiveSync (including proxy servers and client devices) must trust the entire chain of trust for the client certificates (the root certificate of the certification authority, and any intermediate CAs that were used to issue certificates).

For CBA in Outlook on the web, the client certificate needs to be installed on the local computer, device, or on a smart card. For CBA in ActiveSync, the client certificate needs to be installed on the local device. You can automate the installation of certificates on devices by using a mobile device management (MDM) solution like Intune. For more information about Intune, see Overview of Microsoft Intune.

What do you need to know before you begin?

- Estimated time to complete this task: 20 minutes
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "IIS Manager" entry in the Outlook on the web permissions section of the Clients and mobile devices permissions topic.
- To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Step 1: Use the Exchange Management Shell to install the Client

Certificate Mapping Authentication feature on all of your Exchange servers

All Exchange servers that share the same namespace and URLs need to use the same authentication methods. You need to install the Client Certificate Mapping Authentication feature on all of your Exchange servers.

In the Exchange Management Shell, run the following command:

```
Install-WindowsFeature Web-Client-Auth
```

For detailed syntax and parameter information, see Install-WindowsFeature.

Step 2: Use IIS Manager to enable Active Directory Client Certificate Authentication for the Exchange server

1. Open IIS Manager on the Exchange server. An easy way to do this in Windows Server 2012 or later is to press Windows key + Q, type inetmgr, and select **Internet Information Services (IIS) Manager** in the results.

2. Select the server, and verify **Features View** is selected at the bottom of the page.

3. In the **IIS** section, double-click **Authentication**.

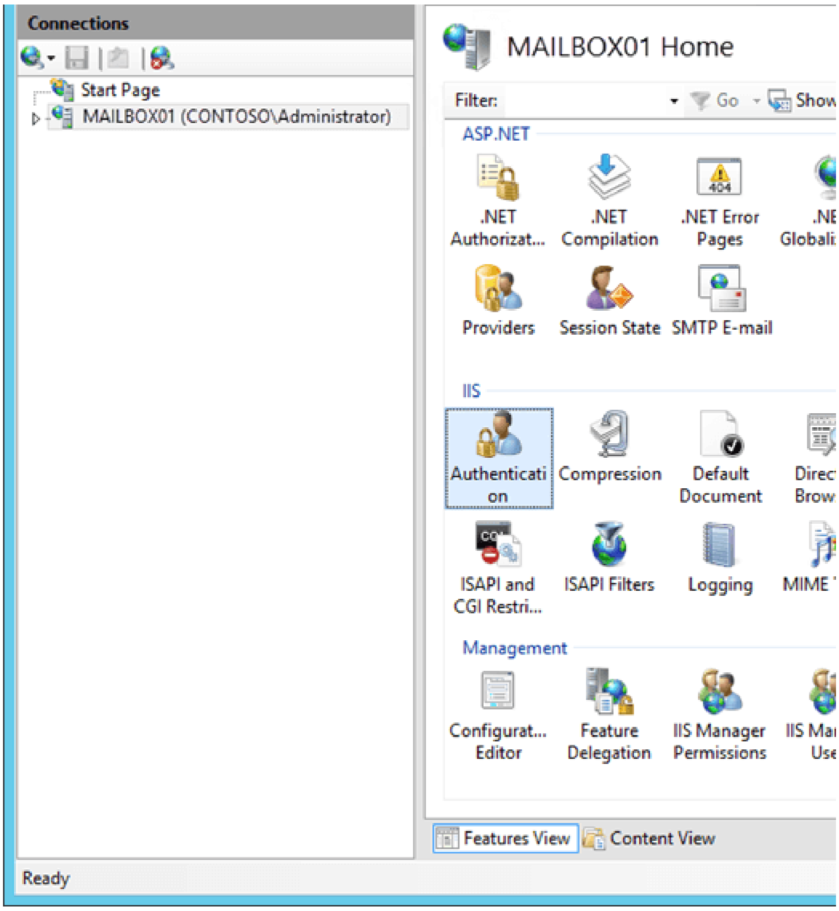

4. On the **Authentication** page that opens, select **Active Directory Client Certificate Authentication** from the list, and in the **Actions** pane, click **Enable**.

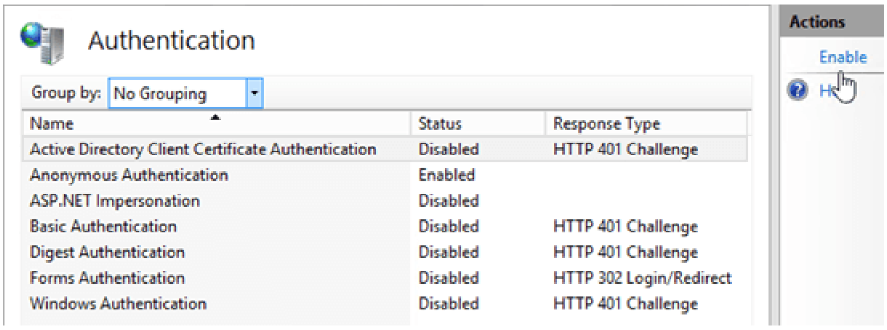

Group by: No Grouping

Name	Status	Response Type
Active Directory Client Certificate Authentication	Disabled	HTTP 401 Challenge
Anonymous Authentication	Enabled	
ASP.NET Impersonation	Disabled	
Basic Authentication	Disabled	HTTP 401 Challenge
Digest Authentication	Disabled	HTTP 401 Challenge
Forms Authentication	Disabled	HTTP 302 Login/Redirect
Windows Authentication	Disabled	HTTP 401 Challenge

You'll see a warning that SSL must be enabled to use Active Directory Client Certificate Mapping.

Step 3: Use IIS Manager to configure the Outlook on the web, Exchange admin center, and ActiveSync virtual directories to require client certificates

Note: You need to *require* client certificates, because *accepting* client certificates (to support both CBA and regular username and password authentication) doesn't work consistently across all types of ActiveSync devices.

1. In IIS Manager, expand the server, expand **Sites**, and then expand **Default Web Site**.

2. Select the **owa** virtual directory, and verify **Features View** is selected at the bottom of the page.

3. In the **IIS** section, double-click **SSL Settings**.

4. On the **SSL Settings** page, verify **Require SSL** is checked, and select the **Client certificates** value **Require**.

5. In the **Actions** pane, click **Apply**.

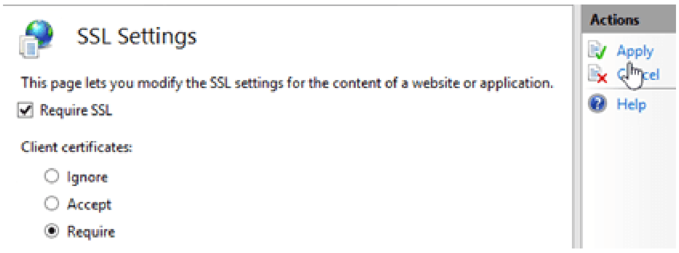

This page lets you modify the SSL settings for the content of a website or application.

☑ Require SSL

Client certificates:

○ Ignore

○ Accept

● Require

6. Select the **Microsoft-Server-ActiveSync** virtual directory.

7. In the **IIS** section, double-click **SSL Settings**.

8. On the **SSL Settings** page, verify **Require SSL** is checked, and select the **Client certificates** value **Require**.

9. In the **Actions** pane, click **Apply**.

Note: Although you can perform these procedures on the command line, the steps might not configure a required registry key. You can use the earlier procedures in IIS Manager (which will definitely set the registry key correctly), or you need to verify that the registry key `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\HTTP\Parameters\SslBindingInfo\0.0.0.0:443` is set to the value `1` after you perform the procedures on the command line.

To perform these procedures on the command line, open an elevated command prompt on the Exchange server (a Command Prompt window you open by selecting **Run as administrator**) and run the following commands:

```
%windir%\system32\inetsrv\appcmd.exe set config "Default Web Site/owa/" -
section:system.webserver/security/access /sslFlags:"Ssl, SslRequireCert" /commit:apphost
```

```
%windir%\system32\inetsrv\appcmd.exe set config "Default Web Site/ecp/" -
section:system.webserver/security/access /sslFlags:"Ssl, SslRequireCert" /commit:apphost
```

```
%windir%\system32\inetsrv\appcmd.exe set config "Default Web Site/Microsoft-Server-ActiveSync/" -
section:system.webserver/security/access /sslFlags:"Ssl, SslRequireCert" /commit:apphost
```

Step 4: Use the Exchange Management Shell to disable authentication other authentication methods on the Outlook on the web, Exchange admin center, and ActiveSync virtual directories

After you require client certificates for authentication, you need to disable all other authentication methods on the Outlook on the web, Exchange admin center (EAC) and ActiveSync virtual directories. By default, only Basic authentication and Forms authentication are enabled.

1. In the Exchange Management Shell, replace *<ServerName>* with the name of your Exchange server, and run the following command to disable all other authentication methods on the Outlook on the web virtual directory:

   ```
   Set-OwaVirtualDirectory "<ServerName>\owa (Default Web Site)" -BasicAuthentication $false -
   WindowsAuthentication $false -DigestAuthentication $false -FormsAuthentication $false -
   AdfsAuthentication $false -OAuthAuthentication $false
   ```

   For detailed syntax and parameter information, see Set-OwaVirtualDirectory.

2. In the Exchange Management Shell, replace *<ServerName>* with the name of your Exchange server, and run the following command to disable all other authentication methods on the EAC virtual directory:

   ```
   Set-EcpVirtualDirectory "<ServerName>\ecp (Default Web Site)" -BasicAuthentication $false -
   WindowsAuthentication $false -DigestAuthentication $false -FormsAuthentication $false -
   AdfsAuthentication $false
   ```

   For detailed syntax and parameter information, see Set-EcpVirtualDirectory.

3. Replace *<ServerName>* with the name of your Exchange server, and run the following command to disable all other authentication methods on the ActiveSync virtual directory:

   ```
   Set-ActiveSyncVirtualDirectory "<ServerName>\Microsoft-Server-ActiveSync (Default Web Site)" -
   BasicAuthEnabled $false -WindowsAuthEnabled $false
   ```

   For detailed syntax and parameter information, see Set-ActiveSyncVirtualDirectory.

Step 5: Use IIS Manager to enable client certificate mapping for the Outlook on the web, Exchange admin center, and ActiveSync virtual directories

> **IMPORTANT**
>
> After you perform this step, running the **Set-ActiveSyncVirtualDirectory** cmdlet might disable the client certificate mapping for ActiveSync.

1. In IIS Manager, expand the server, expand **Sites**, and then expand **Default Web Site**.

2. Select the **owa** virtual directory, and verify **Features View** is selected at the bottom of the page.

3. In the **Management** section, double-click **Configuration Editor**.

4. On the **Configuration Editor** page, click the drop down on **Section**, and navigate to **system.webServer** > **security** > **authentication** > **clientCertificateMappingAuthentication**.

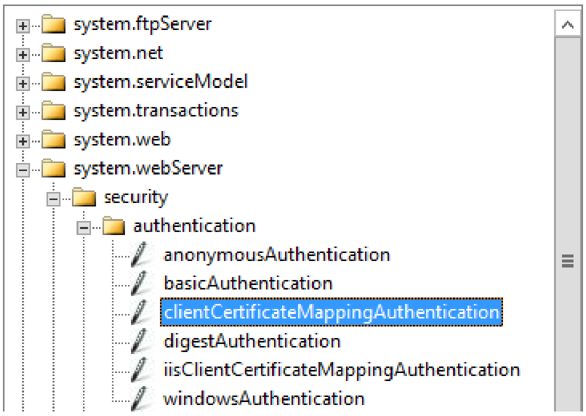

5. Set the **enabled** value to **True**, and in the **Actions** pane, click **Apply**.

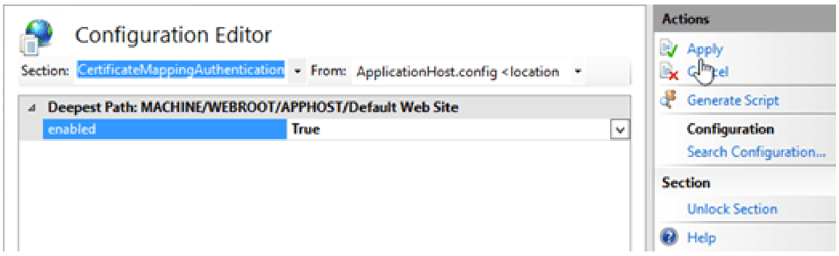
Configuration Editor

Section: CertificateMappingAuthentication From: ApplicationHost.config <location

Deepest Path: MACHINE/WEBROOT/APPHOST/Default Web Site

enabled True

6. Select the **ecp** virtual directory.

7. In the **Management** section, double-click **Configuration Editor**.

8. On the **Configuration Editor** page, click the drop down on **Section**, and navigate to **system.webServer** > **security** > **authentication** > **clientCertificateMappingAuthentication**.

9. Set the **enabled** value to **True**, and in the **Actions** pane, click **Apply**.

10. Select the **Microsoft-Server-ActiveSync** virtual directory.

11. In the **Management** section, double-click **Configuration Editor**.

12. On the **Configuration Editor** page, click the drop down on **Section**, and navigate to **system.webServer** > **security** > **authentication** > **clientCertificateMappingAuthentication**.

13. Set the **enabled** value to **True**, and in the **Actions** pane, click **Apply**.

Note: To perform these procedures on the command line, open an elevated command prompt on the Exchange server and run the following commands:

```
%windir%\system32\inetsrv\appcmd.exe set config "Default Web Site/owa/" -
section:system.webserver/security/authentication/clientCertificateMappingAuthentication /enabled:"True"
/commit:apphost
```

```
%windir%\system32\inetsrv\appcmd.exe set config "Default Web Site/ecp/" -
section:system.webserver/security/authentication/clientCertificateMappingAuthentication /enabled:"True"
/commit:apphost
```

```
%windir%\system32\inetsrv\appcmd.exe set config "Default Web Site/Microsoft-Server-ActiveSync/" -
section:system.webserver/security/authentication/clientCertificateMappingAuthentication /enabled:"True"
/commit:apphost
```

Step 6 (Optional): Add the root certificate of a third-party certification authority to the Enterprise NTAuth store in Active Directory

You only need to perform this step if you aren't using AD CS to issue the client certificates. This setting indicates that the certification authority (CA) is trusted to issue client certificates for Active Directory authentication.

1. Export the CA's root certificate to a Base-64 encoded or DER binary encoded X.509 .cer file. In this example, we'll use C:\Data\CARoot.cer.

2. On any domain member server (for example, a domain controller or an Exchange server), open an elevated command prompt run the following command:

   ```
   %windir%\system32\certutil.exe -enterprise -addstore NTAuth "C:\Data\CARoot.cer"
   ```

   Note that this step requires membership in the **Enterprise Admins** group.

Step 7 (Optional): Use IIS Manager to increase the UploadReadAheadSize value for the Outlook on the web and ActiveSync virtual directories

If your clients receive errors, you might need to increase the **uploadReadAheadSize** values in the IIS metabase to allow for the request headers.

1. In IIS Manager, expand the server, expand **Sites**, and then expand **Default Web Site**.

2. Select the **owa** virtual directory, and verify **Features View** is selected at the bottom of the page.

3. In the **Management** section, double-click **Configuration Editor**.

4. On the **Configuration Editor** page, click the drop down on **Section**, and navigate to **systemwebServer** > **serverRuntime**.

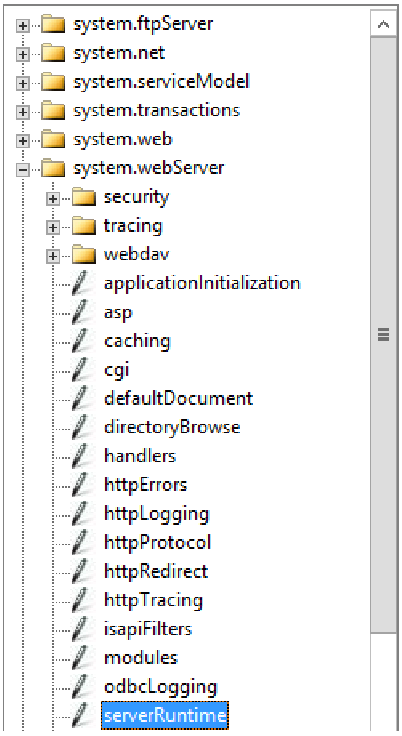

5. Set the **uploadReadAheadSize** value to 49152, and in the **Actions** pane, click **Apply**.

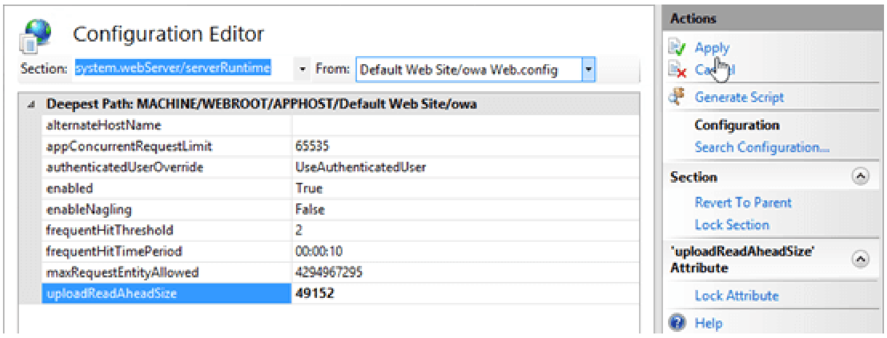

Section: system.webServer/serverRuntime From: Default Web Site/owa Web.config

Deepest Path: MACHINE/WEBROOT/APPHOST/Default Web Site/owa

alternateHostName	
appConcurrentRequestLimit	65535
authenticatedUserOverride	UseAuthenticatedUser
enabled	True
enableNagling	False
frequentHitThreshold	2
frequentHitTimePeriod	00:00:10
maxRequestEntityAllowed	4294967295
uploadReadAheadSize	49152

6. Select the **ecp** virtual directory.

7. In the **Management** section, double-click **Configuration Editor**.

8. On the **Configuration Editor** page, click the drop down on **Section**, and navigate to **systemwebServer** > **serverRuntime**.

9. Set the **uploadReadAheadSize** value to 49152, and in the **Actions** pane, click **Apply**.

10. Select the **Microsoft-Server-ActiveSync** virtual directory.

11. In the **Management** section, double-click **Configuration Editor**.

12. On the **Configuration Editor** page, click the drop down on **Section**, and navigate to **systemwebServer** > **serverRuntime**.

13. Set the **uploadReadAheadSize** value to 49152, and in the **Actions** pane, click **Apply**.

Note: To perform these procedures on the command line, open an elevated command prompt on the Exchange server and run the following commands:

```
%windir%\system32\inetsrv\appcmd.exe set config "Default Web Site/owa/" -
section:system.webserver/serverRuntime /uploadReadAheadSize:49152
```

```
%windir%\system32\inetsrv\appcmd.exe set config "Default Web Site/ecp/" -
section:system.webserver/serverRuntime /uploadReadAheadSize:49152
```

```
%windir%\system32\inetsrv\appcmd.exe set config "Default Web Site/Microsoft-Server-ActiveSync/" -
section:system.webserver/serverRuntime /uploadReadAheadSize:49152
```

Deployment reference

8/3/2020 • 2 minutes to read • Edit Online

Exchange Server readiness checks

Exchange Server editions and versions

Exchange Server language support

Exchange Server storage configuration options

Network ports for clients and mail flow in Exchange

Overview of Exchange services on Exchange servers

Exchange 2019 preferred architecture

Exchange Server readiness checks

8/3/2020 • 2 minutes to read • Edit Online

The topics in this are provide details about the readiness checks that Exchange Server performs when Exchange is installed. Readiness checks ensure that your Active Directory forest and Exchange servers are ready for the version of Exchange that you're installing. Each readiness check topic describes the actions that you can take to resolve issues that are found when the readiness checks are run. You should only perform the steps outlined in a readiness check topic if that readiness check was displayed during setup.

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

AD LDS directory exists in default location [ADAMDataPathExists]

8/3/2020 • 2 minutes to read • Edit Online

Exchange Setup can't continue because the attempt to install Active Directory Lightweight Directory Services (AD LDS) failed.

An older installation of AD LDS exists in the default location. Setup can't perform a new AD LDS install in an existing AD LDS directory structure.

To resolve this issue, remove the existing AD LDS directory and then run Setup again.

For more information about AD LDS directory management, see Administering AD LDS Directory Partitions.

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Duplicate Microsoft Exchange System Objects container exists in Active Directory [AdInitErrorRule]

8/3/2020 • 2 minutes to read • Edit Online

Exchange Setup can't continue because it found a duplicate Microsoft Exchange System Objects container in Active Directory Domain Naming context. When Setup finds a duplicate Microsoft Exchange System Objects container, you need to delete the duplicate container before Setup can continue. Note that running **DomainPrep** again won't fix the problem. You need to find and delete the duplicate Microsoft Exchange System Objects container.

To resolve this issue, do the following steps:

1. Open **Active Directory Users and Computers**. For example:

   - Press Windows key + R, enter **dsc.msc**, and then click **OK**.

   - In **Administrative Tools** > **Active Directory Users and Computers**.

2. In the **Active Directory Users and Computers**, click **View** > **Advanced Features**.

3. Locate the duplicate **Microsoft Exchange System Objects** container.

4. Verify that the duplicate **Microsoft Exchange System Objects** container doesn't contain valid Active Directory objects.

5. Right-click the duplicate **Microsoft Exchange System Objects container**, click **Delete**, and then click **Yes** in the confirmation dialog box.

> **NOTE**
>
> To immediately replicate the change, you need to manually initiate replication between domain controllers.

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Failover Cluster Command Interface Windows feature not installed [RsatClusteringCmdInterfaceInstalled]

8/3/2020 • 2 minutes to read • Edit Online

Microsoft Exchange Server 2016 Setup can't continue because the local computer is missing a required Windows feature. You'll need to install this Windows feature before Exchange 2016 can continue.

Exchange 2016 Setup requires that the **Failover Cluster Command Interface** Windows feature be installed on the computer before installation can continue.

Do the following to install the Windows feature on this computer. If the feature requires a reboot to complete installation, you'll need to exit Exchange 2016 Setup, reboot, and then start Setup again.

> **NOTE**
>
> Additional Windows features or updates might need to be installed before Exchange 2016 Setup can continue. For a complete list of required Windows features and updates, check out Exchange Server prerequisites.

1. Open Windows PowerShell on the local computer.
2. Run the following command to install the required Windows feature.

```
Install-WindowsFeature RSAT-Clustering-CmdInterface
```

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Active Directory does not exist or cannot be contacted [CannotAccessAD]

8/3/2020 • 2 minutes to read • Edit Online

Exchange Setup can't continue because it can't contact a valid Active Directory site. Setup requires that the target server is able to locate the configuration naming context in Active Directory.

To resolve this issue, verify that the account that you're using an Active Directory account to run Setup and then try running Setup again. If this doesn't resolve the issue, follow the guidance about using the support tools in the following topics to further diagnose the problem.

For more information about Active Directory troubleshooting and configuration for Exchange, see the following topics:

- Prepare Active Directory and domains
- Troubleshooting Active Directory Domain Services
- Configuring a Computer for Troubleshooting
- Troubleshooting Active Directory Replication Problems

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

The local computer isn't joined to an Active Directory domain [ComputerNotPartofDomain]

8/3/2020 • 2 minutes to read • Edit Online

Exchange Setup can't continue because it detected that the target server isn't a member of an Active Directory domain. You need to join the target server to an Active Directory domain before you can install the Mailbox server role. You might also see this message if you're using a local computer account instead of a domain user account (with the required permissions) to install Exchange.

For more information, see Exchange Server system requirements

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Installation of the first Exchange server in the organization can't be delegated [DelegatedBridgeheadFirstInstall]

8/3/2020 • 2 minutes to read • Edit Online

Exchange Setup can't continue because this is the first Exchange server in the organization, and the first Exchange server needs to be installed by a member of the Enterprise Admins security group (to create the Exchange Organization container and configure objects in it).

Note: If you haven't already extended the Active Directory schema for Exchange, you need to do one of the following steps:

- A member of the Schema Admins group can extend the Active Directory schema using another computer in the domain before you install Exchange.
- Exchange Setup can extend the schema if your account is a member of the Schema Admins group.

To resolve this issue, run Exchange setup again using an account that's a member of the Enterprise Admins security group (add the current account or use a different account).

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Installation of the first Exchange server in the organization can't be delegated [DelegatedCafeFirstInstall]

8/3/2020 • 2 minutes to read • Edit Online

Exchange Setup can't continue because this is the first Exchange server in the organization, and the first Exchange server needs to be installed by a member of the Enterprise Admins security group (to create the Exchange Organization container and configure objects in it).

Note: If you haven't already extended the Active Directory schema for Exchange, you need to do one of the following steps:

- A member of the Schema Admins group can extend the Active Directory schema using another computer in the domain before you install Exchange.
- Exchange Setup can extend the schema if your account is a member of the Schema Admins group.

To resolve this issue, run Exchange setup again using an account that's a member of the Enterprise Admins security group (add the current account or use a different account).

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Installation of the first Exchange server in the organization can't be delegated [DelegatedClientAccessFirstInstall]

8/3/2020 • 2 minutes to read • Edit Online

Exchange Setup can't continue because this is the first Exchange server in the organization, and the first Exchange server needs to be installed by a member of the Enterprise Admins security group (to create the Exchange Organization container and configure objects in it).

Note: If you haven't already extended the Active Directory schema for Exchange, you need to do one of the following steps:

- A member of the Schema Admins group can extend the Active Directory schema using another computer in the domain before you install Exchange.
- Exchange Setup can extend the schema if your account is a member of the Schema Admins group.

To resolve this issue, run Exchange setup again using an account that's a member of the Enterprise Admins security group (add the current account or use a different account).

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Installation of the first Exchange server in the organization can't be delegated [DelegatedMailboxFirstInstall]

8/3/2020 • 2 minutes to read • Edit Online

Exchange Setup can't continue because this is the first Exchange server in the organization, and the first Exchange server needs to be installed by a member of the Enterprise Admins security group (to create the Exchange Organization container and configure objects in it).

Note: If you haven't already extended the Active Directory schema for Exchange, you need to do one of the following steps:

- A member of the Schema Admins group can extend the Active Directory schema using another computer in the domain before you install Exchange.
- Exchange Setup can extend the schema if your account is a member of the Schema Admins group.

To resolve this issue, run Exchange setup again using an account that's a member of the Enterprise Admins security group (add the current account or use a different account).

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Installation of the first Exchange server in the organization can't be delegated [DelegatedUnifiedMessagingFirstInstall]

8/3/2020 • 2 minutes to read • Edit Online

Exchange 2016 Setup can't continue because this is the first Exchange server in the organization, and the first Exchange server needs to be installed by a member of the Enterprise Admins security group (to create the Exchange Organization container and configure objects in it).

Note: If you haven't already extended the Active Directory schema for Exchange, you need to do one of the following steps:

- A member of the Schema Admins group can extend the Active Directory schema using another computer in the domain before you install Exchange.
- Exchange Setup can extend the schema if your account is a member of the Schema Admins group.

To resolve this issue, run Exchange setup again using an account that's a member of the Enterprise Admins security group (add the current account or use a different account).

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Active Directory functional level isn't Windows Server 2003 or later [ForestLevelNotWin2003Native]

8/3/2020 • 2 minutes to read • Edit Online

Exchange Server 2016 Setup can't continue because the Active Directory forest functional level of the target forest isn't Windows Server 2003 native or later. Before you can install Exchange 2016, you must raise the forest functional level to Windows Server 2003 or later.

For information about how to raise the forest functional level, see Raise the Forest Functional Level.

For more information about Active Directory functional levels, see the following topics:

- What are Active Directory Functional Levels?
- How Active Directory Functional Levels Work

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Cannot write to the Exchange organization container [GlobalServerInstall]

8/3/2020 • 2 minutes to read • Edit Online

Exchange Setup can't continue because the user account doesn't have the permissions that are required to write to the organization container in the Active Directory directory service.

Setup requires that the account you're using to install Exchange has permissions to create and modify objects in Active Directory:

- If this is the first Exchange server in your organizaiton, your account needs to be a member of the Schema Admins security group (to extend the schema) and the Enterprise Admins security group (to prepare Active Directory).
- After you prepare Active Directory for the version of Exchange that you're installing, your account needs to be a member of the Organization Management role group.

For more information, see Prepare Active Directory and domains for Exchange.

To resolve this issue, run Exchange setup again using an account that has the appropriate permissions (grant permissions to the current account or use a different account).

> **IMPORTANT**
>
> Cross-forest installation of Exchange isn't supported. Use an account in the Active Directory forest where you're installing Exchange.

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Global updates required [GlobalUpdateRequired]

8/3/2020 • 2 minutes to read • Edit Online

Exchange Setup can't continue because the user account doesn't have the permissions that are required to write to the organization container in the Active Directory directory service.

Setup requires that the account you're using to install Exchange has permissions to create and modify objects in Active Directory:

- If this is the first Exchange server in your organizaiton, your account needs to be a member of the Schema Admins security group (to extend the schema) and the Enterprise Admins security group (to prepare Active Directory).
- After you prepare Active Directory for the version of Exchange that you're installing, your account needs to be a member of the Organization Management role group.

For more information, see Prepare Active Directory and domains for Exchange.

To resolve this issue, run Setup again using an account that has the appropriate permissions (grant permissions to the current account or use a different account).

> **IMPORTANT**
> Cross-forest installation of Exchange isn't supported. Use an account in the Active Directory forest where you're installing Exchange.

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

The Host record for the local computer cannot be found in the DNS database [HostRecordMissing]

8/3/2020 • 2 minutes to read • Edit Online

Exchange Setup can't continue because the Host (A) record for this computer can't be found in the DNS zone for the domain. Setup requires a valid A record for the server, and Exchange uses email server A records to find the IP address of the next hop to send messages.

To resolve this issue:

- Verify that the local TCP/IP configuration points to the correct DNS server. For more information, see Configure TCP/IP settings.
- Use Nslookup.exe to verify that the Host (A) record exists on the DNS server. For more information, see To verify A resource records exist in DNS.

For information about DNS name resolution, troubleshooting, and A records, see the following:

- Troubleshooting DNS
- Managing resource records

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Installation on domain controllers is not supported with Active Directory split permissions [InstallOnDCInADSplitPermissionMode]

8/3/2020 • 2 minutes to read • Edit Online

Exchange Setup detected that you're installing Exchange on an Active Directory domain controller and one of the following conditions is true:

- The Exchange organization is already configured for Active Directory split permissions.
- You selected the Active Directory split permissions option in Exchange Setup.

Installing Exchange on domain controllers isn't supported when the Exchange organization is configured for Active Directory split permissions. To install Exchange on a domain controller, you need to configure the Exchange organization for Role Based Access Control (RBAC) split permissions or shared permissions.

> **IMPORTANT**
>
> We don't recommend installing Exchange on Active Directory domain controllers. For more information, see Installing Exchange on a domain controller is not recommended [WarningInstallExchangeRolesOnDomainController].

If you want to use Active Directory split permissions, you need install Exchange on a member server.

For more information about split and shared permissions in Exchange 2013 or later, see the following topics:

- Understanding Split Permissions
- Configure Exchange 2013 for Split Permissions
- Configure Exchange 2013 for Shared Permissions

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

The current account isn't logged into an Active Directory domain [LoggedOntoDomain]

8/3/2020 • 2 minutes to read • Edit Online

Exchange Setup can't continue because it detected that the current account isn't logged on to an Active Directory domain. You need to log in using an Active Directory account that has the permissions required to install Exchange.

Setup requires that the account you're using to install Exchange has permissions to create and modify objects in Active Directory:

- If this is the first Exchange server in your organizaiton, your account needs to be a member of the Schema Admins security group (to extend the schema) and the Enterprise Admins security group (to prepare Active Directory).

- After you prepare Active Directory for the version of Exchange that you're installing, your account needs to be a member of the Organization Management role group.

For more information, see Prepare Active Directory and domains for Exchange.

To resolve this issue, run Setup again using an account that has the appropriate permissions (grant permissions to the current account or use a different account).

> **IMPORTANT**
>
> Cross-forest installation of Exchange isn't supported. Use an account in the Active Directory forest where you're installing Exchange.

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

The computer needs to be restarted before Setup can continue [RebootPending]

8/3/2020 • 2 minutes to read • Edit Online

Exchange Setup can't continue because it detected a pending reboot to complete the installation of other programs or Windows updates.

Why is this happening?

When programs and Windows updates are installed, they make changes to files that are stored on your computer. Some programs or updates need to modify or replace files that are currently in use. When this happens, you need to restart the computer before other programs can be installed.

If the installation of a previous program or Windows update didn't complete successfully, Windows and other programs might think a restart is required. You'll continue to see this error each time you run Exchange Setup if this happens (the failed installation can't fix the condition that indicates a restart is required).

How do I fix it?

Typically, you only need to restart the server to get past this error, but you might get this error again after a restart (for example, additional program or Windows updates also require a restart). Try restarting the server again.

If you see this error after you've restarted the server more than two or three times, try reinstalling any programs or Windows updates that you've installed recently. This might allow a failed installation to complete successfully.

If you *still* receive this error after multiple restarts and reinstalling recent programs or Windows updates, we recommend that you contact Microsoft Customer Service and Support. They'll help you find the reason why Windows and other programs think your computer needs to be restarted. To contact Microsoft support, go to Support for business and select **Servers** > **Exchange Server**.

Caution

Although it's tempting, we strongly recommend that you don't attempt to work around this issue by manually deleting or changing registry keys or values. Although you might fix this issue now, manually modifying the registry might cause issues later on. This is especially important if the failed installation was a Windows update.

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

The logged-on user is not a member of the Schema Admins group [SchemaUpdateRequired]

8/3/2020 • 2 minutes to read • Edit Online

Exchange Setup can't continue because the user account isn't a member of the Schema Admins and Enterprise Admins security groups.

Setup requires that the account you're using to install Exchange has permissions to create and modify objects in Active Directory:

- If this is the first Exchange server in your organizaiton, your account needs to be a member of the Schema Admins security group (to extend the schema) and the Enterprise Admins security group (to prepare Active Directory).
- After you prepare Active Directory for the version of Exchange that you're installing, your account needs to be a member of the Organization Management role group.

For more information, see Prepare Active Directory and domains for Exchange.

To resolve this issue, run Exchange setup again using an account that has the appropriate permissions (grant permissions to the current account or use a different account).

> **IMPORTANT**
>
> Cross-forest installation of Exchange isn't supported. Use an account in the Active Directory forest where you're installing Exchange.

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

UCMA 4.0, Core Runtime not installed [UcmaRedistMsi]

8/3/2020 • 2 minutes to read • Edit Online

Exchange 2016 Setup requires the Unified Communications Managed API 4.0 Runtime for Unified Messaging (UM) services on the Mailbox server role. You need to install this update before Exchange 2016 Setup can continue.

Download and install the 64-bit update from Unified Communications Managed API 4.0 Runtime, and then click **Retry** on the **Readiness Checks** page in the Exchange 2016 Setup wizard.

> **NOTE**
>
> If the installation of this update requires a reboot, you'll need to exit Exchange 2016 Setup, reboot, and then start Setup again.

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Cannot remove mailbox database [UnwillingToRemoveMailboxDatabase]

8/3/2020 • 2 minutes to read • Edit Online

Exchange Setup can't continue because it can't remove a user mailbox database from the local server without incurring potential data loss.

Before Exchange Setup removes the Mailbox server role from a server, Setup confirms that all mailbox databases have been removed from the server, or that the mailboxes on the server don't contain active mailboxes. However, user mailboxes might still remain on the server.

To resolve this issue do either of these steps:

- To preserve the mailboxes and their content, move the mailboxes to another server. For instructions, see Mailbox moves in Exchange Server.
- Disable the mailboxes if they're no longer required. For more information, see Disable-Mailbox.
- Remove the mailbox databases if they're no longer required. For instructions, see Manage mailbox databases in Exchange Server.

After you deal with the mailbox databases on the server, run Exchange Setup again.

- For more information about how to identify a mailbox in the database, see Get-Mailbox.

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Installing Exchange on a domain controller is not recommended [WarningInstallExchangeRolesOnDomainController]

8/3/2020 • 2 minutes to read • Edit Online

Exchange Server 2016 or Exchange 2019 Setup has detected that the target computer is an Active Directory domain controller, and we don't recommed installing Exchange on domain controllers.

If you install Exchange on a domain controller, be aware of the following issues:

- Configuring Exchange for Active Directory split permissions isn't supported. For more information about split permissions, see Understanding split permissions.
- The Exchange Trusted Subsystem universal security group (USG) is added to the Domain Admins group. This action grants all Exchange servers domain administrator rights in the domain.
- Exchange Server and Active Directory are both resource-intensive applications. There are performance implications when both applications are running on the same computer.
- The domain controller must be a global catalog server, but Exchange services might not start correctly on a global catalog server.
- System shutdown will take considerably longer if Exchange you don't stop the Exchange services before you shut down or restart the server.
- Demoting the domain controller to a member server isn't supported.
- Running Exchange on a clustered node that's also an Active Directory domain controller isn't supported.

Therefore, we recommend that you install Exchange on a member server, not on a domain controller.

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

KB2619234 update not installed [Win7RpcHttpAssocCookieGuidUpdateNotInstalled]

8/3/2020 • 2 minutes to read • Edit Online

Exchange Server 2016 Setup can't continue because the local computer requires a software update. You'll need to install this update before Exchange 2016 Setup can continue.

Exchange 2016 Setup requires a Windows Server update that allows Outlook Anywhere (RPC over HTTP) to work correctly.

Download and install the 64-bit update from KB2619234, and then click **retry** on the **Readiness Checks** page.

> **NOTE**
>
> If this update requires a reboot to complete installation, you'll need to exit Exchange 2016 Setup, reboot, and then start Setup again.

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Installation of the first Exchange server in the organization can't be delegated [DelegatedFrontendTransportFirstInstall]

8/3/2020 • 2 minutes to read • Edit Online

Exchange Setup can't continue because this is the first Exchange server in the organization, and the first Exchange server needs to be installed by a member of the Enterprise Admins security group (to create the Exchange Organization container and configure objects in it).

Note: If you haven't already extended the Active Directory schema for Exchange, you need to do one of the following steps:

- A member of the Schema Admins group can extend the Active Directory schema using another computer in the domain before you install Exchange.
- Exchange Setup can extend the schema if your account is a member of the Schema Admins group.

To resolve this issue, run Exchange setup again using an account that's a member of the Enterprise Admins security group (add the current account or use a different account).

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

No Exchange 2010 servers detected [NoE14ServerWarning]

8/3/2020 • 2 minutes to read • Edit Online

Microsoft Exchange Server 2016 Setup displayed this warning because no Exchange 2010 servers exist in the organization.

Caution

If you continue with Exchange Server 2016 installation, you won't be able to add Exchange 2010 servers to the organization in the future.

Before deploying Exchange 2016, consider the following factors that may require you to deploy Exchange 2010 servers prior to deploying Exchange 2016:

- **Third-party or in-house developed applications**: Applications developed for earlier versions of Exchange may not be compatible with Exchange 2016. You may need to maintain Exchange 2010 servers to support these applications.
- **Coexistence or migration requirements**: If you plan on migrating mailboxes into your organization, some solutions may require the use of Exchange 2010 servers.

If you decide that you need to deploy Exchange 2010 servers, you need to do so before you deploy Exchange 2016. You need to prepare Active Directory for each Exchange version in the following order:

1. Exchange 2010
2. Exchange 2013 (only required if you're planning to deploy Exchange 2013 at a later date)
3. Exchange 2016

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

The computer needs to be restarted before Setup can continue [PendingRebootWindowsComponents]

8/3/2020 • 2 minutes to read • Edit Online

Exchange Setup can't continue because it detected a pending reboot to complete the installation of other programs or Windows updates.

Why is this happening?

When programs and Windows updates are installed, they make changes to files that are stored on your computer. Some programs or updates need to modify or replace files that are currently in use. When this happens, you need to restart the computer before other programs can be installed.

If the installation of a previous program or Windows update didn't complete successfully, Windows and other programs might think a restart is required. You'll continue to see this error each time you run Exchange Setup if this happens (the failed installation can't fix the condition that indicates a restart is required).

How do I fix it?

Typically, you only need to restart the server to get past this error, but you might get this error again after a restart (for example, additional program or Windows updates also require a restart). Try restarting the server again.

If you see this error after you've restarted the server more than two or three times, try reinstalling any programs or Windows updates that you've installed recently. This might allow a failed installation to complete successfully.

If you *still* receive this error after multiple restarts and reinstalling recent programs or Windows updates, we recommend that you contact Microsoft Customer Service and Support. They'll help you find the reason why Windows and other programs think your computer needs to be restarted. To contact Microsoft support, go to Support for business and select **Servers** > **Exchange Server**.

Caution

Although it's tempting, we strongly recommend that you don't attempt to work around this issue by manually deleting or changing registry keys or values. Although you might fix this issue now, manually modifying the registry might cause issues later on. This is especially important if the failed installation was a Windows update.

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Can't install Exchange 2016 in a forest that contains Exchange 2000 or Exchange 2003 servers. [Exchange2000or2003PresentInOrg]

8/3/2020 • 2 minutes to read • Edit Online

Exchange Setup can't continue because a version of Exchange that's too old for coexistence with the version that you're installing was found in the Active Directory forest. Before you can continue, you need to eliminate all unsupported versions of Exchange from your forest, which might require that you to upgrade to an interim version of Exchange first.

The installation of Exchange Server 2016 or later can't continue because Setup found one or more Exchange 2000 or Exchange 2003 servers in the Active Directory forest. Before you can install Exchange 2016 or later in your organization, you need to remove all Exchange 2000 or Exchange 2003 servers from the forest.

The upgrade path that you need to follow depends on your current version of Exchange. The upgrade paths are described in the following table:

> **NOTE**
>
> When you need to upgrade to an interim version of Exchange, you need to migrate all mailboxes, public folders and other components onto the interim version of Exchange before you decommission and remove the earlier Exchange servers.

CURRENT EXCHANGE VERSION	LATEST EXCHANGE VERSION FOR COEXISTENCE	UPGRADE PATH SUMMARY
Exchange 2000	Exchange 2007	Exchange 2000 > Exchange 2007 > Exchange 2013 > Exchange 2019.
Exchange 2003	Exchange 2010	Exchange 2003 > Exchange 2010 > Exchange 2016 > Exchange 2019.
Exchange 2007	Exchange 2013	Exchange 2007 > Exchange 2013 > Exchange 2019.
Exchange 2010	Exchange 2016	Exchange 2010 > Exchange 2016 > Exchange 2019.

CURRENT EXCHANGE VERSION	LATEST EXCHANGE VERSION FOR COEXISTENCE	UPGRADE PATH SUMMARY
Exchange 2000	Exchange 2007	Exchange 2000 > Exchange 2007 > Exchange 2013 > Exchange 2016.
Exchange 2003	Exchange 2010	Exchange 2003 > Exchange 2010 > Exchange 2016.
Exchange 2007	Exchange 2013	Exchange 2007 > Exchange 2013 > Exchange 2016.

CURRENT EXCHANGE VERSION	LATEST EXCHANGE VERSION FOR COEXISTENCE	UPGRADE PATH SUMMARY
Exchange 2010	Exchange 2016	Exchange 2010 > Exchange 2016.

When upgrading to Exchange 2013 or later, you can use the Exchange Deployment Assistant to help complete your deployment. For more information, see Exchange Deployment Assistant.

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

An incompatible operating system was found [ValidOSVersion]

8/3/2020 • 2 minutes to read • Edit Online

Microsoft Exchange Server 2016 Setup can't continue because it detected an incompatible operating system. You must install a compatible operating system on this computer before you install Exchange 2016. The following table shows the operating systems that are compatible with Exchange 2016.

IMPORTANT

Exchange 2016 doesn't support the Server Core installation option of Windows Server.

Supported operating systems for Exchange 2016

COMPONENT	REQUIREMENT
Mailbox and Edge Transport server roles	Windows Server 2016 Standard or Datacenter* Windows Server 2012 R2 Standard or Datacenter Windows Server 2012 Standard or Datacenter
Management tools	One of the following: • Windows Server 2016 Standard or Datacenter* • Windows Server 2012 R2 Standard or Datacenter • Windows Server 2012 Standard or Datacenter • 64-bit edition of Windows 10 • 64-bit edition of Windows 8.1

* Requires Exchange Server 2016 Cumulative Update 3 or later.

For more information, see Exchange Server system requirements.

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

ExecutionPolicy GPO is defined [PowerShellExecutionPolicyCheckSet]

8/3/2020 • 2 minutes to read • Edit Online

Exchange Setup can't continue because it detected that the **ExecutionPolicy** Group Policy Object (GPO) defines one or both of the following policies:

- **MachinePolicy**
- **UserPolicy**

It doesn't matter how the policies have been defined; it only matters that they have been defined.

Exchange Setup stops and disables the Windows Management Instrumentation (WMI) service. When either of these policies are defined, the WMI service needs to be enabled to run a Windows PowerShell script. If the policies are defined and the WMI service is stopped, Setup will fail and the server will be left in an inconsistent state.

To allow Setup to continue, you need to temporarily remove any definition of **MachinePolicy** or **UserPolicy** in the **ExecutionPolicy** GPO:

```
Set-ItemProperty -Path HKLM:\SOFTWARE\Policies\Microsoft\Windows\PowerShell -Name ExecutionPolicy –Value ""
```

Having problems? Ask for help in the Exchange forums.

Primary DNS Suffix is missing [ms.exch.setupreadiness.FqdnMissing]

8/3/2020 • 2 minutes to read • Edit Online

Exchange Setup can't continue because the primary DNS suffix (for example, contoso.com) hasn't been configured on the target server. Typically, you'll encounter this error when you're trying to install the Edge Transport server role.

To resolve this issue, add a primary DNS suffix on the computer and then run Setup again.

1. Replace <Value> with the DNS suffix you want to use (for example, contoso.com), and run the following command in Winows Powershell on the target server:

```
Set-ItemProperty -Path HKLM:\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters -Name 'NV Domain' -Value
<Value>
```

2. Restart the computer and run Setup again.

> **IMPORTANT**
>
> Changing the computer name or primary DNS suffix after you install Exchange isn't supported.

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

MAPI over HTTP isn't enabled [WarnMapiHttpNotEnabled]

8/3/2020 • 2 minutes to read • Edit Online

Microsoft Exchange Server 2016 Setup displayed this warning because there are servers running Exchange 2016 or later in this organization and MAPI over HTTP isn't enabled.

MAPI over HTTP is the preferred Outlook connectivity method when connecting to servers running Exchange 2016 or later. MAPI over HTTP improves the reliability and stability of the Outlook and Exchange connections by moving the transport layer to the industry-standard HTTP model. This allows a higher level of visibility of transport errors and enhanced recoverability. Additional functionality includes support for an explicit pause-and-resume function. This enables supported clients to change networks or resume from hibernation while maintaining the same server context.

Exchange Setup won't automatically enable MAPI over HTTP to avoid making unexpected changes to client connectivity. However, we recommend that you enable MAPI over HTTP as soon as possible to receive the benefits it provides.

For more information about MAPI over HTTP and how to enable it, see MAPI over HTTP in Exchange Server.

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Can't install Exchange 2016 or later in a forest that contains Exchange 2007 [E16E12CoexistenceMinVersionRequirement]

8/3/2020 • 2 minutes to read • Edit Online

The installation of Exchange Server 2016 or later can't continue because Setup found one or more Exchange 2007 servers in the Active Directory forest. Before you can install Exchange 2016 or later in your organization, you need to remove all Exchange 2007 servers from the forest.

The upgrade steps from Exchange 2007 are:

1. Install Exchange 2013 into your Exchange 2007 organization.
2. Configure Exchange 2013 and Exchange 2007 coexistence.
3. Migrate Exchange 2007 mailboxes, public folders, and other components to Exchange 2013.
4. Decommission and remove all Exchange 2007 servers.
5. Install Exchange 2016 or Exchange 2019 into your Exchange 2013 organization.
6. Configure coexistence with Exchange 2013.
7. Migrate Exchange 2013 mailboxes, public folders, and other components to Exchange 2016 or Exchange 2019.
8. Decommission and remove all Exchange 2013 servers.

The coexistence (and therefore, upgrade) options for Exchange are described in the following table:

CURRENT EXCHANGE VERSION	LATEST EXCHANGE VERSION FOR COEXISTENCE
Exchange 2000	Exchange 2007
Exchange 2003	Exchange 2010
Exchange 2007	Exchange 2013
Exchange 2010	Exchange 2016
Exchange 2013	Exchange 2019

When upgrading to Exchange 2013 or later, you can use the Exchange Deployment Assistant to help complete your deployment. For more information, see Exchange Deployment Assistant.

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Exchange 2010 SP3 RU11 or later is required for coexistence with Exchange 2016 [E16E14CoexistenceMinVersionRequirement]

8/3/2020 • 2 minutes to read • Edit Online

The installation of Exchange Server 2016 can't continue because Setup found one or more Exchange 2010 servers that aren't running the minimum required version of Exchange 2010. Before you can install Exchange 2016 in your organization, all Exchange 2010 servers in the forest need to be running Exchange 2010 Service Pack 3 (SP3) and Update Rollup 11 (RU11) or later. This requirement includes Exchange 2010 Edge Transport servers.

> **IMPORTANT**
>
> After you upgrade your Exchange 2010 Edge Transport servers to Exchange 2010 SP3 RU11 or later, you need to **re-create** the Edge subscription between your Exchange organization and each Edge Transport server (to update the Edge Transport server's Exchange version in Active Directory). For more information about re-creating Edge subscriptions in Exchange 2010, see Managing Edge Subscriptions.

Exchange 2013 CU10 or later is required for coexistence with Exchange 2016 or later [E16E15CoexistenceMinVersionRequirement]

8/3/2020 • 2 minutes to read • Edit Online

The installation of Exchange Server 2016 or later can't continue because Setup found one or more Exchange 2013 servers that aren't running the minimum required version of Exchange 2013. Before you can install Exchange 2016 or later in your organization, all Exchange 2013 servers in the forest need to be running Exchange 2013 Cumulative Update 10 (CU10) or later. This requirement includes Exchange 2013 Edge Transport servers.

> **IMPORTANT**
>
> After you upgrade your Exchange 2013 Edge Transport servers to Exchange 2013 CU10 or later, you need to **`re-create`** the Edge subscription between your Exchange organization and each Edge Transport server (to update the Edge Transport server's Exchange version in Active Directory). For more information about re-creating Edge subscriptions in Exchange 2013, see Manage Edge Subscriptions.

No Exchange 2013 servers detected [NoE15ServerWarning]

8/3/2020 • 2 minutes to read • Edit Online

Microsoft Exchange Server 2016 Setup displayed this warning because no Exchange Server 2013 server roles exist in the organization.

Caution

If you continue with Exchange Server 2016 installation, you won't be able to add Exchange 2013 servers to the organization at a future date.

Before deploying Exchange 2016, consider the following factors that may require you to deploy Exchange 2013 servers prior to deploying Exchange 2016:

- **Third-party or in-house developed applications**: Applications developed for earlier versions of Exchange may not be compatible with Exchange 2016. You may need to maintain Exchange 2013 servers to support these applications.

- **Coexistence or migration requirements**: If you plan on migrating mailboxes into your organization, some solutions may require the use of Exchange 2013 servers.

If you decide that you need to deploy Exchange 2013 servers, you must do so before you deploy Exchange 2016. Active Directory must be prepared for each Exchange version in the following order:

1. Exchange 2013

2. Exchange 2016

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Running "dir" on an ReFS-formatted disk could cause the computer to freeze [Win2k12RefsUpdateNotInstalled]

8/3/2020 • 2 minutes to read • Edit Online

Microsoft Exchange Server 2016 Setup has detected that the computer you're attempting to install Exchange 2016 on doesn't have a recommended Windows update installed. We strongly recommend that you install this Windows update before installing Exchange 2016 to avoid any issues resolved by the update.

Computers running Windows Server 2012 and later support the Resilient File System (ReFS). An issue exists that could cause computers to freeze when the "**dir**" command is run on disks formatted with ReFS.

Download and install the 64-bit update from the following URL, and then click **retry** on the **Readiness Checks** page.

> **NOTE**
>
> If this update requires a reboot to complete installation, you'll need to exit Exchange 2016 Setup, reboot, and then start Setup again.

Microsoft Knowledge Base article KB2894875, Windows 8-based or Windows Server 2012-based computer freezes when you run the "dir" command on an ReFS volume.

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

A Windows Server 2012 update rollup hasn't been installed [Win2k12RollupUpdateNotInstalled]

8/3/2020 • 2 minutes to read • Edit Online

Microsoft Exchange Server 2016 Setup has detected that the computer you're attempting to install Exchange 2016 on doesn't have a recommended Windows update installed. We strongly recommend that you install this Windows update before installing Exchange 2016 to avoid any issues resolved by the update.

A Windows Server 2012 update rollup that resolves several issues, including those that could cause Resilient File System (ReFS)-formatted disks to perform unreliably, hasn't been installed.

Download and install the 64-bit update from the following URL, and then click **retry** on the **Readiness Checks** page.

> **NOTE**
>
> If this update requires a reboot to complete installation, you'll need to exit Exchange 2016 Setup, reboot, and then start Setup again.

Microsoft Knowledge Base article KB2822241, Windows 8 and Windows Server 2012 update rollup: April 2013.

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Disks formatted as ReFS may not perform reliably [Win2k12UrefsUpdateNotInstalled]

8/3/2020 • 2 minutes to read • Edit Online

Microsoft Exchange Server 2016 Setup has detected that the computer you're attempting to install Exchange 2016 on doesn't have a recommended Windows update installed. We strongly recommend that you install this Windows update before installing Exchange 2016 to avoid any issues resolved by the update.

Computers running Windows Server 2012 and later support the Resilient File System (ReFS). An issue in the Virtual Disk Service could cause disks formatted as ReFS to not perform reliably. This could result in data corruption or data loss.

Download and install the 64-bit update from the following URL, and then click **retry** on the **Readiness Checks** page.

> **NOTE**
>
> If this update requires a reboot to complete installation, you'll need to exit Exchange 2016 Setup, reboot, and then start Setup again.

Microsoft Knowledge Base article KB2884597, Virtual Disk Service or applications that use the Virtual Disk Service crash or freeze in Windows Server 2012.

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

KB3206632 security update not installed [Win2k16LSARollupUpdateNotInstalled]

8/3/2020 • 2 minutes to read • Edit Online

Microsoft Exchange Server 2016 Setup can't continue because the local computer requires a software update. You'll need to install this update before Exchange 2016 Setup can continue.

Exchange 2016 Setup requires that the December 13, 2016 (KB3206632) security update be installed on the computer before installation can continue.

Download and install the 64-bit update from the following URL, and then click **retry** on the **Readiness Checks** page.

> **NOTE**
>
> If this update requires a reboot to complete installation, you'll need to exit Exchange 2016 Setup, reboot, and then start Setup again.

December 13, 2016 (KB3206632) security update

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Windows Server Core or Windows Nano Server is installed [IsServerCoreInstalled]

8/3/2020 • 2 minutes to read • Edit Online

Microsoft Exchange Server 2016 Setup can't continue because it detected that the local computer is running Windows Server Core or Windows Nano Server. Exchange 2016 requires that **Windows Server with Desktop Experience** (Windows Server 2016) or **Windows Server with a GUI** (Windows Server 2012 and 2012R2) is installed on the local computer. Before you can install Exchange 2016, you need to do one of the following depending on the version of Windows Server you have installed:

- **Windows Server 2012 and Windows Server 2012 R2**: Run the following command in Windows PowerShell:

  ```
  Install-WindowsFeature Server-Gui-Mgmt-Infra, Server-Gui-Shell -Restart
  ```

- **Windows Server 2016**: Install Windows Server 2016 and choose the **Desktop Experience** installation option. If a computer is running Windows Server 2016 Core or Nano and you want to install Exchange 2016 on it, you'll need to reinstall the operating system and choose the **Desktop Experience** installation option.

For more information, see Exchange Server system requirements.

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Exchange Server editions and versions

8/3/2020 • 2 minutes to read • Edit Online

Exchange Server 2016 and Exchange Server 2019 are available in two server editions:

- **Enterprise Edition**: Can scale up to 100 mounted databases per server.
- **Standard Edition**: Limited to five mounted databases per server.

A *mounted database* is a database that's in use (an active mailbox database that's mounted for use by clients or a passive mailbox database that's mounted in recovery for log replication and replay). While you can create more databases than the described limits, you can only mount the maximum number of databases that are allowed by the edition of Exchange. Note that the recovery database doesn't count towards these limits.

The server editions are defined by a product key. When you enter a valid product key, the supported edition for the server is established. For more information, see Enter your Exchange Server product key.

Notes:

- You can use a valid product key to move from the Trial Edition (evaluation version) of Exchange to either Standard Edition or Enterprise Edition. No loss of functionality occurs after the Trial Edition expires, so you can maintain lab, demo, training, and other non-production environments beyond 120 days without having to reinstall the Trial Edition of Exchange or entering a product key.
- You can use a valid product key to move from Standard Edition to Enterprise Edition.
- You can't use a valid product key to downgrade from Enterprise Edition to Standard Edition or revert to the Trial Edition. You can only do these types of downgrades by uninstalling Exchange, reinstalling Exchange, and entering the correct product key.

Exchange Server versions

For a list of Exchange Server versions and how to download and upgrade to the latest version of Exchange, see the following topics:

- Exchange Server build numbers and release dates
- Install Exchange Mailbox servers using the Setup wizard
- Upgrade Exchange to the latest Cumulative Update

To view the Exchange version and edition information for all Exchange servers in your organization, run the following command in the Exchange Management Shell:

```
Get-ExchangeServer | Format-Table -Auto Name,Edition,AdminDisplayVersion
```

Exchange Server license types

Exchange 2013 and all later versions use a licensing model that's similar to how Exchange 2010 was licensed:

- **Server licenses**: A license must be assigned for each Exchange server. The Server license is sold in two server editions: Standard Edition and Enterprise Edition.
- **Client Access licenses (CALs)**: Exchange also comes in two CAL editions, which are referred to as a

Standard CAL and an Enterprise CAL. You can mix and match the Exchange server editions with the CAL types. For example, you can use Enterprise CALs with Standard Edition or Standard CALs with Enterprise Edition.

For more information about Exchange license types, see Exchange Licensing FAQs.

Exchange Server language support

8/3/2020 • 2 minutes to read • Edit Online

Exchange Server 2016 and Exchange Server 2019 have enhanced language support for both servers and clients. This topic lists the languages that are available for both servers and clients in Exchange 2016 and Exchange 2019.

Supported server languages for Exchange 2016 and Exchange 2019

- Chinese (Simplified)
- Chinese (Traditional)
- English
- French
- German
- Italian
- Japanese
- Korean
- Portuguese
- Russian
- Spanish

Supported client languages for Exchange 2016 and Exchange 2019

- Amharic
- Arabic
- Basque (Basque)
- Bengali (India)
- Bulgarian
- Catalan
- Chinese (Simplified)
- Chinese (Traditional)
- Croatian
- Czech
- Danish
- Dutch
- English
- Estonian

- Filipino (Philippines)
- Finnish
- French
- Galician
- German
- Greek
- Gujarati
- Hebrew
- Hindi
- Hungarian
- Icelandic
- Indonesian
- Italian
- Japanese
- Kannada
- Kazakh
- Kiswahili
- Korean
- Latvian
- Lithuanian
- Malay (Brunei Darussalam)
- Malay (Malaysia)
- Malayalam
- Marathi
- Norwegian (Bokmål)
- Norwegian (Nynorsk)
- Oriya
- Persian
- Polish
- Portuguese (Brazil)
- Portuguese (Portugal)
- Romanian
- Russian

- Serbian (Cyrillic, Serbia)
- Serbian (Latin)
- Slovak
- Slovenian
- Spanish
- Swedish
- Tamil
- Telugu
- Thai
- Turkish
- Ukrainian
- Urdu
- Vietnamese
- Welsh

Exchange Server storage configuration options

8/3/2020 • 21 minutes to read • Edit Online

Understanding the storage options and requirements for Mailbox servers in Exchange Server 2016 and Exchange Server 2019 is an important part of your Mailbox server storage design solution.

Storage architectures

The following table describes supported storage architectures and provides best practice guidance for each type of storage architecture where appropriate.

Supported storage architectures

STORAGE ARCHITECTURE	DESCRIPTION	BEST PRACTICE
Direct-attached storage (DAS)	DAS is a digital storage system directly attached to a server or workstation, without a storage network in between. For example, DAS transports include Serial Attached Small Computer System Interface (SCSI) and Serial Attached Advanced Technology Attachment (ATA).	Not available.
Storage area network (SAN): Internet Small Computer System Interface (iSCSI)	SAN is an architecture to attach remote computer storage devices (such as disk arrays and tape libraries) to servers in such a way that the devices appear as locally attached to the operating system (for example, block storage). iSCSI SANs encapsulate SCSI commands within IP packets and use standard networking infrastructure as the storage transport (for example, Ethernet).	Don't share physical disks backing up Exchange data with other applications. Use dedicated storage networks. Use multiple network paths for stand-alone configurations.
SAN: Fibre Channel	Fibre Channel SANs encapsulate SCSI commands within Fibre Channel packets and generally utilize specialized Fibre Channel networks as the storage transport.	Don't share physical disks backing up Exchange data with other applications. Use multiple Fibre Channel network paths for stand-alone configurations. Follow storage vendor's best practices for tuning Fibre Channel host bus adapters (HBAs), for example, Queue Depth and Queue Target.

A network-attached storage (NAS) unit is a self-contained computer connected to a network, with the sole purpose of supplying file-based data storage services to other devices on the network. The operating system and other software on the NAS unit provide the functionality of data storage, file systems, and access to files, and the management of these functions (for example, file storage).

All storage used by Exchange for storage of Exchange data must be block-level storage because Exchange 2016 doesn't support the use of NAS volumes, other than in the SMB 3.0 scenario outlined in the topic Exchange Server virtualization. Also, in a virtualized environment, NAS storage that's presented to the guest as block-level storage via the hypervisor isn't supported.

Using storage tiers is not recommended, as it could adversely affect system performance. For this reason, do not allow the storage controller to automatically move the most accessed files to "faster" storage.

Physical disk types

The following table provides a list of supported physical disk types and provides best practice guidance for each physical disk type where appropriate.

Supported physical disk types

PHYSICAL DISK TYPE	DESCRIPTION	SUPPORTED OR BEST PRACTICE
Serial ATA (SATA)	SATA is a serial interface for ATA and integrated device electronics (IDE) disks. SATA disks are available in a variety of form factors, speeds, and capacities. In general, choose SATA disks for Exchange 2016 mailbox storage when you have the following design requirements: • High capacity • Moderate performance • Moderate power utilization	Supported: 512-byte sector disks for Windows Server 2008 and Windows Server 2008 R2. In addition, 512e disks are supported for Windows Server 2008 R2 with the following: • The hotfix described in KB982018. • Windows Server 2008 R2 with Service Pack 1 (SP1) and Exchange Server 2010 SP1. Exchange 2013 and later supports native 4-kilobyte (KB) sector disks and 512e disks. Support requires that all copies of a database reside on the same physical disk type. For example, it is not a supported configuration to host one copy of a given database on a 512-byte sector disk and another copy of that same database on a 512e disk or 4K disk. Best practice: Consider enterprise class SATA disks, which generally have better heat, vibration, and reliability characteristics.
Serial Attached SCSI	Serial Attached SCSI is a serial interface for SCSI disks. Serial Attached SCSI disks are available in a variety of form factors, speeds, and capacities. In general, choose Serial Attached SCSI disks for Exchange 2016 mailbox storage when you have the following design requirements: • Moderate capacity • High performance • Moderate power utilization	Supported: 512-byte sector disks for Windows Server 2008 and Windows Server 2008 R2. In addition, 512e disks are supported for Windows Server 2008 R2 with the following: • The hotfix described in KB982018. • Windows Server 2008 R2 SP1 and Exchange Server 2010 SP1. Exchange 2013 and later supports native 4-kilobyte (KB) sector disks and 512e disks. Support requires that all copies of a database reside on the same physical disk type. For example, it is not a supported configuration to host one copy of a given database on a 512-byte sector disk and another copy of that same database on a 512e disk or 4K disk. Best practice: Physical disk-write caching must be disabled when used without a UPS.

PHYSICAL DISK TYPE	DESCRIPTION	SUPPORTED OR BEST PRACTICE
Fibre Channel	Fibre Channel is an electrical interface used to connect disks to Fibre Channel-based SANs. Fibre Channel disks are available in a variety of speeds and capacities. In general, choose Fibre Channel disks for Exchange 2016 mailbox storage when you have the following design requirements: • Moderate capacity • High performance • SAN connectivity	Supported: 512-byte sector disks for Windows Server 2008 and Windows Server 2008 R2. In addition, 512e disks are supported for Windows Server 2008 R2 with the following: • The hotfix described in KB982018. • Windows Server 2008 R2 with Service Pack 1 (SP1) and Exchange Server 2010 SP1. Exchange 2013 and later supports native 4-kilobyte (KB) sector disks and 512e disks. Support requires that all copies of a database reside on the same physical disk type. For example, it is not a supported configuration to host one copy of a given database on a 512-byte sector disk and another copy of that same database on a 512e disk or 4K disk. Best practice: Physical disk-write caching must be disabled when used without a UPS.
Solid-state drive (SSD) (flash disk)	An SSD is a data storage device that uses solid-state memory to store persistent data. An SSD emulates a hard disk drive interface. SSD disks are available in a variety of speeds (different I/O performance capabilities) and capacities. In general, choose SSD disks for Exchange 2016 mailbox storage when you have the following design requirements: • Low capacity • Extremely high performance	Supported: 512-byte sector disks for Windows Server 2008 and Windows Server 2008 R2. In addition, 512e disks are supported for Windows Server 2008 R2 with the following: • The hotfix described in KB982018. • Windows Server 2008 R2 SP1 and Exchange Server 2010 SP1. Exchange 2013 and later supports native 4-kilobyte (KB) sector disks and 512e disks when all copies of a database reside on the same physical disk type. For example, it is not a supported configuration to host one copy of a given database on a 512-byte sector disk and another copy of that same database on a 512e disk or 4K disk. Best practice: Physical disk-write caching must be disabled when used without a UPS. In general, Exchange 2016 Mailbox servers don't require the performance characteristics of SSD storage.

Factors to consider when choosing disk types

There are several trade-offs when choosing disk types for Exchange 2016 storage. The correct disk is one that balances performance (both sequential and random) with capacity, reliability, power utilization, and capital cost. The following table of supported physical disk types provides information to help you when considering these factors.

From a performance perspective, using large, slower disks for Exchange storage is okay, provided the disks can maintain an average read and write latency of 20ms or less under load.

Factors in disk type choice

DISK SPEED (RPM)	DISK FORM FACTOR	INTERFACE OR TRANSPORT	CAPACITY	RANDOM I/O PERFORMANCE	SEQUENTIAL I/O PERFORMANCE	POWER UTILIZATION
5,400	2.5-inch	SATA	Average	Poor	Poor	Excellent
5,400	3.5-inch	SATA	Excellent	Poor	Poor	Above average
7,200	2.5-inch	SATA	Average	Average	Average	Excellent
7,200	2.5-inch	Serial Attached SCSI	Average	Average	Above average	Excellent
7,200	3.5-inch	SATA	Excellent	Average	Above average	Above average
7,200	3.5-inch	Serial Attached SCSI	Excellent	Average	Above average	Above average
7,200	3.5-inch	Fibre Channel	Excellent	Average	Above average	Average
10,000	2.5-inch	Serial Attached SCSI	Below average	Excellent	Above average	Above average
10,000	3.5-inch	SATA	Average	Average	Above average	Above average
10,000	3.5-inch	Serial Attached SCSI	Average	Above average	Above average	Below average
10,000	3.5-inch	Fibre Channel	Average	Above average	Above average	Below average
15,000	2.5-inch	Serial Attached SCSI	Poor	Excellent	Excellent	Average
15,000	3.5-inch	Serial Attached SCSI	Average	Excellent	Excellent	Below average
15,000	3.5-inch	Fibre Channel	Average	Excellent	Excellent	Poor
SSD: enterprise class	Not applicable	SATA, Serial Attached SCSI, Fibre Channel	Poor	Excellent	Excellent	Excellent

Best practices for supported storage configurations

This section provides best practice information about supported disk and array controller configurations. In addition to the commonly used Redundant Array of Indepentdent Disks (RAID), there is also just a bunch of disks (or drives), or JBOD, which refers to a collection of hard disks that have not been configured to act as a redundant array.

RAID is often used to both improve the performance characteristics of individual disks (by striping data across several disks) as well as to provide protection from individual disk failures. With the advancements in Exchange 2016 high availability, RAID is not a required component for Exchange 2016 storage design. However, RAID is still an essential component of Exchange 2016 storage design for standalone servers as well as solutions that require storage fault tolerance.

Operating System, System, or Pagefile Volume

The recommended configuration for an operating system, system or pagefile volume is to utilize RAID technology to protect this data type. The recommended RAID configuration is either RAID-1 or RAID-1/0, however all RAID types are supported.

Separated Mailbox Database and Log Volumes

If you're deploying a standalone Mailbox server role architecture, RAID technology is required for the mailbox database and log volumes. The recommended RAID configuration for mailbox volumes is RAID-1/0 (especially if you're using 5.4K or 7.2K disks); however all RAID types are supported. For log volumes, RAID-1 or RAID-1/0 is the recommended RAID configuration.

When using RAID-5 or RAID-6 configurations for the operating system, pagefile, or Exchange data volumes, note the following:

- RAID-5 configurations, including variations such as RAID-50 and RAID-51, should have no more than 7 disks per array group and array controller high-priority scrubbing and surface scanning enabled.
- RAID-6 configurations should have array controller high-priority scrubbing and surface scanning enabled.

Although JBOD is supported in high availability architectures that have 3 or more highly available database copies, because the log and mailbox database volumes are separated, JBOD is not recommended as a solution.

Mailbox Database and Log Volume Co-Location

Mailbox database and log volume co-location is not recommended in standalone architectures. In high availability architectures, there are two possibilities for this scenario:

1. Single database per volume
2. Multiple databases per volume

Single Database Per Volume

In an Exchange environment, a JBOD storage solution involves having both the database and its associated logs stored on a single disk. To deploy a JBOD solution, you must deploy a minimum of three highly available database copies. Utilizing a single disk is a single point of failure, because when the disk fails, the database copy residing on that disk is lost. Having a minimum of three database copies ensures fault tolerance by having two additional copies in the event that one copy (or one disk) fails. However, placement of three highly available database copies, as well as the use of lagged database copies, can affect storage design. The following table shows guidelines for RAID or JBOD considerations.

RAID or JBOD Considerations

DATACENTER SERVERS	TWO HIGHLY AVAILABLE COPIES (TOTAL)	THREE HIGHLY AVAILABLE COPIES (TOTAL)	TWO OR MORE HIGHLY AVAILABLE COPIES PER DATACENTER	ONE LAGGED COPY	TWO OR MORE LAGGED COPIES PER DATACENTER
Primary datacenter servers	RAID	RAID or JBOD (2 copies)	RAID or JBOD	RAID	RAID or JBOD

DATACENTER SERVERS	TWO HIGHLY AVAILABLE COPIES (TOTAL)	THREE HIGHLY AVAILABLE COPIES (TOTAL)	TWO OR MORE HIGHLY AVAILABLE COPIES PER DATACENTER	ONE LAGGED COPY	TWO OR MORE LAGGED COPIES PER DATACENTER
Secondary datacenter servers	RAID	RAID (1 copy)	RAID or JBOD	RAID	RAID or JBOD

To deploy on JBOD with the primary datacenter servers, you need three or more highly available database copies within the DAG. If mixing lagged copies on the same server hosting highly available database copies (for example, not using dedicated lagged database copy servers), you need at least two lagged database copies.

For the secondary datacenter servers to use JBOD, you should have at least two highly available database copies in the secondary datacenter. The loss of a copy in the secondary datacenter won't result in requiring a reseed across the WAN or having a single point of failure in the event the secondary datacenter is activated. If mixing lagged database copies on the same server hosting highly available database copies (for example, not using dedicated lagged database copy servers), you need at least two lagged database copies.

For dedicated lagged database copy servers, you should have at least two lagged database copies within a datacenter to use JBOD. Otherwise, the loss of disk results in the loss of the lagged database copy, as well as the loss of the protection mechanism.

Multiple Databases Per Volume

Multiple databases per volume is a new JBOD scenario available in Exchange 2016 that allows for active and passive copies (including lagged copies) to be mixed on a single disk, enabling better disk utilization. However, to deploy lagged copies in this manner, automatic lagged copy log file play down must be enabled. The following table shows guidelines for JBOD considerations for multiple databases per volume.

JBOD Considerations

DATACENTER SERVERS	3 OR MORE COPIES (TOTAL)	TWO OR MORE COPIES PER DATACENTER
Primary datacenter servers	JBOD	JBOD
Secondary datacenter servers	N/A	JBOD

The following table provides guidance about storage array configurations for Exchange 2016.

Supported RAID types for the Exchange 2016 Mailbox server role

RAID TYPE	DESCRIPTION	SUPPORTED OR BEST PRACTICE
Disk array RAID stripe size (KB)	The stripe size is the per disk unit of data distribution within a RAID set. Stripe size is also referred to as *block size*.	Best practice: 256 KB or greater. Follow storage vendor best practices.

RAID TYPE	DESCRIPTION	SUPPORTED OR BEST PRACTICE
Storage array cache settings	The cache settings are provided by a battery-backed caching array controller.	Best practice: 100 percent write cache (battery or flash backed cache) for DAS storage controllers in either a RAID or JBOD configuration. 75 percent write cache, 25 percent read cache (battery or flash backed cache) for other types of storage solutions such as SAN. If your SAN vendor has different best practices for cache configuration on their platform, follow the guidance of your SAN vendor.
Physical disk write caching	The settings for the cache are on each individual disk.	Supported: Physical disk write caching must be disabled when used without a UPS.

The following table provides guidance about database and log file choices.

Database and log file choices for the Exchange 2016 Mailbox server role

DATABASE AND LOG FILE OPTIONS	DESCRIPTION	STAND-ALONE: SUPPORTED OR BEST PRACTICE	HIGH AVAILABILITY: SUPPORTED OR BEST PRACTICE
File placement: database per log isolation	Database per log isolation refers to placing the database file and logs from the same mailbox database onto different volumes backed by different physical disks.	Best practice: For recoverability, move database (.edb) file and logs from the same database to different volumes backed by different physical disks.	Supported: Isolation of logs and databases isn't required.
File placement: database files per volume	Database files per volume refers to how you distribute database files within or across disk volumes.	Best practice: Based on your backup methodology.	Supported: When using JBOD, create a single volume with separate directories for database(s) and for log files.
File placement: log streams per volume	Log streams per volume refers to how you distribute database log files within or across disk volumes.	Best practice: Based on your backup methodology.	Supported: When using JBOD, create a single volume with separate directories for database(s) and for log files. Best practice: When using JBOD, leverage multiple databases per volume.
Database size	Database size refers to the disk database (.edb) file size.	Supported: Approximately 16 terabytes. Best practice: • 200 gigabytes (GB) or less. • Provision for 120 percent of calculated maximum database size.	Supported: Approximately 16 terabytes. Best practice: • 2 terabytes or less. • Provision for 120 percent of calculated maximum database size.

DATABASE AND LOG FILE OPTIONS	DESCRIPTION	STAND-ALONE: SUPPORTED OR BEST PRACTICE	HIGH AVAILABILITY: SUPPORTED OR BEST PRACTICE
Log truncation method	Log truncation method is the process for truncating and deleting old database log files. There are two mechanisms: • Circular logging, in which Exchange deletes the logs. • Log truncation, which occurs after a successful full or incremental Volume Shadow Copy Service (VSS) backup.	Best practice: • Use backups for log truncation (for example, circular logging disabled). • Provision for three days of log generation capacity.	Best practice: • Enable circular logging for deployments that use Exchange native data protection features. • Provision for three days beyond replay lag setting of log generation capacity.

The following table provides guidance about Windows disk types.

Windows disk types for the Exchange 2016 Mailbox server role

WINDOWS DISK TYPE	DESCRIPTION	STAND-ALONE: SUPPORTED OR BEST PRACTICE	HIGH AVAILABILITY: SUPPORTED OR BEST PRACTICE
Basic disk	A disk initialized for basic storage is called a basic disk. A basic disk contains basic volumes, such as primary partitions, extended partitions, and logical drives.	Supported. Best practice: Use basic disks.	Supported. Best practice: Use basic disks.
Dynamic disk	A disk initialized for dynamic storage is called a dynamic disk. A dynamic disk contains dynamic volumes, such as simple volumes, spanned volumes, striped volumes, mirrored volumes, and RAID-5 volumes.	Supported.	Supported.

The following table provides guidance on volume configurations.

Volume configurations for the Exchange 2016 Mailbox server role

VOLUME CONFIGURATION	DESCRIPTION	STAND-ALONE: SUPPORTED OR BEST PRACTICE	HIGH AVAILABILITY: SUPPORTED OR BEST PRACTICE
GUID partition table (GPT)	GPT is a disk architecture that expands on the older master boot record (MBR) partitioning scheme. The maximum NTFS formatted partition size is 256 terabytes.	Supported. Best practice: Use GPT partitions.	Supported. Best practice: Use GPT partitions.

VOLUME CONFIGURATION	DESCRIPTION	STAND-ALONE: SUPPORTED OR BEST PRACTICE	HIGH AVAILABILITY: SUPPORTED OR BEST PRACTICE
MBR	An MBR, or partition sector, is the 512-byte boot sector that is the first sector (LBA Sector 0) of a partitioned data storage device such as a hard disk. The maximum NTFS formatted partition size is 2 terabytes.	Supported.	Supported.
Partition alignment	Partition alignment refers to aligning partitions on sector boundaries for optimal performance.	Supported: The Windows Server 2008 R2 and Windows Server 2012 default is 1 megabyte (MB).	Supported: The Windows Server 2008 R2 and Windows Server 2012 default is 1 MB.
Volume path	Volume path refers to how a volume is accessed.	Supported: Drive letter or mount point.Best practice: Mount point host volume must be RAID enabled.	Supported: Drive letter or mount point. Best practice: Mount point host volume must be RAID-enabled.
File system	File system is a method for storing and organizing computer files and the data they contain to make it easy to find and access the files.	Supported: NTFS and ReFS.	Supported: NTFS and ReFS.
NTFS defragmentation	NTFS defragmentation is a process that reduces the amount of fragmentation in Windows file systems. It does this by physically organizing the contents of the disk to store the pieces of each file close together and contiguously.	Supported. Best practice: Not required and not recommended. On Windows Server 2012, we also recommend disabling the automatic disk optimization and defragmentation feature.	Supported. Best practice: Not required and not recommended. On Windows Server 2012, we also recommend disabling the automatic disk optimization and defragmentation feature.
NTFS allocation unit size	NTFS allocation unit size represents the smallest amount of disk space that can be allocated to hold a file.	Supported: All allocation unit sizes. Best practice: 64 KB for both .edb and log file volumes.	Supported: All allocation unit sizes. Best practice: 64 KB for both .edb and log file volumes.
NTFS compression	NTFS compression is the process of reducing the actual size of a file stored on the hard disk.	Supported: Not supported for Exchange database or log files.	Supported: Not supported for Exchange database or log files.

VOLUME CONFIGURATION	DESCRIPTION	STAND-ALONE: SUPPORTED OR BEST PRACTICE	HIGH AVAILABILITY: SUPPORTED OR BEST PRACTICE
NTFS Encrypting File System (EFS)	EFS enables users to encrypt individual files, folders, or entire data drives. Because EFS provides strong encryption through industry-standard algorithms and public key cryptography, encrypted files are confidential even if an attacker bypasses system security.	Supported: Not supported for Exchange database or log files.	Not supported for Exchange database or log files.
Windows BitLocker (volume encryption)	Windows BitLocker is a data protection feature in Windows Server 2008. BitLocker protects against data theft or exposure on computers that are lost or stolen, and it offers more secure data deletion when computers are decommissioned.	Supported: All Exchange database and log files.	Supported: All Exchange database and log files. Windows failover clusters require Windows Server 2008 R2 or Windows Server 2008 R2 SP1. Exchange volumes with Bitlocker enabled are not supported on Windows failover clusters running earlier versions of Windows. For more information about Windows 7 BitLocker encryption, see BitLocker Drive Encryption in Windows 7: Frequently Asked Questions.
Server Message Block (SMB) 3.0	The Server Message Block (SMB) protocol is a network file sharing protocol (on top of TCP/IP or other network protocols) that allows applications on a computer to access files and resources on a remote server. It also allows applications to communicate with any server program that is set up to receive an SMB client request. Windows Server 2012 introduces the new 3.0 version of the SMB protocol with the following features: • SMB Transparent failover • SMB Scaleout • SMB Multichannel • SMB Direct • SMB Encryption • VSS for SMB file shares • SMB Directory Leasing • SMB PowerShell	Limited Support. Supported scenario is a hardware virtualized deployment where the disks are hosted on VHDs on an SMB 3.0 share. These VHDs are presented to the host via a hypervisor. For more information, see Exchange Server virtualization.	Limited Support. Supported scenario is a hardware virtualized deployment where the disks are hosted on VHDs on an SMB 3.0 share. These VHDs are presented to the host via a hypervisor. For more information, see Exchange Server virtualization.

VOLUME CONFIGURATION	DESCRIPTION	STAND-ALONE: SUPPORTED OR BEST PRACTICE	HIGH AVAILABILITY: SUPPORTED OR BEST PRACTICE
Storage Spaces	Storage Spaces is a new storage solution that delivers virtualization capabilities for Windows Server 2012. Storage Spaces allow you to organize physical disks into storage pools, which can be easily expanded by simply adding disks. These disks can be connected either through USB, SATA or SAS. It also utilizes virtual disks (spaces), which behave just like physical disks, with associated powerful capabilities such as thin provisioning, as well as resiliency to failures of underlying physical media. For more information on Storage Spaces, see Storage Spaces Overview.	Supported. Same restrictions as for physical disk types outlined in this topic.	Supported. Same restrictions as for physical disk types outlined in this topic.
Resilient File System (ReFS)	ReFS is a newly engineered file system for Windows Server 2012 that is built on the foundations of NTFS. ReFS maintains high degree of compatibility with NTFS while providing enhanced data verification and auto-correction techniques as well as an integrated end-to-end resiliency to corruptions especially when used in conjunction with the storage spaces feature. For more information on ReFS, see Resilient File System (ReFS) overview: Supported Deployments.	Supported for volumes containing Exchange database files, log files and content indexing files, provided that the following hotfix is installed: Exchange Server 2013 databases become fragmented in Windows Server 2012. Not supported for volumes containing Exchange binaries. Best practice: Data integrity features must be disabled for the Exchange database (.edb) files or the volume that hosts these files. Integrity features can be enabled for volumes containing the content index catalog, provided that the volume does not contain any databases or log files.	Supported for volumes containing Exchange database files, log files and content indexing files, provided that the following hotfix is installed: Exchange Server 2013 databases become fragmented in Windows Server 2012. Not supported for volumes containing Exchange binaries. Best practice: Data integrity features must be disabled for the Exchange database (.edb) files or the volume that hosts these files. Integrity features can be enabled for volumes containing the content index catalog, provided that the volume does not contain any databases or log files.
ReFS allocation unit size	ReFS allocation unit size represents the smallest amount of disk space that can be allocated to hold a file.	Supported: All allocation unit sizes. Best practice: 64 KB for both .edb and log file volumes.	Supported: All allocation unit sizes. Best practice: 64 KB for both .edb and log file volumes.

VOLUME CONFIGURATION	DESCRIPTION	STAND-ALONE: SUPPORTED OR BEST PRACTICE	HIGH AVAILABILITY: SUPPORTED OR BEST PRACTICE
Data De-Duplication	Data deduplication is a technique to optimize storage utilization. It is a method of finding and removing duplication within data without compromising its fidelity or integrity. The goal is to store more data in less space by segmenting files into small variable-sized chunks, identifying duplicate chunks, and maintaining a single copy of each chunk. Data deduplication technologies are typically implemented one of two ways; at the operating system level, or at the storage system level and the operating system is unaware of it being used.	OS Level: Not Supported for Exchange mailbox databases, transport databases, or content index files. Storage System Level: Supported, but falls within the Microsoft third-party storage software solutions support policy. **Note:** OS level dedupe can be used for Exchange database files that are completely offline (used as backups or archives).	OS Level: Not Supported for Exchange mailbox databases, transport databases, or content index files. Storage Level: Supported, but falls within the Microsoft third-party storage software solutions support policy. **Note:** OS level dedupe can be used for Exchange database files that are completely offline (used as backups or archives).

Network ports for clients and mail flow in Exchange

8/3/2020 • 9 minutes to read • Edit Online

This topic provides information about the network ports that are used by Exchange Server 2016 and Exchange Server 2019 for communication with email clients, internet mail servers, and other services that are external to your local Exchange organization. Before we get into that, understand the following ground rules:

- We do not support restricting or altering network traffic between internal Exchange servers, between internal Exchange servers and internal Lync or Skype for Business servers, or between internal Exchange servers and internal Active Directory domain controllers in any and all types of topologies. If you have firewalls or network devices that could potentially restrict or alter this kind of internal network traffic, you need to configure rules that allow free and unrestricted communication between these servers: rules that allow incoming and outgoing network traffic on any port (including random RPC ports) and any protocol that never alter bits on the wire.

- Edge Transport servers are almost always located in a perimeter network, so it's expected that you'll restrict network traffic between the Edge Transport server and the internet, and between the Edge Transport server and your internal Exchange organization. These network ports are described in this topic.

- It's expected that you'll restrict network traffic between external clients and services and your internal Exchange organization. It's also OK if you decide to restrict network traffic between internal clients and internal Exchange servers. These network ports are described in this topic.

Network ports required for clients and services

The network ports that are required for email clients to access mailboxes and other services in the Exchange organization are described in the following diagram and table.

Notes:

- The destination for these clients and services is the Client Access services on a Mailbox server. In Exchange 2016 and Exchange 2019, Client Access (frontend) and backend services are installed together on the same Mailbox server. For more information, see Client Access protocol architecture.

- Although the diagram shows clients and services from the internet, the concepts are the same for internal clients (for example, clients in an accounts forest accessing Exchange servers in a resource forest). Similarly, the table doesn't have a source column because the source could be any location that's external to the Exchange organization (for example, the internet or an accounts forest).

- Edge Transport servers have no involvement in the network traffic that's associated with these clients and services.

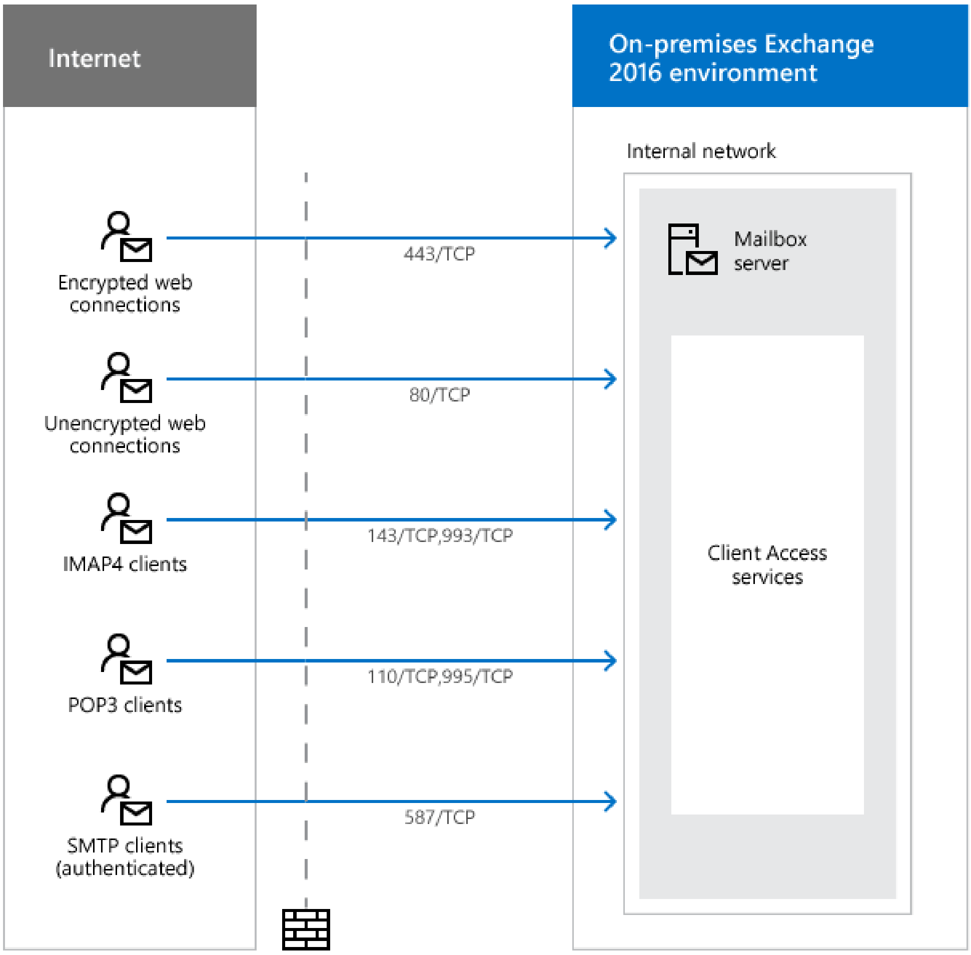

PURPOSE	PORTS	COMMENTS
Encrypted web connections are used by the following clients and services: • Autodiscover service • Exchange ActiveSync • Exchange Web Services (EWS) • Offline address book (OAB) distribution • Outlook Anywhere (RPC over HTTP) • Outlook MAPI over HTTP • Outlook on the web (formerly known as Outlook Web App)	443/TCP (HTTPS)	For more information about these clients and services, see the following topics: • Autodiscover service in Exchange Server • Exchange ActiveSync • EWS reference for Exchange • Offline address books in Exchange Server • Outlook Anywhere • MAPI over HTTP in Exchange Server
Unencrypted web connections are used by the following clients and services: • Internet calendar publishing • Outlook on the web (redirect to 443/TCP) • Autodiscover (fallback when 443/TCP isn't available)	80/TCP (HTTP)	Whenever possible, we recommend using encrypted web connections on 443/TCP to help protect data and credentials. However, you may find that some services must be configured to use unencrypted web connections on 80/TCP to the Client Access services on Mailbox servers. For more information about these clients and services, see the following topics: • Enable Internet Calendar Publishing • Autodiscover service in Exchange Server

PURPOSE	PORTS	COMMENTS
IMAP4 clients	143/TCP (IMAP), 993/TCP (secure IMAP)	IMAP4 is disabled by default. For more information, see POP3 and IMAP4 in Exchange Server. The IMAP4 service in the Client Access services on the Mailbox server proxies connections to the IMAP4 Backend service on a Mailbox server.
POP3 clients	110/TCP (POP3), 995/TCP (secure POP3)	POP3 is disabled by default. For more information, see POP3 and IMAP4 in Exchange Server. The POP3 service in the Client Access services on the Mailbox server proxies connections to the POP3 Backend service on a Mailbox server.
SMTP clients (authenticated)	587/TCP (authenticated SMTP)	The default Received connector named "Client Frontend *<Server name>*" in the Front End Transport service listens for authenticated SMTP client submissions on port 587. **Note**: If you have email clients that are only able to submit authenticated SMTP email on port 25, you can modify the network adapter bindings of the client Receive connector to also listen for authenticated SMTP email submissions on port 25.

Network ports required for mail flow

How mail is delivered to and from your Exchange organization depends on your Exchange topology. The most important factor is whether you have a subscribed Edge Transport server deployed in your perimeter network.

Network ports required for mail flow (no Edge Transport servers)

The network ports that are required for mail flow in an Exchange organization that has only Mailbox servers are described in the following diagram and table.

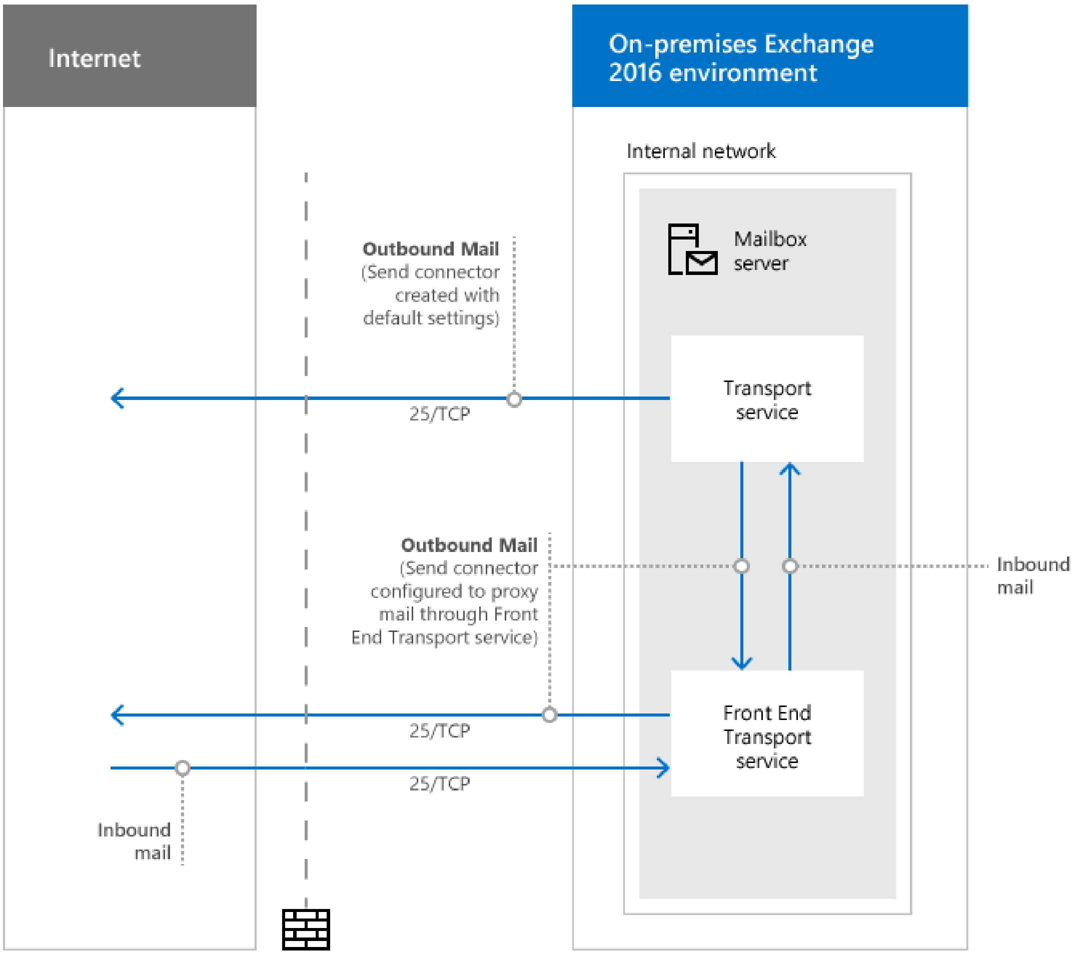
Internet
On-premises Exchange 2016 environment
Internal network
Mailbox server
Outbound Mail (Send connector created with default settings)
25/TCP
Transport service
Outbound Mail (Send connector configured to proxy mail through Front End Transport service)
Inbound mail
25/TCP
25/TCP
Front End Transport service
Inbound mail

PURPOSE	PORTS	SOURCE	DESTINATION	COMMENTS
Inbound mail	25/TCP (SMTP)	Internet (any)	Mailbox server	The default Receive connector named "Default Frontend *<Mailbox server name>*" in the Front End Transport service listens for anonymous inbound SMTP mail on port 25. Mail is relayed from the Front End Transport service to the Transport service on a Mailbox server using the implicit and invisible intra-organization Send connector that automatically routes mail between Exchange servers in the same organization. For more information, see Implicit Send connectors.

PURPOSE	PORTS	SOURCE	DESTINATION	COMMENTS
Outbound mail	25/TCP (SMTP)	Mailbox server	Internet (any)	By default, Exchange doesn't create any Send connectors that allow you to send mail to the internet. You have to create Send connectors manually. For more information, see Create a Send connector to send mail to the internet.
Outbound mail (if proxied through the Front End transport service)	25/TCP (SMTP)	Mailbox server	Internet (any)	Outbound mail is proxied through the Front End Transport service only when a Send connector is configured with **Proxy through Client Access server** in the Exchange admin center or `-FrontEndProxyEnabled $true` in the Exchange Management Shell. In this case, the default Receive connector named "Outbound Proxy Frontend *<Mailbox server name>*" in the Front End Transport service listens for outbound mail from the Transport service on a Mailbox server. For more information, see Configure Send connectors to proxy outbound mail.
DNS for name resolution of the next mail hop (not pictured)	53/UDP,53/TCP (DNS)	Mailbox server	DNS server	See the Name resolution section in this topic.

Network ports required for mail flow with Edge Transport servers

A subscribed Edge Transport server that's installed in your perimeter network affects mail flow in the following ways:

- Outbound mail from the Exchange organization never flows through the Front End Transport service on Mailbox servers. Mail always flows from the Transport service on a Mailbox server in the subscribed Active Directory site to the Edge Transport server (regardless of the version of Exchange on the Edge Transport server).
- Inbound mail flows from the Edge Transport server to a Mailbox server in the subscribed Active Directory

site. Specifically:

- Mail from an Exchange 2013 or later Edge Transport server first arrives at the Front End Transport service before it flows to the Transport service on an Exchange 2016 or Exchange 2019 Mailbox server.
- In Exchange 2016, mail from an Exchange 2010 Edge Transport server always delivers mail directly to the Transport service on an Exchange 2016 Mailbox server. Note that coexistance with Exchange 2010 isn't supported in Exchange 2019.

For more information, see Mail flow and the transport pipeline.

The network ports that are required for mail flow in Exchange organizations that have Edge Transport servers are described in the following diagram and table.

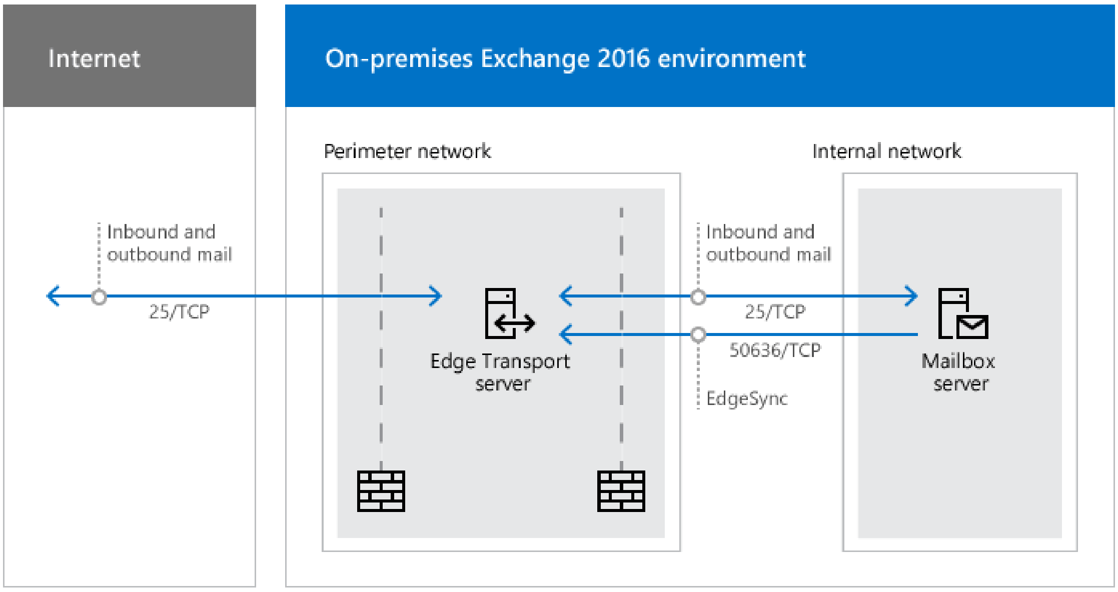

PURPOSE	PORTS	SOURCE	DESTINATION	COMMENTS
Inbound mail - Internet to Edge Transport server	25/TCP (SMTP)	Internet (any)	Edge Transport server	The default Receive connector named "Default internal Receive connector *<Edge Transport server name>*" on the Edge Transport server listens for anonymous SMTP mail on port 25.

PURPOSE	PORTS	SOURCE	DESTINATION	COMMENTS
Inbound mail - Edge Transport server to internal Exchange organization	25/TCP (SMTP)	Edge Transport server	Mailbox servers in the subscribed Active Directory site	The default Send connector named "EdgeSync - Inbound to *<Active Directory site name>*" relays inbound mail on port 25 to any Mailbox server in the subscribed Active Directory site. For more information, see Send connectors created automatically by the Edge Subscription. The default Receive connector named "Default Frontend *<Mailbox server name>*" in the Front End Transport service on the Mailbox server listens for all inbound mail (including mail from Exchange 2013 or later Edge Transport servers) on port 25.

PURPOSE	PORTS	SOURCE	DESTINATION	COMMENTS
Outbound mail - Internal Exchange organization to Edge Transport server	25/TCP (SMTP)	Mailbox servers in the subscribed Active Directory site	Edge Transport servers	Outbound mail always bypasses the Front End Transport service on Mailbox servers. Mail is relayed from the Transport service on any Mailbox server in the subscribed Active Directory site to an Edge Transport server using the implicit and invisible intra-organization Send connector that automatically routes mail between Exchange servers in the same organization. The default Receive connector named "Default internal Receive connector *<Edge Transport server name>*" on the Edge Transport server listens for SMTP mail on port 25 from the Transport service on any Mailbox server in the subscribed Active Directory site.
Outbound mail - Edge Transport server to internet	25/TCP (SMTP)	Edge Transport server	Internet (any)	The default Send connector named "EdgeSync - *<Active Directory site name>* to Internet" relays outbound mail on port 25 from the Edge Transport server to the internet.
EdgeSync synchronization	50636/TCP (secure LDAP)	Mailbox servers in the subscribed Active Directory site that participate in EdgeSync synchronization	Edge Transport servers	When the Edge Transport server is subscribed to the Active Directory site, all Mailbox servers that exist in the site *at the time* participate in EdgeSync synchronization. However, any Mailbox servers that you add later don't *automatically* participate in EdgeSync synchronization.

PURPOSE	PORTS	SOURCE	DESTINATION	COMMENTS
DNS for name resolution of the next mail hop (not pictured)	53/UDP,53/TCP (DNS)	Edge Transport server	DNS server	See the Name resolution section later in this topic.
Open proxy server detection in sender reputation (not pictured)	see comments	Edge Transport server	Internet	By default, sender reputation (the Protocol Analysis agent) uses open proxy server detection as one of the criteria to calculate the sender reputation level (SRL) of the source messaging server. For more information, see Sender reputation and the Protocol Analysis agent. Open proxy server detection uses the following protocols and TCP ports to test source messaging servers for open proxy: • SOCKS4, SOCKS5: 1081, 1080 • Wingate, Telnet, Cisco: 23 • HTTP CONNECT, HTTP POST: 6588, 3128, 80 Also, if your organization uses a proxy server to control outbound internet traffic, you need to define the proxy server name, type, and TCP port that sender reputation requires to access the internet for open proxy server detection. Alternatively, you can disable open proxy server detection in sender reputation. For more information, see Sender reputation procedures.

Name resolution

DNS resolution of the next mail hop is a fundamental part of mail flow in any Exchange organization. Exchange servers that are responsible for receiving inbound mail or delivering outbound mail must be able to resolve both internal and external host names for proper mail routing. And all internal Exchange servers must be able to resolve internal host names for proper mail routing. There are many different ways to design a DNS infrastructure, but the important result is to ensure name resolution for the next hop is working properly for all of your Exchange servers.

Network ports required for hybrid deployments

The network ports that are required for an organization that uses both on-premises Exchange and Microsoft 365 or Office 365 are covered in Hybrid deployment protocols, ports, and endpoints.

Network ports required for Unified Messaging in Exchange 2016

The network ports that are required for Unified Messaging in Exchange 2013 and Exchange 2016 are covered in the topic UM protocols, ports, and services.

Overview of Exchange services on Exchange servers

8/3/2020 • 8 minutes to read • Edit Online

During the installation of Exchange Server 2016 or Exchange Server 2019, Setup runs a set of tasks that install new services in Microsoft Windows. A service is a background process that can be launched during the startup of the server by the Windows Service Control Manager. Services are executable files designed to operate independently and without administrative intervention. A service can run using either a graphical user interface (GUI) mode or a console mode.

All previous versions of Exchange included components that are implemented as services. Each Exchange server role includes services that are part of (or may be needed by) the server role to perform its functions. Note that some services only become active when specific features are used.

The sections in this topic describe the various services that are installed by Exchange 2016 and Exchange 2016 on Mailbox servers and Edge Transport servers. For services that are labeled as optional, you can disable the service if you determine your organization doesn't need the functionality that's provided by the service.

Exchange services on Mailbox servers

The following table describes the Exchange services that are installed on Mailbox servers.

SERVICE NAME	SERVICE SHORT NAME	DESCRIPTION AND DEPENDENCIES	DEFAULT STARTUP TYPE	SECURITY CONTEXT	DEPENDENCIES	REQUIRED OR OPTIONAL
Microsoft Exchange Active Directory Topology	MSExchangeADTopology	Provides Active Directory topology information to Exchange services. If this service is stopped, most Exchange services can't start.	Automatic	Local System	Net.TCP Port Sharing Service	Required

SERVICE NAME	SERVICE SHORT NAME	DESCRIPTION AND DEPENDENCIES	DEFAULT STARTUP TYPE	SECURITY CONTEXT	DEPENDENCIES	REQUIRED OR OPTIONAL
Microsoft Exchange Anti-spam Update	MSExchangeAntispamUpdate	Provides Exchange SmartScreen spam definition updates. **Note**: In November, 2016, Microsoft stopped producing spam definition updates for the SmartScreen filters in Exchange and Outlook. The existing SmartScreen spam definitions were left in place, but their effectiveness will likely degrade over time. For more information, see Deprecating support for SmartScreen in Outlook and Exchange.	Automatic	Local System	Microsoft Exchange Active Directory Topology	Optional
Microsoft Exchange Compliance Audit	MSComplianceAudit	Provides Exchange auditing features.	Automatic	Local System	Microsoft Exchange Active Directory Topology	Required
Microsoft Exchange Compliance Service	MSExchangeCompliance	Provides a host for Exchange compliance services.	Automatic	Local System	Microsoft Exchange Active Directory Topology	Required

SERVICE NAME	SERVICE SHORT NAME	DESCRIPTION AND DEPENDENCIES	DEFAULT STARTUP TYPE	SECURITY CONTEXT	DEPENDENCIES	REQUIRED OR OPTIONAL
Microsoft Exchange DAG Management	MSExchangeDagMgmt	Provides storage and database layout management for Mailbox servers in database availability groups (DAGs).	Automatic	Local System	Microsoft Exchange Active Directory TopologyNet.TCP Port Sharing Service	Required
Microsoft Exchange Diagnostics	MSExchangeDiagnostics	Provides an agent that monitors Exchange server health.	Automatic	Local System	None	Required
Microsoft Exchange EdgeSync	MSExchangeEdgeSync	Replicates configuration and recipient data between the Mailbox server and Active Directory Lightweight Directory Services (AD LDS) on subscribed Edge Transport servers over a secure LDAP channel. If you don't have any subscribed Edge Transport servers, you can disable this service.	Automatic	Local System	Microsoft Exchange Active Directory Topology	Optional
Microsoft Exchange Frontend Transport	MSExchangeFrontEndTransport	Proxies SMTP connections from external hosts to the Microsoft Exchange Transport service on Mailbox servers (the local server or remote servers).	Automatic	Local System	Microsoft Exchange Active Directory Topology	Required

SERVICE NAME	SERVICE SHORT NAME	DESCRIPTION AND DEPENDENCIES	DEFAULT STARTUP TYPE	SECURITY CONTEXT	DEPENDENCIES	REQUIRED OR OPTIONAL
Microsoft Exchange Health Manager	MSExchangeHM	Part of managed availability that monitors the health of key components on the Exchange server.	Automatic	Local System	Windows Event LogWindows Management Instrumentation	Required
Microsoft Exchange Health Manager Recovery	MSExchangeHMRecovery	Part of managed availability that attempts to recover unhealthy components on the Exchange server.	Automatic	Local System	• Windows Event Log • Windows Management Instrumentation	Required

SERVICE NAME	SERVICE SHORT NAME	DESCRIPTION AND DEPENDENCIES	DEFAULT STARTUP TYPE	SECURITY CONTEXT	DEPENDENCIES	REQUIRED OR OPTIONAL
Microsoft Exchange IMAP4	MSExchangeIMAP4	Proxies IMAP4 client connections from the Client Access (frontend) services to the backend IMAP4 service on Mailbox servers. By default, this service isn't running, so IMAP4 clients can't connect to the Exchange server until this service is started. If you don't have any IMAP4 clients, you can disable this service.	Manual	Local System	Microsoft Exchange Active Directory Topology	Optional
Microsoft Exchange IMAP4 Backend	MSExchangeIMAP4BE	Receives proxied IMAP4 client connections from the from the Client Access (frontend) IMAP4 service. By default, this service isn't running, so IMAP4 clients can't connect to the Exchange server until this service is started. If you don't have any IMAP4 clients, you can disable this service.	Manual	Network Service	Microsoft Exchange Active Directory Topology	Optional

SERVICE NAME	SERVICE SHORT NAME	DESCRIPTION AND DEPENDENCIES	DEFAULT STARTUP TYPE	SECURITY CONTEXT	DEPENDENCIES	REQUIRED OR OPTIONAL
Microsoft Exchange Information Store	MSExchangeIS	Manages the mailbox databases on the server. If this service is stopped, mailbox databases on the server are unavailable.	Automatic	Local System	• Microsoft Exchange Active Directory Topology • Remote Procedure Call (RPC) • Server • Windows Event Log • Workstation	Required
Microsoft Exchange Mailbox Assistants	MSExchangeMailboxAssistants	Performs background processing of mailboxes in mailbox databases on the server.	Automatic	Local System	Microsoft Exchange Active Directory Topology	Required
Microsoft Exchange Mailbox Replication	MSExchangeMailboxReplication	Processes mailbox moves and move requests.	Automatic	Local System	Microsoft Exchange Active Directory TopologyNet.TCP Port Sharing Service	Rquired
Microsoft Exchange Mailbox Transport Delivery	MSExchangeDelivery	Receives SMTP messages from the Microsoft Exchange Transport service (on the local or remote Mailbox servers) and delivers them to a local mailbox database using RPC.	Automatic	Network Service	Microsoft Exchange Active Directory Topology	Required

SERVICE NAME	SERVICE SHORT NAME	DESCRIPTION AND DEPENDENCIES	DEFAULT STARTUP TYPE	SECURITY CONTEXT	DEPENDENCIES	REQUIRED OR OPTIONAL
Microsoft Exchange Mailbox Transport Submission	MSExchangeSubmission	Receives RPC messages from a local mailbox database, and submits them over SMTP to the Microsoft Exchange Transport service (on the local or remote Mailbox servers).	Automatic	Local System	Microsoft Exchange Active Directory Topology	Required
Microsoft Exchange Notifications Broker	MSExchangeNotificationsBroker	Provides Exchange notifications to local and remote Exchange processes.	Automatic	Local System	Microsoft Exchange Active Directory Topology • Net.TCP Port Sharing Service	Required
Microsoft Exchange POP3	MSExchangePOP3	Proxies POP3 client connections from the Client Access (frontend) services to the backend IMAP4 service on Mailbox servers. By default, this service isn't running, so POP3 clients can't connect to the Exchange server until this service is started.	Manual	Network Service	Microsoft Exchange Active Directory Topology	Optional

SERVICE NAME	SERVICE SHORT NAME	DESCRIPTION AND DEPENDENCIES	DEFAULT STARTUP TYPE	SECURITY CONTEXT	DEPENDENCIES	REQUIRED OR OPTIONAL
Microsoft Exchange POP3 Backend	MSExchangePOP3BE	Receives proxied POP3 client connections from the from the Client Access (frontend) POP3 service. By default, this service isn't running, so POP3 clients can't connect to the Exchange server until this service is started.	Manual	Network Service	Microsoft Exchange Active Directory Topology	Optional
Microsoft Exchange Replication Service	MSExchangeRepl	Provides replication functionality for mailbox databases in a database availability groups (DAGs).	Automatic	Local System	Microsoft Exchange Active Directory Topology	Required
Microsoft Exchange RPC Client Access	MSExchangeRPC	Manages client RPC connections for Exchange.	Automatic	Network Service	Microsoft Exchange Active Directory Topology	Required
Microsoft Exchange Search	MSExchangeFastSearch	Provides indexing of mailbox content, which improves the performance of content search.	Automatic	Local System	Microsoft Exchange Active Directory Topology	Required
Microsoft Exchange Search Host Controller	HostControllerService	Provides deployment and management services for applications on the local Exchange server.	Automatic	Local System	HTTP Service	Requirerd

SERVICE NAME	SERVICE SHORT NAME	DESCRIPTION AND DEPENDENCIES	DEFAULT STARTUP TYPE	SECURITY CONTEXT	DEPENDENCIES	REQUIRED OR OPTIONAL
Microsoft Exchange Server Extension for Windows Server Backup	WSBExchange	Enables Windows Server Backup to back and restore Exchange server data.	Manual	Local System	None	Optional
Microsoft Exchange Service Host	MSExchangeServiceHost	Provides a service host for Exchange components that don't have their own services.	Automatic	Local System	Microsoft Exchange Active Directory Topology	Required
Microsoft Exchange Throttling	MSExchangeThrottling	Provides user workload management that limits the rate of user operations (formerly known as user throttling).	Automatic	Network Service	Microsoft Exchange Active Directory Topology	Required
Microsoft Exchange Transport	MSExchangeTransport	Provides SMTP server and transport stack.	Automatic	Network Service	Microsoft Exchange Active Directory Topology • Microsoft Filtering Management Service	Required
Microsoft Exchange Transport Log Search	MSExchangeTransportLogSearch	Provides remote search capability for transport log files (for example, message tracking).	Automatic	Local System	Microsoft Exchange Active Directory Topology	Optional

SERVICE NAME	SERVICE SHORT NAME	DESCRIPTION AND DEPENDENCIES	DEFAULT STARTUP TYPE	SECURITY CONTEXT	DEPENDENCIES	REQUIRED OR OPTIONAL
Microsoft Exchange Unified Messaging (Exchange 2016 only)	MSExchangeUM	Provides Unified Messaging (UM) features: allows voice and fax messages to be stored in Exchange 2016 and gives users telephone access to email, voice mail, calendar, contacts, or an auto attendant. If this service is stopped, Unified Messaging isn't available. If you don't use UM in Exchange 2016, you can disable this service.	Automatic	Local System	• CNG Key Isolation • Microsoft Exchange Active Directory Topology	Optional
Microsoft Exchange Unified Messaging Call Router (Exchange 2016 only)	MSExchangeUMCR	Redirects UM client connections from the Client Access (frontend) services to the backend Unified Messaging service on Exchange 2016 Mailbox servers. If you don't use UM in Exchange 2016, you can disable this service.	Automatic	Local System	• CNG Key Isolation • Microsoft Exchange Active Directory Topology	Optional

Exchange services on Edge Transport servers

The following table describes the Exchange services that are installed on Edge Transport servers.

SERVICE NAME	SERVICE SHORT NAME	DESCRIPTION	DEFAULT STARTUP TYPE	SECURITY CONTEXT	DEPENDENCIES	REQUIRED OR OPTIONAL
Microsoft Exchange ADAM	ADAM_MSExchange	Stores configuration data and recipient data on the Edge Transport server. This service represents the named instance of the Active Directory Lightweight Directory Services (AD LDS) that's automatically created by Exchange Setup.	Automatic	Network Service	COM+ Event System	Required

SERVICE NAME	SERVICE SHORT NAME	DESCRIPTION	DEFAULT STARTUP TYPE	SECURITY CONTEXT	DEPENDENCIES	REQUIRED OR OPTIONAL
Microsoft Exchange Anti-spam Update	MSExchangeAntispamUpdate	Provides Exchange SmartScreen spam definition updates. **Note**: In November, 2016, Microsoft stopped producing spam definition updates for the SmartScreen filters in Exchange and Outlook. The existing SmartScreen spam definitions were left in place, but their effectiveness will likely degrade over time. For more information, see Deprecating support for SmartScreen in Outlook and Exchange.	Automatic	Local System	Microsoft Exchange ADAM	Optional
Microsoft Exchange Credential Service	MSExchangeEdgeCredential	Monitors credential changes in Active Directory Lightweight Directory Services (AD LDS) and installs the changes on the Edge Transport server.	Automatic	Local System	Microsoft Exchange ADAM	Required

SERVICE NAME	SERVICE SHORT NAME	DESCRIPTION	DEFAULT STARTUP TYPE	SECURITY CONTEXT	DEPENDENCIES	REQUIRED OR OPTIONAL
Microsoft Exchange Diagnostics	MSExchangeDiagnostics	Provides an agent that monitors Exchange server health.	Automatic	Local System	None	Required
Microsoft Exchange Health Manager	MSExchangeHM	Part of managed availability that monitors the health of key components on the Exchange server.	Automatic	Local System	• Windows Event Log • Windows Management Instrumentation	Required
Microsoft Exchange Health Manager Recovery	MSExchangeHMRecovery	Part of managed availability that attempts to recover unhealthy components on the Exchange server.	Automatic	Local System	Windows Event LogWindows Management Instrumentation	Required
Microsoft Exchange Service Host	MSExchangeServiceHost	Provides a service host for Exchange components that don't have their own services.	Automatic	Local System	Microsoft Exchange ADAM	Required
Microsoft Exchange Transport	MSExchangeTransport	Provides SMTP server and transport stack.	Automatic	Network Service	Microsoft Exchange ADAM	Required
Microsoft Exchange Transport Log Search	MSExchangeTransportLogSearch	Provides remote search capability for transport log files (for example, message tracking).	Automatic	Local System	Microsoft Exchange ADAM	Optional

Exchange 2019 preferred architecture

8/3/2020 • 16 minutes to read • Edit Online

With each new release of Exchange Server for our on-premises customers we update our Preferred Architecture and discuss what changes we would like our customers to be aware of. Exchange Server 2013 brought us the first of the Preferred Architectures in modern Exchange history and was then followed with a refresh for Exchange Server 2016 by providing refinements for the changes that came with the 2016 release. With this update for Exchange Server 2019 we will iterate on the previous PA to take advantage of new technologies and improvements.

The preferred architecture

The PA is the Exchange Server Engineering Team's best practice recommendation for what we believe is the optimum deployment architecture for Exchange Server 2019 in an on-premises environment.

While Exchange 2019 offers a wide variety of architectural choices for on-premises deployments, the architecture discussed here is the most scrutinized one. While there are other supported deployment architectures, they are not our recommended practice.

Following the PA helps customers become a member of a community of organizations with similar Exchange Server deployments. This allows easier knowledge sharing and facilitates a more rapid response to unforeseen circumstances. Our own support organization is well aware what an Exchange Server PA deployment should look like and prevents them from spending lengthy cycles learning and understanding a customer's highly custom environment before working with them towards a support case resolution.

The PA is designed with several business requirements in mind, such as the requirement that the architecture be able to:

- Include both high availability within the datacenter, and site resilience between datacenters
- Support multiple copies of each database, thereby allowing for quick activation
- Reduce the cost of the messaging infrastructure
- Increase availability by optimizing around failure domains and reducing complexity

The specific prescriptive nature of the PA means of course that not every customer will be able to deploy it word for word. For example, not all our customers have multiple datacenters. Some of our customers may have different business requirements or internal policies they must adhere to which necessitate a different deployment architecture. If you fall into those categories, and you want to deploy Exchange on-premises, there are still advantages to adhering as closely as possible to the PA and deviate only where your requirements or policies force you to differ. Alternatively, you can always consider Microsoft 365 or Office 365 where you no longer must deploy or manage a large number of servers.

The PA removes complexity and redundancy where necessary to drive the architecture to a predictable recovery model: when a failure occurs, another copy of the affected database is activated.

The PA covers the following four areas of focus:

1. Namespace design
2. Site resilient datacenter pair design
3. Server design

4. Database availability group design

For Exchange Server 2019 we have no changes in three of the four categories from the Exchange Server 2016 Preferred Architecture. The areas of Namespace design, Datacenter design, and DAG design are receiving no major changes. We have been very pleased with customer deployments that closely followed the Exchange Server 2016 PA and see no need to deviate from the recommendations in those areas.

The most noteworthy changes in the Exchange Server 2019 PA focus on the area of Server design due to some new and exciting technologies.

Namespace design

In the Namespace Planning and Load Balancing Principles articles for Exchange Server 2016, Ross Smith IV outlined the various configuration choices that were available with Exchange 2016 and these concepts continue to apply for Exchange Server 2019. For the namespace, the choices are to either deploy a bound namespace (having a preference for the users to operate out of a specific datacenter) or an unbound namespace (having the users connect to any datacenter without preference).

The recommended approach is to utilize the unbounded model, deploying a single Exchange namespace per client protocol for the site resilient datacenter pair (where each datacenter is assumed to represent its own Active Directory site - see more details on that below). For example:

- For the Autodiscover service: autodiscover.contoso.com
- For HTTP clients: mail.contoso.com
- For IMAP clients: imap.contoso.com
- For SMTP clients: smtp.contoso.com

Each Exchange namespace is load balanced across both datacenters in a layer 7 configuration that does not leverage session affinity, resulting in fifty percent of traffic being proxied between datacenters. Traffic is equally distributed across the datacenters in the site resilient pair, via round robin DNS, geo-DNS, or other similar solutions. From our perspective, the simpler solution is the least complex and easier to manage, so our recommendation is to leverage round robin DNS.

One caution we have for customers is to ensure you assign a low TTL (time to live) value for any DNS record associated with your Exchange architecture. In the event if a full datacenter outage when using round robin DNS you must maintain the ability to quickly update your DNS records to remove the IP addresses from the offline datacenter so they are not returned for DNS queries. For example, if your DNS records have a longer TTL value of 24 hours it may take up to a day for downstream DNS caches to properly update. If you do not perform this step you may find some clients are unable to properly transition to the still available IP addresses in your remaining datacenter. Do not forget to add the IP addresses back to your DNS records when your previously offline datacenter is recovered and ready to host services once again.

Datacenter affinity is required for the Office Online Server farms, thus a namespace is deployed per datacenter with the load balancer utilizing layer 7, and maintaining session affinity via cookie-based persistence.

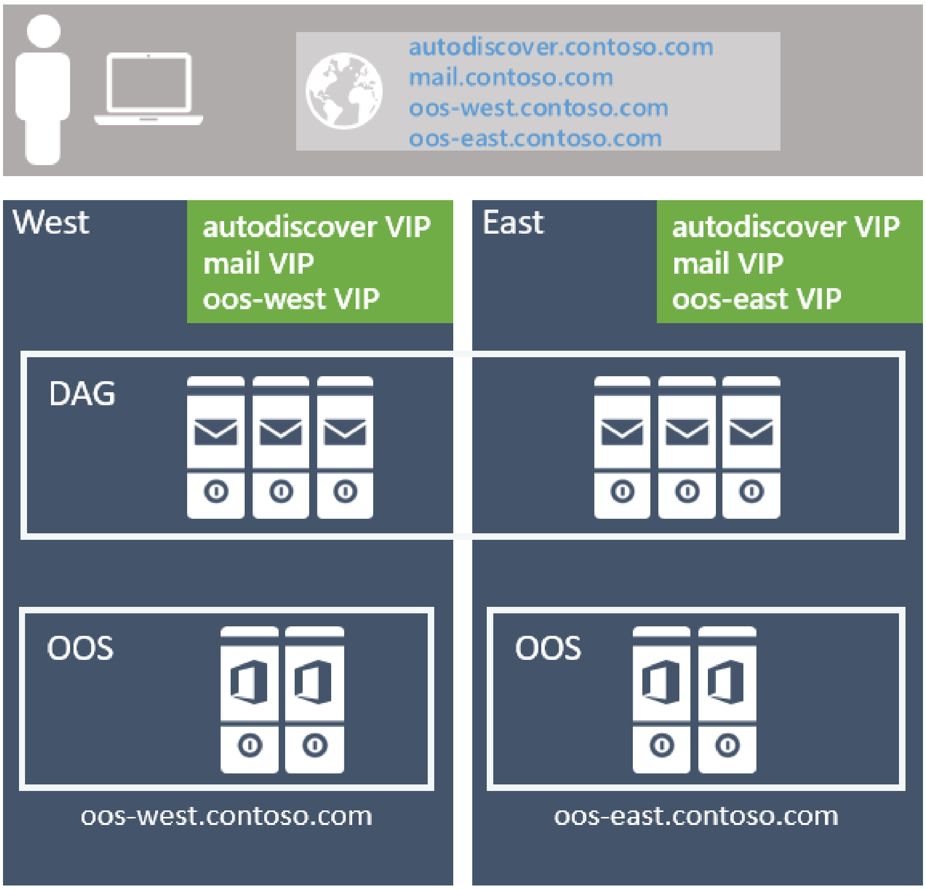

In the event that you have multiple site resilient datacenter pairs in your environment, you will need to decide if you want to have a single worldwide namespace, or if you want to control the traffic to each specific datacenter by using regional namespaces. Your decision depends on your network topology and the associated cost with using an unbound model; for example, if you have datacenters located in North America and South Africa, the network link between these regions might not only be costly, but it might also have high latency, which can introduce user pain and operational issues. In that case, it makes sense to deploy a bound model with a separate namespace for each region. However, options like geographical DNS offer you the ability to deploy a single unified namespace, even when you have costly network links; geo-DNS allows you to have your users directed to the closest datacenter based on their client's IP address.

Site resilient datacenter pair design

To achieve a highly available *and* site resilient architecture, you must have two or more datacenters that are well-connected (ideally, you want a low round-trip network latency, otherwise replication and the client experience are adversely affected). In addition, the datacenters should be connected via redundant network paths supplied by different operating carriers.

While we support stretching an Active Directory site across multiple datacenters, for the PA we recommend that each datacenter be its own Active Directory site. There are two reasons:

1. Transport site resilience via Shadow redundancy in Exchange Server and Safety Net in Exchange Server can only be achieved when the DAG has members located in more than one Active Directory site.

2. Active Directory has published guidance that states that subnets should be placed in different Active Directory sites when the round trip latency is greater than 10ms between the subnets.

Server design

In the PA, all servers are physical servers and use locally attached storage. Physical hardware is deployed rather than virtualized hardware for two reasons:

1. The servers are scaled to use 80% of resources during the worst-failure mode.
2. Virtualization comes with a slight performance penalty as well as adding an additional layer of management and complexity, which introduces additional recovery modes that do not add value, particularly since Exchange Server natively provides the same functionality.

Commodity servers

Commodity server platforms are used in the PA. Current commodity platforms are and include:

- 2U, dual socket servers with up to 48 physical processor cores (an increase from 24 cores in Exchange 2016)
- Up to 256GB of memory (an increase from 192GB in Exchange 2016)
- A battery-backed write cache controller
- 12 or more drive bays within the server chassis
- The ability to mix traditional rotating platter storage (HDD) and solid-state storage (SSD) within the same chassis.

Scale Theory

It is important to note even though we have increased the allowed processor and memory capacity in Exchange Server 2019 the Exchange Server PG's recommendation remains to scale out rather than up. Scaling out vs up means we would much rather see you deploy a larger number of servers with slightly less resources per server rather than a smaller number of very dense servers using maximum resources and populated with large numbers of mailboxes. By locating a reasonable number of mailboxes within a server, you lessen the impact of any planned or unplanned outage as well as reduce the risk of discovering other system bottlenecks.

An increase in system resources should not result in the assumption you will see linear performance gains in Exchange Server 2019 using the maximum allowed resources when comparing it to Exchange 2016's maximum allowed resources. Each new version of Exchange brings new processes and updates which in turn make it difficult to compare a current version to prior version. Please follow any and all sizing guidance from Microsoft when determining your server design.

Storage

Additional drive bays may be directly attached per-server depending on the number of mailboxes, mailbox size, and the server's resource scalability.

Each server houses a single RAID1 disk pair for the operating system, Exchange binaries, protocol/client logs, and the transport database.

The remaining storage is configured as JBOD (Just a Bunch of Disks). Please be aware some hardware storage controllers may require each disk to each be configured as a single-disk RAID0 group for write caching to be utilized. Consult with your hardware manufacturer to confirm the proper configuration for your system that guarantees write-cache will be utilized.

New to the Exchange Server 2019 PA is the recommendation of having two classes of storage for everything not already located on the RAID1 disk pair previously mentioned.

Traditional storage class

This storage class contains Exchange Server database files and Exchange Server transaction log files. These disks are large capacity 7.2K RPM serially attached SCSI (SAS) disks. While SATA disks are also available we observe better IO and a lower annualized failure rate using the SAS equivalent.

To ensure the capacity and IO of each disk is used as efficiently as possible, up to four database copies are deployed per-disk. The normal run-time copy layout ensures that there is no more than a single active database copy per disk.

At least one disk in the traditional storage disk pool is reserved as a hot spare. AutoReseed is enabled and quickly restores database redundancy after a disk failure by activating the hot spare and initiating database copy reseeds.

Solid state storage class

This storage class contains Exchange 2019's new MetaCache Database (MCDB) files. These solid-state drives may come in different form factors such as but not limited to traditional 2.5"/3.5" SAS connected or M.2 PCIe connected drives.

Customers should expect to deploy roughly 5-10% additional storage as solid state storage. For example, if a single server was expected to hold 28TB of mailbox database files on traditional storage, then an additional 1.4-2.8TB TB of solid-state storage would also be recommended as additional storage for the same server.

Traditional and solid-state disks should be deployed in a 3:1 ratio where possible. For every three traditional disks within the server a single solid-state disk will be deployed and hold the MCDBs for all DBs within the three associated traditional disks. This recommendation limits the failure domain a solid-state drive failure can impose on a system. When an SSD fails, Exchange 2019 will failover all database copies using that SSD for their MCDB to another DAG node with healthy MCDB resources for the affected database. Limiting the number of database failovers reduce the chance of impacting users if many more databases were sharing a smaller number of solid-state drives.

In the event of a solid-state drive failure Exchange High Availability services will attempt to mount the affected databases on different DAG nodes where a healthy MCDB for each affected database still exists. If for some reason no healthy MCDBs exist for one of the affected databases, then Exchange High Availability services will leave the local affected database copy running without the performance benefits of the MCDB.

For example, if a customer were to deploy a system capable of holding 20 drives it may have a layout like the following.

- 2 HDDs for OS mirror, Exchange Binaries, and Transport Database
- 12 HDDs for Exchange Database storage
- 1 HDD as the AutoReseed spare
- 4 SSDs for Exchange MCDBs that provide between 5-10% of the cumulative database storage capacity.
- Optionally a customer may elect to add a spare SSD or a second AutoReseed drive.

This can be visualized as the following;

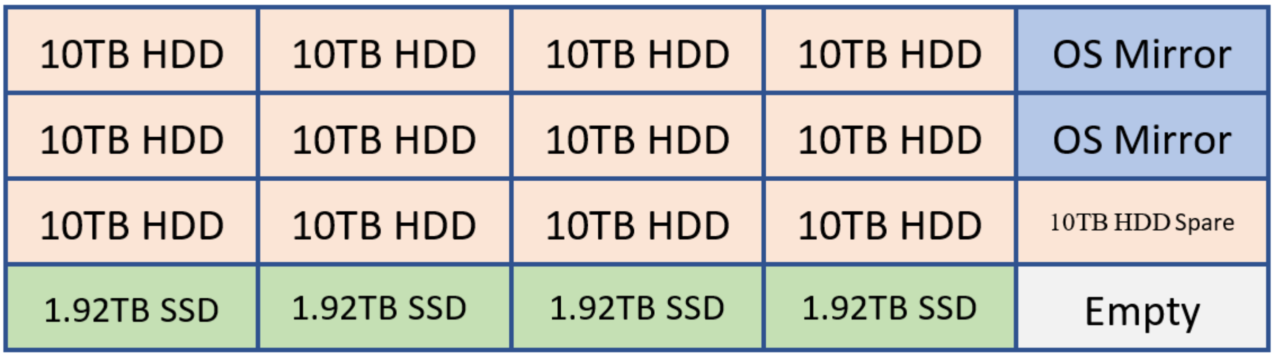

10TB HDD	10TB HDD	10TB HDD	10TB HDD	OS Mirror
10TB HDD	10TB HDD	10TB HDD	10TB HDD	OS Mirror
10TB HDD	10TB HDD	10TB HDD	10TB HDD	10TB HDD Spare
1.92TB SSD	1.92TB SSD	1.92TB SSD	1.92TB SSD	Empty

In the example above, we have 120 TB of Exchange database storage and 7.68 TB of MCDB storage which is roughly 6.4% the traditional database storage space. With this amount of MCBD storage we are perfectly aligned within the guidance of 5-10%. Each of the 10 TB drives will hold four database copies and each MCDB drive would hold twelve MCDBs.

Common storage settings

Whether Traditional or Solid-State, all disks that houses an Exchange data are formatted with ReFS (with the integrity feature disabled) and the DAG is configured such that AutoReseed formats the disks with ReFS:

```
Set-DatabaseAvailabilityGroup -Identity <DAGIdentity> -FileSystem ReFS
```

BitLocker is used to encrypt each disk, thereby providing data encryption at rest and mitigating concerns around data theft or disk replacement. For more information, see Enabling BitLocker on Exchange Servers.

Database availability group design

Within each site resilient datacenter pair, you will have one or more DAGs. It is not recommended to stretch a DAG across more than two datacenters.

DAG configuration

As with the namespace model, each DAG within the site resilient datacenter pair operates in an unbound model with active copies distributed equally across all servers in the DAG. This model:

1. Ensures that each DAG member's full stack of services (client connectivity, replication pipeline, transport, etc.) is being validated during normal operations.

2. Distributes the load across as many servers as possible during a failure scenario, thereby only incrementally increasing resource use across the remaining members within the DAG.

Each datacenter is symmetrical, with an equal number of DAG members in each datacenter. This means that each DAG has an even number of servers and uses a witness server for quorum maintenance.

The DAG is the fundamental building block in Exchange 2019. With respect to DAG size, a DAG with a greater number of participating member nodes provides more redundancy and resources. Within the PA, the goal is to deploy DAGs with a greater number of member nodes, typically starting with an eight-member DAG and increasing the number of servers as required to meet your requirements. You should only create new DAGs when scalability introduces concerns over the existing database copy layout.

DAG network design

The PA leverages a single, non-teamed network interface for both client connectivity and data replication. A single network interface is all that is needed because ultimately our goal is to achieve a standard recovery model regardless of the failure - whether a server failure occurs, or a network failure occurs, the result is the same: a database copy is activated on another server within the DAG. This architectural change simplifies the network stack and obviates the need to manually eliminate heartbeat cross-talk.

Witness server placement

The placement of the witness server determines whether the architecture can provide automatic datacenter failover capabilities or whether it will require a manual activation to enable service in the event of a site failure.

If your organization has a third location with a network infrastructure that is isolated from network failures that affect the site resilient datacenter pair in which the DAG is deployed, then the recommendation is to deploy the DAG's witness server in that third location. This configuration gives the DAG the ability to automatically failover databases to the other datacenter in response to a datacenter-level failure event, regardless of which datacenter has the outage.

If your organization does not have a third location, consider placing the server witness in Azure; alternatively, place the witness server in one of the datacenters within the site resilient datacenter pair. If you have multiple DAGs within the site resilient datacenter pair, then place the witness server for all DAGs in the same datacenter (typically the datacenter where the majority of the users are physically located). Also, make sure the Primary Active Manager (PAM) for each DAG is also located in the same datacenter.

Exchange Server 2019 and all earlier versions do not support the use of the Cloud Witness feature first introduced in Windows Server 2016 Failover Cluster.

Data resiliency

Data resiliency is achieved by deploying multiple database copies. In the PA, database copies are distributed across the site resilient datacenter pair, thereby ensuring that mailbox data is protected from software, hardware and even datacenter failures.

Each database has four copies, with two copies in each datacenter, which means at a minimum, the PA requires four servers. Out of these four copies, three of them are configured as highly available. The fourth copy (the copy with the highest Activation Preference number) is configured as a lagged database copy. Due to the server design, each copy of a database is isolated from its other copies, thereby reducing failure domains and increasing the overall availability of the solution as discussed in DAG: Beyond the "A".

The purpose of the lagged database copy is to provide a recovery mechanism for the rare event of system-wide, catastrophic logical corruption. It is not intended for individual mailbox recovery or mailbox item recovery.

The lagged database copy is configured with a seven day ReplayLagTime. In addition, the Replay Lag Manager is also enabled to provide dynamic log file play down for lagged copies when availability is compromised due to the loss of non-lagged copies.

By using the lagged database copy in this manner, it is important to understand that the lagged database copy is not a guaranteed point-in-time backup. The lagged database copy will have an availability threshold, typically around 90%, due to periods where the disk containing a lagged copy is lost due to disk failure, the lagged copy becoming an HA copy (due to automatic play down), as well as, the periods where the lagged database copy is re-building the replay queue.

To protect against accidental (or malicious) item deletion, Single Item Recover or In-Place Hold technologies are used, and the Deleted Item Retention window is set to a value that meets or exceeds any defined item-level recovery SLA.

With all of these technologies in play, traditional backups are unnecessary; as a result, the PA leverages Exchange Native Data Protection.

Office Online Server design

At a minimum, you will want to deploy an Office Online Server (OOS) farm with at least two OOS nodes in each datacenter that hosts Exchange 2019 servers. Each Office Online Server should have at least 8 processor cores, 32GB of memory and at least 40GB of space dedicated for log files. Exchange 2019 mailbox servers should be configured to rely on the local OOS farm in their datacenter to ensure the lowest possible latency and highest possible bandwidth between the servers to render file content to users.

Summary

Exchange Server 2019 continues to improve upon the investments introduced in previous versions of Exchange as well as introduces additional technologies originally invented for use in Microsoft 365 and Office 365.

By aligning with the Preferred Architecture you will take advantage of these changes and provide the best on-premises user experience possible. You will continue the tradition of having a highly reliable, predictable, and resilient Exchange deployment.

Exchange Server permissions

8/3/2020 • 12 minutes to read • Edit Online

Microsoft Exchange Server includes a large set of predefined permissions, based on the Role Based Access Control (RBAC) permissions model, which you can use right away to easily grant permissions to your administrators and users. You can use the permissions features in Exchange Server so that you can get your new organization up and running quickly.

Role-based permissions

In Exchange Server, the permissions that you grant to administrators and users are based on management roles. A role defines the set of tasks that an administrator or user can perform. For example, a management role called `Mail Recipients` defines the tasks that someone can perform on a set of mailboxes, contacts, and distribution groups. When a role is assigned to an administrator or user, that person is granted the permissions provided by the role.

There are two types of roles, administrative roles and end-user roles:

- **Administrative roles**: These roles contain permissions that can be assigned to administrators or specialist users using role groups that manage a part of the Exchange organization, such as recipients, servers, or databases.

- **End-user roles**: These roles, assigned using role assignment policies, enable users to manage aspects of their own mailbox and distribution groups that they own. End-user roles begin with the prefix `My`.

Roles give permissions to perform tasks to administrators and users by making cmdlets available to those who are assigned the roles. Because the Exchange admin center (EAC) and the Exchange Management Shell use cmdlets to manage Exchange, granting access to a cmdlet gives the administrator or user permission to perform the task in each of the Exchange management interfaces.

Role groups and role assignment policies

Roles grant permissions to perform tasks in Exchange Server, but you need an easy way to assign them to administrators and users. Exchange Server provides you with the following to help you do that:

- **Role groups**: Role groups enable you to grant permissions to administrators and specialist users.

- **Role assignment policies**: Role assignment policies enable you to grant permissions to end users to change settings on their own mailbox or distribution groups that they own.

For more information about role groups and role assignment policies, see the following sections.

Role groups

Every administrator that manages Exchange Server needs to be assigned at least one or more roles. Administrators might have more than one role because they may perform job functions that span multiple areas in Exchange. For example, one administrator might manage both recipients and Exchange servers. In this case, that administrator might be assigned both the `Mail Recipients` and `Exchange Servers` roles.

To make it easier to assign multiple roles to an administrator, Exchange Server includes role groups. Role groups are special universal security groups (USGs) used by Exchange Server that can contain Active Directory users, USGs, and other role groups. When a role is assigned to a role group, the permissions granted by the role are granted to all the members of the role group. This enables you to assign many roles to many role group members at once. Role groups typically encompass broader management areas, such as recipient management. They're used only with administrative roles, and not end-user roles.

> **NOTE**
>
> It's possible to assign a role directly to a user or USG without using a role group. However, that method of role assignment is an advanced procedure and isn't covered in this topic. We recommend that you use role groups to manage permissions.

The following figure shows the relationship between users, role groups, and roles.

Roles, role groups, and role group members

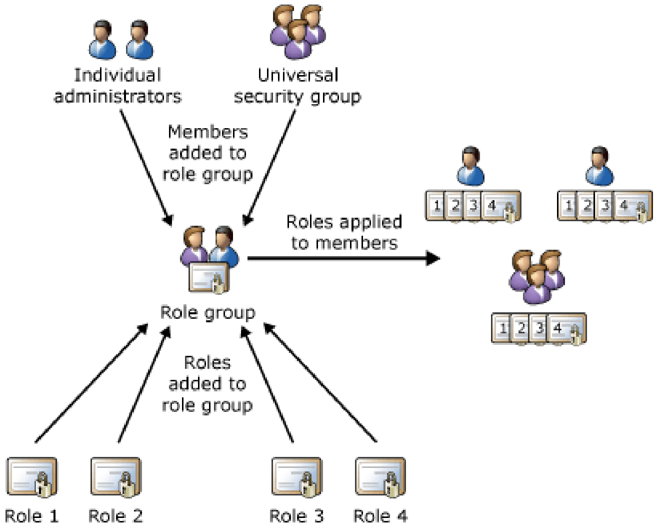

Exchange Server includes several built-in role groups, each one providing permissions to manage specific areas in Exchange Server. Some role groups may overlap with others. The following table lists each role group with a description of its use. If you want to see the roles assigned to each role group, click the name of the role group in the "Role group" column, and then open the "Management Roles Assigned to This Role Group" section.

> **IMPORTANT**
>
> If an administrator is a member of more than one role group, Exchange Server grants the administrator all of the permissions provided by the role groups he or she is a member of.

Built-in role groups

ROLE GROUP	DESCRIPTION
Organization Management	Administrators who are members of the Organization Management role group have administrative access to the entire Exchange Server organization and can perform almost any task against any Exchange Server object, with some exceptions, such as the `Discovery Management` role. **Important**: Because the Organization Management role group is a powerful role, only users or USGs that perform organizational-level administrative tasks that can potentially impact the entire Exchange organization should be members of this role group.
View-Only Organization Management	Administrators who are members of the View Only Organization Management role group can view the properties of any object in the Exchange organization.

ROLE GROUP	DESCRIPTION
Recipient Management	Administrators who are members of the Recipient Management role group have administrative access to create or modify Exchange Server recipients within the Exchange Server organization.
UM Management	Administrators who are members of the UM Management role group can manage features in the Exchange organization such as Unified Messaging (UM) service configuration, UM properties on mailboxes, UM prompts, and UM auto attendant configuration. (**Note**: UM is not available on Exchange 2019.)
Help Desk	The Help Desk role group, by default, enables members to view and modify the Outlook on the web (formerly known as Outlook Web App) options of any user in the organization. These options might include modifying the user's display name, address, and phone number. They don't include options that aren't available in Outlook on the web options, such as modifying the size of a mailbox or configuring the mailbox database on which a mailbox is located.
Hygiene Management	Administrators who are members of the Hygiene Management role group can configure the antivirus and antispam features of Exchange Server. Third-party programs that integrate with Exchange Server can add service accounts to this role group to grant those programs access to the cmdlets required to retrieve and configure the Exchange configuration.
Records Management	Users who are members of the Records Management role group can configure compliance features, such as retention policy tags, message classifications, and mail flow rules (also known as transport rules).
Discovery Management	Administrators or users who are members of the Discovery Management role group can perform searches of mailboxes in the Exchange organization for data that meets specific criteria and can also configure legal holds on mailboxes.
Public Folder Management	Administrators who are members of the Public Folder Management role group can manage public folders on servers running Exchange Server.
Server Management	Administrators who are members of the Server Management role group can configure server-specific configuration of transport, Unified Messaging, client access, and mailbox features such as database copies, certificates, transport queues and Send connectors, virtual directories, and client access protocols. (**Note**: UM is not available on Exchange 2019.)
Delegated Setup	Administrators who are members of the Delegated Setup role group can deploy servers running Exchange Server that have been previously provisioned by a member of the Organization Management role group.

ROLE GROUP	DESCRIPTION
Compliance Management	Users who are members of the Compliance Management role group can configure and manage Exchange compliance settings in accordance with their organization's policy.

If you work in a small organization that has only a few administrators, you might only ever use the Organization Management role group, and none of the others. If you work in a larger organization, you might have administrators who perform specific tasks administering Exchange, such as recipient or server management. In those cases, you might add one administrator to the Recipient Management role group, and another administrator to the Server Management role group. Those administrators can then manage their specific areas of Exchange Server but won't have permissions to manage areas they're not responsible for.

If you can't find a built-in role group that fits the jobs your administrators need to do, you can create role groups and add roles to them. For more information, see Work with role groups later in this topic.

Role assignment policies

Exchange Server provides role assignment policies so that you can control what settings your users can configure on their own mailboxes and on distribution groups they own. These settings include their display name, contact information, voice mail settings, and distribution group membership.

Your Exchange Server organization can have multiple role assignment policies that provide different levels of permissions for the different types of users in your organizations. Some users can be allowed to change their address or create distribution groups, while others can't. It all depends on the role assignment policy associated with their mailbox. Role assignment policies are added directly to mailboxes, and each mailbox can only be associated with one role assignment policy at a time.

Of the role assignment policies in your organization, one is marked as default. The default role assignment policy is associated with new mailboxes that aren't explicitly assigned a specific role assignment policy when they're created. The default role assignment policy should contain the permissions that should be applied to the majority of your mailboxes.

Permissions are added to role assignment policies using end-user roles. End-user roles begin with `My` and grant permissions for users to manage only their mailbox or distribution groups they own. They can't be used to manage any other mailbox. Only end-user roles can be assigned to role assignment policies.

When an end-user role is assigned to a role assignment policy, all of the mailboxes associated with that role assignment policy receive the permissions granted by the role. This enables you to add or remove permissions to sets of users without having to configure individual mailboxes. The following figure shows:

- End-user roles are assigned to role assignment policies. Role assignment policies can share the same end-user roles.
- Role assignment policies are associated with mailboxes. Each mailbox can only be associated with one role assignment policy.
- After a mailbox is associated with a role assignment policy, the end-user roles are applied to that mailbox. The permissions granted by the roles are granted to the user of the mailbox.

Roles, role assignment policies, and mailboxes

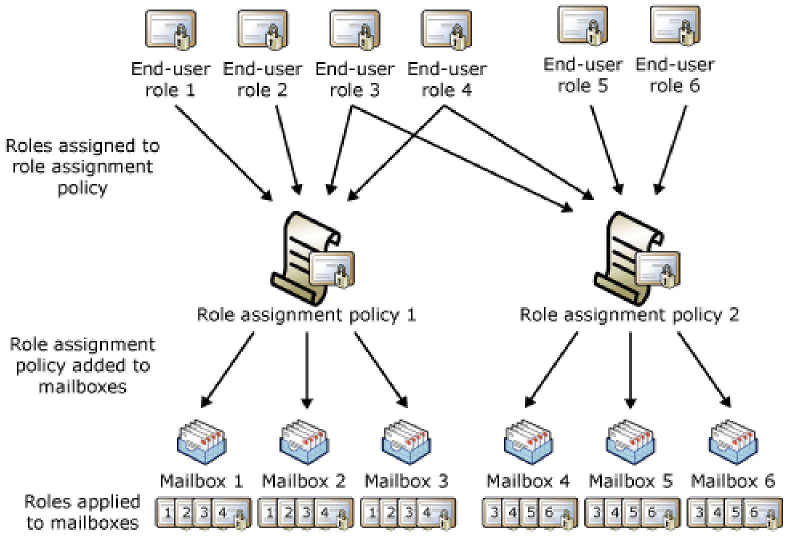

The Default Role Assignment Policy role assignment policy is included with Exchange Server. As the name implies, it's the default role assignment policy. If you want to change the permissions provided by this role assignment policy, or if you want to create role assignment policies, see Work with role assignment policies later in this topic.

Work with role groups

To manage your permissions using role groups in Exchange Server, we recommend that you use the Exchange admin center (EAC). When you use the EAC to manage role groups, you can add and remove roles and members, create role groups, and copy role groups with a few clicks of your mouse. The EAC provides simple dialog boxes, such as the **new role group** dialog box, shown in the following figure, to perform these tasks.

New role group dialog box in the EAC

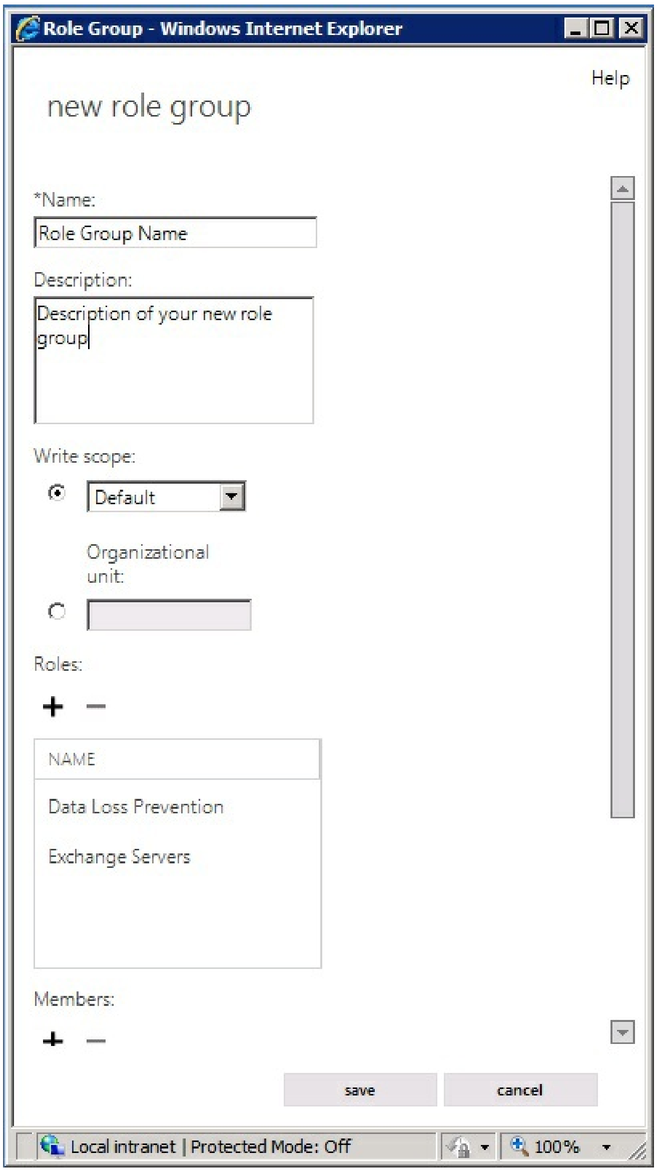
Role Group - Windows Internet Explorer
Help
save
cancel
Local intranet | Protected Mode: Off
100%

new role group

*Name: Role Group Name

Description: Description of your new role group

Write scope:

● Default

○ Organizational unit:

Roles:

NAME
Data Loss Prevention
Exchange Servers

Members:

If none of the role groups included with Exchange Server have the permissions you need, you can use the EAC to create a role group and add the roles that have the permissions you need. For your new role group, you'll need to:

1. Choose a name for your role group.

2. Select the roles you want to add to the role group.

3. Add members to the role group.

4. Save the role group.

After you create the role group, you manage it like any other role group.

If there's an existing role group that has some, but not all, of the permissions you need, you can copy it and then make changes to create a role group. Copying an existing role group lets you make changes to it without affecting the original role group. As part of copying the role group, you can add a new name and description, add and remove roles to and from the new role group, and add new members. When you create or copy a role group, you use the same dialog box that's shown in the preceding figure.

Existing role groups can also be modified. You can add and remove roles from existing role groups, and add and remove members from it at the same time, using an EAC dialog box similar to the one in the preceding figure. By adding and removing roles to and from role groups, you turn on and off administrative features for members of that role group.

NOTE

Although you can change which roles are assigned to built-in role groups, we recommend that you copy built-in role groups, modify the role group copy, and then add members to the role group copy.

Work with role assignment policies

To manage the permissions that you grant end users to manage their own mailbox in Exchange Server, we recommend that you use the EAC. When you use the EAC to manage end-user permissions, you can add roles, remove roles, and create role assignment policies with a few clicks of your mouse. The EAC provides simple dialog boxes, such as the **role assignment policy** dialog box, shown in the following figure, to perform these tasks.

Role assignment policy dialog box in the EAC

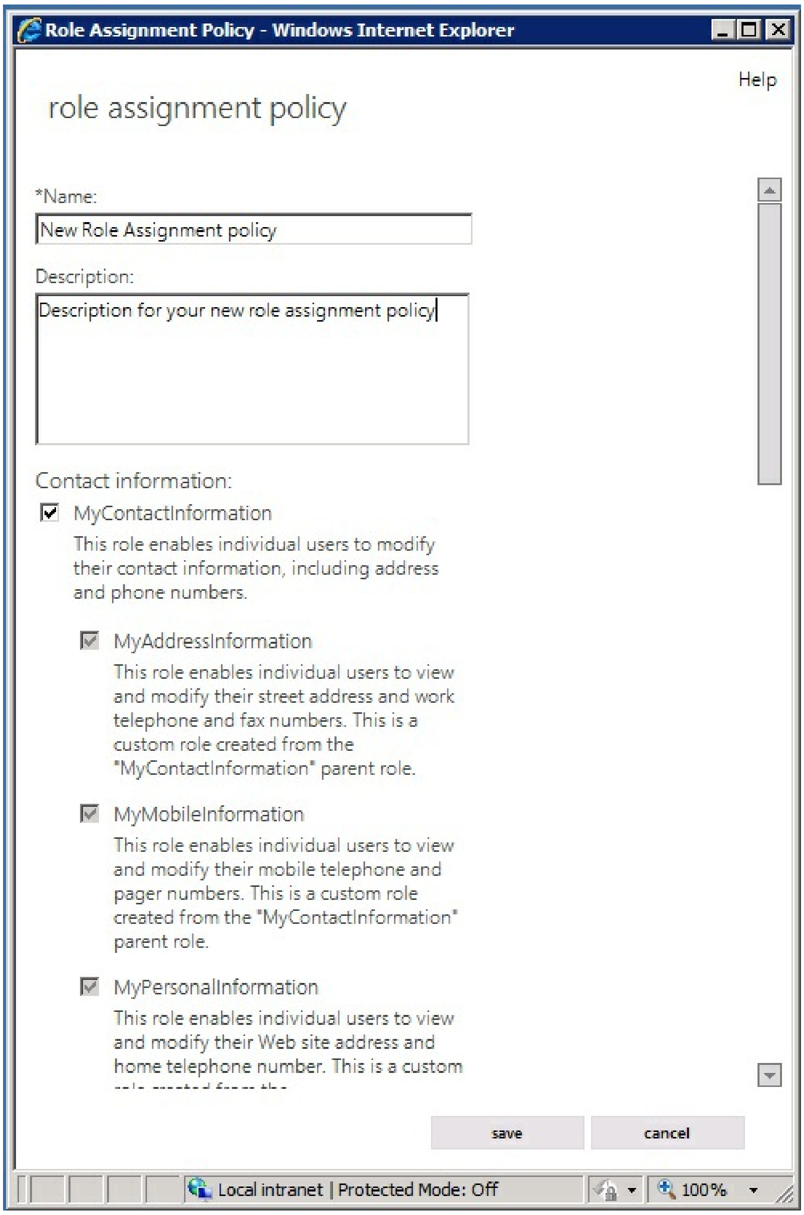
Role Assignment Policy - Windows Internet Explorer
Help
save
cancel
Local intranet | Protected Mode: Off
100%

role assignment policy

*Name:

New Role Assignment policy

Description:

Description for your new role assignment policy

Contact information:

☑ MyContactInformation

This role enables individual users to modify their contact information, including address and phone numbers.

☑ MyAddressInformation

This role enables individual users to view and modify their street address and work telephone and fax numbers. This is a custom role created from the "MyContactInformation" parent role.

☑ MyMobileInformation

This role enables individual users to view and modify their mobile telephone and pager numbers. This is a custom role created from the "MyContactInformation" parent role.

☑ MyPersonalInformation

This role enables individual users to view and modify their Web site address and home telephone number. This is a custom role created from the

Exchange Server includes a role assignment policy named Default Role Assignment Policy. This role assignment policy enables users whose mailboxes are associated with it to do the following:

- Join or leave distribution groups that allow members to manage their own membership.
- View and modify basic mailbox settings on their own mailbox, such as Inbox rules, spelling behavior, junk mail settings, and Microsoft ActiveSync devices.
- Modify their contact information, such as work address and phone number, mobile phone number, and

pager number.

- Create, modify, or view text message settings.
- View or modify voice mail settings.
- View and modify their marketplace apps.
- Create team mailboxes and connect them to Microsoft SharePoint lists.

If you want to add or remove permissions from the Default Role Assignment Policy or any other role assignment policy, you can use the EAC. When you open the role assignment policy in the EAC, select the check box next to the roles you want to assign to it or clear the check box next to the roles you want to remove. The change you make to the role assignment policy is applied to every mailbox associated with it.

If you want to assign different end-user permissions to the various types of users in your organization, you can create role assignment policies. You can specify a new name for the role assignment policy, and then select the roles you want to assign to the role assignment policy. After you create a role assignment policy, you can associate it with mailboxes using the EAC.

If you want to change which role assignment policy is the default, you needs to use the Exchange Management Shell. When you change the default role assignment policy, any mailboxes that are created will be associated with the new default role assignment policy if one wasn't explicitly specified. The role assignment policy associated with existing mailboxes doesn't change when you select a new default role assignment policy.

Notes:

- If you select a check box for a role that has child roles, the check boxes for the child roles are also selected. If you clear the check box for a role with child roles, the check boxes for the child roles are also cleared.
- For detailed steps about how to create role assignment policies or make changes to existing role assignment policies, see the following topics:
  - Manage role assignment policies
  - Change the assignment policy on a mailbox

Manage role groups

8/3/2020 • 24 minutes to read • Edit Online

A management role group is a universal security group (USG) used in the Role Based Access Control (RBAC) permissions model in Exchange Server. A management role group simplifies the assignment of management roles to a group of users. All members of a role group are assigned the same set of roles. For more information about role groups in Exchange Server, see Understanding Management Role Groups.

For additional management tasks related to role groups, see Permissions.

What do you need to know before you begin?

- Estimated time to complete each procedure: 5 to 10 minutes
- To open the EAC, see Exchange admin center in Exchange Server. To open the Exchange Management Shell, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Role groups" entry in the Role management permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Create a role group

If you want to customize the permissions that you can assign to a group of users, create a new custom management role group.

Use the EAC to create a role group

1. In the Exchange admin center (EAC), navigate to **Permissions** > **Admin Roles** and then click **Add** +.
2. In the **New role group** window, provide a name for the new role group.
3. You can either select the roles that you want to be assigned to the role group and the members you want to be added to the role group now, or you can do this at another time.
4. Select the write scope that you want to apply to the new role group.
5. Click **Save** to create the role group.

Use the Exchange Management Shell to create a role group

To create a role group, see the Examples section in New-RoleGroup.

How do you know this worked?

To verify that you have successfully created a role group, do the following:

1. In the EAC, navigate to **Permissions** > **Admin Roles**.

2. Verify that the new role group appears in the role group list and then select it.

3. Verify that members, assigned roles, and scope that you specified on the new role group are listed in the role group details pane.

Copy a role group

Use the EAC to copy a role group

If you have a role group that contains the permissions you want to grant to users, but you want to apply a different management scope, or remove or add one or two management roles without having to add all the other roles manually, you can copy the existing role group.

> **IMPORTANT**
>
> You can't use the EAC to copy a role group if you've used the Exchange Management Shell to configure multiple management role scopes or exclusive scopes on the role group. If you've configured multiple scopes or exclusive scopes on the role group, you must use the Exchange Management Shell procedures later in this topic to copy the role group. For more information about management role scopes, see Understanding Management Role Scopes.

1. In the EAC, navigate to **Permissions** > **Admin Roles**.

2. Select the role group you want to copy and then click **Copy**.

3. In the **New role group** window, provide a name for the new role group.

4. Review the roles that have been copied to the new role group. Add or remove roles as necessary.

5. Review the write scope, and change it as necessary.

6. Review the members that have been copied to the new role group. Add or remove members as necessary.

7. Click **Save** to create the role group.

Use the Exchange Management Shell to copy a role group with no scope

1. Store the role group that you want to copy in a variable using the following syntax.

   ```
   $RoleGroup = Get-RoleGroup <name of role group to copy>
   ```

2. Create the new role group, and also add members to the role group and specify who can delegate the new role group to other users, using the following syntax.

   ```
   New-RoleGroup <name of new role group> -Roles $RoleGroup.Roles -Members <member1, member2, member3...>
   -ManagedBy <user1, user2, user3...>
   ```

   For example, the following commands copy the Organization Management role group, and name the new role group "Limited Organization Management". It adds the members Isabelle, Carter, and Lukas and can be delegated by Jenny and Katie.

   ```
   $RoleGroup = Get-RoleGroup "Organization Management"
   New-RoleGroup "Limited Organization Management" -Roles $RoleGroup.Roles -Members Isabelle, Carter,
   Lukas -ManagedBy Jenny, Katie
   ```

After the new role group is created, you can add or remove roles, change the scope of role assignments on the role, and more.

For detailed syntax and parameter information, see Get-RoleGroup and New-RoleGroup.

Use the Exchange Management Shell to copy a role group with a custom scope

1. Store the role group that you want to copy in a variable using the following syntax.

   ```
   $RoleGroup = Get-RoleGroup <name of role group to copy>
   ```

2. Create the new role group with a custom scope using the following syntax.

   ```
   New-RoleGroup <name of new role group> -Roles $RoleGroup.Roles -CustomRecipientWriteScope <recipient
   scope name> -CustomConfigWriteScope <configuraiton scope name>
   ```

For example, the following commands copy the Organization Management role group and create a new role group called Vancouver Organization Management with the Vancouver Users recipient scope and Vancouver Servers configuration scope.

```
$RoleGroup = Get-RoleGroup "Organization Management"
New-RoleGroup "Vancouver Organization Management" -Roles $RoleGroup.Roles -CustomRecipientWriteScope
"Vancouver Users" -CustomConfigWriteScope "Vancouver Servers"
```

You can also add members to the role group when you create it by using the *Members* parameter as shown in Use the Exchange Management Shell to create a role assignment with no scope earlier in this topic. For more information about management scopes, see Understanding Management Scopes.

After the new role group is created, you can add or remove roles, change the scope of role assignments on the role, and perform other tasks.

For detailed syntax and parameter information, see Get-RoleGroup and New-RoleGroup.

Use the Exchange Management Shell to copy a role group with an OU scope

1. Store the role group that you want to copy in a variable using the following syntax.

   ```
   $RoleGroup = Get-RoleGroup <name of role group to copy>
   ```

2. Create the new role group with a custom scope using the following syntax.

   ```
   New-RoleGroup <name of new role group> -Roles $RoleGroup.Roles -RecipientOrganizationalUnitScope <OU
   name>
   ```

For example, the following commands copy the Recipient Management role group and create a new role group called Toronto Recipient Management that allows management of only users in the Toronto Users OU.

```
$RoleGroup = Get-RoleGroup "Recipient Management"
New-RoleGroup "Toronto Recipient Management" -Roles $RoleGroup.Roles -RecipientOrganizationalUnitScope
"contoso.com/Toronto Users"
```

You can also add members to the role group when you create it by using the *Members* parameter as shown in Use the Exchange Management Shell to create a role assignment with no scope earlier in this topic. For more information about management scopes, see Understanding Management Scopes.

After the new role group is created, you can add or remove roles, change the scope of role assignments on the role, and more.

For detailed syntax and parameter information, see Get-RoleGroup and New-RoleGroup.

How do you know this worked?

To verify that you have successfully copied a role group, do the following:

1. In the EAC, navigate to **Permissions** > **Admin Roles**.

2. Verify that the copied role group appears in the role group list, and then select it.

3. Verify that members, assigned roles, and scope that you specified on the copied role group are listed in the role group details pane.

Remove a role group

If you no longer need a role group you created, you can remove it. When you remove a role group, the management role assignments between the role group and the management roles are deleted. The management roles aren't deleted. If a user depended on the role group for access to a feature, the user will no longer have access to the feature. You can't remove built-in role groups.

Use the EAC to remove a role group

1. In the EAC, navigate to **Permissions** > **Admin Roles**.

2. Select the role group you want to remove and then click **Delete** 🗑.

3. Verify that you want to remove the selected role group, and if so, respond **Yes** to the warning.

Use the Exchange Management Shell to remove a role group

To remove a role group, see the Examples section in Remove-RoleGroup.

View role groups

You can view either a list of role groups or the detailed information about a specific role group that exists in your organization.

Use the EAC to view a list of role groups and role group details

1. In the EAC, navigate to **Permissions** > **Admin Roles**. All of the role groups in your organization are listed here.

2. Select a role group to view the members, assigned roles, and scope that are configured on the role group.

Use the Exchange Management Shell to view a list of role groups and role group details

To view a list of role groups, see the Examples section in Get-RoleGroup.

Add a role to a role group

Adding a management role to a role group is the best and simplest way to grant permissions to a group of administrators or specialist users. If you want to give users that are members of a role group the ability to manage a feature, you add the management role that manages the feature to the role group. After the role is added, the members of the role group are granted the permissions provided by the role.

Use the EAC to add a management role to a role group

> **IMPORTANT**
>
> You can't use the EAC to add roles to a role group if you've used the Exchange Management Shell to configure multiple management role scopes or exclusive scopes on the role group. If you've configured multiple scopes or exclusive scopes on the role group, you must use the Exchange Management Shell procedures later in this topic to add roles to the role group. For more information about management role scopes, see Understanding Management Role Scopes.

1. In the EAC, navigate to **Permissions** > **Admin Roles**.

2. Select the role group you want to add a role to, and then click **Edit** ✎.

3. In the **Roles** section, select the roles you want to add to the role group.

4. When you've finished adding roles to the role group, click **Save**.

Use the Exchange Management Shell to create a role assignment with no scope

You can create a role assignment with no scope between a role and a role group. When you do this, the implicit read and implicit write scopes of the role apply.

Use the following syntax to assign a role without any scope to a role group. A role assignment name is created automatically if you don't specify one.

```
New-ManagementRoleAssignment -SecurityGroup <role group name> -Role <role name>
```

This example assigns the Transport Rules management role to the Seattle Compliance role group.

```
New-ManagementRoleAssignment -SecurityGroup "Seattle Compliance" -Role "Transport Rules"
```

For detailed syntax and parameter information, see New-ManagementRoleAssignment.

Use the Exchange Management Shell to create a role assignment with a predefined scope

If a predefined scope meets your business requirements, you can apply that scope to the role assignment rather than create a new one. For a list of predefined scopes and their descriptions, see Understanding Management Role Scopes.

For more information about role assignments, see Understanding Management Role Assignments.

Use the following syntax to assign a role to a role group with a predefined scope. A role assignment name is created automatically if you don't specify one.

```
New-ManagementRoleAssignment -SecurityGroup <role group name> -Role <role name> -RecipientRelativeWriteScope <
MyGAL | MyDistributionGroups | Organization | Self >
```

This example assigns the Message Tracking role to the Enterprise Support role group and applies the Organization predefined scope.

```
New-ManagementRoleAssignment -SecurityGroup "Enterprise Support" -Role "Message Tracking" -
RecipientRelativeWriteScope Organization
```

For detailed syntax and parameter information, see New-ManagementRoleAssignment.

Use the Exchange Management Shell to create a role assignment with a recipient filter-based scope

If you created a recipient filter-based scope, you need to include the scope in the command used to assign the role to a role group by using the *CustomRecipientWriteScope* parameter.

You can also include a configuration write scope when you create a role assignment that has a recipient write scope.

For more information about role assignments and scopes, see the following topics:

- Understanding Management Role Assignments
- Understanding Management Role Scopes

Use the following syntax to assign a role to a role group with a recipient filter-based scope. A role assignment name is created automatically if you don't specify one.

```
New-ManagementRoleAssignment -SecurityGroup <role group name> -Role <role name> -CustomRecipientWriteScope
<role scope name>
```

This example assigns the Message Tracking role to the Seattle Recipient Admins role group and applies the Seattle Recipients scope.

```
New-ManagementRoleAssignment -SecurityGroup "Seattle Recipient Admins" -Role "Message Tracking" -
CustomRecipientWriteScope "Seattle Recipients"
```

For detailed syntax and parameter information, see New-ManagementRoleAssignment.

Use the Exchange Management Shell to create a role assignment with a configuration scope

If you created a server or database configuration filter or list-based scope, you need to include the scope in the command used to assign the role to a role group by using the *CustomConfigWriteScope* parameter.

You can also include a recipient write scope when you create a role assignment that has a configuration write scope.

For more information about role assignments and management scopes, see the following topics:

- Understanding Management Role Assignments
- Understanding Management Role Scopes

Use the following syntax to assign a role to a role group with a configuration scope. A role assignment name is created automatically if you don't specify one.

```
New-ManagementRoleAssignment -SecurityGroup <role group name> -Role <role name> -CustomConfigWriteScope <role
scope name>
```

This example assigns the Databases role to the Seattle Server Admins role group and applies the Seattle Servers scope.

```
New-ManagementRoleAssignment -SecurityGroup "Seattle Server Admins" -Role "Databases" -CustomConfigWriteScope
"Seattle Servers"
```

For detailed syntax and parameter information, see New-ManagementRoleAssignment.

Use the Exchange Management Shell to create a role assignment with an OU scope

If you want to scope a role's write scope to an OU, you can specify the OU in the *RecipientOrganizationalUnitScope* parameter directly.

For more information about role assignments and management scopes, see the following topics:

- Understanding Management Role Assignments
- Understanding Management Role Scopes

Use the following command to assign a role to a role group and restrict the write scope of a role to a specific OU. A role assignment name is created automatically if you don't specify one.

```
New-ManagementRoleAssignment -SecurityGroup <role group name> -Role <role name> -
RecipientOrganizationalUnitScope <OU>
```

This example assigns the Mail Recipients role to the Seattle Recipient Admins role group and scopes the assignment to the Sales\Users OU in the Contoso.com domain.

```
New-ManagementRoleAssignment -SecurityGroup "Seattle Recipient Admins" -Role "Mail Recipients" -
RecipientOrganizationalUnitScope contoso.com/sales/users
```

For detailed syntax and parameter information, see New-ManagementRoleAssignment.

How do you know this worked?

To verify that you have successfully added roles to a role group, do the following:

1. In the EAC, navigate to **Permissions** > **Admin Roles**.
2. Select the role group you added roles to. In the role group details pane, verify that the roles that you added are listed.

Remove a role from a role group

Removing a role from a management role group is the best and simplest way to revoke permissions granted to a group of administrators or specialist users. If you don't want administrators or specialist users to have permissions to manage a feature, you remove the management role from the management role group that manages the permissions. After the role is removed, the members of the role group will no longer have permissions to manage the feature.

> **NOTE**
>
> Some role groups, such as the Organization Management role group, restrict what roles can be removed from a role group. For more information, see Understanding Management Role Groups. > If an administrator is a member of another role group that contains management roles that grants permissions to manage the feature, you need to either remove the administrator from the other role groups, or remove the role that grants permissions to manage the feature from the other role groups.

Use the EAC to remove a management role from a role group

> **IMPORTANT**
>
> You can't use the EAC to remove roles from a role group if you've used the Exchange Management Shell to configure multiple scopes or exclusive scopes on the role group. If you've configured multiple scopes or exclusive scopes on the role group, you must use the Exchange Management Shell procedures later in this topic to remove roles from the role group. For more information about management role scopes, see Understanding Management Role Scopes.

1. In the EAC, navigate to **Permissions** > **Admin Roles**.
2. Select the role group you want to remove a role from, and then click **Edit** ✏.

3. In the **Roles** section, select the roles you want to remove from the role group.

4. When you've finished removing roles from the role group, click **Save**.

Use the Exchange Management Shell to remove a role from a role group

You can remove roles from role groups by retrieving the associated management role assignment using the **Get-ManagementRoleAssignment** cmdlet and then piping the role assignment returned to the **Remove-ManagementRoleAssignment** cmdlet. Unless you want to remove both delegating and regular role assignments at the same time, specify the *Delegating* parameter to specify whether you want to remove regular or delegating role assignments.

For more information about regular and delegating role assignments, see Understanding Management Role Assignments.

This procedure uses pipelining. For more information about pipelining, see about_Pipelines.

To remove a role from a role group, use the following syntax.

```
Get-ManagementRoleAssignment -RoleAssignee <role group name> -Role <role name> -Delegating <$true | $false> |
Remove-ManagementRoleAssignment
```

This example removes the Distribution Groups role, which enables administrators to manage distribution groups, from the Seattle Recipient Administrators role group. Because we want to remove the role assignment that provides permissions to manage distribution groups, the *Delegating* parameter is set to `$False`, which returns only regular role assignments.

```
Get-ManagementRoleAssignment -RoleAssignee "Seattle Recipient Administrators" -Role "Distribution Groups" -
Delegating $false | Remove-ManagementRoleAssignment
```

For detailed syntax and parameter information, see Remove-ManagementRoleAssignment.

How do you know this worked?

To verify that you have successfully removed roles from a role group, do the following:

1. In the EAC, navigate to **Permissions** > **Admin Roles**.

2. Select the role group you removed roles from. In the role group details pane, verify that the roles that you removed are no longer listed.

Change a role group's scope

The management role assignments between a role group and a role contain management scopes, which determine what objects are made available to members of that role group. By changing the write scope on a role group, you can change what objects are made available to role group members to create, change, or remove. You can't change the read scope on a role group.

Exchange Server includes scopes that are applied by default to role assignments when no custom scopes are created. If you want to use a custom scope with a role assignment on a role group, you must create one first. For more information about creating custom scopes, which is an advanced task, see Create a Regular or Exclusive Scope.

For more information about management role scopes and assignments in Exchange Server, see the following topics:

- Understanding Management Role Scopes

- Understanding Management Role Assignments

Use the EAC to change the scope on a role group

When you use the EAC to change the scope on a role group, you're actually changing the scope on all the role assignments between the role group and each of the management roles assigned to the role group. If you want to change the scope on specific role assignments, you must use the Exchange Management Shell procedures later in this topic.

> **IMPORTANT**
>
> You can't use the EAC to manage scopes on role assignments between roles and a role group if you've used the Exchange Management Shell to configure multiple scopes or exclusive scopes on those role assignments. If you've configured multiple scopes or exclusive scopes on those role assignments, you must use the Exchange Management Shell procedures later in this topic to manage scopes. For more information about management role scopes, see Understanding Management Role Scopes.

1. In the EAC, navigate to **Permissions** > **Admin Roles**.
2. Select the role group you want to change the scope on, and then click **Edit** ✎.
3. Select one of the two following **Write scope** options:
   - A write scope from the drop-down box, where you can select either the default write scope or a custom write scope.
   - **Organizational unit**: Select this option and provide an organizational unit (OU) if you want to scope this role group to an OU.
4. Click **Save** to save the changes to the role group.

Use the Exchange Management Shell to change the scope of all role assignments on a role group at the same time

Role assignments between the role group and the roles assigned to it can use the implicit scope obtained from the roles themselves, the same custom scope, or different custom scopes. For more information about role assignments, see Understanding Management Role Assignments.

The scopes on the role assignments are managed using the **Set-ManagementRoleAssignment** cmdlet. You can't manage scopes using the **Set-RoleGroup** cmdlet.

To change the scope of all the role assignments between a role group and a set of management roles at the same time, you need to first retrieve the role assignments on the role group, and then set the new scope on each of the assignments. You can do this by using the **Get-ManagementRoleAssignment** cmdlet to retrieve the role assignments, and then pipe them to the **Set-ManagementRoleAssignment** cmdlet.

This procedure uses the concepts of pipelining and the *WhatIf* switch. For more information, see the following topics:

- about_Pipelines
- WhatIf, Confirm, and ValidateOnly Switches

To set the scope on all of the role assignments on a role group at the same time, use the following syntax.

```
Get-ManagementRoleAssignment -RoleAssignee <name of role group> | Set-ManagementRoleAssignment -
CustomRecipientWriteScope <recipient scope name> -CustomConfigWriteScope <configuration scope name> -
RecipientRelativeScopeWriteScope < MyDistributionGroups | Organization | Self> -ExclusiveRecipientWriteScope
<exclusive recipient scope name> -ExclusiveConfigWriteScope <exclusive configuration scope name> -
RecipientOrganizationalUnitScope <organizational unit>
```

You use only the parameters you need to configure the scope you want to use. For example, if you want to change the recipient scope for all role assignments on the Sales Recipient Management role group to Direct Sales Employees, use the following command.

```
Get-ManagementRoleAssignment -RoleAssignee "Sales Recipient Management" | Set-ManagementRoleAssignment -
CustomRecipientWriteScope "Direct Sales Employees"
```

> **NOTE**
>
> You can use the *WhatIf* switch to verify that only the role assignments you want to change are changed. Run the preceding command with the *WhatIf* switch to verify the results, and then remove the *WhatIf* switch to apply the changes.

For more information about changing management role assignments, see Change a Role Assignment.

For detailed syntax and parameter information, see Get-ManagementRoleAssignment.

Use the Exchange Management Shell to change the scope of individual role assignments on a role group

Role assignments between the role group and the roles assigned to it can use the implicit scope obtained from the roles themselves, the same custom scope, or different custom scopes. For more information about role assignments, see Understanding Management Role Assignments.

The scopes on the role assignments are managed using the **Set-ManagementRoleAssignment** cmdlet. You can't manage scopes using the **Set-RoleGroup** cmdlet.

This procedure uses the concepts of pipelining and the **Format-List** cmdlet. For more information, see the following topics:

- about_Pipelines
- Working with Command Output

To change the scope on a role assignment between a role group and a management role, you first find the name of the role assignment, and then set the scope on the role assignment.

1. To find the names of all the role assignments on a role group, use the following command. By piping the management role assignments to the **Format-List** cmdlet, you can view the full name of the assignment.

   ```
   Get-ManagementRoleAssignment -RoleAssignee <role group name> | Format-List Name
   ```

2. Find the name of the role assignment you want to change. Use the name of the role assignment in the next step.

3. To set the scope on an individual assignment, use the following syntax.

   ```
   Set-ManagementRoleAssignment <role assignment name> -CustomRecipientWriteScope <recipient scope name> -
   CustomConfigWriteScope <configuration scope name> -RecipientRelativeScopeWriteScope <
   MyDistributionGroups | Organization | Self> -ExclusiveRecipientWriteScope <exclusive recipient scope
   name> -ExclusiveConfigWriteScope <exclusive configuration scope name> -RecipientOrganizationalUnitScope
   <organizational unit>
   ```

You use only the parameters you need to configure the scope you want to use. For example, if you want to change the recipient scope for the Mail Recipients_Sales Recipient Management role assignment to All Sales Employees, use the following command.

```
Set-ManagementRoleAssignment "Mail Recipients_Sales Recipient Management" -CustomRecipientWriteScope "All
Sales Employees"
```

For more information about changing management role assignments, see Change a Role Assignment.

For detailed syntax and parameter information, see Set-ManagementRoleAssignment.

How do you know this worked?

To verify that you have successfully changed the scope of a role assignment on a role group, do the following:

- If you used the EAC to configure the scope on the role group, do the following:

  1. In the EAC, navigate to **Permissions** > **Admin Roles**. All the role groups in your organization are listed here.

  2. Select a role group to view the scope that's configured on the role group.

- If you used the Exchange Management Shell to configure the scope on the role group, do the following:

  1. Run the following command in the Exchange Management Shell.

     ```
     Get-ManagementRoleAssignment -RoleAssignee <role group name> | Format-Table *WriteScope
     ```

  2. Verify that the write scope on the role assignments has been changed to the scope you specified.

Add or remove a role group delegate

Role group delegates are users or universal security groups (USGs) that can add or remove members from a role group or change the properties of a role group. By adding or removing role group delegates, you can control who is allowed to manage a role group.

> **IMPORTANT**
>
> After you add a delegate to a role group, the role group can only be managed by the delegates on the role group, or by users who are assigned, either directly or indirectly, the Role Management management role. > If a user is assigned, either directly or indirectly, the Role Management role and isn't added as a delegate of the role group, the user must use the *BypassSecurityGroupManagerCheck* switch on the **Add-RoleGroupMember**, **Remove-RoleGroupMember**, **Update-RoleGroupMember**, and **Set-RoleGroup** cmdlets to manage a role group.

> **NOTE**
>
> You can't use the EAC to add a delegate to a role group.

Use the Exchange Management Shell to add a delegate to a role group

To change the list of delegates on a role group, you use the *ManagedBy* parameter on the **Set-RoleGroup** cmdlet. The *ManagedBy* parameter overwrites the entire delegate list on the role group. If you want to add delegates to the role group rather than replace the entire list of delegates, use the following steps:

1. Store the role group in a variable using the following command.

   ```
   $RoleGroup = Get-RoleGroup <role group name>
   ```

2. Add the delegate to the role group stored in the variable using the following command.

```
$RoleGroup.ManagedBy += (Get-User <user to add>).Identity
```

> **NOTE**
>
> Use the **Get-Group** cmdlet if you want to add a USG.

3. Repeat Step 2 for each delegate you want to add.

4. Apply the new list of delegates to the actual role group using the following command.

```
Set-RoleGroup <role group name> -ManagedBy $RoleGroup.ManagedBy
```

This example adds the user David Strome as a delegate on the Organization Management role group.

```
$RoleGroup = Get-RoleGroup "Organization Management"
$RoleGroup.ManagedBy += (Get-User "David Strome").Identity
Set-RoleGroup "Organization Management" -ManagedBy $RoleGroup.ManagedBy
```

For detailed syntax and parameter information, see Set-RoleGroup.

Use the Exchange Management Shell to remove a delegate from a role group

To change the list of delegates on a role group, you use the *ManagedBy* parameter on the **Set-RoleGroup** cmdlet. The *ManagedBy* parameter overwrites the entire delegate list on the role group. If you want to remove delegates from the role group rather than replace the entire list of delegates, use the following steps:

1. Store the role group in a variable using the following command.

```
$RoleGroup = Get-RoleGroup <role group name>
```

2. Remove the delegate from the role group stored in the variable using the following command.

```
$RoleGroup.ManagedBy -= (Get-User <user to remove>).Identity
```

> **NOTE**
>
> Use the **Get-Group** cmdlet if you want to remove a USG.

3. Repeat Step 2 for each delegate you want to remove.

4. Apply the new list of delegates to the actual role group using the following command.

```
Set-RoleGroup <role group name> -ManagedBy $RoleGroup.ManagedBy
```

This example removes the user David Strome as a delegate on the Organization Management role group.

```
$RoleGroup = Get-RoleGroup "Organization Management"
$RoleGroup.ManagedBy -= (Get-User "David Strome").Identity
Set-RoleGroup "Organization Management" -ManagedBy $RoleGroup.ManagedBy
```

For detailed syntax and parameter information, see Set-RoleGroup.

How do you know this worked?

To verify that you have successfully changed the delegate list on a role group, do the following:

1. In the Exchange Management Shell, run the following command.

   ```
   Get-RoleGroup <role group name> | Format-List ManagedBy
   ```

2. Verify that the delegates listed on the *ManagedBy* property include only the delegates that should be able to manage the role group.

Manage role group members

8/3/2020 • 3 minutes to read • Edit Online

To learn about role groups in Exchange Server, see Understanding Management Role Groups.

For additional management tasks related to role groups, see Permissions.

What do you need to know before you begin?

- Estimated time to complete each procedure: 5 minutes
- To open the EAC, see Exchange admin center in Exchange Server. To open the Exchange Management Shell, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Role groups" entry in the Role management permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Add members to a role group

To give a user the permissions that are granted by a role group, you need to add the user, or a universal security group (USG), or another role group that the user is a member of, as a member of the role group.

Use the EAC to add members to a role group

1. In the Exchange admin center (EAC), navigate to **Permissions** > **Admin Roles**.
2. Select the role group you want to add members to, and then click **Edit** ✏.
3. In the **Members** section, click **Add** +.
4. Select the users, USGs, or other role groups you want to add to the role group, click **Add**, and then click **OK**.
5. Click **Save** to save the changes to the role group.

Use the Exchange Management Shell to add members to a role group

To add a role group member, see the Examples section in Add-RoleGroupMember.

To add multiple role group members or to replace the role group membership entirely, see the Examples section in Update-RoleGroupMember.

How do you know this worked?

To verify that you have successfully added one or more members to a role group, do the following:

1. In the EAC, navigate to **Permissions** > **Admin Roles**.
2. Select the role group you added members to.

3. In the role group details pane, verify that the members you added are listed.

Remove members from a role group

To remove the permissions granted by a role group from a user, you need to remove the user, or the universal security group (USG) the user is a member of, from the role group's membership.

Use the EAC to remove members from a role group

1. In the EAC, navigate to **Permissions** > **Admin Roles**.

2. Select the role group you want to remove members from, and then click **Edit** ✎.

3. In the **Members** section, select the members you want to remove, click **Remove** −, and then click **Save**.

Use the Exchange Management Shell to remove members from a role group

To remove a role group member, see the Examples section in Remove-RoleGroupMember.

To remove multiple role group members or to replace the role group membership entirely, see the Examples section in Update-RoleGroupMember.

How do you know this worked?

To verify that you have successfully removed one or more members to a role group, do the following:

1. In the EAC, navigate to **Permissions** > **Admin Roles**.

2. Select the role group you removed members from.

3. In the role group details pane, verify that the members you removed are no longer listed.

View the members of a role group

The members of a role group are granted the permissions provided by the management roles assigned to the role group. You can view the members of a role group to see which users, universal security groups (USG), or other role groups are granted permissions by the role group you specify.

Use the EAC to view the members of a role group

1. In the EAC, navigate to **Permissions** > **Admin Roles**.

2. Select the role group you want to view the members of.

3. In the role group details pane, view the members in the role group details pane.

Use the Exchange Management Shell to view the members of a role group

To view the members of a role group, see the "Examples" section in Get-RoleGroupMember.

Manage role assignment policies

8/3/2020 • 9 minutes to read • Edit Online

If you want to customize the permissions that you assign to a group of end users, create a new custom management role assignment policy. The assignment policy you create can be customized to suit your end user's specific requirements. For more information about assignment policies in Exchange Server, see Understanding Management Role Assignment Policies.

Looking for other management tasks related to managing permissions? Check out Permissions.

What do you need to know before you begin?

- Estimated time to complete each procedure: 5 minutes
- To open the EAC, see Exchange admin center in Exchange Server. To open the Exchange Management Shell, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Assignment policies" entry in the Role management permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Add an assignment policy

After you've created the new assignment policy, you assign users to it. For more information, see Change the assignment policy on a mailbox.

Use the EAC to create a new assignment policy

> **NOTE**
>
> You can only create explicit assignment policies using the Exchange admin center (EAC). If you want to create a new default assignment policy, you must use the Exchange Management Shell. For more information, see the "Use the Exchange Management Shell to create a default assignment policy" section later in this topic.

1. In the EAC, navigate to **Permissions** > **User Roles** and then click **Add** ＋.
2. In the role assignment policy window, provide a name for the new assignment policy.
3. Select the check box next to the role or roles you want to add to the assignment policy. You can select multiple roles, including end-user roles you've added. If you select a role that has child roles, the child roles are automatically selected.
4. Click **Save** to save the changes to the assignment policy.

Use the Exchange Management Shell to create an explicit assignment policy

To create an explicit assignment policy that can be manually assigned to mailboxes, use the following syntax.

```
New-RoleAssignmentPolicy <assignment policy name> -Roles <roles to assign>
```

This example creates the explicit assignment policy Limited Mailbox Configuration and assigns the `MyBaseOptions`, `MyAddressInformation`, and `MyDisplayName` roles to it.

```
New-RoleAssignmentPolicy "Limited Mailbox Configuration" -Roles MyBaseOptions, MyAddressInformation,
MyDisplayName
```

For detailed syntax and parameter information, see New-RoleAssignmentPolicy.

Use the Exchange Management Shell to create a default assignment policy

To create a default assignment policy assigned to new mailboxes, use the following syntax.

```
New-RoleAssignmentPolicy <assignment policy name> -Roles <roles to assign> -IsDefault
```

This example creates the default assignment policy Limited Mailbox Configuration and assigns the `MyBaseOptions`, `MyAddressInformation`, and `MyDisplayName` roles to it.

```
New-RoleAssignmentPolicy "Limited Mailbox Configuration" -Roles MyBaseOptions, MyAddressInformation,
MyDisplayName -IsDefault
```

For detailed syntax and parameter information, see New-RoleAssignmentPolicy.

Remove an assignment policy

If you no longer need a management role assignment policy, you can remove it.

What do you need to know before you begin?

- All users assigned the assignment policy must be changed to another assignment policy. For more information about how to change an assignment policy on a mailbox, see Change the assignment policy on a mailbox.

- All the management role assignments between the assignment policy and the assigned management roles must be removed. For more information about how to remove a role assignment from an assignment policy, see the Remove a role from an assignment policy section later in this topic.

- If you want to remove a default assignment policy, it must be the last assignment policy in the Exchange Server organization.

Use the EAC to remove an assignment policy

1. In the EAC, navigate to **Permissions** > **User Roles**.

2. Select the assignment policy you want to remove, and then click **Delete** 🗑.

Use the Exchange Management Shell to remove an assignment policy

To remove an assignment policy, use the following syntax.

```
Remove-RoleAssignmentPolicy <role assignment policy>
```

This example removes the New York Temporary Users assignment policy.

```
Remove-RoleAssignmentPolicy "New York Temporary Users"
```

For detailed syntax and parameter information, see Remove-RoleAssignmentPolicy.

View a list of assignment policies or assignment policy details

You can view management role assignment policies in a variety of ways, depending on the information you want and whether you're using the EAC or the Exchange Management Shell.

In the EAC, you can view the list of assignment policies and the roles assigned to them. In the Exchange Management Shell, you can view all the assignment policies in your organization, list the mailboxes assigned a specific policy, and more.

Use the EAC to view a list of assignment policies

1. In the EAC, navigate to **Permissions** > **User Roles**. All of the assignment policies in the organization are listed here.

2. To view the details of a specific assignment policy, select the assignment policy you want to view. The description and the roles assigned to the assignment policy are displayed in the details pane.

Use the Exchange Management Shell to view a list of assignment policies

You can view a list of all the assignment policies in your organization by not specifying any assignment policies when you run the **Get-RoleAssignmentPolicy** cmdlet.

This procedure makes use of pipelining and the **Format-Table** cmdlet. For more information about these concepts, see the following topics:

- about_Pipelines

- Working with Command Output

To return a list of all assignment policies in your organization, use the following command.

```
Get-RoleAssignmentPolicy
```

To return a list of specific properties for all the assignment policies in your organization, you can pipe the results to the **Format-Table** cmdlet and specify the properties you want in the list of results. Use the following syntax.

```
Get-RoleAssignmentPolicy | Format-Table <property 1>, <property 2...>
```

This example returns a list of all the assignment policies in your organization and includes the **Name** and **IsDefault** properties.

```
Get-RoleAssignmentPolicy | Format-Table Name, IsDefault
```

For detailed syntax and parameter information, see Get-Mailbox or Get-RoleAssignmentPolicy.

Use the Exchange Management Shell to view the details of a single assignment policy

You can view the details of a specific assignment policy by using the **Get-RoleAssignmentPolicy** cmdlet and piping the output to the **Format-List** cmdlet.

This procedure makes use of pipelining and the **Format-List** cmdlet. For more information about these concepts, see the following topics:

- about_Pipelines

- Working with Command Output

To view the details of a specific assignment policy, use the following syntax.

```
Get-RoleAssignmentPolicy <assignment policy name> | Format-List
```

This example views the details about the Redmond Users - no Text Messaging assignment policy.

```
Get-RoleAssignmentPolicy "Redmond Users - no Text Messaging" | Format-List
```

For detailed syntax and parameter information, see Get-Mailbox or Get-RoleAssignmentPolicy.

Use the Exchange Management Shell to find the default assignment policy

You can find the default assignment policy by piping the output of the **Get-RoleAssignmentPolicy** cmdlet to the **Where** cmdlet. With the **Where** cmdlet, filter the data returned to display only the assignment policy that has its *IsDefault* property set to `$True`.

This procedure makes use of pipelining and the **Where** cmdlet. For more information about these concepts, see the following topics:

- about_Pipelines
- Working with Command Output

This example returns the default assignment policy.

```
Get-RoleAssignmentPolicy | Where {$_.IsDefault -eq $True}
```

For detailed syntax and parameter information, see Get-Mailbox or Get-RoleAssignmentPolicy.

Use the Exchange Management Shell to view mailboxes that are assigned a specific policy

You can find all the mailboxes assigned a specific assignment policy by piping the output of the **Get-Mailbox** cmdlet to the **Where** cmdlet. With the **Where** cmdlet, filter the data returned to display only the mailboxes that have their *RoleAssignmentPolicy* property set to the assignment policy name you specify.

This procedure makes use of pipelining and the **Where** cmdlet. For more information about these concepts, see the following topics:

- about_Pipelines
- Working with Command Output

Use the following syntax.

```
Get-Mailbox | Where {$_.RoleAssignmentPolicy -Eq "<role assignment policy>"}
```

This example finds all the mailboxes assigned the policy Vancouver End Users.

```
Get-Mailbox | Where {$_.RoleAssignmentPolicy -Eq "Vancouver End Users"}
```

For detailed syntax and parameter information, see Get-Mailbox or Get-RoleAssignmentPolicy.

Change the default assignment policy

You can change the management role assignment policy assigned to new mailboxes that are created. Changing the default role assignment policy doesn't change the assignment policy assigned to existing mailboxes. To change the assignment policy assigned to existing mailboxes, see Change the assignment policy on a mailbox.

> **NOTE**
>
> You can't use the EAC to change the default assignment policy. You need to use the Exchange Management Shell.

Use the Exchange Management Shell to change the default assignment policy

To change the default assignment policy, use the following syntax.

```
Set-RoleAssignmentPolicy <assignment policy name> -IsDefault
```

This example sets the Vancouver End Users assignment policy as the default assignment policy.

```
Set-RoleAssignmentPolicy "Vancouver End Users" -IsDefault
```

> **IMPORTANT**
>
> New mailboxes are assigned the default assignment policy even if the policy hasn't been assigned management roles. Mailboxes assigned assignment policies with no assigned management roles can't access any mailbox configuration features in Outlook on the web.

For detailed syntax and parameter information, see Set-RoleAssignmentPolicy.

Add a role to an assignment policy

Use the EAC to add a role to an assignment policy

1. In the EAC, navigate to **Permissions** > **User Roles**.

2. Select the assignment policy you want to add one or more roles to, and then click **Edit** ✎.

3. Select the check box next to the role or roles you want to add to the assignment policy. You can select multiple roles, including end-user roles you've added. If you select a role that has child roles, the child roles are automatically selected.

4. Click **Save** to save the changes to the assignment policy.

Use the Exchange Management Shell to add a role to an assignment policy

To create a management role assignment between a role and an assignment policy, use the following syntax.

```
New-ManagementRoleAssignment -Name <role assignment name> -Role <role name> -Policy <assignment policy name>
```

This example creates the role assignment Seattle Users - Voicemail between the MyVoicemail role and the Seattle Users assignment policy.

```
New-ManagementRoleAssignment -Name "Seattle Users - Voicemail" -Role MyVoicemail -Policy "Seattle Users"
```

For detailed syntax and parameter information, see New-ManagementRoleAssignment.

Remove a role from an assignment policy

If you don't want end users to have permissions to manage certain features of their mailbox or distribution group, you can remove the management role that grants the permissions from the management role assignment policy to which the user is assigned. If other users are assigned the same assignment policy, they also lose the ability to manage that feature.

Use the EAC to remove a role from an assignment policy

1. In the EAC, navigate to **Permissions** > **User Roles**.

2. Select the assignment policy you want to remove one or more roles from, and then click **Edit** ✎ .

3. Clear the check box next to the role or roles you want to remove from the assignment policy. If you clear the check box for a role that has child roles, the check boxes for the child roles are also cleared.

4. Click **Save** to save the changes to the assignment policy.

Use the Exchange Management Shell to remove a role from an assignment policy

You can remove roles from assignment policies by retrieving the associated management role assignment using the **Get-ManagementRoleAssignment** cmdlet and then piping the role assignment returned to the **Remove-ManagementRoleAssignment** cmdlet.

For more information about regular and delegating role assignments, see Understanding Management Role Assignments.

This procedure uses pipelining. For more information about pipelining, see about_Pipelines.

To remove a role from an assignment policy, use the following syntax.

```
Get-ManagementRoleAssignment -RoleAssignee <assignment policy name> -Role <role name> | Remove-
ManagementRoleAssignment
```

This example removes the MyVoicemail management role, which enables users to manage their voice mail options, from the Seattle Users assignment policy.

```
Get-ManagementRoleAssignment -RoleAssignee "Seattle Users" -Role MyVoicemail | Remove-ManagementRoleAssignment
```

For detailed syntax and parameter information, see Remove-ManagementRoleAssignment.

Change the assignment policy on a mailbox

8/3/2020 • 2 minutes to read • Edit Online

When you change a mailbox's assignment policy, the change takes effect as soon as the user refreshes the connection, such as the next time they log into their mailbox or open the mailbox options page. For more information about assignment policies in Exchange Server, see Understanding Management Role Assignment Policies.

Looking for other management tasks related to permissions? Check out Permissions.

What do you need to know before you begin?

- Estimated time to complete each procedure: 5 minutes
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Role groups" entry in the Role management permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to change the assignment policy on a mailbox

1. In the Exchange admin center (EAC), navigate to **Recipients** > **Mailboxes**.
2. Select the user or resource mailbox you want to change the assignment policy on and then click **Edit** ✎.
3. Select **Mailbox Features**.
4. In the **Role assignment policy** list, select the assignment policy you want to assign to the mailbox and then click **Save**.

Use the Exchange Management Shell to change the assignment policy on a mailbox

To change the assignment policy that's assigned to a mailbox, use the following syntax.

```
Set-Mailbox <mailbox alias or name> -RoleAssignmentPolicy <assignment policy>
```

This example sets the assignment policy to Engineering Users on the mailbox Brian.

```
Set-Mailbox Brian -RoleAssignmentPolicy "Engineering Users"
```

Use the Exchange Management Shell to change the assignment policy on a group of mailboxes assigned a specific assignment policy

NOTE

You can't use the EAC to change the assignment policy on a group of mailboxes all at once.

This procedure makes use of pipelining, the **Where** cmdlet, and the *WhatIf* parameter. For more information about these concepts, see the following topics:

- about_Pipelines
- Working with Command Output
- WhatIf, Confirm, and ValidateOnly Switches

If you want to change the assignment policy for a group of mailboxes that are assigned a specific policy, use the following syntax.

```
Get-Mailbox | Where {$_.RoleAssignmentPolicy -Eq "<assignment policy to find>"} | Set-Mailbox -
RoleAssignmentPolicy <assignment policy to set>
```

This example finds all the mailboxes assigned to the Redmond Users - No Voicemail assignment policy and changes the assignment policy to Redmond Users - Voicemail Enabled.

```
Get-Mailbox | Where {$_.RoleAssignmentPolicy -Eq "Redmond Users - No Voicemail"} | Set-Mailbox -
RoleAssignmentPolicy "Redmond Users - Voicemail Enabled"
```

This example includes the *WhatIf* parameter so that you can see all the mailboxes that would be changed without committing any changes.

```
Get-Mailbox | Where {$_.RoleAssignmentPolicy -Eq "Redmond Users - No Voicemail"} | Set-Mailbox -
RoleAssignmentPolicy "Redmond Users - Voicemail Enabled" -WhatIf
```

For detailed syntax and parameter information, see Get-Mailbox or Set-Mailbox.

Feature permissions

8/3/2020 • 2 minutes to read • Edit Online

Permissions in Microsoft Exchange Server are managed using the Role Based Access Control (RBAC) permissions model. The following topics identify the management role groups required to administer the features associated with each functional area in Exchange Server.

- Role management permissions
- Messaging policy and compliance in Exchange Server
- Antispam and antimalware permissions
- Mail flow permissions
- Recipients Permissions
- Email address and address book permissions
- Sharing and collaboration permissions
- Clients and mobile devices permissions
- Unified Messaging permissions
- High availability and site resilience permissions
- Exchange infrastructure and PowerShell permissions
- Server health and performance permissions

Role management permissions

8/3/2020 • 2 minutes to read • Edit Online

The permissions required to perform tasks to configure management roles vary depending on the procedure being performed or the cmdlet you want to run. For more information about management roles, see Understanding Management Roles.

To find out what permissions you need to perform the procedure or run the cmdlet, do the following:

1. In the table below, find the feature that is most related to the procedure you want to perform or the cmdlet you want to run.

2. Next, look at the permissions required for the feature. You must be assigned one of those role groups, an equivalent custom role group, or an equivalent management role. You can also click on a role group to see its management roles. If a feature lists more than one role group, you only need to be assigned one of the role groups to use the feature. For more information about role groups and management roles, see Understanding Role Based Access Control.

3. Now, run the **Get-ManagementRoleAssignment** cmdlet to look at the role groups or management roles assigned to you to see if you have the permissions that are necessary to manage the feature.

> **NOTE**
>
> You must be assigned the Role Management management role to run the **Get-ManagementRoleAssignment** cmdlet. If you don't have permissions to run the **Get-ManagementRoleAssignment** cmdlet, ask your Exchange administrator to retrieve the role groups or management roles assigned to you.

If you want to delegate the ability to manage a feature to another user, see Delegate role assignments.

Role management permissions

You can use the features in the following table to manage the management role groups, roles, assignment policies, assignments, scopes that define the permissions you can apply to administrators, and end users. Users who are assigned the View-Only Management role group can view the configuration of the features in the following table. For more information, see View-only Organization Management.

FEATURE	PERMISSIONS REQUIRED
Management roles	Organization Management
Unscoped management roles	Unscoped Role Management management role
Role groups	Organization Management
Assignment policies	Organization Management
Role assignments	Organization Management
Management scopes	Organization Management

FEATURE	PERMISSIONS REQUIRED
Management role entries	Organization Management
Legacy permissions	Organization Management
Active Directory split permissions	Organization Management **Important**: To run the `setup.exe` command with the *PrepareAD* and *ActiveDirectorySplitPermissions* parameters, the account you use must be a member of the Schema Admins and Enterprise Administrators groups.

Messaging policy and compliance permissions in Exchange Server

8/3/2020 • 3 minutes to read • Edit Online

The permissions required to configure messaging policy and compliance vary depending on the procedure being performed or the cmdlet you want to run. For more information about messaging policy and compliance, see Messaging policy and compliance in Exchange Server.

To find out what permissions you need to perform the procedure or run the cmdlet, do the following:

1. In the table below, find the feature that is most related to the procedure you want to perform or the cmdlet you want to run.

2. Next, look at the permissions required for the feature. You must be assigned one of those role groups, an equivalent custom role group, or an equivalent management role. You can also click on a role group to see its management roles. If a feature lists more than one role group, you only need to be assigned one of the role groups to use the feature. For more information about role groups and management roles, see Understanding Role Based Access Control.

3. Now, run the **Get-ManagementRoleAssignment** cmdlet to look at the role groups or management roles assigned to you to see if you have the permissions that are necessary to manage the feature.

   > **NOTE**
   >
   > You must be assigned the Role Management management role to run the **Get-ManagementRoleAssignment** cmdlet. If you don't have permissions to run the **Get-ManagementRoleAssignment** cmdlet, ask your Exchange administrator to retrieve the role groups or management roles assigned to you.

If you want to delegate the ability to manage a feature to another user, see Delegate role assignments.

Messaging policy and compliance permissions

You can use the features in the following table to configure messaging policy and compliance features. The role groups that are required to configure each feature are listed.

Users who are assigned the View-Only Management role group can view the configuration of the features in the following table. For more information, see View-Only Organization Management.

FEATURE	PERMISSIONS REQUIRED
Data loss prevention (DLP)	Compliance Management
Delete mailbox content (using the Search-Mailbox cmdlet with the *DeleteContent* switch)	Discovery Management **and** Mailbox Import Export Role **Note**: By default, the Mailbox Import Export role isn't assigned to any role group. You can assign a management role to a built-in or custom role group, a user, or a universal security group. Assigning a role to a role group is recommended. For more information, see Add a role to a role group.

FEATURE	PERMISSIONS REQUIRED
Discovery mailboxes - Create	Organization Management Recipient Management
Information Rights Management (IRM) configuration	Compliance Management Organization Management
In-Place Archive	Organization Management Recipient Management
In-Place Archive - Test connectivity	Organization Management Server Management
In-Place eDiscovery	Discovery Management **Note**: By default, the Discovery Management role group doesn't have any members. No users, including administrators, have the required permissions to search mailboxes. For more information, see Assign eDiscovery permissions in Exchange Server.
In-Place Hold	Discovery Management Organization Management **Notes**: • To create a query-based In-Place Hold, a user requires both the Mailbox Search and Legal Hold roles to be assigned directly or via membership in a role group that has both roles assigned. To create an In-Place Hold without using a query, which places all mailbox items on hold, you must have the Legal Hold role assigned. The Discovery Management role group is assigned both roles. • The Organization Management role group is assigned the Legal Hold role. Members of the Organization Management role group can place an In-Place Hold on all items in a mailbox, but can't create a query-based In-Place Hold.
Journaling	Organization Management Records Management
Litigation Hold	Organization Management
Mailbox audit logging	Organization Management Records Management
Message classifications	Organization Management
Messaging records management	Compliance Management Organization Management Records Management
Retention policies - Apply	Organization Management Recipient Management Records Management
Retention policies - Create	See the entry for Messaging records management

FEATURE	PERMISSIONS REQUIRED
Mail flow rules (also known as transport rules)	Organization Management Records Management

Antispam and antimalware permissions

8/3/2020 • 2 minutes to read • Edit Online

The permissions required to perform tasks related to antispam and antimalware vary depending on the procedure being performed or the cmdlet you want to run. For more information about transport features, see Mail flow and the transport pipeline.

This topic lists the permissions required to manage the mail flow features in Microsoft Exchange Server.

To find out what permissions you need to perform the procedure or run the cmdlet, do the following:

1. In the table below, find the feature that is most related to the procedure you want to perform or the cmdlet you want to run.

2. Next, look at the permissions required for the feature. You must be assigned one of those role groups, an equivalent custom role group, or an equivalent management role. You can also click on a role group to see its management roles. If a feature lists more than one role group, you only need to be assigned one of the role groups to use the feature. For more information about role groups and management roles, see Understanding Role Based Access Control.

3. Now, run the **Get-ManagementRoleAssignment** cmdlet to look at the role groups or management roles assigned to you to see if you have the permissions that are necessary to manage the feature.

   > **NOTE**
   >
   > You must be assigned the Role Management management role to run the **Get-ManagementRoleAssignment** cmdlet. If you don't have permissions to run the **Get-ManagementRoleAssignment** cmdlet, ask your Exchange administrator to retrieve the role groups or management roles assigned to you.

If you want to delegate the ability to manage a feature to another user, see Delegate role assignments.

> **NOTE**
>
> Some features that you want to manage might exist on Edge Transport servers. To manage features on Edge Transport servers, you need to become a member of the Local Administrators group on the Edge Transport server you want to manage. Edge Transport servers don't use Role Based Access Control (RBAC). Features that can be managed on Edge Transport servers have Edge Transport Local Administrator in the "Permissions required" column in the table below.

> **NOTE**
>
> Some features may require that you have local administrator permissions on the server you want to manage. To manage these features, you must be a member of the Local Administrators group on that server.

Antispam and Anti-Malware Permissions

You can use the features in the following tables to configure antispam and antimalware settings in your organization. The permissions that are required to configure each feature are listed.

Users who are assigned the View Only Management role group can view the configuration of the features shown in the following table. For more information, see View Only Organization Management.

FEATURE	PERMISSIONS REQUIRED
Anti-malware	Organization Management Hygiene Management
Antispam features	Organization Management Hygiene Management
Antispam features - Edge Transport	Edge Transport server local administrator

Mail flow permissions

8/3/2020 • 3 minutes to read • Edit Online

The permissions required to perform tasks related to mail flow vary depending on the procedure being performed or the cmdlet you want to run. For more information about transport features, see Mail flow and the transport pipeline.

This topic lists the permissions required to manage the mail flow features in Exchange Server 2016 and Exchange Server 2019. For information about how Microsoft 365 or Office 365 permissions relate to Exchange permissions, see About admin roles.

To find out what permissions you need to perform the procedure or run the cmdlet, do the following:

1. In the table below, find the feature that is most related to the procedure you want to perform or the cmdlet you want to run.

2. Next, look at the permissions required for the feature. You must be assigned one of those role groups, an equivalent custom role group, or an equivalent management role. You can also click a role group to see its management roles. If a feature lists more than one role group, you need to be assigned to only one of the role groups to use the feature. For more information about role groups and management roles, see Understanding Role Based Access Control.

3. Now, run the **Get-ManagementRoleAssignment** cmdlet to look at the role groups or management roles assigned to you to see if you have the permissions that are necessary to manage the feature.

   > **NOTE**
   > You must be assigned the Role Management management role to run the **Get-ManagementRoleAssignment** cmdlet. If you don't have permissions to run the **Get-ManagementRoleAssignment** cmdlet, ask your Exchange administrator to retrieve the role groups or management roles assigned to you.

If you want to delegate the ability to manage a feature to another user, see Delegate role assignments.

> **NOTE**
> Some features that you want to manage might exist on Edge Transport servers. To manage features on Edge Transport servers, you need to become a member of the Local Administrators group on the Edge Transport server you want to manage. Edge Transport servers don't use Role Based Access Control (RBAC). Features that can be managed on Edge Transport servers have Edge Transport Local Administrator in the "Permissions required" column in the table below.

> **NOTE**
> Some features may require that you have local administrator permissions on the server you want to manage. To manage these features, you must be a member of the Local Administrators group on that server.

Mail flow permissions

You can use the features in the following tables to configure mail flow settings in the Front End Transport, Mailbox Transport, and Transport services on Mailbox servers, and on Edge Transport servers. The permissions that are required to configure each feature are listed.

Users who are assigned the View Only Management role group can view the configuration of the features shown in the following table. For more information, see View Only Organization Management.

Mailbox servers

FEATURE	PERMISSIONS REQUIRED
Accepted domains	Organization Management
Active Directory site and site link management	Organization Management
Antispam features	Organization Management Hygiene Management
Antispam updates	Organization Management Hygiene Management
Certificate management	Organization Management
Delivery Agent connectors	Organization Management Server Management
DSNs	Organization Management
EdgeSync	Organization Management
Foreign connectors	Organization Management
Front End Transport service	Organization Management Server Management Hygiene Management
Journaling	Organization Management Records Management
Mailbox access	Organization Management
Mailbox junk email configuration	Organization Management Records Management Recipient Management Help Desk
Mailbox Transport service	Organization Management Server Management Hygiene Management
MailTips	Organization Management
Message classifications	Organization Management Records Management
Message tracking	Organization Management Records Management Recipient Management

FEATURE	PERMISSIONS REQUIRED
Moderated transport	Organization Management Recipient Management
Queues	Organization Management Server Management
Receive connectors	Organization Management Server Management Hygiene Management
Remote domains	Organization Management
SafeList aggregation	Organization Management Records Management
Send connectors	Organization Management
Shadow redundancy	Organization Management
Testing mail flow	Organization Management Server Management
Testing mail flow rule (also known as transport rule) processing	Organization Management
Transport agents	Organization Management Records Management
Transport configuration	Organization Management
Transport logs	Organization Management Server Management
Mail flow rules (also known as transport rules)	Organization Management Records Management
Transport service	Organization Management Server Management Hygiene Management
X.400 domains	Organization Management

Edge Transport servers

FEATURE	PERMISSIONS REQUIRED
Accepted domains - Edge Transport	Edge Transport server local administrator
Address Rewriting - Edge Transport	Edge Transport server local administrator
Edge Transport server	Edge Transport server local administrator

FEATURE	PERMISSIONS REQUIRED
EdgeSync - Edge Transport	Edge Transport server local administrator
Queues - Edge Transport	Edge Transport server local administrator
Receive connectors - Edge Transport	Edge Transport server local administrator
Send connectors - Edge Transport	Edge Transport server local administrator
Transport configuration - Edge Transport	Edge Transport server local administrator
Transport logs - Edge Transport	Edge Transport server local administrator
Mail flow rules (also known as transport rules) - Edge Transport	Edge Transport server local administrator

Recipients Permissions

8/3/2020 • 4 minutes to read • Edit Online

The permissions required to perform tasks to manage recipients vary depending on the procedure being performed or the cmdlet you want to run.

To find out what permissions you need to perform the procedure or run the cmdlet, do the following:

1. In the table below, find the feature that is most related to the procedure you want to perform or the cmdlet you want to run.
2. Next, look at the permissions required for the feature. You must be assigned one of those role groups, an equivalent custom role group, or an equivalent management role. You can also click on a role group to see its management roles. If a feature lists more than one role group, you need to be assigned to only one of the role groups to use the feature. For more information about role groups and management roles, see Understanding Role Based Access Control.
3. Now, run the **Get-ManagementRoleAssignment** cmdlet to look at the role groups or management roles assigned to you to see if you have the permissions that are necessary to manage the feature.

> **NOTE**
>
> You must be assigned the Role Management management role to run the **Get-ManagementRoleAssignment** cmdlet. If you don't have permissions to run the **Get-ManagementRoleAssignment** cmdlet, ask your Exchange administrator to retrieve the role groups or management roles assigned to you.

If you want to delegate the ability to manage a feature to another user, see Delegate role assignments.

Mailbox server permissions

Users who are assigned the View-Only Management role group can view the configuration of the features in the following table. For more information, see View-only Organization Management.

FEATURE	PERMISSIONS REQUIRED
Calendar repair, server configuration	Organization Management Server Management
Delegating Mailbox servers	Organization Management
Email address policies	Organization Management Server Management
Exchange Search	Organization Management View-Only Organization Management Server Management

FEATURE	PERMISSIONS REQUIRED
Exchange Search - diagnostics	Organization Management View-Only Organization Management Support Diagnostics role **Note:**: The Support Diagnostics role isn't assigned to a role group. For more information, see Add a role to a role group.
Group metrics	Organization Management Server Management
Import Export	Mailbox Import Export role **Note:**: The Mailbox Import Export role isn't assigned to a role group. For more information, see Mailbox Import Export Role.
Mailbox Assistants	Organization Management Server Management
Mailbox moves	Organization Management Recipient Management
Mailbox recovery	Organization Management
Mailbox repair request	Organization Management Server Management Recipient Management
Mailbox restore request	Organization Management
Mailbox server configuration	Organization Management Server Management
Manage Exchange Search Indexer service on a Mailbox server	Local Administrator on the Mailbox server
MAPI connectivity	Organization Management Server Management
OAB virtual directories	Organization Management Server Management
Remove store mailbox	Organization Management Server Management

Calendar and sharing permissions

Users who are assigned the View-Only Management role group can view the configuration of the features in the following table. For more information, see View-only Organization Management.

FEATURE	PERMISSIONS REQUIRED

FEATURE	PERMISSIONS REQUIRED
Calendar configuration	Organization Management Recipient Management Help Desk
Calendar diagnostics	Organization Management Records Management Hygiene Management Compliance Management Help Desk
Calendar processing	Organization Management Recipient Management Help Desk
Notifications	Organization Management Recipient Management
Organization relationships	Organization Management
Sharing policies	Organization Management

Resource mailbox configuration permissions

Users who are assigned the View-Only Management role group can view the configuration of the features in the following table. For more information, see View-only Organization Management.

FEATURE	PERMISSIONS REQUIRED
Booking policies	Organization Management Recipient Management Help Desk
Delegation	Organization Management Recipient Management
Resource mailbox schema configuration	Organization Management

Mailbox database permissions

Users who are assigned the View-Only Management role group can view the configuration of the features in the following table. For more information, see View-only Organization Management.

FEATURE	PERMISSIONS REQUIRED
Mailbox databases	Organization Management Server Management

Recipient provisioning permissions

This table contains the various permissions that are required to manage recipients.

Users who are assigned the View-Only Management role group can view the configuration of the features in the following table. For more information, see View-only Organization Management.

FEATURE	PERMISSIONS REQUIRED
Address list, GAL	Organization Management
Antispam	Organization Management Recipient Management
Apps for Outlook	Organization Management View-Only Organization Management Help Desk
Applying sharing policies	Organization Management Recipient Management
Arbitration	Organization Management
Archive connectivity	Organization Management View-Only Organization Management Server Management
Assigning offline address books	Organization Management Recipient Management
Automatic replies	Organization Management Recipient Management Help Desk
Calendar configuration	Organization Management Recipient Management
Calendar repair	Organization Management Recipient Management
Contact aggregation settings	Organization Management Recipient Management View-Only Organization Management
Convert mailboxes	Organization Management Recipient Management
Disconnected mailboxes	Organization Management Recipient Management Help Desk
Distribution groups	Organization Management Recipient Management
Dynamic distribution groups	Organization Management Recipient Management
Email addresses	Organization Management Recipient Management UM Management

FEATURE	PERMISSIONS REQUIRED
Inbox rules	Organization Management Recipient Management Help Desk
Mail contacts	Organization Management Recipient Management
Mail tips	Organization Management Recipient Management
Mail user	Organization Management Recipient Management
Mailbox folder permissions	Organization Management Recipient Management Help Desk
Mailbox folders	Organization Management Recipient Management
MAPI connectivity	Organization Management
Message configuration	Organization Management Recipient Management Help Desk
Message quotas	Organization Management Recipient Management
Moderation	Organization Management Recipient Management
Permissions and delegation	Organization Management
Archive mailboxes	Organization Management Recipient Management
Recipient data properties	Organization Management Recipient Management
Remote mailboxes	Organization Management Recipient Management
Retention and legal holds	Organization Management Recipient Management Records Management
Send As	Organization Management Recipient Management
Spelling configuration	Organization Management Recipient Management Help Desk

FEATURE	PERMISSIONS REQUIRED
Unified Messaging (in Exchange 2016; not available in Exchange 2019)	Organization Management UM Management
User mailboxes	Organization Management Recipient Management
User photos	Organization Management Recipient Management Help Desk

Mailbox move and migration permissions

The table contains the permissions that are required to move on-premises mailboxes to different domains or forests and to migrate on-premises mailboxes to and from your cloud-based organization.

FEATURE	PERMISSIONS REQUIRED
Mailbox moves (local or cross-forest)	Organization Management Recipient Management
Mailbox moves (hybrid deployment)	Organization Management Recipient Management
Migration (on-boarding and off-boarding from the cloud)	Organization Management Recipient Management

Email address and address book permissions

8/3/2020 • 2 minutes to read • Edit Online

The permissions required to configure email address and address book features vary depending on the procedure being performed or the cmdlet you want to run. For more information about email addresses and address books, see Email addresses and address books in Exchange Server.

To find out what permissions you need to perform the procedure or run the cmdlet, do the following:

1. In the table below, find the feature that is most related to the procedure you want to perform or the cmdlet you want to run.

2. Next, look at the permissions required for the feature. You must be assigned one of those role groups, an equivalent custom role group, or an equivalent management role. You can also click on a role group to see its management roles. If a feature lists more than one role group, you only need to be assigned one of the role groups to use the feature. For more information about role groups and management roles, see Understanding Role Based Access Control.

3. Now, run the **Get-ManagementRoleAssignment** cmdlet to look at the role groups or management roles assigned to you to see if you have the permissions that are necessary to manage the feature.

> **NOTE**
> You must be assigned the Role Management management role to run the **Get-ManagementRoleAssignment** cmdlet. If you don't have permissions to run the **Get-ManagementRoleAssignment** cmdlet, ask your Exchange administrator to retrieve the role groups or management roles assigned to you.

If you want to delegate the ability to manage a feature to another user, see Delegate role assignments.

Email address and address book permissions

Users who are assigned the View-Only Management role group can view the configuration of the features in the following table. For more information, see View-only Organization Management.

FEATURE	PERMISSIONS REQUIRED
Address book policies	Organization Management
Address lists	Organization Management
Email address policies	Organization Management
Details templates	Organization Management
Global address lists	Organization Management
Offline address books	Organization Management
Offline address book connectivity	Organization Management

Sharing and collaboration permissions

8/3/2020 • 2 minutes to read • Edit Online

The permissions required to configure sharing and collaboration features vary depending on the procedure being performed or the cmdlet you want to run. For more information about sharing and collaboration, see Collaboration and Sharing.

To find out what permissions you need to perform the procedure or run the cmdlet, do the following:

1. In the table below, find the feature that is most related to the procedure you want to perform or the cmdlet you want to run.
2. Next, look at the permissions required for the feature. You must be assigned one of those role groups, an equivalent custom role group, or an equivalent management role. You can also click on a role group to see its management roles. If a feature lists more than one role group, you only need to be assigned one of the role groups to use the feature. For more information about role groups and management roles, see Understanding Role Based Access Control.
3. Now, run the **Get-ManagementRoleAssignment** cmdlet to look at the role groups or management roles assigned to you to see if you have the permissions that are necessary to manage the feature.

> **NOTE**
>
> You must be assigned the Role Management management role to run the **Get-ManagementRoleAssignment** cmdlet. If you don't have permissions to run the **Get-ManagementRoleAssignment** cmdlet, ask your Exchange administrator to retrieve the role groups or management roles assigned to you.

If you want to delegate the ability to manage a feature to another user, see Delegate role assignments.

Sharing and collaboration feature permissions

You can use the features in the following table to configure sharing and collaboration features. The role groups that are required to configure each feature are listed.

Users who are assigned the View-Only Management role group can view the configuration of the features in the following table. For more information, see View-only Organization Management.

FEATURE	PERMISSIONS REQUIRED
Partner applications - configure	Organization Management
Public folders, mail-enabled	Organization Management Recipient Management
Public folders	Organization Management Public Folder Management
Site mailboxes	Organization Management Recipient Management

FEATURE	PERMISSIONS REQUIRED
Site mailbox provisioning policy	Organization Management Recipient Management

Clients and mobile devices permissions

8/3/2020 • 5 minutes to read • Edit Online

The permissions required to perform tasks for clients and mobile devices vary depending on the procedure being performed or the cmdlet you want to run. For more information about client and mobile device features, see Clients and mobile.

To find out what permissions you need to perform the procedure or run the cmdlet, do the following:

1. In the table below, find the feature that is most related to the procedure you want to perform or the cmdlet you want to run.

2. Next, look at the permissions required for the feature. You must be assigned one of those role groups, an equivalent custom role group, or an equivalent management role. You can also click on a role group to see its management roles. If a feature lists more than one role group, you only need to be assigned one of the role groups to use the feature. For more information about role groups and management roles, see Understanding Role Based Access Control.

3. Now, run the **Get-ManagementRoleAssignment** cmdlet to look at the role groups or management roles assigned to you to see if you have the permissions that are necessary to manage the feature.

   > **NOTE**
   > You must be assigned the Role Management management role to run the **Get-ManagementRoleAssignment** cmdlet. If you don't have permissions to run the **Get-ManagementRoleAssignment** cmdlet, ask your Exchange administrator to retrieve the role groups or management roles assigned to you.

If you want to delegate the ability to manage a feature to another user, see Delegate role assignments.

> **NOTE**
> Some features may require that you have local administrator permissions on the server you want to manage. To manage these features, you must be a member of the Local Administrators group on that server.

Client Access service permissions

You can configure any of the following features for the Client Access service.

Users who are assigned the View-Only Management role group can view the configuration of the features in the following table. For more information, see View-only Organization Management.

FEATURE	PERMISSIONS REQUIRED
Client Access service array settings	Organization Management Server Management
Client Access service settings	Server Management
Client Access service email channel settings	Organization Management Server Management

FEATURE	PERMISSIONS REQUIRED
Client Access user settings	Server Management
Client Access virtual directory settings	Organization Management Server Management
RPC Client Access settings	Organization Management Server Management View-Only Organization Management
Push notification proxy settings	Organization Management Recipient Management
OAuth authentication redirection settings	Organization Management

Exchange ActiveSync permissions

You can configure any of the following for Exchange ActiveSync.

Users who are assigned the View-Only Management role group can view the configuration of the features in the following table. For more information, see View-only Organization Management.

FEATURE	PERMISSIONS REQUIRED
Exchange ActiveSync Autoblock settings	Organization Management
Exchange ActiveSync mailbox policy settings	Organization Management Server Management
Exchange ActiveSync server settings	Organization Management Server Management
Exchange ActiveSync settings	Organization Management Server Management
Exchange ActiveSync user settings	Recipient Management
Exchange ActiveSync virtual directory settings	Organization Management Server Management
Mobile device mailbox policy settings	Organization Management Server Management
Mobile device user settings	Organization Management Server Management Recipient Management

Autodiscover permissions

You can configure the following for the Autodiscover service.

Users who are assigned the View-Only Management role group can view the configuration of the features in the following table. For more information, see View-only Organization Management.

FEATURE	PERMISSIONS REQUIRED
Autodiscover service configuration settings	Organization Management Server Management View-Only Organization Management Delegated Setup Hygiene Management
Autodiscover virtual directory settings	Organization Management Server Management

Availability service permissions

You can configure the following for the Availability service.

Users who are assigned the View-Only Management role group can view the configuration of the features in the following table. For more information, see View-only Organization Management.

FEATURE	PERMISSIONS REQUIRED
Availability service address space settings	Organization Management View-Only Organization Management
Availability service configuration settings	Organization Management Server Management View-Only Organization Management

Client throttling permissions

You can configure the following for client throttling.

Users who are assigned the View-Only Management role group can view the configuration of the features in the following table. For more information, see View-only Organization Management.

FEATURE	PERMISSIONS REQUIRED
Client throttling settings	Organization Management View-Only Organization Management

Exchange Web Services permissions

You can configure the following for Web Services virtual directories.

Users who are assigned the View-Only Management role group can view the configuration of the features in the following table. For more information, see View-only Organization Management.

FEATURE	PERMISSIONS REQUIRED
Exchange Web Services virtual directory settings	Organization Management Server Management

FEATURE	PERMISSIONS REQUIRED
Test Exchange Web Services	Organization Management Server Management
Test Outlook Web Services	Organization Management

Outlook Anywhere permissions

You can configure and manage the following settings for Outlook Anywhere.

Users who are assigned the View-Only Management role group can view the configuration of the features in the following table. For more information, see View-only Organization Management.

FEATURE	PERMISSIONS REQUIRED
Outlook Anywhere configuration (enable, disable, change, view)	Organization Management Server Management View-Only Organization Management Delegated Setup Hygiene Management
RPC over HTTP Proxy component	Local Server Administrator
Test Outlook Anywhere connectivity	Organization Management View-Only Organization Management Server Management

Outlook on the web permissions

You can use the following features to view Outlook on the web settings, control security and user access to Outlook on the web, and test Outlook on the web connectivity.

Users who are assigned the View-Only Management role group can view the configuration of the features in the following table. For more information, see View-only Organization Management.

FEATURE	PERMISSIONS REQUIRED
Graphics editor	Local Server Administrator
IIS Manager	Local Server Administrator
ISA Server 2006	ISA Server Enterprise Administrator
Outlook on the web mailbox policies	Organization Management Recipient Management
Outlook on the web virtual directories	Organization Management Server Management
Registry Editor	Local Server Administrator
S/MIME configuration	Organization Management

FEATURE	PERMISSIONS REQUIRED
Text editor	Local Server Administrator
View Outlook on the web mailbox policies	Organization Management Recipient Management View-Only Organization Management Delegated Setup Hygiene Management

POP3 and IMAP4 permissions

You can configure the following for POP3 and IMAP4.

Users who are assigned the View-Only Management role group can view the configuration of the features in the following table. For more information, see View-only Organization Management.

FEATURE	PERMISSIONS REQUIRED
IMAP4 settings	Organization Management Server Management View-Only Organization Management
POP3 settings	Organization Management Server Management View-Only Organization Management
Test IMAP4 settings	Organization Management Server Management View-Only Organization Management
Test POP3 settings	Organization Management Server Management View-Only Organization Management

Windows PowerShell virtual directory permissions

You can configure the following for Windows PowerShell.

Users who are assigned the View-Only Management role group can view the configuration of the features in the following table. For more information, see View-only Organization Management.

FEATURE	PERMISSIONS REQUIRED
Test Windows PowerShell	Organization Management
Windows PowerShell settings	Organization Management

Text Messaging permissions

You can configure the following for text messaging.

Users who are assigned the View-Only Management role group can view the configuration of the features in the following table. For more information, see View-only Organization Management.

FEATURE	PERMISSIONS REQUIRED
Text messaging notification settings	Recipient Management
Text messaging settings	Recipient Management
Text messaging user settings	Recipient Management

Unified Messaging permissions

8/3/2020 • 2 minutes to read • Edit Online

The permissions required to manage Unified Messaging services and features on Exchange 2016 Mailbox servers vary depending on the procedure being performed or the cmdlet you want to run.

> **NOTE**
>
> Unified Messaging is not available in Exchange 2019.

To find out what permissions you need to perform the procedure or run the cmdlet, do the following:

1. In the table below, find the feature that is most related to the procedure you want to perform or the cmdlet you want to run.

2. Next, look at the permissions required for the feature. You must be assigned one of those role groups, an equivalent custom role group, or an equivalent management role. You can also click on a role group to see its management roles. If a feature lists more than one role group, you only need to be assigned one of the role groups to use the feature. For more information about role groups and management roles, see Understanding Role Based Access Control.

3. Now, run the **Get-ManagementRoleAssignment** cmdlet to look at the role groups or management roles assigned to you to see if you have the permissions that are necessary to manage the feature.

   > **NOTE**
   >
   > You must be assigned the Role Management management role to run the **Get-ManagementRoleAssignment** cmdlet. If you don't have permissions to run the **Get-ManagementRoleAssignment** cmdlet, ask your Exchange administrator to retrieve the role groups or management roles assigned to you.

If you want to delegate the ability to manage a feature to another user, see Delegate role assignments.

UM component permissions

You can configure settings for the UM components and features in the following table.

Users who are assigned the View-Only Management role group can view the configuration of the features in the following table. For more information, see View-only Organization Management.

FEATURE	PERMISSIONS REQUIRED
UM auto attendants	Organization Management Unified Messaging Management
UM call answering rules	Organization Management Unified Messaging Management
UM call data and summary reports	Organization Management Unified Messaging Management

FEATURE	PERMISSIONS REQUIRED
UM Call Router service (front-end)	Organization Management Unified Messaging Management
UM dial plans	Organization Management Unified Messaging Management
UM hunt groups	Organization Management Unified Messaging Management
UM IP gateways	Organization Management Unified Messaging Management
UM mailbox policies	Organization Management Unified Messaging Management
UM mailboxes	Organization Management Unified Messaging Management
UM prompts	Organization Management Unified Messaging Management
UM service (back-end)	Organization Management Server Management

High availability and site resilience permissions

8/3/2020 • 2 minutes to read • Edit Online

The permissions required to configure high availability vary depending on the procedure being performed or the cmdlet you want to run. For more information about high availability, see High availability and site resilience.

To find out what permissions you need to perform the procedure or run the cmdlet, do the following:

1. In the table below, find the feature that is most related to the procedure you want to perform or the cmdlet you want to run.

2. Next, look at the permissions required for the feature. You must be assigned one of those role groups, an equivalent custom role group, or an equivalent management role. You can also click on a role group to see its management roles. If a feature lists more than one role group, you only need to be assigned one of the role groups to use the feature. For more information about role groups and management roles, see Understanding Role Based Access Control.

3. Now, run the **Get-ManagementRoleAssignment** cmdlet to look at the role groups or management roles assigned to you to see if you have the permissions that are necessary to manage the feature.

> **NOTE**
>
> You must be assigned the Role Management management role to run the **Get-ManagementRoleAssignment** cmdlet. If you don't have permissions to run the **Get-ManagementRoleAssignment** cmdlet, ask your Exchange administrator to retrieve the role groups or management roles assigned to you.

If you want to delegate the ability to manage a feature to another user, see Delegate role assignments.

Database availability group permissions

You can use the features in the following table to add, remove, and configure settings for database availability groups (DAGs).

Users who are assigned the View-Only Management role group can view the configuration of the features in the following table. For more information, see View-only Organization Management.

FEATURE	PERMISSIONS REQUIRED
Database availability group membership	Organization Management Database Availability Groups
Database availability group properties	Organization Management Database Availability Groups
Database availability groups	Organization Management Database Availability Groups
Database availability networks	Organization Management Database Availability Groups

Mailbox database copy permissions

You can use the features in the following table to add, remove, update, and activate mailbox database copies.

FEATURE	PERMISSIONS REQUIRED
Database switchover	Organization Management Database Copies
Mailbox database copies	Organization Management Database Copies
Server switchover	Organization Management Database Copies
Update a mailbox database copy	Organization Management Database Copies

Exchange infrastructure and PowerShell permissions

8/3/2020 • 3 minutes to read • Edit Online

The permissions required to perform tasks to configure various components of Exchange Server depend on the procedure being performed or the cmdlet you want to run. See each of the sections in this topic for more information about their respective features.

To find out what permissions you need to perform the procedure or run the cmdlet, do the following:

1. In the table below, find the feature that is most related to the procedure you want to perform or the cmdlet you want to run.
2. Next, look at the permissions required for the feature. You must be assigned one of those role groups, an equivalent custom role group, or an equivalent management role. You can also click on a role group to see its management roles. If a feature lists more than one role group, you only need to be assigned one of the role groups to use the feature. For more information about role groups and management roles, see Understanding Role Based Access Control.
3. Now, run the **Get-ManagementRoleAssignment** cmdlet to look at the role groups or management roles assigned to you to see if you have the permissions that are necessary to manage the feature.

   > **NOTE**
   >
   > You must be assigned the Role Management management role to run the **Get-ManagementRoleAssignment** cmdlet. If you don't have permissions to run the **Get-ManagementRoleAssignment** cmdlet, ask your Exchange administrator to retrieve the role groups or management roles assigned to you.

If you want to delegate the ability to manage a feature to another user, see Delegate role assignments.

> **NOTE**
>
> Some features may require that you have local administrator permissions on the server you want to manage. To manage these features, you must be a member of the Local Administrators group on that server.

Exchange infrastructure permissions

The following table lists the permissions required to perform tasks that configure general Exchange settings.

Users who are assigned the View-Only Management role group can view the configuration of the features in the following table. For more information, see View-only Organization Management.

FEATURE	PERMISSIONS REQUIRED
Administrator audit logging	Organization Management Records Management
Exchange admin center configuration settings	View-Only Organization Management
Exchange admin center connectivity	Organization Management Server Management

FEATURE	PERMISSIONS REQUIRED
Exchange server configuration settings	Organization Management Server Management
Exchange Help settings	Organization Management
Message categories	Organization Management Hygiene Management Recipient Management Help Desk
Product key	Organization Management
Test system health	Organization Management Server Management
View-only administrator audit logging	Organization Management Records Management **Note**: You can also manually assign the View-Only Audit Logs management role to a management role group. For more information, see View-Only Audit Logs.
Write to audit log	Users that are members of any role group or assigned any management role can write to the administrator audit log.

Exchange PowerShell infrastructure permissions

The following table lists the permissions required to perform tasks that configure features that control how the Exchange Management Shell runs.

Users who are assigned the View-Only Management role group can view the configuration of the features in the following table. For more information, see View-only Organization Management.

FEATURE	PERMISSIONS REQUIRED
Active Directory Domain Services server settings	Organization Management Server Management Recipient Management UM Management
Cmdlet extension agents	Organization Management
PowerShell virtual directories	Organization Management Server Management
PowerShell and WinRM installation	Local Server Administrator
Remote PowerShell	Organization Management

Federation and certificates permissions

The following table lists permissions required for performing tasks related to federation trusts, OAuth configuration, certificate management, and hybrid deployment configuration.

Users who are assigned the View-Only Management role group can view the configuration of the features in the following table. For more information, see View-only Organization Management.

FEATURE	PERMISSIONS REQUIRED
Certificate management	Organization Management Server Management
Federation trusts, OAuth	Organization Management
Test federation trusts, OAuth	Organization Management View-Only Organization Management Server Management
Hybrid deployment configuration	Organization Management
Intra-Organization connectors	Organization Management Recipient Management Records Management

Server health and performance permissions

8/3/2020 • 2 minutes to read • Edit Online

The permissions required to perform tasks to configure various components of Exchange Server depend on the procedure being performed or the cmdlet you want to run. See each of the sections in this topic for more information about their respective features.

To find out what permissions you need to perform the procedure or run the cmdlet, do the following:

1. In the table below, find the feature that is most related to the procedure you want to perform or the cmdlet you want to run.

2. Next, look at the permissions required for the feature. You must be assigned one of those role groups, an equivalent custom role group, or an equivalent management role. You can also click on a role group to see its management roles. If a feature lists more than one role group, you only need to be assigned one of the role groups to use the feature. For more information about role groups and management roles, see Understanding Role Based Access Control.

3. Now, run the **Get-ManagementRoleAssignment** cmdlet to look at the role groups or management roles assigned to you to see if you have the permissions that are necessary to manage the feature.

   > **NOTE**
   >
   > You must be assigned the Role Management management role to run the **Get-ManagementRoleAssignment** cmdlet. If you don't have permissions to run the **Get-ManagementRoleAssignment** cmdlet, ask your Exchange administrator to retrieve the role groups or management roles assigned to you.

If you want to delegate the ability to manage a feature to another user, see Delegate role assignments.

> **NOTE**
>
> Some features may require that you have local administrator permissions on the server you want to manage. To manage these features, you must be a member of the Local Administrators group on that server.

Exchange workload management permissions

The following table lists the permissions required to perform tasks that manage the health and performance of your Exchange Server organization. For more information, see User workload management in Exchange Server.

Users who are assigned the View-Only Management role group can view the configuration of the features in the following table. For more information, see View-only Organization Management.

FEATURE	PERMISSIONS REQUIRED
User throttling	Organization Management Recipient Management View-Only Organization Management
Exchange workload throttling	Organization Management View-Only Organization Management

Exchange event log permissions

The following table lists the permissions required to perform tasks that manage Exchange event log settings.

Users who are assigned the View-Only Management role group can view the configuration of the features in the following table. For more information, see View-only Organization Management.

FEATURE	PERMISSIONS REQUIRED
Exchange event log management	Organization Management Server Management View-Only Organization Management UM Management

Split permissions in Exchange Server

8/3/2020 • 19 minutes to read • Edit Online

Organizations that separate the management of Exchange Server 2016 and Exchange Server 2019 objects and Active Directory objects use what's called a *split permissions* model. Split permissions enable organizations to assign specific permissions and related tasks to specific groups within the organization. This separation of work helps to maintain standards and workflows, and helps to control change in the organization.

The highest level of split permissions is the separation of Exchange management and Active Directory management. Many organizations have two groups: administrators that manage the organization's Exchange infrastructure, including servers and recipients, and administrators that manage the Active Directory infrastructure. This is an important separation for many organizations because the Active Directory infrastructure often spans many locations, domains, services, applications, and even Active Directory forests. Active Directory administrators must ensure that changes made to Active Directory don't negatively impact any other services. As a result, typically only a small group of administrators is allowed to manage that infrastructure.

At the same time, the infrastructure for Exchange, including servers and recipients, can also be complex and require specialized knowledge. Additionally, Exchange stores extremely confidential information about the business of the organization. Exchange administrators can potentially access this information. By limiting the number of Exchange administrators, the organization limits who can make changes to Exchange configuration and who can access sensitive information.

Split permissions typically make a distinction between the creation of security principals in Active Directory, such as users and security groups, and the subsequent configuration of those objects. This helps to reduce the chance of unauthorized access to the network by controlling who can create objects that grant access to it. Most often only Active Directory administrators can create security principals while other administrators, such as Exchange administrators, can manage specific attributes on existing Active Directory objects.

To support the varying needs to separate the management of Exchange and Active Directory, Exchange lets you choose whether you want a shared permissions model or a split permissions model. Exchange offers two types of split permissions models: RBAC and Active Directory. Exchange defaults to a shared permissions model.

Explanation of Role Based Access Control and Active Directory

To understand split permissions, you need to understand how the Role Based Access Control (RBAC) permissions model in Exchange works with Active Directory. The RBAC model controls who can perform what actions, and on which objects those actions can be performed. For more information about the various components of RBAC that are discussed in this topic, see Exchange Server permissions.

All tasks that are performed on Exchange objects must be done through the Exchange Management Shell or the Exchange admin center (EAC) interface. Both of these management tools use RBAC to authorize all tasks that are performed.

RBAC is a component that exists on every Exchange server. RBAC checks whether the user performing an action is authorized to do so:

- If the user isn't authorized to perform the action, RBAC doesn't allow the action to proceed.
- If the user is authorized to perform the action, RBAC checks whether the user is authorized to perform the action against the specific object being requested:
- If the user is authorized, RBAC allows the action to proceed.

- If the user isn't authorized, RBAC doesn't allow the action to proceed.

If RBAC allows an action to proceed, the action is performed in the context of the Exchange Trusted Subsystem and not the user's context. The Exchange Trusted Subsystem is a highly privileged universal security group (USG) that has read/write access to every Exchange-related object in the Exchange organization. It's also a member of the Administrators local security group and the Exchange Windows Permissions USG, which enables Exchange to create and manage Active Directory objects.

> **WARNING**
>
> Don't make any manual changes to the membership of the Exchange Trusted Subsystem security group. Also, don't add it to or remove it from object access control lists (ACLs). By making changes to the Exchange Trusted Subsystem USG yourself, you could cause irreparable damage to your Exchange organization.

It's important to understand that it doesn't matter what Active Directory permissions a user has when using the Exchange management tools. If the user is authorized, via RBAC, to perform an action in the Exchange management tools, the user can perform the action regardless of his or her Active Directory permissions. Conversely, if a user is an Enterprise Admin in Active Directory but isn't authorized to perform an action, such as creating a mailbox, in the Exchange management tools, the action won't succeed because the user doesn't have the required permissions according to RBAC.

> **IMPORTANT**
>
> Although the RBAC permissions model doesn't apply to the Active Directory Users and Computers management tool, Active Directory Users and Computers can't manage the Exchange configuration. S, although a user may have access to modify some attributes on Active Directory objects, such as the display name of a user, the user must use the Exchange management tools, and therefore must be authorized by RBAC, to manage Exchange attributes.

Shared permissions

The shared permissions model is the default model for Exchange. You don't need to change anything if this is the permissions model you want to use. This model doesn't separate the management of Exchange and Active Directory objects from within the Exchange management tools. It allows administrators using the Exchange management tools to create security principals in Active Directory.

The following table shows the roles that enable the creation of security principals in Exchange and the management role groups they're assigned to by default.

MANAGEMENT ROLE	ROLE GROUP
Mail Recipient Creation role	Organization Management Recipient Management
Security Group Creation and Membership role	Organization Management

Only role groups, users, or USGs that are assigned the Mail Recipient Creation role can create security principals such as Active Directory users. By default, the Organization Management and Recipient Management role groups are assigned this role. Therefore members of these role groups can create security principals.

Only role groups, users, or USGs that are assigned the Security Group Creation and Membership role can create security groups or manage their memberships. By default, only the Organization Management role group is assigned this role. Therefore only members of the Organization Management role group can create or manage the membership of security groups.

You can assign the Mail Recipient Creation role and the Security Group Creation and Membership role to other role groups, users, or USGs if you want other users to be able to create security principals.

To enable the management of existing security principals in Exchange, the Mail Recipients role is assigned to the Organization Management and Recipient Management role groups by default. Only role groups, users, or USGs that are assigned the Mail Recipients role can manage existing security principals. If you want other role groups, users, or USGs to be able to manage existing security principals, you must assign the Mail Recipients role to them.

For more information about how to add roles to role groups, users, or USGs, see the following topics:

- Manage role groups
- Add a role to a user or USG

If you switched to a split permissions model and want to change back to a shared permissions model, see Configure Exchange Server for shared permissions.

Split permissions

If your organization separates Exchange management and Active Directory management, you need to configure Exchange to support the split permissions model. When configured correctly, only the administrators who you want to create security principals, such as Active Directory administrators, will be able to do so and only Exchange administrators will be able to modify the Exchange attributes on existing security principals. This splitting of permissions also falls roughly along the lines of the domain and configuration partitions in Active Directory. Partitions are also called naming contexts. The domain partition stores the users, groups, and other objects for a specific domain. The configuration partition stores the forest-wide configuration information for the services that used Active Directory, such as Exchange. Data that's stored in the domain partition is typically managed by Active Directory administrators, although objects may contain Exchange-specific attributes that can be managed by Exchange administrators. Data that's stored in the configuration partition is managed by the administrators for each respective service that stores data in this partition. For Exchange, this is typically Exchange administrators.

Exchange supports the two following types of split permissions:

- **RBAC split permissions**: Permissions to create security principals in the Active Directory domain partition are controlled by RBAC. Only Exchange servers, services, and those who are members of the appropriate role groups can create security principals.
- **Active Directory split permissions**: Permissions to create security principals in the Active Directory domain partition are completely removed from any Exchange user, service, or server. No option is provided in RBAC to create security principals. Creation of security principals in Active Directory must be performed using Active Directory management tools.

> **IMPORTANT**
>
> In coexistence scenarios, Active Directory split permissions configuration also applies to any Exchange 2010 or later servers in the organization.

If your organization chooses to use a split permissions model instead of shared permissions, we recommend that you use the RBAC split permissions model. The RBAC split permissions model provides significantly more flexibility while providing the nearly same administration separation as Active Directory split permissions, with the exception that Exchange servers and services can create security principals in the RBAC split permissions model.

You're asked whether you want to enable Active Directory split permissions during Setup. If you choose to enable Active Directory split permissions, you can only change to shared permissions or RBAC split permissions by rerunning Setup and disabling Active Directory split permissions. This choice applies to all Exchange 2010 or later servers in the organization.

The following sections describe RBAC and Active Directory split permissions in more detail.

RBAC split permissions

The RBAC security model modifies the default management role assignments to separate who can create security principals in the Active Directory domain partition from those who administer the Exchange organization data in the Active Directory configuration partition. Security principals, such as users with mailboxes and distribution groups, can be created by administrators who are members of the Mail Recipient Creation and Security Group Creation and Membership roles. These permissions remain separate from the permissions required to create security principals outside of the Exchange management tools. Exchange administrators who aren't assigned the Mail Recipient Creation or Security Group Creation and Membership roles can still modify Exchange-related attributes on security principals. Active Directory administrators also have the option of using the Exchange management tools to create Active Directory security principals.

Exchange servers and the Exchange Trusted Subsystem also have permissions to create security principals in Active Directory on behalf of users and third-party programs that integrate with RBAC.

RBAC split permissions is a good choice for your organization if the following are true:

- Your organization doesn't require that security principal creation be performed using only Active Directory management tools and only by users who are assigned specific Active Directory permissions.
- Your organization allows services, such as Exchange servers, to create security principals.
- You want to simplify the process required to create mailboxes, mail-enabled users, distribution groups, and role groups by allowing their creation from within the Exchange management tools.
- You want to manage the membership of distribution groups and role groups within the Exchange management tools.
- You have third-party programs that require that Exchange servers be able to create security principals on their behalf.

If your organization requires a complete separation of Exchange and Active Directory administration where no Active Directory administration can be performed using Exchange management tools or by Exchange services, see the Active Directory Split Permissions section later in this topic.

Switching from shared permissions to RBAC split permissions is a manual process where you remove the permissions required to create security principals from the role groups that are granted them by default. The following table shows the roles that enable the creation of security principals in Exchange and the management role groups they're assigned to by default.

MANAGEMENT ROLE	ROLE GROUP
Mail Recipient Creation role	Organization Management Recipient Management
Security Group Creation and Membership role	Organization Management

By default, members of the Organization Management and Recipient Management role groups can create security principals. You must transfer the ability to create security principals from the built-in role groups to a new role group that you create.

To configure RBAC split permissions, you must do the following:

1. Disable Active Directory split permissions if it's enabled.

2. Create a role group, which will contain the Active Directory administrators that will be able to create security principals.

3. Create regular and delegating role assignments between the Mail Recipient Creation role and the new role group.

4. Create regular and delegating role assignments between the Security Group Creation and Membership role and the new role group.

5. Remove the regular and delegating management role assignments between the Mail Recipient Creation role and the Organization Management and Recipient Management role groups.

6. Remove the regular and delegating role assignments between the Security Group Creation and Membership role and the Organization Management role group.

After completing these steps, only members of the new role group that you create will be able to create security principals, such as mailboxes. The new group will only be able to create the objects. It won't be able to configure the Exchange attributes on the new object. An Active Directory administrator, who is a member of the new group, will need to create the object, and then an Exchange administrator will need to configure the Exchange attributes on the object. Exchange administrators won't be able to use the following cmdlets:

- **New-Mailbox**
- **New-MailContact**
- **New-MailUser**
- **New-RemoteMailbox**
- **Remove-Mailbox**
- **Remove-MailContact**
- **Remove-MailUser**
- **Remove-RemoteMailbox**

Exchange administrators will, however, be able to create and manage Exchange-specific objects, such as mail flow rules (also known as transport rules), distribution groups, and so on and manage Exchange-related attributes on any object.

Additionally, the associated features in the EAC and Outlook on the web (formerly known as Outlook Web App), such as the New Mailbox Wizard, will also no longer be available or will generate an error if you try to use them.

If you want the new role group to also be able to manage the Exchange attributes on the new object, the Mail Recipients role also needs to be assigned to the new role group.

For more information about configuring a split permissions model, see Configure Exchange 2013 for split permissions.

Active Directory split permissions

With Active Directory split permissions, the creation of security principals in the Active Directory domain partition, such as mailboxes and distribution groups, must be performed using Active Directory management tools. Several changes are made to the permissions granted to the Exchange Trusted Subsystem and Exchange servers to limit what Exchange administrators and servers can do. The following changes in functionality occur when you enable Active Directory split permissions:

- Creation of mailboxes, mail-enabled users, distribution groups, and other security principals is removed from the Exchange management tools.

- Adding and removing distribution group members can't be done from the Exchange management tools.
- All permissions granted to the Exchange Trusted Subsystem and Exchange servers to create security principals are removed.
- Exchange servers and the Exchange management tools can only modify the Exchange attributes of existing security principals in Active Directory.

For example, to create a mailbox with Active Directory split permissions enabled, a user must first be created using Active Directory tools by a user with the required Active Directory permissions. Then, the user can be mailbox-enabled using the Exchange management tools. Only the Exchange-related attributes of the mailbox can be modified by Exchange administrators using the Exchange management tools.

Active Directory split permissions is a good choice for your organization if the following are true:

- Your organization requires that security principals be created using only the Active Directory management tools or only by users who are granted specific permissions in Active Directory.
- You want to completely separate the ability to create security principals from those who manage the Exchange organization.
- You want to perform all distribution group management, including creation of distribution groups and adding and removing members of those groups, using Active Directory management tools.
- You don't want Exchange servers, or third-party programs that use Exchange on their behalf, to create security principals.

Notes:

- Switching to Active Directory split permissions is a choice that you can make when you install Exchange by using the Setup wizard or the */ActiveDirectorySplitPermissions* command line switch with Setup.exe (and you must always specify the */PrepareAD* switch along with the */ActiveDirectorySplitPermissions* switch).
- You can also enable or disable Active Directory split permissions **after** you've installed Exchange by rerunning Setup.exe from the command line. To enable Active Directory split permissions, use the value `/ActiveDirectorySplitPermissions:True`. To disable it, use the value `/ActiveDirectorySplitPermissions:False`.
- If you have multiple domains within the same forest, you must also do one of the following steps:
  - Specify the */PrepareAllDomains* switch when you apply Active Directory split permissions.
  - Run Setup.exe with the */PrepareDomain* switch in each domain. You must prepare every domain that contains Exchange servers, mail-enabled objects, or global catalog servers that could be accessed by an Exchange server.
- You can't enable Active Directory split permissions if you've installed Exchange 2010 or later on a domain controller.
- After you enable or disable Active Directory split permissions, we recommend that you restart the Exchange servers in your organization to force them to pick up the new Active Directory access token with the updated permissions.

Exchange achieves Active Directory split permissions by removing permissions and membership from the Exchange Windows Permissions security group. This security group, in shared permissions and RBAC split permissions, is given permissions to many non-Exchange objects and attributes throughout Active Directory. By removing the permissions and membership to this security group, Exchange administrators and services are prevented from creating or modifying those non-Exchange Active Directory objects.

For a list of changes that occur to the Exchange Windows Permissions security group and other Exchange components when you enable or disable Active Directory split permissions, see the following table.

NOTE

Role assignments to role groups that enable Exchange administrators to create security principals are removed when Active Directory split permissions is enabled. This is done to remove access to cmdlets that would otherwise generate an error when they're run because they don't have permissions to create the associated Active Directory object.

ACTION	CHANGES MADE BY EXCHANGE
Enable Active Directory split permissions during first Exchange Server installation	The following actions happen when you enable Active Directory split permissions either through the Setup wizard or by running Setup.exe with the */PrepareAD* and `/ActiveDirectorySplitPermissions:true` command line switches: • An organizational unit (OU) named **Microsoft Exchange Protected Groups** is created. • The **Exchange Windows Permissions** security group is created in the **Microsoft Exchange Protected Groups** OU. • The **Exchange Trusted Subsystem** security group isn't added to the **Exchange Windows Permissions** security group. • Creation of non-delegating management role assignments to management roles with the following management role types is skipped: `MailRecipientCreation` and `SecurityGroupCreationandMembership`. • Access control entries (ACEs) that would have been assigned to the **Exchange Windows Permissions** security group aren't added to the Active Directory domain object. If you run Setup.exe with the */PrepareAllDomains* or */PrepareDomain* switch, the following actions happen in each child domain that's prepared: • All ACEs assigned to the **Exchange Windows Permissions** security group are removed from the domain object. • ACEs are set in each domain with the exception of any ACEs assigned to the **Exchange Windows Permissions** security group.

ACTION	CHANGES MADE BY EXCHANGE
Switch from shared permissions or RBAC split permissions to Active Directory split permissions	The following actions happen when you run the setup.exe command with the */PrepareAD* and `/ActiveDirectorySplitPermissions:true` command line switches: • An OU named **Microsoft Exchange Protected Groups** is created. • The **Exchange Windows Permissions** security group is moved to the **Microsoft Exchange Protected Groups** OU. • The **Exchange Trusted Subsystem** security group is removed from the **Exchange Windows Permissions** security group. • Any non-delegating role assignments to management roles with the following role types are removed: `MailRecipientCreation` and `SecurityGroupCreationandMembership`. • All ACEs assigned to the **Exchange Windows Permissions** security group are removed from the domain object. If you run Setup.exe with either the */PrepareAllDomains* or */PrepareDomain* switch, the following actions happen in each child domain that's prepared: • All ACEs assigned to the **Exchange Windows Permissions** security group are removed from the domain object. • ACEs are set in each domain with the exception of any ACEs assigned to the **Exchange Windows Permissions** security group.

ACTION	CHANGES MADE BY EXCHANGE
Switch from Active Directory split permissions to shared permissions or RBAC split permissions	The following actions happen when you run Setup.exe with the */PrepareAD* and `/ActiveDirectorySplitPermissions:false` switches: • The **Exchange Windows Permissions** security group is moved to the **Microsoft Exchange Security Groups** OU. • The **Microsoft Exchange Protected Groups** OU is removed. • The **Exchange Trusted Subsystems** security group is added to the **Exchange Windows Permissions** security group. •ACEs are added to the domain object for the **Exchange Windows Permissions** security group. If you run setup with either the */PrepareAllDomains* or */PrepareDomain* switch, the following actions happen in each child domain that's prepared: • ACEs are added to the domain object for the **Exchange Windows Permissions** security group. • ACEs are set in each domain including ACEs assigned to the **Exchange Windows Permissions** security group. Role assignments to the Mail Recipient Creation and Security Group Creation and Membership roles aren't automatically created when switching from Active Directory split to shared permissions. If delegating role assignments were customized prior to Active Directory split permissions being enabled, those customizations are left intact. To create role assignments between these roles and the Organization Management role group, see Configure Exchange Server for shared permissions.

After you enable Active Directory split permissions, the following cmdlets are no longer available:

- **New-Mailbox**
- **New-MailContact**
- **New-MailUser**
- **New-RemoteMailbox**
- **Remove-Mailbox**
- **Remove-MailContact**
- **Remove-MailUser**
- **Remove-RemoteMailbox**

After you enable Active Directory split permissions, the following cmdlets are accessible but you can't use them to create distribution groups or modify distribution group membership:

- **Add-DistributionGroupMember**
- **New-DistributionGroup**
- **Remove-DistributionGroup**
- **Remove-DistributionGroupMember**

- **Update-DistributionGroupMember**

Some cmdlets, although still available, may offer only limited functionality when used with Active Directory split permissions. This is because they may allow you to configure recipient objects that are in the domain Active Directory partition and Exchange configuration objects that are in the configuration Active Directory partition. They may also allow you to configure Exchange-related attributes on objects stored in the domain partition. Attempts to use the cmdlets to create objects, or modify non-Exchange-related attributes on objects, in the domain partition will result in an error. For example, the **Add-ADPermission** cmdlet will return an error if you attempt to add permissions to a mailbox. However, the **Add-ADPermission** cmdlet will succeed if you configure permissions on a Receive connector. This is because a mailbox is stored in the domain partition while Receive connectors are stored in the configuration partition.

Additionally, the associated features in the EAC and Outlook on the web, such as the New Mailbox wizard, will also no longer be available or will generate an error if you try to use them.

Exchange administrators will, however, be able to create and manage Exchange-specific objects, such as mail flow rules, and so on.

For more information about configuring an Active Directory split permissions model, see Configure Exchange 2013 for split permissions.

Configure Exchange Server for split permissions

8/3/2020 • 8 minutes to read • Edit Online

Split permissions enable two separate groups, such as Active Directory administrators and Exchange administrators, to manage their respective services, objects, and attributes. Active Directory administrators manage security principals, such as users, that provide permissions to access an Active Directory forest. Exchange administrators manage the Exchange-related attributes on Active Directory objects and Exchange-specific object creation and management.

Exchange Server 2016 and Exchange Server 2019 offer the following types of split permissions models:

- **RBAC split permissions**: Permissions to create security principals in the Active Directory domain partition are controlled by Role Based Access Control (RBAC). Only those who are members of the appropriate role groups can create security principals.

- **Active Directory split permissions**: Permissions to create security principals in the Active Directory domain partition are completely removed from any Exchange user, service, or server. No option is provided in RBAC to create security principals. Creation of security principals in Active Directory must be performed using Active Directory management tools.

The model that you choose depends on the structure and needs of your organization. Choose the procedure that follows that's applicable to the model you want to configure. We recommend that you use the RBAC split permissions model. The RBAC split permissions model provides significantly more flexibility while providing the same administration separation as Active Directory split permissions.

For more information about shared and split permissions, see Split permissions in Exchange Server.

For more information about management role groups, management roles, and regular and delegating management role assignments, see the following topics:

- Understanding Role Based Access Control

- Understanding management role groups

- Understanding management roles

- Understanding management role assignments

What do you need to know before you begin?

- Estimated time to complete each procedure: 5 minutes

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Active Directory split permissions" entry in the Role management permissions topic.

- The permissions model that you select will be applied to all Exchange 2010 or later servers in your organization.

- To download the latest version of Exchange, see Updates for Exchange Server.

- To open the Exchange Management Shell, see Open the Exchange Management Shell.

TIP

Having problems? Ask for help in the Exchange Server forums.

Switch to RBAC split permissions

After you've switched to RBAC split permissions, only Active Directory administrators will be able to create Active Directory security principals. This means that Exchange administrators won't be able to use the following cmdlets:

- **New-Mailbox**
- **New-MailContact**
- **New-MailUser**
- **New-RemoteMailbox**
- **Remove-Mailbox**
- **Remove-MailContact**
- **Remove-MailUser**
- **Remove-RemoteMailbox**

Exchange administrators will only be able to manage the Exchange attributes on existing Active Directory security principals. However, They will be able to create and manage Exchange-specific objects, such as mail flow rules (also known as transport rules) and distribution groups. For more information, see the "RBAC Split Permissions" section in Split permissions in Exchange Server.

To configure Exchange for split permissions, you must assign the Mail Recipient Creation role and the Security Group Creation and Membership role to a role group that contains members that are Active Directory administrators. You must then remove the assignments between those roles and any role group or universal security group (USG) that contains Exchange administrators.

To configure RBAC split permissions, do the following steps:

1. If your organization is currently configured for Active Directory split permissions, do the following steps:

   a. On the target server, open File Explorer, right-click on the Exchange ISO image file, and then select **Mount**. Note the virtual DVD drive letter that's assigned.

   b. Open a Windows Command Prompt window. For example:

   - Press the Windows key + 'R' to open the **Run** dialog, type cmd.exe, and then press **OK**.
   - Press **Start**. In the **Search** box, type **Command Prompt**, then in the list of results, select **Command Prompt**.

   c. In the Command Prompt window, run the following command to disable Active Directory split permissions:

   ```
   Setup.exe /IAcceptExchangeServerLicenseTerms /PrepareAD /ActiveDirectorySplitPermissions:false
   ```

   d. Restart all Exchange servers in your organization or wait for the Active Directory access token to replicate to all of your Exchange servers.

2. Do the following steps in the Exchange Management Shell:

a. Create a role group for the Active Directory administrators. In addition to creating the role group, the command creates regular role assignments between the new role group and the Mail Recipient Creation role and the Security Group Creation and Membership role.

```
New-RoleGroup "Active Directory Administrators" -Roles "Mail Recipient Creation", "Security Group
Creation and Membership"
```

> **NOTE**
>
> If you want members of this role group to be able to create role assignments, include the Role Management role. You don't have to add this role now. However, if you ever want to assign either the Mail Recipient Creation role or Security Group Creation and Membership role to other role assignees, the Role Management role must be assigned to this new role group. The steps that follow configure the Active Directory Administrators role group as the only role group that can delegate these roles.

b. Create delegating role assignments between the new role group and the Mail Recipient Creation role and Security Group Creation and Membership role by running the following commands:

```
New-ManagementRoleAssignment -Role "Mail Recipient Creation" -SecurityGroup "Active Directory
Administrators" -Delegating
New-ManagementRoleAssignment -Role "Security Group Creation and Membership" -SecurityGroup
"Active Directory Administrators" -Delegating
```

c. Add members to the new role group by running the following command:

```
Add-RoleGroupMember "Active Directory Administrators" -Member <user to add>
```

d. Replace the delegate list on the new role group so that only members of the role group can add or remove members by running the following command:

```
Set-RoleGroup "Active Directory Administrators" -ManagedBy "Active Directory Administrators"
```

> **IMPORTANT**
>
> Members of the Organization Management role group, or those who are assigned the Role Management role, either directly or through another role group or USG, can bypass this delegate security check. If you want to prevent any Exchange administrator from adding himself or herself to the new role group, you must remove the role assignment between the Role Management role and any Exchange administrator and assign it to another role group.

e. Find all of the regular and delegating role assignments to the Mail Recipient Creation role by running the following command:

```
Get-ManagementRoleAssignment -Role "Mail Recipient Creation" | Format-Table Name, Role,
RoleAssigneeName -Auto
```

f. Remove all of the regular and delegating role assignments to the Mail Recipient Creation role that aren't associated with the new role group or any other role groups, USGs, or direct assignments you want to keep by running the following command:

```
Remove-ManagementRoleAssignment <Mail Recipient Creation role assignment to remove>
```

> **NOTE**
>
> If you want to remove all of the regular and delegating role assignments to the Mail Recipient Creation role on any role assignee other than the Active Directory Administrators role group, use the following command. The *WhatIf* switch lets you see what role assignments will be removed. Remove the *WhatIf* switch and run the command again to remove the role assignments.

```
Get-ManagementRoleAssignment -Role "Mail Recipient Creation" | Where { $_.RoleAssigneeName -NE
"Active Directory Administrators" } | Remove-ManagementRoleAssignment -WhatIf
```

g. Find all of the regular and delegating role assignments to the Security Group Creation and Membership role by running the following command.

```
Get-ManagementRoleAssignment -Role "Security Group Creation and Membership" | Format-Table Name,
Role, RoleAssigneeName -Auto
```

h. Remove all of the regular and delegating role assignments to the Security Group Creation and Membership role that aren't associated with the new role group or any other role groups, USGs, or direct assignments you want to keep by running the following command:

```
Remove-ManagementRoleAssignment <Security Group Creation and Membership role assignment to
remove>
```

> **NOTE**
>
> You can use the same command in the preceding Note to remove all of the regular and delegating role assignments to the Security Group Creation and Membership role on any role assignee other than the Active Directory Administrators role group, as shown in this example.

```
Get-ManagementRoleAssignment -Role "Security Group Creation and Membership" | Where {
$_.RoleAssigneeName -NE "Active Directory Administrators" } | Remove-ManagementRoleAssignment -
WhatIf
```

For detailed syntax and parameter information, see the following topics:

- New-RoleGroup
- New-ManagementRoleAssignment
- Add-RoleGroupMember
- Set-RoleGroup
- Get-ManagementRoleAssignment
- Remove-ManagementRoleAssignment

Switch to Active Directory split permissions

You can configure your Exchange organization for Active Directory split permissions. Active Directory split permissions completely remove the permissions that allow Exchange administrators and servers from creating security principals in Active Directory or modifying non-Exchange attributes on those objects. When you are done, only Active Directory administrators will be able to create Active Directory security principals. This means that Exchange administrators won't be able to use the following cmdlets:

- Add-DistributionGroupMember
- New-DistributionGroup
- New-Mailbox
- New-MailContact
- New-MailUser
- New-RemoteMailbox
- Remove-DistributionGroup
- Remove-DistributionGroupMember
- Remove-Mailbox
- Remove-MailContact
- Remove-MailUser
- Remove-RemoteMailbox
- Update-DistributionGroupMember

Exchange administrators and servers will only be able to manage the Exchange attributes on existing Active Directory security principals. However, they will be able to create and manage Exchange-specific objects, such as transport rules and Unified Messaging dial plans.

> **WARNING**
> After you enable Active Directory split permissions, Exchange administrators and servers will no longer be able to create security principals in Active Directory, and they won't be able to manage distribution group membership. These tasks must be performed using Active Directory management tools with the required Active Directory permissions. Before you make this change, you should understand the impact it will have on your administration processes and third-party applications that integrate with Exchange and the RBAC permissions model.
>
> For more information, see the "Active Directory split permissions" section in Split permissions in Exchange Server.

To switch from shared or RBAC split permissions to Active Directory split permissions, do the following steps:

1. On the target server, open File Explorer, right-click on the Exchange ISO image file, and then select **Mount**. Note the virtual DVD drive letter that's assigned.

2. In a Windows Command Prompt window, run the following command to enable Active Directory split permissions:

   ```
   Setup.exe /IAcceptExchangeServerLicenseTerms /PrepareAD /ActiveDirectorySplitPermissions:true
   ```

3. If you have multiple Active Directory domains in your organization, you must either run `Setup.exe /IAcceptExchangeServerLicenseTerms /PrepareDomain` in each child domain that contains Exchange servers or objects or run `Setup.exe /IAcceptExchangeServerLicenseTerms /PrepareAllDomains` from a site that has an Active Directory server from every domain.

4. Restart all Exchange servers in your organization or wait for the Active Directory access token to replicate to all of you Exchange 2013 servers.

Configure Exchange Server for shared permissions

8/3/2020 • 7 minutes to read • Edit Online

If you've never configured your organization for split permissions, you don't need to perform this procedure. Exchange Server 2016 and Exchange Server 2019 are configured for shared permissions by default.

Shared permissions enable you, as an Exchange administrator, to create Active Directory security principals, such as users, and then configure them as Exchange recipients. Unlike split permissions, which separate management tasks between groups of Exchange administrators and Active Directory administrators, there's no separation of tasks with shared permissions.

For more information about shared and split permissions, see Split permissions in Exchange Server.

You can configure your Exchange organization for shared permissions if you've previously set your organization for split permissions. The procedure to switch to shared permissions is different depending on whether you're currently using Role Based Access Control (RBAC) split permissions or Active Directory split permissions. Choose the procedure that follows that's applicable to your current configuration. If the following are true, your organization is using Active Directory split permissions:

- The Microsoft Exchange Protected Groups organizational unit (OU) exists.
- The Exchange Windows Permissions security group is located in the Microsoft Exchange Protected Groups OU.
- The Exchange Trusted Subsystem security group is a member of the Exchange Windows Permissions security group.
- There are no regular management role assignments to the Mail Recipient Creation role or Security Group Creation and Membership role.

For more information about management role groups, management roles, and regular and delegating management role assignments, see the following topics:

- Understanding Role Based Access Control
- Understanding management role groups
- Understanding management roles
- Understanding management role assignments

What do you need to know before you begin?

- Estimated time to complete each procedure: 5 minutes
- Procedures in this topic require specific permissions. See each procedure for its permissions information.
- The Exchange organization must currently be configured for RBAC or Active Directory split permissions.
- The permissions model that you select will be applied to all Exchange 2010 or later servers in your organization.
- You must have permissions to delegate the Mail Recipient Creation management role and the Security Group Creation and Membership management role to the Organization Management management role group or another role group that's assigned the Mail Recipients role.

- To download the latest version of Exchange on the target computer, see Updates for Exchange Server.
- To open the Exchange Management Shell, see Open the Exchange Management Shell.

> **TIP**
>
> Having problems? Ask for help in the Exchange Server forums.

Switch from RBAC split permissions to shared permissions

You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Role groups" entry in the Role management permissions topic.

To switch from RBAC split permissions to Exchange shared permissions, you must assign the Mail Recipient Creation role and the Security Group Creation and Membership role to a role group that's also assigned the Mail Recipients role and has Exchange administrators as members. In the default shared permissions configuration, the Organization Management role group contains each of these roles. Because of this, the Organization Management role group is in this procedure.

Configure shared permissions

To configure shared permissions on the Organization Management role group, do the following steps using an account that has permissions to delegate role assignments for the Mail Recipient Creation role and the Security Group Creation and Membership role:

1. Add delegating role assignments for the Mail Recipient Creation role and Security Group Creation and Membership role to the Organization Management role group using the following commands.

   ```
   New-ManagementRoleAssignment -Role "Mail Recipient Creation" -SecurityGroup "Organization Management" -
   Delegating
   New-ManagementRoleAssignment -Role "Security Group Creation and Membership" -SecurityGroup
   "Organization Management" -Delegating
   ```

   > **NOTE**
   >
   > The role group (in this procedure, the Active Directory Administrators role group) that has delegating role assignments for the Mail Recipient Creation role and Security Group Creation and Membership role must be assigned the Role Management role to run the **New-ManagementRoleAssignment** cmdlet. The role assignee that can delegate the Role Management role must assign that role to the Active Directory Administrators role group.

2. Add regular role assignments for the Mail Recipient Creation role to the Organization Management and Recipient Management role groups using the following commands.

   ```
   New-ManagementRoleAssignment -Role "Mail Recipient Creation" -SecurityGroup "Organization Management"
   New-ManagementRoleAssignment -Role "Security Group Creation and Membership" -SecurityGroup "Recipient
   Management"
   ```

3. Add a regular role assignment for the Security Group Creation and Membership role to the Organization Management role group using the following command.

   ```
    New-ManagementRoleAssignment -Role "Security Group Creation and Membership" -SecurityGroup
   "Organization Management"
   ```

For detailed syntax and parameter information, see New-ManagementRoleAssignment.

Remove permissions from Active Directory administrators (Optional)

You can optionally remove the permissions granted to Active Directory administrators if you no longer want them to be able to create or manage Active Directory objects using the Exchange management tools. If you want to remove permissions from Active Directory administrators, perform this procedure.

> **NOTE**
>
> Although you can remove permissions for Active Directory administrators to manage Active Directory objects using the Exchange management tools, Active Directory administrators can continue to manage Active Directory objects using Active Directory management tools, if their Active Directory permissions allow it. They won't, however, be able to manage Exchange-specific attributes on Active Directory objects. For more information, see Split permissions in Exchange Server.

To remove Exchange-related split permissions from Active Directory administrators, do the following steps:

1. Remove the regular and delegating role assignments that assign the Mail Recipient Creation role to the role group or universal security group (USG) that contains the Active Directory administrators as members using the following command. This command uses the Active Directory Administrators role group as an example. The *WhatIf* switch lets you see what role assignments will be removed. Remove the *WhatIf* switch, and run the command again to remove the role assignments.

   ```
   Get-ManagementRoleAssignment -Role "Mail Recipient Creation" | Where { $_.RoleAssigneeName -EQ "Active Directory Administrators" } | Remove-ManagementRoleAssignment -WhatIf
   ```

2. Remove the regular and delegating role assignments that assign the Security Group Creation and Membership role to the role group or USG that contains the Active Directory administrators as members using the following command. This command uses the Active Directory Administrators role group as an example. The *WhatIf* switch lets you see what role assignments will be removed. Remove the *WhatIf* switch, and run the command again to remove the role assignments.

   ```
   Get-ManagementRoleAssignment -Role "Security Group Creation and Membership" | Where { $_.RoleAssigneeName -EQ "Active Directory Administrators" } | Remove-ManagementRoleAssignment -WhatIf
   ```

3. Optional. If you want to remove all Exchange permissions from the Active Directory administrators, you can remove the role group or USG in which they're members. For more information about how to remove a role group, see Manage role groups.

For detailed syntax and parameter information, see Get-ManagementRoleAssignment or Remove-ManagementRoleAssignment.

Switch from Active Directory split permissions to shared permissions

You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Active Directory split permissions" entry in the Role management permissions topic.

To switch from Active Directory split permissions to Exchange shared permissions, you must rerun Exchange Setup to disable Active Directory split permissions in the Exchange organization, and then create role assignments between a role group and the Mail Recipient Creation role and Security Group Creation and Membership role. In the default shared permissions configuration, the Organization Management role group contains each of these roles. Because of this, the Organization Management role group is in this procedure.

IMPORTANT

The Setup.exe command in this procedure makes changes to Active Directory. You must use an account that has the permissions required to make these changes. This account might not be the same account that has permissions to create role assignments using the **New-ManagementRoleAssignment** cmdlet. Use the account, or accounts, with the permissions necessary to successfully complete each step in this procedure.

To switch from Active Directory split permissions to shared permissions, do the following steps:

1. On the target server, open File Explorer, right-click on the Exchange ISO image file that you downloaded, and then select **Mount**. Note the virtual DVD drive letter that's assigned.
2. Open a Windows Command Prompt window. For example:
   - Press the Windows key + 'R' to open the **Run** dialog, type cmd.exe, and then press **OK**.
   - Press **Start**. In the **Search** box, type **Command Prompt**, then in the list of results, select **Command Prompt**.
3. In the Command Prompt window, run the following command:

```
Setup.exe /IAcceptExchangeServerLicenseTerms /PrepareAD /ActiveDirectorySplitPermissions:false
```

4. In the Exchange Management Shell, run the following commands to add regular role assignments between the Mail Recipient Creation role and Security Group Creation and Management role and the Organization Management and Recipient Management role groups.

```
New-ManagementRoleAssignment "Mail Recipient Creation_Organization Management" -Role "Mail Recipient
Creation" -SecurityGroup "Organization Management"
New-ManagementRoleAssignment "Security Group Creation and Membership_Org Management" -Role "Security
Group Creation and Membership" -SecurityGroup "Organization Management"
New-ManagementRoleAssignment "Mail Recipient Creation_Recipient Management" -Role "Mail Recipient
Creation" -SecurityGroup "Recipient Management"
```

5. Restart all Exchange servers in your organization.

For detailed syntax and parameter information, see New-ManagementRoleAssignment.

Recipients

8/3/2020 • 15 minutes to read • Edit Online

The people and resources that send and receive messages are the core of any messaging and collaboration system. In an Exchange organization, these people and resources are referred to as *recipients*. A recipient is any mail-enabled object in Active Directory to which Microsoft Exchange can deliver or route messages.

Exchange recipient types

Exchange includes several explicit recipient types. Each recipient type is identified in the Exchange admin center (EAC) and has a unique value in the *RecipientTypeDetails* property in the Exchange Management Shell. The use of explicit recipient types has the following benefits:

- At a glance, you can differentiate between various recipient types.
- You can search and sort by each recipient type.
- You can more easily perform bulk management operations for selected recipient types.
- You can more easily view recipient properties because the EAC uses the recipient types to render different property pages. For example, the resource capacity is displayed for a room mailbox, but isn't present for a user mailbox.

The following table lists the available recipient types. All these recipient types are discussed in more detail later in this topic.

RECIPIENT TYPE	DESCRIPTION
Dynamic distribution group	A distribution group that uses recipient filters and conditions to derive its membership at the time messages are sent.
Equipment mailbox	A resource mailbox that's assigned to a resource that's not location-specific, such as a portable computer, projector, microphone, or a company car. Equipment mailboxes can be included as resources in meeting requests, providing a simple and efficient way of using resources for your users.
Linked mailbox	A mailbox that's assigned to an individual user in a separate, trusted forest.
Mail contact	A mail-enabled Active Directory contact that contains information about people or organizations that exist outside the Exchange organization. Each mail contact has an external email address. All messages sent to the mail contact are routed to this external email address.
Mail forest contact	A mail contact that represents a recipient object from another forest. Mail forest contacts are typically created by Microsoft Identity Integration Server (MIIS) synchronization. **Note**: Mail forest contacts are read-only recipient objects that are updated only through MIIS or similar custom synchronization. You can't use the EAC or the Exchange Management Shell to remove or modify a mail forest contact.

RECIPIENT TYPE	DESCRIPTION
Mail user	A mail-enabled Active Directory user that represents a user outside the Exchange organization. Each mail user has an external email address. All messages sent to the mail user are routed to this external email address. A mail user is similar to a mail contact, except that a mail user has Active Directory logon credentials and can access resources.
Mail-enabled non-universal group	A mail-enabled Active Directory global or local group object. Mail-enabled non-universal groups were discontinued in Exchange Server 2007 and can exist only if they were migrated from Exchange 2003 or earlier versions of Exchange. You can't use Exchange Server 2013 to create non-universal distribution groups.
Mail-enabled public folder	An Exchange public folder that's configured to receive messages.
Distribution groups	A distribution group is a mail-enabled Active Directory distribution group object that can be used only to distribute messages to a group of recipients.
Mail-enabled security group	A mail-enabled security group is an Active Directory universal security group object that can be used to assign access permissions to resources in Active Directory and can also be used to distribute messages.
Microsoft Exchange recipient	A special recipient object that provides a unified and well-known message sender that differentiates system-generated messages from other messages. It replaces the System Administrator sender used for system-generated messages in earlier versions of Exchange.
Room mailbox	A resource mailbox that's assigned to a meeting location, such as a conference room, auditorium, or training room. Room mailboxes can be included as resources in meeting requests, providing a simple and efficient way of organizing meetings for your users.
Shared mailbox	A mailbox that's not primarily associated with a single user and is generally configured to allow access for multiple users.
Site mailbox	A mailbox comprised of an Exchange mailbox to store email messages and a SharePoint site to store documents. Users can access both email messages and documents using the same client interface. For more information, see Site mailboxes.
User mailbox	A mailbox that's assigned to an individual user in your Exchange organization. It typically contains messages, calendar items, contacts, tasks, documents, and other important business data.
Microsoft 365 or Office 365 mailbox	In hybrid deployments, a Microsoft 365 or Office 365 mailbox consists of a mail user that exists in Active Directory on-premises and an associated cloud mailbox that exists in Exchange Online.

RECIPIENT TYPE	DESCRIPTION
Linked user	A linked user is a user whose mailbox resides in a different forest than the forest in which the user resides.

Mailboxes

Mailboxes are the most common recipient type used by information workers in an Exchange organization. Each mailbox is associated with an Active Directory user account. The user can use the mailbox to send and receive messages, and to store messages, appointments, tasks, notes, and documents. Mailboxes are the primary messaging and collaboration tool for the users in your Exchange organization.

Mailbox components

Each mailbox consists of an Active Directory user and the mailbox data that's stored in the Exchange mailbox database (as shown in the following figure). All configuration data for the mailbox is stored in the Exchange attributes of the Active Directory user object. The mailbox database contains the actual data that's in the mailbox associated with the user account.

> **IMPORTANT**
>
> When you create a mailbox for a new or existing user, the Exchange attributes required for a mailbox are added to the user object in Active Directory. The associated mailbox data isn't created until the mailbox either receives a message or the user signs in to it.

Mailbox components

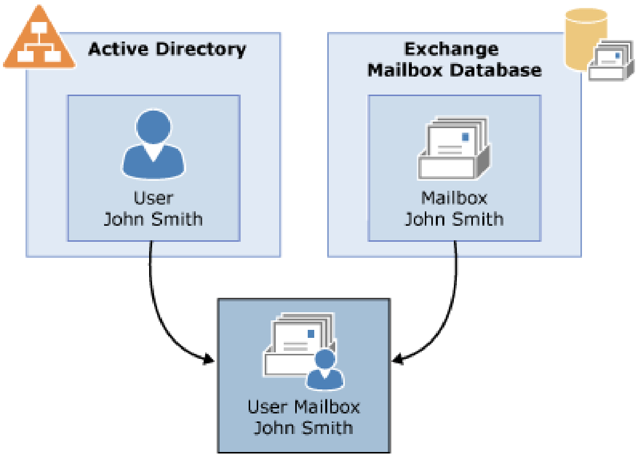

Caution

If you remove a mailbox, the mailbox data stored in the Exchange mailbox database is marked for deletion and the associated user account is also deleted from Active Directory. To retain the user account and delete only the mailbox data, you must disable the mailbox.

Mailbox types

Exchange supports the following mailbox types:

- **User mailboxes**: User mailboxes are assigned to individual users in your Exchange organization. User mailboxes provide your users with a rich collaboration platform. Users can send and receive messages, manage their contacts, schedule meetings, and maintain a task list. They can also have voice mail messages delivered to their mailboxes. User mailboxes are the most commonly used mailbox type and are typically the mailbox type assigned to users in your organization.
- **Linked mailboxes**: Linked mailboxes are mailboxes that are accessed by users in a separate, trusted forest. Linked mailboxes may be necessary for organizations that deploy Exchange in a resource forest. The resource forest scenario allows an organization to centralize Exchange in a single forest, while

allowing access to the Exchange organization with user accounts in one or more trusted forests.

As stated earlier, every mailbox must have a user account associated with it. However, the user account that accesses the linked mailbox doesn't exist in the forest where Exchange is deployed. Therefore, a disabled user account that exists in the same forest as Exchange is associated with each linked mailbox. The following figure illustrates the relationship between the linked user account used to access the linked mailbox and the disabled user account in the Exchange resource forest associated with the linked mailbox.

Linked mailbox

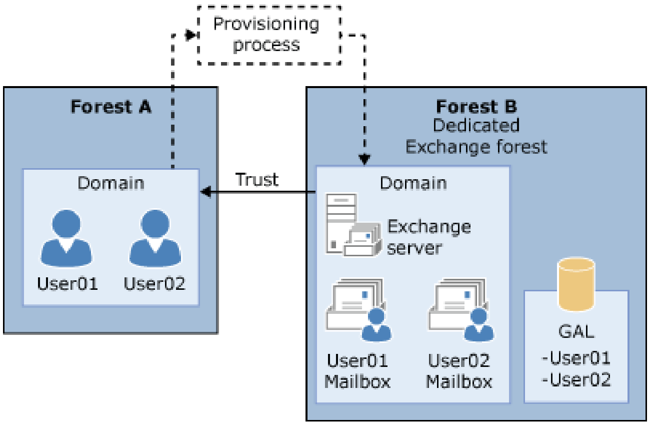

- **Microsoft 365 or Office 365 mailboxes**: When you create a Microsoft 365 or Office 365 mailbox in Exchange Online in a hybrid deployment, the mail user is created in Active Directory on-premises. Directory synchronization, if it's configured, automatically synchronizes this new user object to Microsoft 365 or Office 365, where it's converted to a cloud mailbox in Exchange Online. You can create Microsoft 365 or Office 365 mailboxes as regular user mailboxes, resource mailboxes for meeting rooms and equipment, and shared mailboxes.

- **Shared mailboxes**: Shared mailboxes aren't primarily associated with individual users and are generally configured to allow access by multiple users.

  Although it's possible to assign additional users the logon access permissions to any mailbox type, shared mailboxes are dedicated for this functionality. The Active Directory user associated with a shared mailbox must be a disabled account. After you create a shared mailbox, you must assign permissions to all users that require access to the shared mailbox.

- **Resource mailboxes**: Resource mailboxes are special mailboxes designed to be used for scheduling resources. Like all mailbox types, a resource mailbox has an associated Active Directory user account, but it must be a disabled account. The following are the types of resource mailboxes:

  - **Room mailboxes**: These mailboxes are assigned to meeting locations, such as conference rooms, auditoriums, and training rooms.

  - **Equipment mailboxes**: These mailboxes are assigned to resources that aren't location-specific, such as portable computers, projectors, microphones, or company cars.

    You can include both types of resource mailboxes in meeting requests, providing a simple and efficient way for your users to use resources. You can configure resource mailboxes to automatically process incoming meeting requests based on the resource booking policies that are defined by the resource owners. For example, you can configure a conference room to automatically accept incoming meeting requests except recurring meetings, which can be subject to approval by the resource owner.

System mailboxes

System mailboxes are created by Exchange in the root domain of the Active Directory forest during installation. Users or administrators can't sign in to these mailboxes. System mailboxes are created for Exchange features such as Unified Messaging (UM), migration, message approval, and In-Place eDiscovery. This table lists information about system mailboxes as they're displayed in Active Directory.

> **NOTE**
>
> Unified Messaging is not available in Exchange 2019

MAILBOX	NAME
Organization	SystemMailbox {bb558c35-97f1-4cb9-8ff7-d53741dc928c}
Message approval	SystemMailbox {1f05a927-*xxxx-xxxx-xxxx-xxxxxxxxxxxx*} where *xxxx-xxxx-xxxx-xxxxxxxxxxxx* is a randomly assigned and unique GUID for each Exchange forest
UM data storage	SystemMailbox {e0dc1c29-89c3-4034-b678-e6c29d823ed9} This mailbox exists in Exchange 2016, not in Exchange 2019
Discovery	DiscoverySearchMailbox {D919BA05-46A6-415f-80AD-7E09334BB852}
Federated email	FederatedEmail.4c1f4d8b-8179-4148-93bf-00a95fa1e042
Migration	Migration.8f3e7716-2011-43e4-96b1-aba62d229136

If you want to decommission the last Mailbox server in your Exchange organization, you should first disable these system mailboxes by using the Disable-Mailbox cmdlet. When you decommission a Mailbox server that contains these system mailboxes, you should move the system mailboxes to another Mailbox server to make sure that you don't lose functionality.

Planning for mailboxes

Mailboxes are created in mailbox databases on Exchange servers that have the Mailbox server role installed. To help provide a reliable and effective platform for your mailbox users, detailed planning for the deployment of Mailbox servers and databases is essential. To learn more about planning for Mailbox servers and databases, see Planning and deployment.

Distribution groups

Distribution groups are mail-enabled Active Directory group objects that are primarily used for distributing messages to multiple recipients. Any recipient type can be a member of a distribution group.

> **IMPORTANT**
>
> Note the terminology differences between Active Directory and Exchange. In Active Directory, a distribution group refers to any group that doesn't have a security context, whether it's mail-enabled or not. In Exchange, all mail-enabled groups are referred to as distribution groups, whether they have a security context or not.

Exchange supports the following types of distribution groups:

- **Distribution groups**: These are Active Directory universal distribution group objects that are mail-enabled. They can be used only to distribute messages to a group of recipients.
- **Mail-enabled security groups**: These are Active Directory universal security group objects that are mail-enabled. They can be used to assign access permissions to resources in Active Directory and can also be used to distribute messages.

- **Mail-enabled non-universal groups**: These are Active Directory global or local group objects that are mail-enabled. You can create or mail-enable only universal distribution groups. You may have mail-enabled groups that were migrated from previous versions of Exchange that aren't universal groups. These groups can still be managed by using the EAC or the Exchange Management Shell.

> **NOTE**
>
> To convert a domain-local or a global group to a universal group, you can use the Set-Group cmdlet in the Exchange Management Shell.

Dynamic distribution groups

Dynamic distribution groups are distribution groups whose membership is based on specific recipient filters rather than a defined set of recipients.

Unlike regular distribution groups, the membership list for dynamic distribution groups is calculated each time a message is sent to them, based on the filters and conditions that you specify. When an email message is sent to a dynamic distribution group, it's delivered to all recipients in the organization that match the criteria defined for that dynamic distribution group.

> **IMPORTANT**
>
> A dynamic distribution group includes any recipient in Active Directory that has attributes that match the group's filter at the time a message is sent. If a recipient's properties are modified to match the group's filter, that recipient could inadvertently become a group member and start receiving messages that are sent to the dynamic distribution group. Well-defined, consistent account provisioning processes can reduce the chances of this issue occurring.

To help you create recipient filters for dynamic distribution groups, you can use precanned filters. A *precanned filter* is a commonly used filter that you can use to meet a variety of recipient-filtering criteria. You can use these filters to specify the recipient types that you want to include in a dynamic distribution group. In addition, you can also specify a list of conditions that the recipients must meet. You can create precanned conditions based on the following properties:

- Custom attributes 1-15
- State or province
- Company
- Department
- Recipient container

You can also specify conditions based on recipient properties other than those previously listed. To do this, you must use the Exchange Management Shell to create a custom query for the dynamic distribution group. Remember that the filter and condition settings for dynamic distribution groups that have custom recipient filters can be managed only by using the Exchange Management Shell. For an example of how to create a dynamic distribution group by using a custom query, see Manage dynamic distribution groups.

Mail contacts

Mail contacts typically contain information about people or organizations that exist outside your Exchange organization. Mail contacts can appear in your organization's shared address book (also called the global address list or GAL) and other address lists, and can be added as members to distribution groups. Each contact has an external email address, and all email messages that are sent to a contact are automatically forwarded to that address. Contacts are ideal for representing people external to your Exchange organization (in the shared address book) who don't need access to any internal resources. The following are mail contact types:

- **Mail contacts**: These are mail-enabled Active Directory contacts that contain information about people or organizations that exist outside your Exchange organization.
- **Mail forest contacts**: These represent recipient objects from another forest. These contacts are typically created by directory synchronization. Mail forest contacts are read-only recipient objects that can be updated or removed only by means of synchronization. You can't use Exchange management interfaces to modify or remove a mail forest contact.

Mail users

Mail users are similar to mail contacts. Both have external email addresses, both contain information about people outside your Exchange organization, and both can be displayed in the shared address book and other address lists. However, unlike a mail contact, mail users have Active Directory logon credentials and can access resources to which they are assigned permissions.

If a person external to your organization requires access to resources on your network, you should create a mail user instead of a mail contact. For example, you may want to create mail users for short-term consultants who require access to your server infrastructure, but who will use their own external addresses.

Another scenario is to create mail users in your organization for users who you don't want to maintain an Exchange mailbox. For example, after an acquisition, the acquired company may maintain their separate messaging infrastructure, but may also need access to resources on your network. For those users, you may want to create mail users instead of mailbox users.

> **NOTE**
>
> In the EAC, you use the **Recipients** > **Contacts** page to create and manage mail users. There isn't a separate page for mail users.

Mail-enabled public folders

Public folders are intended to serve as a repository for information shared among many users. Mail-enabling a public folder provides an extra level of functionality to users. In addition to being able to post messages to the folder, users can send email messages to, and sometimes receive email messages from, the public folder. Each mail-enabled folder has an object in Active Directory that stores its email address, address book name, and other mail-related attributes.

You can manage public folders by using either the EAC or the Exchange Management Shell. For more information about managing public folders, see Public folders.

Microsoft Exchange recipient

The Microsoft Exchange recipient is a special recipient object that provides a unified and well-known message sender that differentiates system-generated messages from other messages. It replaces the System Administrator sender that was used for system-generated messages in earlier versions of Exchange.

The Microsoft Exchange recipient isn't a typical recipient object, such as a mailbox, mail user, or mail contact, and it isn't managed by using the typical recipient tools. However, you can use the Set-OrganizationConfig cmdlet in the Exchange Management Shell to configure the Microsoft Exchange recipient.

> **NOTE**
>
> When system-generated messages are sent to an external sender, the Microsoft Exchange recipient isn't used as the sender of the message. Instead, the email address specified by the *ExternalPostmasterAddress* parameter in the Set-TransportConfig cmdlet is used.

Recipients documentation

The following table contains links to topics that will help you learn about and manage Exchange recipients.

TOPIC	DESCRIPTION
Create user mailboxes in Exchange Server	Learn how to create user mailboxes using the Exchange admin center or the Exchange Management Shell.
Manage user mailboxes	Learn how to create user mailboxes, change mailbox properties, and bulk-edit selected properties for multiple mailboxes.
Manage distribution groups	Learn how to create and manage distribution groups, and create a group naming policy for your organization.
Manage dynamic distribution groups	Learn how to create dynamic distribution groups and manage dynamic distribution group properties, such as using custom attributes and other properties to determine group membership.
Manage mail contacts	Learn how to create and manage mail contacts.
Manage mail users	Learn how to create and manage mail users.
Create and manage room mailboxes	Learn how to create room mailboxes and manage room mailbox properties, such as enabling recurring meetings and configuring booking and scheduling options.
Manage equipment mailboxes	Learn how to create equipment mailboxes, configure booking and scheduling options, and manage other mailbox properties.
Disconnected mailboxes	Learn about the two types of disconnected mailboxes and how to work with them.
Custom attributes	Learn how to add information about a recipient by using custom attributes.
Filters in recipient Shell commands	Learn how to use precanned or custom filters with commands to filter a set of recipients.
Manage permissions for recipients	Learn how to use the EAC or the Exchange Management Shell to assign permissions to users and groups.
Automatic Mailbox Distribution	Learn about how automatic mailbox distribution works and how to control which mailbox databases are selected for new and moved mailboxes.

Create user mailboxes in Exchange Server

8/3/2020 • 8 minutes to read • Edit Online

User mailboxes are Exchange mailboxes that are associated with people, typically one mailbox per person. Each user mailbox has an associated Active Directory account that gives the person access to the mailbox to send and receive email messages, and create meetings and appointments.

When you create a new user mailbox in Exchange, you also create the corresponding Active Directory user at the same time. Or, you can create a new mailbox for an existing Active Directory account that doesn't have an associated mailbox. This is known as *mailbox-enabling* an existing user.

You can create user mailboxes in Exchange Server by using the Exchange admin center (EAC) or the Exchange Management Shell. The following table describes some of the important properties for user mailboxes.

PROPERTY	REQUIRED OR OPTIONAL	DESCRIPTION
Alias	Optional	The Exchange alias (also known as the *mail nickname*) for the mailbox. The maximum length is 64 characters. Valid characters are letters, numbers and ASCII text characters that are allowed in email addresses. For example, periods are allowed, but each period must be surrounded by other valid characters (for example, pilar.pinilla). The alias value is used to generate the primary email address (*<alias>*@*<domain>*). If you don't specify an alias value, the username part of the account name (user principal name) is used. The alias value must be unique. **Note**: Don't use apostrophes (') or quotation marks (") in the alias. Although these characters are allowed, they might cause problems later.
Display name	EAC: Required Exchange Management Shell: Optional	Identifies the mailbox in the EAC, and in address lists in Outlook and Outlook on the web (formerly known as Outlook Web App). The maximum length is 256 characters. Spaces and other text characters are allowed. In the EAC, the display name is populated by the values that you enter for the first name, middle initial, and last name, but you can specify a custom value. In the Exchange Management Shell, if you don't specify a value for the display name, the value of the **Name** property is used. The display name value doesn't need to be unique, but having multiple mailboxes with the same display name would be confusing.

PROPERTY	REQUIRED OR OPTIONAL	DESCRIPTION
Name	Required	Specifies the name of the object in Active Directory. Only administrators see this value in Exchange or Active Directory management tools. The maximum length is 64 characters. Spaces and other text characters are allowed. The name value must be unique.

What do you need to know before you begin?

- Estimated time to complete each user mailbox task: 2 to 5 minutes.
- For more information about the EAC, see Exchange admin center in Exchange Server. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Recipient Provisioning Permissions" section in the Recipients Permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Create user mailboxes

The procedures in this section describe how to create a new mailbox and the associated Active Directory user account.

Use the EAC to create user mailboxes

1. In the EAC, go to **Recipients** > **Mailboxes**.
2. Click **New** (+) and then select **User mailbox**.

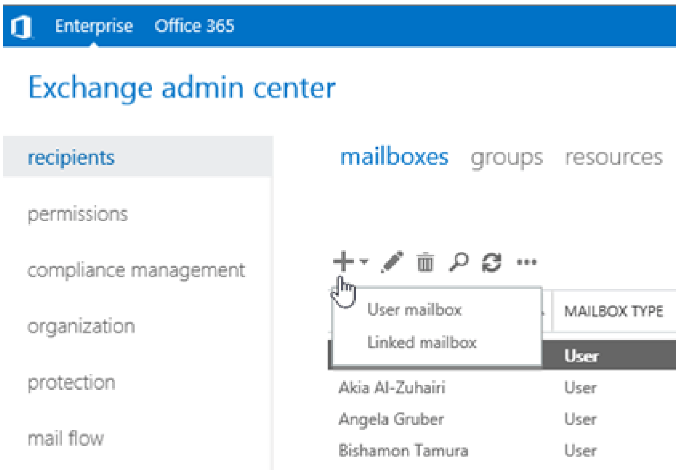

Note: A linked mailbox is a local mailbox that's associated with a user account in a different (trusted) Active Directory forest. For more information, see Manage linked mailboxes.

3. On the **New user mailbox** page, configure the following settings. Settings marked with an asterisk (*) are required.

   - **Alias**
   - **Existing user** or **New user**: Select **New user**.
   - **First name**
   - **Initials**
   - **Last name**
   - *** Display name**: By default, this field is populated with the names you enter in the **First name**, **Initials**, and **Last name** fields, but you can override it. The maximum length is 256 characters.
   - *** Name**: By default, this field is populated with the names you enter in the **First name**, **Initials**, and **Last name** field, but you can override it. The maximum length is 64 characters, and the value must be unique in your organization.
   - **Organizational unit**: Typically, the default location for the user account is the Users container. To change it, click **Browse** and select the OU or container where you want to create the account.
   - *** User logon name**: This is the Active Directory user account that's created and associated with the mailbox.

   Notes:

   - Don't use apostrophes (') or quotation marks ("). Although these characters are allowed, they might cause problems later (for example, assigning access permissions to the mailbox).
   - If this value is different than the **Alias** value, the user's email address and account name will be different (important if the email domain and the Active Directory domain are the same).
   - *** New Password**: Verify the value complies with your organization's password length, complexity, and history requirements.
   - *** Confirm password**
   - **Require password change on next logon**: Select this check box to force the user to change the initial password when they first sign in to the mailbox.

4. You can click **Save** to create the mailbox and the associated Active Directory user account, or you can click **More options** to configure the following additional settings:

   - **Mailbox database**: Click **Browse** to select the mailbox database that holds the mailbox.
   - **Create an on-premises archive mailbox for this user**: Select this check box to create an archive mailbox for the mailbox, and then click **Browse** to select the mailbox database that holds the archive mailbox. Items are automatically moved from the primary mailbox to the archive based on the retention policy settings. For more information, see In-Place Archiving in Exchange Server.
   - **Address book policy**: ABPs define a global address list (GAL), an offline address book (OAB), a room list, and a set of address lists. An ABP gives the user access to a customized GAL in Outlook and Outlook on the web. For more information, see Address book policies in Exchange Server.

   When you're finished, click **Save**.

Use the Exchange Management Shell to create user mailboxes

To create a user mailbox in the Exchange Management Shell, use the following syntax:

```
New-Mailbox -Name <Name> -UserPrincipalName <UPN> -Password (ConvertTo-SecureString -String '<Password>' -
AsPlainText -Force) [-Alias <Alias>] [-FirstName <FirstName>] [-LastName <LastName>] [-DisplayName
<DisplayName>] -[OrganizationalUnit <OU>]
```

This example creates a new mailbox and Active Directory user account for Pilar Pinilla with the following settings:

- **Required parameters**:
  - *Name*: Pilar Pinilla. This value is also used for the display name, because we aren't using the *DisplayName* parameter.
  - *UserPrincipalName*: The Active Directory account name is `pilarp@contoso.com`.
  - *Password*: `Pa$$word1`
- **Optional parameters**:
  - *FirstName*: Pilar
  - *LastName*: Pinilla
  - The alias value is `pilarp` because we aren't using the *Alias* parameter, and `pilarp` is taken from the *UserPrincipalName* parameter value.

```
New-Mailbox -Name "Pilar Pinilla" -UserPrincipalName pilarp@contoso.com -Password (ConvertTo-SecureString -
String 'Pa$$word1' -AsPlainText -Force) -FirstName Pilar -LastName Pinilla
```

For detailed syntax and parameter information, see New-Mailbox.

How do you know this worked?

To verify that you've successfully created a user mailbox, use either of the following procedures:

- In the EAC, go to **Recipients** > **Mailboxes**, and verify the mailbox is displayed in the list.
- In the Exchange Management Shell, replace *<Name>* with the *Name* parameter value that you used, and run the following command:

  ```
  Get-Mailbox -Identity <Name> | Format-List Name,DisplayName,Alias,PrimarySmtpAddress,Database
  ```

Create mailboxes for existing user accounts

When you mailbox-enable a user account, you can only select existing Active Directory users that aren't already mail-enabled (no mail users or accounts that already have an associated mailbox).

Use the EAC to create mailboxes for existing user accounts

1. In the EAC, go to **Recipients** > **Mailboxes**.
2. Click **New** (+) and then select **User mailbox**.

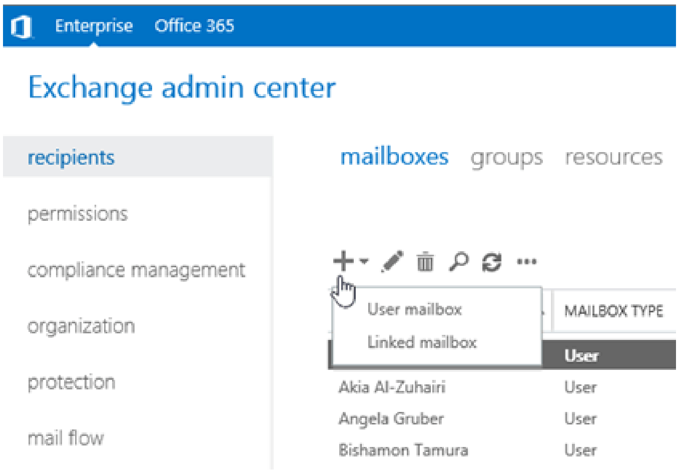

3. On the **New user mailbox** page, configure the following settings.

   - **Alias**: This setting is optional.

     Notes:

     - Don't use apostrophes (') or quotation marks ("). Although these characters are allowed, they might cause problems later.

     - If this value is different than the username part of the user principal name, the user's email address and account name will be different (important if the email domain and the Active Directory domain are the same).

   - **Existing user** or **New user**: Verify **Existing user** is selected, and then click **Browse** to select an available account.

4. You can click **Save** to create the mailbox, or you can click **More options** to configure the following additional settings:

   - **Mailbox database**: Click **Browse** to select the mailbox database that holds the mailbox.

   - **Create an on-premises archive mailbox for this user**: Select this check box to create an archive mailbox for the mailbox, and then click **Browse** to select the mailbox database that holds the archive mailbox. Items are automatically moved from the primary mailbox to the archive based on the retention policy settings. For more information, see In-Place Archiving in Exchange Server.

   - **Address book policy**: ABPs define a global address list (GAL), an offline address book (OAB), a room list, and a set of address lists. An ABP gives the user access to a customized GAL in Outlook and Outlook on the web. For more information, see Address book policies in Exchange Server.

   When you're finished, click **Save**.

Use the Exchange Management Shell to create mailboxes for existing user accounts

To create a mailbox for an existing user account, use the following syntax:

```
Enable-Mailbox -Identity <Account> [-Alias <Alias>] [-DisplayName <DisplayName>] [-Database <Database>]
```

This example creates a mailbox in the mailbox database named UsersMailboxDatabase for the existing user named Kathleen Reiter, whose account name (user principal name) is kreiter@contoso.com.

- Because we aren't using the *Alias* parameter, the alias value is `kreiter`.

- Because we aren't using the *DisplayName* parameter, the value of the **name** attribute in Active Directory is used as the display name.

```
Enable-Mailbox -Identity kreiter@contoso.com -Database UsersMailboxDatabase
```

This example finds all user accounts that aren't mail-enabled and that aren't system accounts (the **userPrincipalName** attribute isn't blank), and then creates mailboxes for those accounts.

```
Get-User -RecipientTypeDetails User -Filter "UserPrincipalName -ne `$null" -ResultSize unlimited | Enable-Mailbox
```

For detailed syntax and parameter information, see Enable-Mailbox and Get-User.

How do you know this worked?

To verify that you've successfully created a mailbox for an existing user, use either of the following procedures:

- In the EAC, go to **Recipients** > **Mailboxes** and verify the mailbox is displayed in the list.
- In the Exchange Management Shell, replace *<Name>* with the name attribute of the user, and run the following command:

  ```
  Get-Mailbox -Identity <Name> | Format-List Name,DisplayName,Alias,PrimarySmtpAddress,Database
  ```

Manage user mailboxes

8/3/2020 • 23 minutes to read • Edit Online

After you create a user mailbox, you can make changes and set additional properties by using the EAC or the Exchange Management Shell.

You can also change properties for multiple user mailboxes at the same time. For more information, see Use the EAC to bulk edit user mailboxes.

What do you need to know before you begin?

- Estimated time to complete each user mailbox task: 2 to 5 minutes.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Recipient Provisioning Permissions" section in the Recipients Permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Change user mailbox properties

Use the EAC to change user mailbox properties

1. In the EAC, navigate to **Recipients** > **Mailboxes**.
2. In the list of user mailboxes, click the mailbox that you want to change the properties for, and then click **Edit** ✏.
3. On the mailbox properties page, click one of the following sections to view or change properties.
    - General
    - Mailbox Usage
    - Contact Information
    - Organization
    - Email Address
    - Mailbox Features
    - Member Of
    - MailTip
    - Mailbox Delegation

General

Use the **General** section to view or change basic information about the user.

- **First name**, **Initials**, **Last name**

- *** Name**: This is the name that's listed in Active Directory. If you change this name, it can't exceed 64 characters.

- *** Display name**: This name appears in your organization's address book, on the To: and From: lines in email, and in the Mailbox list. This name can't contain empty spaces before or after the display name.

- *** Alias**: This specifies the email alias for the user. The user's alias is the portion of the email address on the left side of the at (@) symbol. It must be unique in the forest.

- *** User logon name**: This is the name that the user uses to sign in to their mailbox and to log on to the domain. Typically the user logon name consists of the user's alias on the left side of the @ symbol, and the domain name in which the user account resides on the right side of the @ symbol.

- **Require password change on next logon**: Select this check box if you want the user to reset their password the next time they sign in to their mailbox.

- **Hide from address lists**: Select this check box to prevent the recipient from appearing in the address book and other address lists that are defined in your Exchange organization. After you select this check box, users can still send messages to the recipient by using the email address.

Click **More options** to view or change these additional properties:

- **Organizational unit**: This read-only box displays the organizational unit (OU) that contains the user account. You have to use Active Directory Users and Computers to move the user account to a different OU.

- **Mailbox database**: This read-only box displays the name of the mailbox database that hosts the mailbox. To move the mailbox to a different database, select it in the mailbox list, and then click **Move mailbox to another database** in the Details pane.

- **Custom attributes**: This section displays the custom attributes defined for the user mailbox. To specify custom attribute values, click **Edit**. You can specify up to 15 custom attributes for the recipient.

Mailbox Usage

Use the **Mailbox Usage** section to view or change the mailbox storage quota and deleted item retention settings for the mailbox. These settings are configured by default when the mailbox is created. They use the values that are configured for the mailbox database and apply to all mailboxes in that database. You can customize these settings for each mailbox instead of using the mailbox database defaults.

- **Last logon**: This read-only box displays the last time that the user signed in to their mailbox.

- **Mailbox usage**: This area shows the total size of the mailbox and the percentage of the total mailbox quota that has been used.

> **NOTE**
>
> To obtain the information that's displayed in the previous two boxes, the EAC queries the mailbox database that hosts the mailbox. If the EAC is unable to communicate with the Exchange store that contains the mailbox database, these boxes will be blank. A warning message is displayed if the user hasn't signed in to the mailbox for the first time.

Click **More options** to view or change the mailbox storage quota and the deleted item retention settings for the mailbox.

- **Storage quota settings**: To customize these settings for the mailbox and not use the mailbox database defaults, click **Customize the settings for this mailbox**, type a new value, and then click **Save**.

  The value range for any of the storage quota settings is from 0 through 2047 gigabytes (GB).

    - **Issue a warning at (GB)**: This box displays the maximum storage limit before a warning is issued to the user. If the mailbox size reaches or exceeds the value specified, Exchange sends a warning message to the user.

    - **Prohibit send at (GB)**: This box displays the *prohibit send* limit for the mailbox. If the mailbox size reaches or exceeds the specified limit, Exchange prevents the user from sending new messages and displays a descriptive error message.

    - **Prohibit send and receive at (GB)**: This box displays the *prohibit send and receive* limit for the mailbox. If the mailbox size reaches or exceeds the specified limit, Exchange prevents the mailbox user from sending new messages and won't deliver any new messages to the mailbox. Any messages sent to the mailbox are returned to the sender with a descriptive error message.

- **Deleted item retention settings**: To customize these settings for the mailbox and not use the mailbox database defaults, click **Customize the settings for this mailbox**, type a new value, and then click **Save**.

    - **Keep deleted items for (days)**: This box displays the length of time that deleted items are retained before they are permanently deleted and can't be recovered by the user. When the mailbox is created, this value is based on the deleted item retention settings configured for the mailbox database. By default, a mailbox database is configured to retain deleted items for 14 days. The value range for this property is from 0 through 24855 days.

    - **Don't permanently delete items until the database is backed up**: Select this check box to prevent mailboxes and email messages from being deleted until after the mailbox database on which the mailbox is located has been backed up.

Contact Information

Use the **Contact Information** section to view or change the user's contact information. The information on this page is displayed in the address book. Click **More options** to display additional boxes.

> **TIP**
>
> You can use the **State/Province** box to create recipient conditions for dynamic distribution groups, email address policies, or address lists.

Mailbox users can use Outlook or Outlook on the web (formerly known as Outlook Web App) to view and change their own contact information. But they can't change the information in the **Notes** and **Web page** boxes.

Organization

Use the **Organization** section to record detailed information about the user's role in the organization. This information is displayed in the address book. Also, you can create a virtual organization chart that is accessible from email clients such as Outlook.

- **Title**: Use this box to view or change the recipient's title.

- **Department**: Use this box to view or change the department in which the user works. You can use this box to create recipient conditions for dynamic distribution groups, email address policies, or address lists.

- **Company**: Use this box to view or change the company for which the user works. You can use this box to create recipient conditions for dynamic distribution groups, email address policies, or address lists.

- **Manager**: To add a manager, click **Browse**. In **Select Manager**, select a person, and then click **OK**.

- **Direct reports**: You can't modify this box. A *direct report* is a user who reports to a specific manager. If you've specified a manager for the user, that user appears as a direct report in the details of the manager's mailbox. For example, Kari manages Chris and Kate, so Kari's mailbox is specified in the **Manager** box of Chris's mailbox and Kate's mailbox, and Chris and Kate appear in the **Direct reports** box in the properties of Kari's mailbox.

Email Address

Use the **Email Address** section to view or change the email addresses associated with the user mailbox. This includes the user's primary SMTP address and any associated proxy addresses. The primary SMTP address (also known as the *default reply address*) is displayed in bold text in the address list, with the uppercase **SMTP** value in the **Type** column.

- **Add**: Click **Add** + to add a new email address for this mailbox. Select one of following address types:
  - **SMTP**: This is the default address type. Click this button and then type the new SMTP address in the *** Email address** box.
  - **EUM**: An EUM (Exchange Unified Messaging) address is used by the Microsoft Exchange Unified Messaging service in Exchange 2016 to locate UM-enabled users within an Exchange organization. EUM addresses consist of the extension number and the UM dial plan for the UM-enabled user. Click this button and type the extension number in the **Address/Extension** box. Then click **Browse** and select a dial plan for the user. (**Note**: Unified Messaging is not available in Exchange 2019.)
  - **Custom address type**: Click this button and type one of the supported non-SMTP email address types in the *** Email address** box.

    > **NOTE**
    >
    > With the exception of X.400 addresses, Exchange doesn't validate custom addresses for proper formatting. You must make sure that the custom address you specify complies with the format requirements for that address type.

  - **Make this the reply address**: In Exchange Online, you can select this check box to make the new email address the primary SMTP address for the mailbox. This check box isn't available in the EAC in Exchange Server.
- **Automatically update email addresses based on the email address policy applied to this recipient**: Select this check box to have the recipient's email addresses automatically updated based on changes made to email address policies in your organization. This box is selected by default.

Mailbox Features

Use the **Mailbox Features** section to view or change the following mailbox features and settings:

- **Sharing policy**: This box shows the sharing policy applied to the mailbox. A sharing policy controls how users in your organization can share calendar and contact information with users outside your Exchange organization. The Default Sharing Policy is assigned to mailboxes when they are created. To change the sharing policy that's assigned to the user, select a different one from the drop-down list.
- **Role assignment policy**: This box shows the role assignment policy assigned to the mailbox. The role assignment policy specifies the role-based access control (RBAC) roles that are assigned to the user and control what specific mailbox and distribution group configuration settings users can modify. To change the role assignment policy that's assigned to the user, select a different one from the drop-down list.
- **Retention policy**: This box shows the retention policy assigned to the mailbox. A retention policy is a group of retention tags that are applied to the user's mailbox. They allow you to control how long to keep items in users' mailboxes and define what action to take on items that have reached a certain age. A retention policy isn't assigned to mailboxes when they are created. To assign a retention policy to the user, select one from the drop-down list.
- **Address book policy**: This box shows the address book policy applied to the mailbox. An address book policy allows you to segment users into specific groups to provide customized views of the address book.

To apply or change the address book policy applied to the mailbox, select one from the drop-down list.

- **Unified Messaging (not available in Exchange 2019)**: This feature is disabled by default. When you enable Unified Messaging (UM) in Exchange 2016, the user will be able to use your organization's UM features and a default set of UM properties are applied to the user. Click **Enable** to enable UM for the mailbox. For information about how to enable UM, see Enable a User for Unified Messaging.

  > **NOTE**
  >
  > A UM dial plan and a UM mailbox policy must exist before you can enable UM.

- **Mobile Devices**: Use this section to view and change the settings for Exchange ActiveSync, which is enabled by default. Exchange ActiveSync enables access to an Exchange mailbox from a mobile device. Click **Disable Exchange ActiveSync** to disable this feature for the mailbox.

- **Outlook Web App**: This feature is enabled by default. Outlook on the web enables access to an Exchange mailbox from a web browser. Click **Disable** to disable Outlook on the web for the mailbox. Click **Edit details** to add or change an Outlook on the web mailbox policy for the mailbox.

- **IMAP**: This feature is enabled by default. Click **Disable** to disable IMAP for the mailbox.

- **POP3**: This feature is enabled by default. Click **Disable** to disable POP3 for the mailbox.

- **MAPI**: This feature is enabled by default. MAPI enables access to an Exchange mailbox from a MAPI client such as Outlook. Click **Disable** to disable MAPI for the mailbox.

- **Litigation hold**: This feature is disabled by default. Litigation hold preserves deleted mailbox items and records changes made to mailbox items. Deleted items and all instances of changed items are returned in a discovery search. Click **Enable** to put the mailbox on litigation hold. If the mailbox is on litigation hold, click **Disable** to remove the litigation hold. Mailboxes on litigation hold are inactive mailboxes and can't be deleted. To delete the mailbox, remove the litigation hold. If the mailbox is on litigation hold, click **Edit details** to view and change the following litigation hold settings:

  - **Hold date**: This read-only box indicates the date and time when the mailbox was put on litigation hold.

  - **Put on hold by**: This read-only box indicates the user who put the mailbox on litigation hold.

  - **Note**: Use this box to notify the user about the litigation hold, explain why the mailbox is on litigation hold, or provide additional guidance to the user, such as informing them that the litigation hold won't affect their day-to-day use of email.

  - **URL**: Use this box to provide a URL to a website that provides information or guidance about the litigation hold on the mailbox.

    > **NOTE**
    >
    > The text from these boxes appears in the user's mailbox only if they are using Outlook 2010 or later versions. It doesn't appear in Outlook on the web or other email clients. To view the text from the Note and URL boxes in Outlook, click the **File** tab, and on the **Info** page, under **Account Settings**, you'll see the litigation hold comment.

- **Archiving**: If an archive mailbox doesn't exist for the user, this feature is disabled. To enable an archive mailbox, click **Enable**. If the user has an archive mailbox, the size of the archive mailbox and usage statistics are displayed. Click **Edit details** to view and change the following archive mailbox settings:

  - **Status**: This read-only box indicates whether an archive mailbox exists.

    - **Database**: This read-only box shows the name of the mailbox database that hosts the archive mailbox.
    - **Name**: Type the name of the archive mailbox in this box. This name is displayed under the folder list in Outlook or Outlook on the web.
    - **Archive quota (GB)**: This box shows the total size of the archive mailbox. To change the size, type a new value in the box or select a value from the drop-down list.
    - **Issue warning at (GB)**: This box shows the maximum storage limit for the archive mailbox before a warning is issued to the user. If the archive mailbox size reaches or exceeds the value specified, Exchange sends a warning message to the user. To change this limit, type a new value in the box or select a value from the drop-down list.
- **Delivery Options**: Use to forward email messages sent to the user to another recipient and to set the maximum number of recipients that the user can send a message to. Click **View details** to view and change these settings.
    - **Forwarding address**: Select the **Enable forwarding** check box and then click **Browse** to display the **Select Mail User and Mailbox** page. Use this page to select a recipient to whom you want to forward all email messages that are sent to this mailbox.
    - **Deliver message to both forwarding address and mailbox**: Select this check box so that messages will be delivered to both the forwarding address and the user's mailbox.
    - **Recipient limit**: This setting controls the maximum number of recipients the user can send a message to. Select the **Maximum recipients** check box to limit the number of recipients allowed in the To:, Cc:, and Bcc: boxes of an email message and then specify the maximum number of recipients. For on-premises Exchange organizations, the recipient limit is unlimited.
- **Message Size Restrictions**: These settings control the size of messages that the user can send and receive. Click **View details** to view and change maximum size for sent and received messages.
    - **Sent messages**: To specify a maximum size for messages sent by this user, select the **Maximum message size (KB)** check box and type a value in the box. The message size must be between 0 and 2,097,151 KB. If the user sends a message larger than the specified size, the message will be returned to the user with a descriptive error message.
    - **Received messages**: To specify a maximum size for messages received by this user, select the **Maximum message size (KB)** check box and type a value in the box. The message size must be between 0 and 2,097,151 KB. If the user receives a message larger than the specified size, the message will be returned to the sender with a descriptive error message.
- **Message Delivery Restrictions**: These settings control who can send email messages to this user. Click **View details** to view and change these restrictions.
    - **Accept messages from**: Use this section to specify who can send messages to this user.
    - **All senders**: Select this option to specify that the user can accept messages from all senders. This includes both senders in your Exchange organization and external senders. This option is selected by default. This option includes external users only if you clear the **Require that all senders are authenticated** check box. If you select this check box, messages from external users will be rejected.
    - **Only senders in the following list**: Select this option to specify that the user can accept messages only from a specified set of senders in your Exchange organization. Click **Add** + to display the **Select Recipients** page, which displays a list of all recipients in your Exchange organization. Select the recipients you want, add them to the list, and then click **OK**. You can also search for a specific recipient by typing the recipient's name in the search box and then clicking **Search** 🔎.

  - **Require that all senders are authenticated**: Select this option to prevent anonymous users from sending messages to the user.

  - **Reject messages from**: Use this section to block people from sending messages to this user.

  - **No senders**: Select this option to specify that the mailbox won't reject messages from any senders in the Exchange organization. This option is selected by default.

  - **Senders in the following list**: Select this option to specify that the mailbox will reject messages from a specified set of senders in your Exchange organization. Click **Add** + to display the **Select Recipients** page, which displays a list of all recipients in your Exchange organization. Select the recipients you want, add them to the list, and then click **OK**. You can also search for a specific recipient by typing the recipient's name in the search box and then clicking **Search** 🔎.

Member Of

Use the **Member Of** section to view a list of the distribution groups or security groups to which this user belongs. You can't change membership information on this page. Note that the user may match the criteria for one or more dynamic distribution groups in your organization. However, dynamic distribution groups aren't displayed on this page because their membership is calculated each time they are used.

MailTip

Use the **MailTip** section to add a MailTip to alert users of potential issues if they send a message to this recipient. A MailTip is text that is displayed in the InfoBar when this recipient is added to the To, Cc, or Bcc boxes of a new email message.

> **NOTE**
>
> MailTips can include HTML tags, but scripts aren't allowed. The length of a custom MailTip can't exceed 175 displayed characters. HTML tags aren't counted in the limit.

Mailbox Delegation

Use the **Mailbox Delegation** section to assign permissions to other users (also called *delegates*) to allow them to sign in to the user's mailbox or send messages on behalf of the user. You can assign the following permissions:

- **Send As**: This permission allows users other than the mailbox owner to use the mailbox to send messages. After this permission is assigned to a delegate, any message that a delegate sends from this mailbox will appear as if it was sent by the mailbox owner. However, this permission doesn't allow a delegate to sign in to the user's mailbox.

- **Send on Behalf Of**: This permission also allows a delegate to use this mailbox to send messages. However, after this permission is assigned to a delegate, the **From:** address in any message sent by the delegate indicates that the message was sent by the delegate on behalf of the mailbox owner.

- **Full Access**: This permission allows a delegate to sign in to the user's mailbox and view the contents of the mailbox. However, after this permission is assigned to a delegate, the delegate can't send messages from the mailbox. To allow a delegate to send email from the user's mailbox, you still have to assign the delegate the Send As or the Send on Behalf Of permission.

To assign permissions to delegates, click **Add** + under the appropriate permission to display a page that displays a list of all recipients in your Exchange organization that can be assigned the permission. Select the recipients you want, add them to the list, and then click **OK**. You can also search for a specific recipient by typing the recipient's name in the search box and then clicking **Search** 🔎.

Use the Exchange Management Shell to change user mailbox properties

Use the **Get-Mailbox** and **Set-Mailbox** cmdlets to view and change properties for user mailboxes. One advantage of using the Exchange Management Shell is the ability to change the properties for multiple mailboxes.

For information about what parameters correspond to mailbox properties, see the following topics:

- Get-Mailbox
- Set-Mailbox

Here are some examples of using the Exchange Management Shell to change user mailbox properties.

This example shows how to forward Pat Coleman's email messages to Sunil Koduri's (sunilk@contoso.com) mailbox.

```
Set-Mailbox -Identity patc -DeliverToMailboxAndForward $true -ForwardingAddress sunilk@contoso.com
```

This example uses the **Get-Mailbox** command to find all user mailboxes in the organization, and then uses the **Set-Mailbox** command to set the recipient limit to 500 recipients allowed in the To:, Cc:, and Bcc: boxes of an email message.

```
Get-Mailbox -ResultSize unlimited -Filter "RecipientTypeDetails -eq 'UserMailbox'" | Set-Mailbox -
RecipientLimits 500
```

This example uses the **Get-Mailbox** command to find all the mailboxes in the Marketing organizational unit, and then uses the **Set-Mailbox** command to configure these mailboxes. The custom warning, prohibit send, and prohibit send and receive limits are set to 200 megabytes (MB), 250 MB, and 280 MB respectively, and the mailbox database's default limits are ignored. This command can be used to configure a specific set of mailboxes to have larger or smaller limits than other mailboxes in the organization.

```
Get-Mailbox -OrganizationalUnit "Marketing" | Set-Mailbox -IssueWarningQuota 209715200 -ProhibitSendQuota
262144000 -ProhibitSendReceiveQuota 293601280 -UseDatabaseQuotaDefaults $false
```

This example uses the **Get-Mailbox** cmdlet to find all users in the Customer Service department, and then uses the **Set-Mailbox** cmdlet to change the maximum message size for sending messages to 2 MB.

```
Get-Mailbox -Filter "Department -eq 'Customer Service'" | Set-Mailbox -MaxSendSize 2097152
```

This example sets the MailTip translation in French and Chinese.

```
Set-Mailbox john@contoso.com -MailTipTranslations ("FR: C'est la langue française", "CHT: 這是漢語語言")
```

How do you know this worked?

To verify that you've successfully changed properties for a user mailbox, do the following:

- In the EAC, select the mailbox and then click **Edit** ✎ to view the property or feature that you changed. Depending on the property that you changed, it might be displayed in the Details pane for the selected mailbox.

- In the Exchange Management Shell, use the **Get-Mailbox** cmdlet to verify the changes. One advantage of using the Exchange Management Shell is that you can view multiple properties for multiple mailboxes. In the example above where the recipient limit was changed, run the following command to verify the new value.

  ```
  Get-Mailbox -ResultSize unlimited -Filter "RecipientTypeDetails -eq 'UserMailbox'" | Format-List
  Name,RecipientLimits
  ```

For the example above where the message limits were changed, run this command.

```
Get-Mailbox -OrganizationalUnit "Marketing" | Format-List
Name,IssueWarningQuota,ProhibitSendQuota,ProhibitSendReceiveQuota,UseDatabaseQuotaDefaults
```

Bulk edit user mailboxes

You can use the EAC to change the properties for multiple user mailboxes. When you select two or more user mailboxes from the mailbox list in the EAC, the properties that can be bulk edited are displayed in the Details pane. When you change one of these properties, the change is applied to all selected mailboxes.

Here's a list of the user mailbox properties and features that can be bulk edited. Note that not all properties in each area are available to be changed.

- **Contact Information**: Change shared properties such as street, postal code, and city name.
- **Organization**: Change shared properties such as department name, company name, and the manager that the selected users report to.
- **Custom attributes**: Change or add values for custom attributes 1 - 15.
- **Mailbox quota**: Change the mailbox quota values and the retention period for deleted items.
- **Email connectivity**: Enable or disable Outlook on the web, POP3, IMAP, MAPI, and Exchange ActiveSync.
- **Archive**: Enable or disable the archive mailbox.
- **Retention policy, role assignment policy, and sharing policy**: Update the settings for each of these mailbox features.
- **Move mailboxes to another database**: Move the selected mailboxes to a different database.
- **Delegate permissions**: Assign permissions to users or groups that allow them to open or send messages from other mailboxes. You can assign Full, Send As and Send on Behalf permissions to users or groups. Check out Manage permissions for recipients for more details.

> **NOTE**
>
> The estimated time to complete this task is 2 minutes, but may take longer if you change multiple properties or features.

Use the EAC to bulk edit user mailboxes

1. In the EAC, navigate to **Recipients** > **Mailboxes**.
2. In the list of mailboxes, select two or more mailboxes.

   > **TIP**
   >
   > You can select multiple adjacent mailboxes by holding down the Shift key and clicking the first mailbox, and then clicking the last mailbox you want to edit. You can also select multiple non-adjacent mailboxes by holding down the Ctrl key and clicking each mailbox that you want to edit.

3. In the Details pane, under **Bulk Edit**, select the mailbox properties or feature that you want to edit.
4. Make the changes on the properties page and then save your changes.

How do you know this worked?

To verify that you've successfully bulk edited user mailboxes, do one of the following:

- In the EAC, select each of the mailboxes that you bulk edited and then click **Edit** ✏ to view the property or feature that you changed.

- In the Exchange Management Shell, use the **Get-Mailbox** cmdlet to verify the changes. One advantage of using the Exchange Management Shell is that you can view multiple properties for multiple mailboxes. For example, say you used the bulk edit feature in the EAC to enable the archive mailbox and assign a retention policy to all users in your organization. To verify these changes, you could run the following command:

  ```
  Get-Mailbox -ResultSize unlimited -Filter "RecipientTypeDetails -eq 'UserMailbox'" | Format-List
  Name,ArchiveDatabase,RetentionPolicy
  ```

  For more information about the available parameters for the **Get-Mailbox** cmdlet, see Get-Mailbox.

Add or remove email addresses for a mailbox

8/3/2020 • 6 minutes to read • Edit Online

You can use the EAC or the Exchange Management Shell to add or remove an email address for a user mailbox. You can configure more than one email address for the same mailbox. The additional addresses are called *proxy addresses*. A proxy address lets a user receive email that's sent to a different email address. Any email message sent to the user's proxy address is delivered to their primary email address, which is also known as the *primary SMTP address* or the *default reply address*.

> **NOTE**
>
> The procedures in this topic show how to add or remove email addresses for a user mailbox. You can use similar procedures to add or remove email addresses for other recipient types.

For additional management tasks related to managing recipients, see the "Recipients documentation" table in Recipients.

What do you need to know before you begin?

- Estimated time to complete each procedure: 2 minutes.
- To open the EAC, see Exchange admin center in Exchange Server. To open the Exchange Management Shell, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Recipient Provisioning Permissions" section in the Recipients Permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Add an email address to a user mailbox

Use the EAC to add an email address

1. In the EAC, navigate to **Recipients** > **Mailboxes**.
2. In the list of user mailboxes, click the mailbox that you want to add an email address to, and then click **Edit** ✎.
3. On the mailbox properties page, click **Email Address**.

   > **NOTE**
   >
   > On the **Email Address** page, the primary SMTP address is displayed in bold text in the address list, with the uppercase **SMTP** value in the **Type** column.

4. Click **Add** +, and then click **SMTP** to add an SMTP email address to this mailbox.

> **NOTE**
>
> SMTP is the default email address type. You can also add Exchange Unified Messaging (EUM) addresses or custom addresses to a mailbox in Exchange 2016. For more information, see "Change user mailbox properties" in the Manage user mailboxes topic. (**Note**: Unified Messaging is not available in Exchange 2019.)

5. Type the new SMTP address in the **Email address** box, and then click **OK**.

   The new address is displayed in the list of email addresses for the selected mailbox.

6. Click **Save** to save the change.

Use the Exchange Management Shell to add an email address

The email addresses associated with a mailbox are contained in the *EmailAddresses* property for the mailbox. Because it can contain more than one email address, the *EmailAddresses* property is known as a *multivalued* property. The following examples show different ways to modify a multivalued property.

This example shows how to add an SMTP address to the mailbox of Dan Jump.

```
Set-Mailbox "Dan Jump" -EmailAddresses @{add="dan.jump@northamerica.contoso.com"}
```

This example shows how to add multiple SMTP addresses to a mailbox.

```
Set-Mailbox "Dan Jump" -EmailAddresses @{add="dan.jump@northamerica.contoso.com","danj@tailspintoys.com"}
```

For more information about how to use this method of adding and removing values for multivalued properties, see Modifying Multivalued Properties.

This example shows another way to add email addresses to a mailbox by specifying all addresses associated with the mailbox. In this example, danj@tailspintoys.com is the new email address that you want to add. The other two email addresses are existing addresses. The address with the case-sensitive qualifier `SMTP` is the primary SMTP address. You have to include all email addresses for the mailbox when you use this command syntax. If you don't, the addresses specified in the command will overwrite the existing addresses.

```
Set-Mailbox "Dan Jump" -EmailAddresses
"SMTP:dan.jump@contoso.com","dan.jump@northamerica.contoso.com","danj@tailspintoys.com"
```

For detailed syntax and parameter information, see Set-Mailbox.

How do you know this worked?

To verify that you've successfully added an email address to a mailbox, do one of the following:

- In the EAC, navigate to **Recipients** > **Mailboxes**, click the mailbox, and then click **Edit** ✎.

- On the mailbox properties page, click **Email Address**.

- In the list of email addresses for the mailbox, verify that the new email address is included.

Or

- Run the following command in the Exchange Management Shell.

```
Get-Mailbox <identity> | Format-List EmailAddresses
```

- Verify that the new email address is included in the results.

Remove an email address from a user mailbox

Use the EAC to remove an email address

1. In the EAC, navigate to **Recipients** > **Mailboxes**.
2. In the list of user mailboxes, click the mailbox that you want to remove an email address from, and then click **Edit** ✏.
3. On the mailbox properties page, click **Email Address**.
4. In the list of email addresses, select the address you want to remove, and then click **Remove** −.
5. Click **Save** to save the change.

Use the Exchange Management Shell to remove an email address

This example shows how to remove an email address from the mailbox of Janet Schorr.

```
Set-Mailbox "Janet Schorr" -EmailAddresses @{remove="janets@corp.contoso.com"}
```

This example shows how to remove multiple addresses from a mailbox.

```
Set-Mailbox "Janet Schorr" -EmailAddresses @{remove="janet.schorr@corp.contoso.com","janets@tailspintoys.com"}
```

For more information about how to use this method of adding and removing values for multivalued properties, see Modifying Multivalued Properties.

You can also remove an email address by omitting it from the command to set email addresses for a mailbox. For example, let's say Janet Schorr's mailbox has three email addresses: janets@contoso.com (the primary SMTP address), janets@corp.contoso.com, and janets@tailspintoys.com. To remove the address janets@corp.contoso.com, you would run the following command.

```
Set-Mailbox "Janet Schorr" -EmailAddresses "SMTP:janets@contoso.com","janets@tailspintoys.com"
```

Because janets@corp.contoso.com was omitted in the previous command, it's removed from the mailbox.

For detailed syntax and parameter information, see Set-Mailbox.

How do you know this worked?

To verify that you've successfully removed an email address from a mailbox, do one of the following:

- In the EAC, navigate to **Recipients** > **Mailboxes**, click the mailbox, and then click **Edit** ✏.
- On the mailbox properties page, click **Email Address**.
- In the list of email addresses for the mailbox, verify that the email address isn't included.

Or

- Run the following command in the Exchange Management Shell.

```
Get-Mailbox <identity> | Format-List EmailAddresses
```

- Verify that the email address isn't included in the results.

Use the Exchange Management Shell to add email addresses to multiple mailboxes

You can add a new email address to multiple mailboxes at one time by using the Exchange Management Shell and a comma separated values (CSV) file.

This example imports data from C:\Users\Administrator\Desktop\AddEmailAddress.csv, which has the following format.

```
Mailbox,NewEmailAddress
Dan Jump,danj@northamerica.contoso.com
David Pelton,davidp@northamerica.contoso.com
Kim Akers,kima@northamerica.contoso.com
Janet Schorr,janets@northamerica.contoso.com
Jeffrey Zeng,jeffreyz@northamerica.contoso.com
Spencer Low,spencerl@northamerica.contoso.com
Toni Poe,tonip@northamerica.contoso.com
...
```

Run the following command to use the data in the CSV file to add the email address to each mailbox specified in the CSV file.

```
Import-CSV "C:\Users\Administrator\Desktop\AddEmailAddress.csv" | foreach {Set-Mailbox $_.Mailbox -
EmailAddresses @{add=$_.NewEmailAddress}}
```

> **NOTE**
>
> The column names in the first row of this CSV file (`Mailbox,NewEmailAddress`) are arbitrary. Whatever you use for column names, make sure you use the same column names in the Exchange Management Shell command.

How do you know this worked?

To verify that you've successfully added an email address to multiple mailboxes, do one of the following:

- In the EAC, navigate to **Recipients** > **Mailboxes**, click a mailbox that you added the address to, and then click **Edit** ✏.

- On the mailbox properties page, click **Email Address**.

- In the list of email addresses for the mailbox, verify that the new email address is included.

Or

- Run the following command in the Exchange Management Shell, using the same CSV file that you used to add the new email address.

  ```
  Import-CSV "C:\Users\Administrator\Desktop\AddEmailAddress.csv" | foreach {Get-Mailbox $_.Mailbox |
  Format-List Name,EmailAddresses}
  ```

- Verify that the new email address is included in the results for each mailbox.

Configure email forwarding for a mailbox

8/3/2020 • 3 minutes to read • Edit Online

Email forwarding lets you to set up a mailbox to forward email messages sent to a user's mailbox to another user's mailbox in or outside of your organization.

Use the Exchange admin center and the Exchange Management Shell

You can use either the Exchange admin center (EAC) or Exchange Management Shell to set up email forwarding.

You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Recipient Provisioning Permissions" entry in the Recipients Permissions topic.

Use the Exchange admin center to set up email forwarding

1. In the Exchange admin center, navigate to **Recipients** > **Mailboxes**.

2. In the list of user mailboxes, click or tap the mailbox that you want to set up mail forwarding for, and then click or tap **Edit** ✏.

3. On the mailbox properties page, click **Mailbox Features**.

4. Under **Mail Flow**, select **View details** to view or change the setting for forwarding email messages.

   On this page, you can set the maximum number of recipients that the user can send a message to. The recipient limit is unlimited by default. If you want to specify a limit, click the **Maximum recipients** check box and then type the limit in the text box beneath the check box.

5. Check the **Enable forwarding** check box, and then click or tap **Browse**.

6. On the **Select Recipient** page, select a user you want to forward all email to. Select the **Deliver message to both forwarding address and mailbox** check box if you want both the recipient and the forwarding email address to get copies of the emails sent. Click or tap **OK**, and then click or tap **Save**.

> **NOTE**
>
> What if you want to forward emails to an email address outside your organization? You can use the Exchange Management Shell to do this. See the following example in "Use the Exchange Management Shell to configure mail forwarding".

Use the Exchange Management Shell to set up mail forwarding

Haven't used Exchange Management Shell much? Check out the Exchange Management Shell topic to learn more. Take a look at the Get-Mailbox and Set-Mailbox topics for more details on the cmdlets used here.

This example delivers email to the mailbox of Douglas Kohn and also forwards all mail sent to Douglas Kohn to an external email address, douglaskohn.parents@fineartschool.net.

```
Set-Mailbox -Identity "Douglas Kohn" -DeliverToMailboxAndForward $true -ForwardingSMTPAddress
"douglaskohn.parents@fineartschool.net"
```

This example forwards all email sent to the mailbox of Ken Sanchez, an employee of Contoso Suites, to one of his coworkers, pilarp@contoso.com.

```
Set-Mailbox -Identity "Ken Sanchez" -ForwardingAddress "pilarp@contoso.com"
```

For detailed syntax and parameter information, see Set-Mailbox.

How do you know this worked?

To make sure that you've successfully set up email forwarding, do one of the following:

1. In the Exchange admin center, go to **Recipients** > **Mailboxes**.

2. In the list of user mailboxes, click or tap the mailbox that you configured email forwarding for, and then click **Edit** ✎.

3. On the mailbox properties page, click or tap **Mailbox Features**.

4. Under **Mail Flow**, click or tap **View details** to view the mail forwarding settings.

Or

Run the following command in the Exchange Management Shell.

```
Get-Mailbox <identity> | Format-List ForwardingSMTPAddress,DeliverToMailboxandForward
```

Make sure that the forwarding address is listed in the *ForwardingSMTPAddress* parameter. Also, if the *DeliverToMailboxAndForward* parameter is set to `$true`, messages will be delivered to the mailbox and to the forwarding address. If the parameter is set to `$false`, messages are delivered only to the forwarding address.

End users

Check out the following topics on how to forward your email to another email address by using Outlook and Outlook Web App.

- Forward email to another email account

- Manage email messages by using rules

Additional information

For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Configure message delivery restrictions for a mailbox

8/3/2020 • 4 minutes to read • Edit Online

You can use the EAC or the Exchange Management Shell to place restrictions on whether messages are delivered to individual recipients. Message delivery restrictions are useful to control who can send messages to users in your organization. For example, you can configure a mailbox to accept or reject messages sent by specific users or to accept messages only from users in your Exchange organization.

The message delivery restrictions covered in this topic apply to all recipient types. To learn more about the different recipient types, see Recipients.

For additional management tasks related to recipients, see the following topics:

- Manage user mailboxes
- Manage distribution groups
- Manage dynamic distribution groups
- Manage mail users
- Manage mail contacts

What do you need to know before you begin?

- Estimated time to complete: 5 minutes.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Recipient Provisioning Permissions" section in the Recipients Permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to place message delivery restrictions

1. In the EAC, navigate to **Recipients** > **Mailboxes**.
2. In the list of user mailboxes, click the mailbox that you want to set up message delivery restrictions for, and then click **Edit** ✎.
3. On the mailbox properties page, click **Mailbox Features**.
4. Under **Message Delivery Restrictions**, click **View details** to view and change the following delivery restrictions:
   - **Accept messages from**: Use this section to specify who can send messages to this user.
   - **All senders**: This option specifies that the user can accept messages from all senders. This includes

both senders in your Exchange organization and external senders. This is the default option. It includes external users only if you clear the **Require that all senders are authenticated** check box. If you select this check box, messages from external users will be rejected.

- **Only senders in the following list**: This option specifies that the user can accept messages only from a specified set of senders in your Exchange organization. Click **Add** + to display a list of all recipients in your Exchange organization. Select the recipients you want, add them to the list, and then click **OK**. You can also search for a specific recipient by typing the recipient's name in the search box and then clicking **Search** 🔍.

- **Require that all senders are authenticated**: This option prevents anonymous users from sending messages to the user. This includes external users that are outside of your Exchange organization.

- **Reject messages from**: Use this section to block people from sending messages to this user.

- **No senders**: This option specifies that the mailbox won't reject messages from any senders in the Exchange organization. This is the default option.

- **Senders in the following list**: This option specifies that the mailbox will reject messages from a specified set of senders in your Exchange organization. Click **Add** + to display a list of all recipients in your Exchange organization. Select the recipients you want, add them to the list, and then click **OK**. You can also search for a specific recipient by typing the recipient's name in the search box and then clicking **Search** 🔍.

5. Click **OK** to close the **Message Delivery Restrictions** page, and then click **Save** to save your changes.

Use the Exchange Management Shell to place message delivery restrictions

The following examples show how to use the Exchange Management Shell to configure message delivery restrictions for a mailbox. For other recipient types, use the corresponding **Set-** cmdlet with the same parameters.

This example configures the mailbox of Robin Wood to accept messages only from the users Lori Penor, Jeff Phillips, and members of the distribution group Legal Team 1.

```
Set-Mailbox -Identity "Robin Wood" -AcceptMessagesOnlyFrom "Lori Penor","Jeff Phillips" -
AcceptMessagesOnlyFromDLMembers "Legal Team 1"
```

NOTE

If you're configuring a mailbox to accept messages only from individual senders, you have to use the *AcceptMessagesOnlyFrom* parameter. If you're setting up a mailbox to accept messages only from senders that are members of a specific distribution group, use the *AcceptMessagesOnlyFromDLMembers* parameter.

This example adds the user named David Pelton to the list of users whose messages will be accepted by the mailbox of Robin Wood.

```
Set-Mailbox -Identity "Robin Wood" -AcceptMessagesOnlyFrom @{add="David Pelton"}
```

This example configures the mailbox of Robin Wood to require all senders to be authenticated. This means the mailbox will only accept messages sent by other users in your Exchange organization.

```
Set-Mailbox -Identity "Robin Wood" -RequireSenderAuthenticationEnabled $true
```

This example configures the mailbox of Robin Wood to reject messages from the users Joe Healy, Terry Adams, and members of the distribution group Legal Team 2.

```
Set-Mailbox -Identity "Robin Wood" -RejectMessagesFrom "Joe Healy","Terry Adams" -RejectMessagesFromDLMembers
"Legal Team 2"
```

This example configures the mailbox of Robin Wood to also reject messages sent by members of the group Legal Team 3.

```
Set-Mailbox -Identity "Robin Wood" -RejectMessagesFromDLMembers @{add="Legal Team 3"}
```

> **NOTE**
>
> If you're setting up a mailbox to reject messages from individual senders, you have to use the *RejectMessagesFrom* parameter. If you're setting up a mailbox to reject messages from senders that are members of a specific distribution group, use the *RejectMessagesFromDLMembers* parameter.

For detailed syntax and parameter information related to placing delivery restrictions for different types of recipients, see the following topics:

- Set-DistributionGroup
- Set-DynamicDistributionGroup
- Set-Mailbox
- Set-MailContact
- Set-MailUser

How do you know this worked?

To verify that you've successfully placed message delivery restrictions for a user mailbox, do one the following:

1. In the EAC, navigate to **Recipients** > **Mailboxes**.
2. In the list of user mailboxes, click the mailbox that you want to verify the message delivery restrictions for, and then click **Edit** ✎.
3. On the mailbox properties page, click **Mailbox Features**.
4. Under **Message Delivery Restrictions**, click **View details** to verify the delivery restrictions for the mailbox.

Or

Run the following command in the Exchange Management Shell.

```
Get-Mailbox <identity> | Format-List
AcceptMessagesOnlyFrom,AcceptMessagesOnlyFromDLMembers,RejectMessagesFrom,RejectMessagesFromDLMembers,RequireS
enderAuthenticationEnabled
```

Configure message size limits for a mailbox

8/3/2020 • 2 minutes to read • Edit Online

Message size limits control the size of messages that a user can send and receive. By default, when a mailbox is created, there isn't a size limit for sent and received messages.

Keep in mind that there are other settings in an Exchange organization that determine the maximum message size a mailbox can send and receive (for example, the maximum message size configured on a Mailbox server). To learn more about the message size restrictions in Exchange, including the types of message size limits, their scope, and the order of precedence, see Message size and recipient limits in Exchange Server.

For additional management tasks related to user mailboxes, see Manage user mailboxes.

> **NOTE**
>
> You can also control the size of messages sent and received by mail users and from shared mailboxes.

What do you need to know before you begin?

- Estimated time to complete: 2 minutes.
- To open the EAC, see Exchange admin center in Exchange Server. To open the Exchange Management Shell, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Recipient Provisioning Permissions" section in the Recipients Permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to set message size limits

1. In the EAC, navigate to **Recipients** > **Mailboxes**.
2. In the list of user mailboxes, click the mailbox that you want to change the message size limits for, and then click **Edit** ✏.
3. On the mailbox properties page, click **Mailbox Features**.
4. Under **Message Size Restrictions**, click **View details** to view and change the following message size limits:
    - **Sent messages**: To set a maximum size for messages sent by this user, select the **Maximum message size (KB)** check box and type a value in the box. The message size must be between 0 and 2,097,151 KB. If the user sends a message larger than the specified size, the message will be returned to the user with a descriptive error message.

   - **Received messages**: To set a maximum size for messages received by this user, select the **Maximum message size (KB)** check box and type a value in the box. The message size must be between 0 and 2,097,151 KB. If the user receives a message larger than the specified size, the message will be returned to the sender with a descriptive error message.

5. Click **OK**, and then click **Save** to save your changes.

Use the Exchange Management Shell to configure message size limits

This example sets the maximum size for sent messages to 25 MB and the maximum size for received messages to 35 MB for the mailbox of Debra Garcia.

```
Set-Mailbox -Identity "Debra Garcia" -MaxSendSize 25mb -MaxReceiveSize 35mb
```

For detailed syntax and parameter information, see Set-Mailbox.

How do you know this worked?

To verify that you've successfully set up message size limits for a mailbox, do one of the following:

1. In the EAC, navigate to **Recipients** > **Mailboxes**.

2. In the list of user mailboxes, click the mailbox that you want to verify the message size limits for, and then click **Edit** ✎.

3. On the mailbox properties page, click **Mailbox Features**.

4. Under **Message Size Restrictions**, click **View details** to verify the message size limits for the mailbox.

   Or

   Run the following command in the Exchange Management Shell.

   ```
   Get-Mailbox -Identity <Identity> | Format-List MaxSendSize,MaxReceiveSize
   ```

Configure storage quotas for a mailbox

8/3/2020 • 3 minutes to read • Edit Online

You can use the Exchange admin center (EAC) or the Exchange Management Shell to customize the mailbox storage quotas for specific mailboxes. Storage quotas let you control the size of mailboxes and manage the growth of mailbox databases. When a mailbox reaches or exceeds a specified storage quota, Exchange sends a descriptive notification to the mailbox owner.

Storage quotas are typically configured on a per-database basis. This means that the quotas configured for a mailbox database apply to all mailboxes in that database. For more information about managing per-database mailbox settings, see Manage mailbox databases in Exchange Server.

This topic shows you how to customize storage settings for a specific mailbox instead of using the storage settings from the mailbox database. For additional management tasks related to user mailboxes, see Manage user mailboxes.

What do you need to know before you begin?

- Estimated time to complete: 2 minutes.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Recipient Provisioning Permissions" section in the Recipients Permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to set storage quotas for a mailbox

1. In the EAC, navigate to **Recipients** > **Mailboxes**.
2. In the list of user mailboxes, click the mailbox that you want to change the storage quotas for, and then click **Edit** ✎.
3. On the mailbox properties page, click **Mailbox Usage**, and then click **More options**.
4. Click **Customize the settings for this mailbox**, and then set the following boxes. The value range for any of the storage quota settings is from 0 through 2047 gigabytes (GB).

   - **Issue a warning at (GB)**: This box displays the maximum storage limit before a warning is issued to the user. If the mailbox size reaches or exceeds the value specified, Exchange sends a warning message to the user.

> **IMPORTANT**
>
> The message associated with the **Issue warning** quota won't be sent to the user unless the value of this setting is greater than 50% of the value specified in the **Prohibit send** quota. For example, if you set the **Prohibit send** quota to 8 MB, you must set the **Issue warning** quota to at least 4 MB. If you don't, the **Issue warning** quota message won't be sent.

- **Prohibit send at (GB)**: If the mailbox size reaches or exceeds the specified limit, Exchange prevents the user from sending new messages and displays a descriptive error message.

- **Prohibit send and receive at (GB)**: If the mailbox size reaches or exceeds the specified limit, Exchange prevents the mailbox user from sending new messages and won't deliver any new messages to the mailbox. Any messages sent to the mailbox are returned to the sender with a descriptive error message.

5. Click **Save** to save your changes.

Use the Exchange Management Shell to configure storage quotas for a mailbox

This example sets the issue warning, prohibit send, and prohibit send and receive quotas for Joe Healy's mailbox to 24.5 GB, 24.75 GB, and 25 GB respectively.

> **NOTE**
>
> To ensure that the custom settings for the mailbox are used rather than the mailbox database defaults, you must set the *UseDatabaseQuotaDefaults* parameter to `$false`.

```
Set-Mailbox -Identity "Joe Healy" -IssueWarningQuota 24.5gb -ProhibitSendQuota 24.75gb -
ProhibitSendReceiveQuota 25gb -UseDatabaseQuotaDefaults $false
```

This example sets the issue warning, prohibit send, and prohibit send and receive quotas for Ayla Kol's mailbox to 900 megabytes (MB), 950 MB, and 1 GB respectively, and configures the mailbox to use custom settings.

```
Set-Mailbox -Identity "Ayla Kol" -IssueWarningQuota 900mb -ProhibitSendQuota 950mb -ProhibitSendReceiveQuota
1gb -UseDatabaseQuotaDefaults $false
```

For detailed syntax and parameter information, see Set-Mailbox.

How do you know this worked?

To verify that you've successfully set the storage quotas for a mailbox, do one of the following:

1. In the EAC, navigate to **Recipients** > **Mailboxes**.

2. In the list of user mailboxes, click the mailbox that you want to verify the storage quotas for, and then click **Edit** ✎.

3. On the mailbox properties page, click **Mailbox Usage**, and then click **More options**.

4. Verify that **Customize the settings for this mailbox** is selected.

5. Verify the storage quota settings.

Or

In the Exchange Management Shell, replace <Identity> with the name, email address or alias of the mailbox, and run the following command:

```
Get-Mailbox <Identity> | Format-List
IssueWarningQuota,ProhibitSendQuota,ProhibitSendReceiveQuota,UseDatabaseQuotaDefaults
```

Configure Deleted Item retention and Recoverable Items quotas

8/3/2020 • 3 minutes to read • Edit Online

When a user deletes items from the Deleted Items default folder by using the Delete, Shift+Delete, or **Empty Deleted Items Folder** actions, the items are moved to the **Recoverable Items\Deletions** folder. The duration that deleted items remain in this folder is based on the deleted item retention settings configured for the mailbox database or the mailbox. By default, a mailbox database is configured to retain deleted items for 14 days, and the recoverable items warning quota and recoverable items quota are set to 20 gigabytes (GB) and 30 GB respectively.

> **NOTE**
>
> Before the retention time for deleted items elapses,Outlook and Outlook on the web users can recover deleted items by using the Recover Deleted Items feature. To learn more about these features, see the "Recover deleted items" topic for Outlook for Windows or Outlook on the web.

You can use the Exchange Management Shell to configure deleted item retention settings and recoverable items quotas for a mailbox or mailbox database. Deleted item retention settings are ignored when a mailbox is placed on In-Place Hold or litigation hold.

To learn more about deleted item retention, the Recoverable Items folder, In-Place Hold, and litigation hold, see Recoverable Items folder in Exchange Server.

What do you need to know before you begin?

- Estimated time to completion: 5 minutes.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Messaging records management" entry in the Messaging policy and compliance permissions in Exchange Server topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Configure deleted item retention for a mailbox

Use the Exchange admin center (EAC) to configure deleted item retention for a mailbox

1. Navigate to **Recipients** > **Mailboxes**.
2. In the list view, select a mailbox, and then click **Edit** ✎.
3. On the mailbox property page, click **Mailbox usage**, click **More options**, and then select one of the following:
   - **Use the default retention settings from the mailbox database**: Use the deleted item retention

setting that's configured for the mailbox database.

- **Customize the settings for this mailbox**: Configure deleted item retention settings for the mailbox.

- ***Keep deleted items for (days)**: Displays the length of time that deleted items are retained before they're permanently deleted and can't be recovered by the user. When the mailbox is created, this value is based on the deleted item retention settings configured for the mailbox database. By default, a mailbox database is configured to retain deleted items for 14 days. The value range for this property is from 0 through 24,855 days.

- **Don't permanently delete items until the database is backed up**: Check this box to prevent mailboxes and email messages from being deleted until after the mailbox database on which the mailbox is located has been backed up.

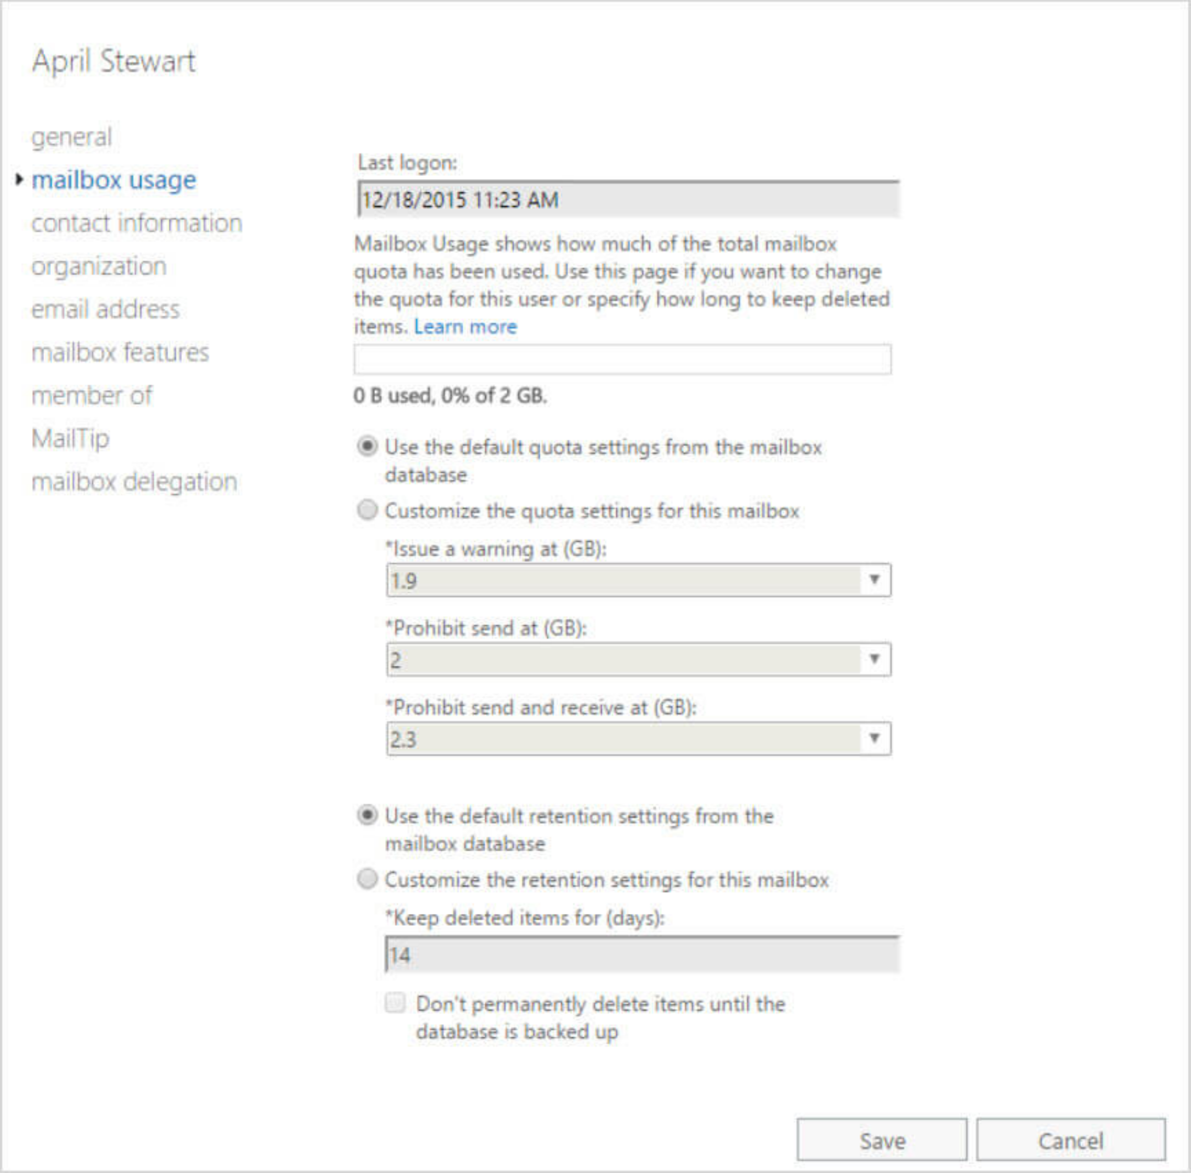

April Stewart

Last logon:
12/18/2015 11:23 AM

Mailbox Usage shows how much of the total mailbox quota has been used. Use this page if you want to change the quota for this user or specify how long to keep deleted items. Learn more

0 B used, 0% of 2 GB.

● Use the default quota settings from the mailbox database

○ Customize the quota settings for this mailbox

*Issue a warning at (GB):
1.9

*Prohibit send at (GB):
2

*Prohibit send and receive at (GB):
2.3

● Use the default retention settings from the mailbox database

○ Customize the retention settings for this mailbox

*Keep deleted items for (days):
14

□ Don't permanently delete items until the database is backed up

Use the Exchange Management Shell to configure deleted item retention for a mailbox

This example configures April Stewart's mailbox to retain deleted items for 30 days.

```
Set-Mailbox -Identity - "April Stewart" -RetainDeletedItemsFor 30
```

For detailed syntax and parameter information, see Set-Mailbox.

Use the Exchange Management Shell to configure recoverable items quotas for a mailbox

> **NOTE**
>
> You can't use the EAC to configure recoverable items quotas for a mailbox.

This example configures a recoverable items warning quota of 12 GB and a recoverable items quota of 15 GB for April Stewart's mailbox.

```
Set-Mailbox -Identity "April Stewart" -RecoverableItemsWarningQuota 12GB -RecoverableItemsQuota 15GB -
UseDatabaseQuotaDefaults $false
```

> **NOTE**
>
> To configure a mailbox to use different recoverable items quotas than the mailbox database in which it resides, you must set the *UseDatabaseQuotaDefaults* parameter to `$false`.

For detailed syntax and parameter information, see Set-Mailbox.

Use the Exchange Management Shell to configure deleted item retention for a mailbox database

> **NOTE**
>
> You can't use the EAC to configure deleted item retention for a mailbox database.

This example configures a deleted item retention period of 10 days for the mailbox database MDB2.

```
Set-MailboxDatabase -Identity MDB2 -DeletedItemRetention 10
```

For detailed syntax and parameter information, see Set-MailboxDatabase.

Use the Exchange Management Shell to configure recoverable items quotas for a mailbox database

> **NOTE**
>
> You can't use the EAC to configure recoverable items quotas for a mailbox database

This example configures a recoverable items warning quota of 15 GB and a recoverable items quota of 20 GB on mailbox database MDB2.

```
Set-MailboxDatabase -Identity MDB2 -RecoverableItemsWarningQuota 15GB -RecoverableItemsQuota 20GB
```

For detailed syntax and parameter information, see Set-MailboxDatabase.

Convert a mailbox in Exchange Server

8/3/2020 • 3 minutes to read • Edit Online

In Exchange Server 2013 or later, converting a mailbox from one type of mailbox to another is mostly unchanged from the experience in Exchange 2010. You still need to use the **Set-Mailbox** cmdlet in the Exchange Management Shell to do the conversion.

You can convert the following mailboxes to a different type:

- User mailbox to room or equipment mailbox
- User mailbox to shared mailbox
- Shared mailbox to user mailbox
- Shared mailbox to room or equipment mailbox
- Room or equipment mailbox to user mailbox
- Room or equipment mailbox to shared mailbox

> **NOTE**
> If your organization uses a hybrid Exchange environment, you need to manage your mailboxes by using the on-premises Exchange management tools. To convert a mailbox in a hybrid environment, you might need to move the mailbox back to on-premises Exchange, convert the mailbox type, and then move it back to Microsoft 365 or Office 365.

What do you need to know before you begin?

- Estimated time to complete: 5 minutes.
- Room, equipment, and shared mailboxes have associated user accounts in Active Directory, but the accounts are disabled. When you convert one of these mailbox types to a regular (user) mailbox, you need to specify a password that satisfies the length and complexity requirements for your organization.

  Overwriting an existing password requires the Reset Password role, which isn't assigned to any role groups by default. To assign the role to a role group that you belong to, see Add a role to a role group. Note that changes in permission require you to log off and log on for the changes to take effect.
- When you convert a regular (user) mailbox to a room, equipment, or shared mailbox, the associated account is disabled.

  For room mailboxes, you can enable the associated user account, which also requires you to specify a password (which requires the Reset Password role). You need to enable the room mailbox user account for features like the Skype for Business Room System.
- To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Recipient Provisioning Permissions" section in the Recipients Permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard

shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to convert a mailbox

To convert a mailbox to a different type, use this syntax:

```
Set-Mailbox -Identity <MailboxIdentity> -Type <Regular | Room | Equipment | Shared> [-Password (ConvertTo-SecureString -String '<Password>' -AsPlainText -Force)] [-EnableRoomMailboxAccount <$true | $false>] [-RoomMailboxPassword (ConvertTo-SecureString -String '<Password>' -AsPlainText -Force)] [-ResetPasswordOnNextLogon <$true | $false>]
```

This example converts the shared mailbox named Marketing Dept 01 to a user mailbox with the new password P@ssw0rd25, and the requirement to change the password the next time the user logs in to the mailbox.

```
Set-Mailbox -Identity "Marketing Dept 01" -Type Regular -Password (ConvertTo-SecureString -String 'P@ssw0rd25' -AsPlainText -Force) -ResetPasswordOnNextLogon $true
```

This example converts the user mailbox named Conference Room 01 to a room mailbox.

```
Set-Mailbox -Identity "Conference Room 01" -Type Room
```

This is the same example, but the user account for the room mailbox is enabled, and the password is P@ssw0rd25

```
Set-Mailbox -Identity "Conference Room 01" -Type Room -EnableRoomMailboxAccount $true -RoomMailboxPassword (ConvertTo-SecureString -String 'P@ssw0rd25' -AsPlainText -Force)
```

Note: Even when you convert a user mailbox with a known password to a room mailbox, you still need to use the *RoomMailboxPassword* parameter to specify a password.

For detailed syntax and parameter information, see Set-Mailbox.

How do you know this worked?

To verify that you've successfully converted a mailbox, replace *<MailboxIdentity>* with the name, alias, or email address of the mailbox, and run this command in the Exchange Management Shell to verify the property values:

```
Get-Mailbox -Identity <MailboxIdentity> | Format-List Name,RecipientTypeDetails,UserPrincipalName,AccountDisabled
```

For detailed syntax and parameter information, see Get-Mailbox.

Enable or disable single item recovery for a mailbox

8/3/2020 • 2 minutes to read • Edit Online

You can use the Exchange Management Shell to enable or disable single item recovery on a mailbox. In Exchange Server, single item recovery is disabled when a mailbox is created. If single item recovery is enabled, messages that are purged (hard-deleted) by the user are retained in the Recoverable Items folder of the mailbox until the deleted item retention period expires. This lets an administrator recover messages purged by the user before the deleted item retention period expires. Also, if a message is changed by a user or a process, copies of the original item are also retained when single item recovery is enabled.

What do you need to know before you begin?

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Retention and legal holds" entry in the Recipients Permissions topic.
- You can't use the Exchange admin center (EAC) to enable or disable single item recovery.
- In Exchange Server, the mailbox uses the deleted item retention settings of the mailbox database, by default. The deleted item retention period for a mailbox database is set to 14 days, but you can override the default by configuring this setting on a per-mailbox basis. For details, see Configure Deleted Item retention and Recoverable Items quotas.

Enable single item recovery

This example enables single item recovery for the mailbox of April Summers.

```
Set-Mailbox -Identity "April Summers" -SingleItemRecoveryEnabled $true
```

This example enables single item recovery for the mailbox of Pilar Pinilla and sets the number of days that deleted items are retained to 30 days.

```
Set-Mailbox -Identity "Pilar Pinilla" -SingleItemRecoveryEnabled $true -RetainDeletedItemsFor 30
```

This example enables single item recovery for all user mailboxes in the organization.

```
Get-Mailbox -ResultSize unlimited -Filter "RecipientTypeDetails -eq 'UserMailbox'" | Set-Mailbox -
SingleItemRecoveryEnabled $true
```

This example enables single item recovery for all user mailboxes in the organization and sets the number of days that deleted items are retained to 30 days

```
Get-Mailbox -ResultSize unlimited -Filter "RecipientTypeDetails -eq 'UserMailbox'" | Set-Mailbox -
SingleItemRecoveryEnabled $true -RetainDeletedItemsFor 30
```

For detailed syntax and parameter information, see Set-Mailbox.

Disable single item recovery

You might need to disable single item recovery for a user's mailbox. For example, before you can use `Search-`

Mailbox -DeleteContent to permanently delete content from a mailbox, you have to disable single item recovery. For more information, see Search for and delete messages in Exchange Server.

This example disables single item recovery for the mailbox of Ayla Kol.

```
Set-Mailbox -Identity "Ayla Kol" -SingleItemRecoveryEnabled $false
```

How do you know this worked?

To verify that you've enabled single item recovery for a mailbox and display the value for how long deleted items will be retained (in days), run the following command.

```
Get-Mailbox <Name> | Format-List SingleItemRecoveryEnabled,RetainDeletedItemsFor
```

You can use this same command to verify that single item recovery is disabled for a mailbox.

More information

- To learn more about single item recovery, see Recoverable Items folder in Exchange Server. To recover messages purged by the user before the deleted item retention period expires, see Recover deleted messages in a user's mailbox .

- If a mailbox is placed on In-Place Hold or Litigation Hold, messages in the Recoverable Items folder are retained until the hold duration expires. If the hold duration is unlimited, then items are retained until the hold is removed or the hold duration is changed.

Recover deleted messages in a user's mailbox

8/3/2020 • 6 minutes to read • Edit Online

Administrators can search for items that are purged (hard-deleted) by a user by using the Recover Deleted Items feature in Outlook or Outlook on the web. They can also search for items deleted by an automated process, such as the retention policy assigned to user mailboxes. In these situations, the purged items can't be recovered by a user. But administrators can recover purged messages if the deleted item retention period for the item hasn't expired.

> **NOTE**
>
> In addition to using this procedure to search for and recover deleted items, you can also use this procedure to search for items residing in other folders in the mailbox and to delete items from the source mailbox (also known as *search and purge*).

What you need to know before you begin?

- Procedures in this topic require specific permissions. See each procedure for its permissions information.
- Single item recovery should be enabled for a mailbox before the item you want to recover is deleted. In Exchange Server, single item recovery is disabled when a mailbox is created. For more information, see Enable or disable single item recovery for a mailbox.
- To search for and recover items, you need the following information:
  - **Source mailbox**: The mailbox being searched.
  - **Target mailbox**: The discovery mailbox in which messages will be recovered. Exchange Server Setup creates a default discovery mailbox. In Exchange Online, a discovery mailbox is also created by default. If required, you can create additional discovery mailboxes. For details, see Create a Discovery Mailbox.

    > **NOTE**
    >
    > When using the **Search-Mailbox** cmdlet, you can also specify a target mailbox that isn't a discovery mailbox. However, you can't specify the same mailbox as the source and target mailbox.

  - **Search criteria**: Criteria include sender or recipient, or keywords (words or phrases) in the message.

Step 1: Search for and recover missing items

You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "In-Place eDiscovery" entry in the Messaging policy and compliance permissions in Exchange Server topic.

NOTE

You can use In-Place eDiscovery in the Exchange admin center (EAC) to search for missing items. However, when using the EAC, you can't restrict the search to the Recoverable Items folder. All messages matching your search parameters will be returned even if they're not deleted. After they're recovered to the specified discovery mailbox, you may need to review the search results and remove unnecessary messages before recovering the remaining messages to the user's mailbox or exporting them to a .pst file. For details about how to use the EAC to perform an In-Place eDiscovery search, see Create an In-Place eDiscovery search in Exchange Server.

The first step in the recovery process is to search for messages in the source mailbox. Use one of the following methods to search a user mailbox and copy messages to a discovery mailbox.

Use the Exchange Management Shell to search for messages

This example searches for messages in April Stewart's mailbox that meet the following criteria:

- Sender: Ken Kwok
- Keyword: Seattle

```
Search-Mailbox "April Stewart" -SearchQuery "from:'Ken Kwok' AND seattle" -TargetMailbox "Discovery Search
Mailbox" -TargetFolder "April Stewart Recovery" -LogLevel Full
```

NOTE

When using the **Search-Mailbox** cmdlet, you can scope the search by using the *SearchQuery* parameter to specify a query formatted using Keyword Query Language (KQL). You can also use the *SearchDumpsterOnly* switch to search only items in the Recoverable Items folder.

For detailed syntax and parameter information, see Search-Mailbox.

How do you know this worked?

To verify that you have successfully searched the messages you want to recover, log on to the discovery mailbox you selected as the target mailbox and review the search results.

Step 2: Restore recovered items

You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "In-Place eDiscovery" entry in the Messaging policy and compliance permissions in Exchange Server topic.

NOTE

You can't use the EAC to restore recovered items.

After messages have been recovered to a discovery mailbox, you can restore them to the user's mailbox by using the **Search-Mailbox** cmdlet. In Exchange Server, you can also use the **New-MailboxExportRequest** and **New-MailboxImportRequest** cmdlets to export the messages to or import the messages from a .pst file.

Use the Exchange Management Shell to restore messages

This example restores messages to April Stewart's mailbox and deletes them from the Discovery Search Mailbox.

```
Search-Mailbox "Discovery Search Mailbox" -SearchQuery "from:'Ken Kwok' AND seattle" -TargetMailbox "April
Stewart" -TargetFolder "Recovered Messages" -LogLevel Full -DeleteContent
```

For detailed syntax and parameter information, see Search-Mailbox.

How do you know this worked?

To verify that you have successfully recovered messages to the user's mailbox, have the user review messages in the target folder you specified in the above command.

Use the Exchange Management Shell to export and import messages from a .pst file

In Exchange Server, you can export contents from a mailbox to a .pst file and import the contents of a .pst file to a mailbox. To learn more about mailbox import and export, see Mailbox imports and exports in Exchange Server. You can't perform this task in Exchange Online.

This example uses the following settings to export messages from the folder April Stewart Recovery in the Discovery Search Mailbox to a .pst file:

- **Mailbox**: Discovery Search Mailbox
- **Source folder**: April Stewart Recovery
- **ContentFilter**: April travel plans
- **PST file path**: \MYSERVER\HelpDeskPst\AprilStewartRecovery.pst

```
New-MailboxExportRequest -Mailbox "Discovery Search Mailbox" -SourceRootFolder "April Stewart Recovery" -
ContentFilter "Subject -eq 'April travel plans'" -FilePath \\MYSERVER\HelpDeskPst\AprilStewartRecovery.pst
```

For detailed syntax and parameter information, see New-MailboxExportRequest.

This example uses the following settings to import messages from a .pst file to the folder Recovered By Helpdesk in April Stewart's mailbox:

- **Mailbox**: April Stewart
- **Target folder**: Recovered By Helpdesk
- **PST file path**: \MYSERVER\HelpDeskPst\AprilStewartRecovery.pst

```
New-MailboxImportRequest -Mailbox "April Stewart" -TargetRootFolder "Recovered By Helpdesk" -FilePath
\\MYSERVER\HelpDeskPst\AprilStewartRecovery.pst
```

For detailed syntax and parameter information, see New-MailboxImportRequest.

How do you know this worked?

To verify that you have successfully exported messages to a .pst file, use Outlook to open the .pst file and inspect its contents. To verify that you have successfully imported messages from the .pst file, have the user inspect the contents of the target folder you specified in the above command.

More information

- The ability to recover deleted items is enabled by *single item recovery*, which lets an administrator recover a message that's been purged by a user or by retention policy as long as the deleted item retention period hasn't expired for that item. To learn more about single item recovery, see Recoverable Items folder in Exchange Server.

- In Exchange Server, a mailbox database is configured to retain deleted items for 14 days, by default. You can configure deleted item retention settings for a mailbox or mailbox database. For more information, see:
  - Configure Deleted Item retention and Recoverable Items quotas
- Users can recover a deleted item if it hasn't been purged and if the deleted item retention period for that item hasn't expired. If users need to recover deleted items from the Recoverable Items folder, point them to the following topics:
  - Recover deleted items in Outlook for Windows
  - Recover deleted items or email in Outlook on the web
- This topic shows you how to use the **Search-Mailbox** cmdlet to search for and recover missing items. If you use this cmdlet, you can search only one mailbox at a time. If you want to search multiple mailboxes at the same time, you can use In-Place eDiscovery in Exchange Server in the Exchange admin center (EAC) or the New-ComplianceSearch cmdlet in Windows PowerShell.
- In addition to using this procedure to search for and recover deleted items, you can also use a similar procedure to search for items in user mailboxes and then delete those items from the source mailbox. For more information, see Search for and delete messages in Exchange Server.

Manage linked mailboxes

8/3/2020 • 25 minutes to read • Edit Online

Linked mailboxes may be necessary for organizations that deploy Exchange in a *resource forest*. The resource forest scenario lets an organization centralize Exchange in a single forest, while allowing access to the Exchange organization with user accounts that are located in one or more trusted forests (called *account forests*). The user account that accesses the linked mailbox doesn't exist in the forest where Exchange is deployed. Therefore, a disabled user account that exists in the same forest as Exchange is created and associated with the corresponding linked mailbox.

The following figure illustrates the relationship between the linked user account used to access the linked mailbox (located in the account forest) and the disabled user account in the Exchange resource forest that's associated with the linked mailbox.

Linked mailboxes

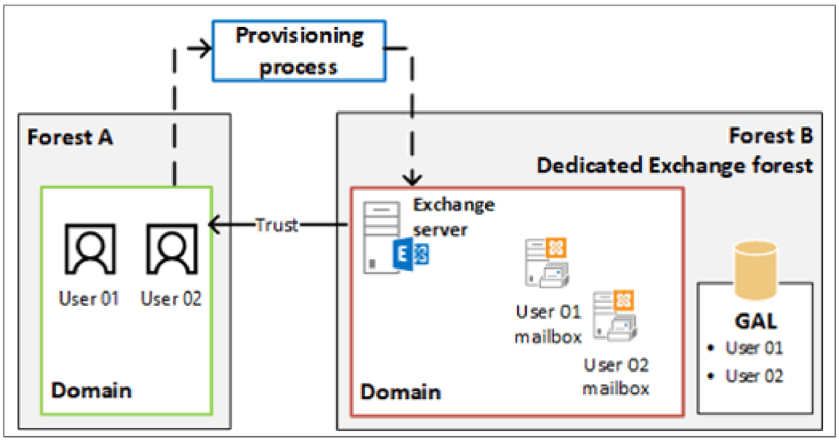

> **NOTE**
>
> A trust between the Exchange forest and at least one account forest must be set up before you can create linked mailboxes. At a minimum, you must set up a one-way, outgoing trust so that the Exchange forest trusts the account forest. For more information, see Learn more about setting up a forest trust to support linked mailboxes.

What do you need to know before you begin?

- Estimated time to complete: 2 to 5 minutes.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Recipient Provisioning Permissions" section in the Recipients Permissions topic.
- A user account (called the *linked master account*) must exist in the account forest before you can create a linked mailbox. This is because the linked mailbox is associated with a user in the account forest.
- If you've configured a one-way outgoing trust where the Exchange forest trusts the account forest, you'll need administrator credentials in the account forest to create a linked mailbox.

  To create a linked mailbox without being prompted for administrator credentials in the account forest, you have to create a two-way trust, or create another one-way outgoing trust where the account forest also trusts the Exchange forest. This step also requires administrator credentials in the account forest.

- You can complete this procedure in the Exchange admin center (EAC) or use the Exchange Management Shell.

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Create a linked mailbox

Use the EAC to create a linked mailbox

1. In the EAC, navigate to **Recipients** > **Mailboxes**.

2. Click **New** > **Linked mailbox**.

3. On the **New linked mailbox** page, in the **Trusted forest or domain** box, select the name of the account forest that contains the user account that you're creating the linked mailbox for. Click **Next**.

4. If your organization has configured a one-way outgoing trust where the Exchange forest trusts the account forest, you're prompted for administrator credentials in the account forest so that you can gain access to a domain controller in the trusted forest. Type the username and password for an administrator account in the account forest, and then click **Next**.

   > **NOTE**
   >
   > You won't be prompted for administrator credentials if you've created a two-way trust or have created another one-way outgoing trust where the account forest trusts the Exchange forest.

5. Complete the following boxes on the **Select linked master account** page.

   - **Linked domain controller**: Select a domain controller in the account forest. Exchange will connect to this domain controller to retrieve the list of user accounts in the account forest so that you can select the linked master account.

   - **Linked master account**: Click **Browse**, select a user account in the account forest, and then click **OK**. The new linked mailbox will be associated with this account.

6. Click **Next** and complete the following boxes on the **Enter general information** page.

   - *** Name**: Use this box to type a name for the user. This is the name used as the display name in the EAC and your organization's address book, and the name that's listed in Active Directory. This name is required.

   - **Organizational unit**: You can select an organizational unit (OU) other than the default (which is the recipient scope). If the recipient scope is set to the forest, the default value is set to the Users container in the Active Directory domain that contains the computer on which the EAC is running. If the recipient scope is set to a specific domain, the Users container in that domain is selected by default. If the recipient scope is set to a specific OU, that OU is selected by default.

     To select a different OU, click **Browse**. The dialog box displays all OUs in the Exchange forest that are within the specified scope. Select the OU you want, and then click **OK**.

   - *** User logon name**: Use this box to type the user logon name, which is required to create a linked mailbox. Type the username here. This name will be used in the left portion of the email address for the linked mailbox if you don't specify an alias.

> **NOTE**
>
> Because the user account that is created in the Exchange forest is disabled when you create a linked mailbox, the user doesn't use the user logon name to sign in to the linked mailbox. They sign in using their credentials from the account forest.

7. Click **More options** to configure the following boxes. Otherwise, skip to Step 8 to save the new linked mailbox.

   - **Alias**: Type the alias, which specifies the email alias for the linked mailbox. The user's alias is the portion of the email address on the left side of the at (@) symbol. It must be unique in the forest.

     > **NOTE**
     >
     > If you leave this box blank, the value from the username portion of the **User Logon Name** is used for the email alias.

   - **First name**, **Initials**, **Last name**

   - **Mailbox database**: Use this option to specify a mailbox database instead of allowing Exchange to choose a database for you. Click **Browse** to open the **Select Mailbox Database** dialog box. This dialog box lists all the mailbox databases in your Exchange organization. By default, the mailbox databases are sorted by name. You can also click the title of the corresponding column to sort the databases by server name or version. Select the mailbox database you want to use, and then click **OK**.

   - **Address book policy**: Use this option to specify an address book policy (ABP) for the linked mailbox. An ABP contains a global address list (GAL), an offline address book (OAB), a room list, and a set of address lists. When assigned to users, an ABP provides them with access to a customized GAL in Outlook and Outlook on the web (formerly known as Outlook Web App). To learn more, see Address book policies in Exchange Server.

     In the drop-down list, select the policy that you want associated with this mailbox.

8. When you're finished, click **Save** to create the new linked mailbox.

Use the Exchange Management Shell to create a linked mailbox

This example creates a linked mailbox for Ayla Kol in the CONTOSO Exchange resource forest. The FABRIKAM domain is in the account forest. The administrator account FABRIKAM \administrator is used to access the linked domain controller.

```
New-Mailbox -Name "Ayla Kol" -LinkedDomainController "DC1_FABRIKAM" -LinkedMasterAccount " FABRIKAM\aylak" -
OrganizationalUnit Users -UserPrincipalName aylak@contoso.com -LinkedCredential:(Get-Credential
FABRIKAM\administrator)
```

For syntax and parameter information, see New-Mailbox.

How do you know this worked?

To verify that you've successfully created a linked mailbox, do one of the following:

- In the EAC, navigate to **Recipients** > **Mailboxes**. The new linked mailbox is displayed in the mailbox list. Under **Mailbox Type**, the type is **Linked**.

- In the Exchange Management Shell, run the following command to display information about the new linked mailbox.

```
Get-Mailbox <Name> | Format-List Name,RecipientTypeDetails,IsLinked,LinkedMasterAccount
```

Change linked mailbox properties

After you create a linked mailbox, you can make changes and set additional properties by using the EAC or the Exchange Management Shell.

You can also change properties for multiple linked mailboxes at the same time. For more information, see Bulk edit user mailboxes.

> **IMPORTANT**
>
> The estimated time to complete this task will vary based on the number of properties you want to view or change.

Use the EAC to change linked mailbox properties

1. In the EAC, navigate to **Recipients** > **Mailboxes**.

2. In the list of mailboxes, click the linked mailbox that you want to change the properties for, and then click **Edit** ✎.

3. On the mailbox properties page, click one of the following sections to view or change properties.

   - General
   - Mailbox Usage
   - Email Address
   - Mailbox Features
   - Member Of
   - MailTip

General

Use the **General** section to view or change basic information about the user.

- *** Linked mailbox name**: This is the name that's listed in Active Directory. If you change this name, it can't exceed 64 characters.

- *** Display name**: This name appears in your organization's address book, on the To: and From: lines in email, and in the Mailboxes list in the EAC. This name can't contain empty spaces before or after the display name.

- *** User logon name**: For user mailboxes, this is the name that the user uses to sign in to their mailbox and to log on to the domain. For linked mailboxes, the corresponding user account that is created in the Exchange forest when the linked mailbox was created is disabled. The user uses their credentials from the account forest to sign in to the linked mailbox.

  If you change this name, it must be unique in your organization.

- **Linked master account**: This read-only box displays the user (in the format domain\username format) from the account forest that is associated with the linked mailbox. To change the linked master account associated with the linked mailbox, you have to use the **Set-Mailbox** cmdlet in the Exchange Management Shell. If you change the linked master account, the user will have to use the credentials for the new linked master account to sign in to the linked mailbox. For the command syntax to change the linked master account, see Use the Exchange management Shell to change linked mailbox properties.

- **Hide from address lists**: Select this check box to prevent the linked mailbox from appearing in the address book and other address lists that are defined in your Exchange organization. After you select this check box, users can still send messages to this user by using the email address.

Click **More options** to view or change these additional properties:

- **Organizational unit**: This read-only box displays the organizational unit (OU) that contains the user account. You have to use Active Directory Users and Computers to move the user account to a different OU.
- **Mailbox database**: This read-only box displays the name of the mailbox database that hosts the mailbox. To move the mailbox to a different database, select it in the mailbox list, and then click **Move mailbox to a different database** in the Details pane.
- *** Alias** This specifies the email alias for the linked mailbox. The alias is the portion of the email address on the left side of the at (@) symbol. It must be unique in the forest.
- **First name**, **Initials**, **Last name**
- **Custom attributes**: This section displays the custom attributes defined for the linked mailbox. To specify custom attribute values, click **Edit** ✎. You can specify up to 15 custom attributes for the recipient.

Mailbox Usage

Use the **Mailbox Usage** section to view or change the mailbox storage quota and deleted item retention settings for the linked mailbox. These settings are configured by default when the linked mailbox is created. They use the values that are configured for the mailbox database and apply to all mailboxes in that database. You can customize these settings for each mailbox instead of using the mailbox database defaults.

- **Last logon**: This read-only box displays the last time that the user signed in to the mailbox.
- **Mailbox usage**: This area shows the total size of the mailbox and the percentage of the total mailbox quota that has been used.

> **NOTE**
>
> To obtain the information that's displayed in the previous two boxes, the EAC queries the mailbox database that hosts the mailbox. If the EAC can't communicate with the Exchange store that contains the mailbox database, these boxes will be blank. A warning message is displayed if the user hasn't signed in to the mailbox for the first time.

Click **More options** to view or change the mailbox storage quota and the deleted item retention settings for the mailbox.

- **Storage quota settings**: To customize these settings for the mailbox and not use the mailbox database defaults, click **Customize settings for this mailbox**, type a new value, and then click **Save**.

  The value range for any of the storage quota settings is from 0 through 2047 gigabytes (GB).

  - **Issue a warning at (GB)**: This box displays the maximum storage limit before a warning is issued to the user. If the mailbox size reaches or exceeds the value specified, Exchange sends a warning message to the user.
  - **Prohibit send at (GB)**: This box displays the *prohibit send* limit for the mailbox. If the mailbox size reaches or exceeds the specified limit, Exchange prevents the user from sending new messages and displays a descriptive error message.
  - **Prohibit send and receive at (GB)**: This box displays the *prohibit send and receive* limit for the mailbox. If the mailbox size reaches or exceeds the specified limit, Exchange prevents the mailbox user

from sending new messages and won't deliver any new messages to the mailbox. Any messages sent to the mailbox are returned to the sender with a descriptive error message.

- **Deleted item retention settings**: To customize these settings for the mailbox and not use the mailbox database defaults, click **Customize settings for this mailbox**, type a new value, and then click **Save**.
  - **Keep deleted items for (days)**: This box displays the length of time that deleted items are retained before they're permanently deleted and can't be recovered by the user. When the mailbox is created, this length of time is based on the deleted item retention settings configured for the mailbox database. By default, a mailbox database is configured to retain deleted items for 14 days. The value range for this property is from 0 through 24855 days.
  - **Don't permanently delete items until the database is backed up**: Select this check box to prevent mailboxes and email messages from being deleted until after the mailbox database on which the mailbox is located has been backed up.

Email Address

Use the **Email address** section to view or change the email addresses associated with the linked mailbox. This includes the user's primary SMTP addresses and any associated proxy addresses. The primary SMTP address (also known as the *default reply address*) is displayed in bold text in the address list, with the uppercase **SMTP** value in the **Type** column.

- **Add**: Click **Add** + to add a new email address for this mailbox. Select one of following address types:
  - **SMTP**: This is the default address type. Click this radio button and then type the new SMTP address in the *** Email address** box.
  - **EUM**: An EUM (Exchange Unified Messaging) address is used by the Exchange Unified Messaging service in Exchange 2016 to locate UM-enabled users within an Exchange organization. EUM addresses consist of the extension number and the UM dial plan for the UM-enabled user. Click this radio button and type the extension number in the **Address/Extension** box. Then click **Browse** and select a dial plan for the user. (**Note**: Unified Messaging is not available in Exchange 2019.)
  - **Custom address type**: Click this button and type one of the supported non-SMTP email address types in the *** Email address** box.

    > **NOTE**
    >
    > With the exception of X.400 addresses, Exchange doesn't validate custom addresses for proper formatting. You must make sure that the custom address you specify complies with the format requirements for that address type.

- **Automatically update email addresses based on the email address policy applied to this recipient**: Select this check box if you want the recipient's email addresses to be updated automatically when changes are made to email address policies in your organization. This box is selected by default.

Mailbox Features

Use the **Mailbox Features** section to view or change the following mailbox features and settings:

- **Sharing policy**: This box shows the sharing policy applied to the mailbox. A sharing policy controls how users in your organization can share calendar and contact information with users outside your Exchange organization. The Default Sharing Policy is assigned to mailboxes when they are created. To change the sharing policy that's assigned to the user, select a different one from the drop-down list.
- **Role assignment policy**: This box shows the role assignment policy assigned to the mailbox. The role assignment policy specifies the role-based access control (RBAC) roles that are assigned to the user and controls which mailbox and distribution group configuration settings users can modify. To change the role

assignment policy that's assigned to the user, select a different one from the drop-down list.

- **Retention policy**: This box shows the retention policy assigned to the mailbox. A retention policy is a group of retention tags that are applied to the user's mailbox. The tags allow you to control how long to keep items in users' mailboxes and define which action to take on items that have reached a certain age. A retention policy isn't assigned to mailboxes when they are created. To assign a retention policy to the user, select one from the drop-down list.

- **Address Book policy**: This box shows the address book policy applied to the mailbox. An address book policy allows you to segment users into specific groups to provide customized views of the address book. To apply or change the address book policy that's applied to the mailbox, select one from the drop-down list.

- **Unified Messaging**: This feature is disabled by default. When you enable Unified Messaging (UM) in Exchange 2016, the user will be able to use your organization's UM features and a default set of UM properties are applied to the user. Click **Enable** to enable UM for the mailbox. For information about how to enable UM, see Enable a User for Unified Messaging. (**Note**: Unified Messaging is not available in Exchange 2019.)

  > **NOTE**
  >
  > A UM dial plan and a UM mailbox policy must exist before you can enable UM.

- **Mobile Devices**: Use this section to view and change the settings for Exchange ActiveSync, which is enabled by default. Exchange ActiveSync enables access to an Exchange mailbox from a mobile device. Click **Disable Exchange ActiveSync** to disable this feature for the mailbox.

- **Outlook Web App**: This feature is enabled by default. Outlook on the web provides access to an Exchange mailbox via a web browser. Click **Disable** to disable Outlook on the web for the mailbox. Click **Edit details** to add or change an Outlook on the web mailbox policy for the mailbox.

- **IMAP**: This feature is enabled by default. Click **Disable** to disable IMAP for the mailbox.

- **POP3**: This feature is enabled by default. Click **Disable** to disable POP3 for the mailbox.

- **MAPI**: This feature is enabled by default. MAPI enables access to an Exchange mailbox from a MAPI client such as Outlook. Click **Disable** to disable MAPI for the mailbox.

- **Litigation hold**: This feature is disabled by default. Litigation hold preserves deleted mailbox items and records changes made to mailbox items. Deleted items and all instances of changed items are returned in a discovery search. Click **Enable** to put the mailbox on litigation hold. If the mailbox is on litigation hold, click **Disable** to remove the litigation hold. If the mailbox is on litigation hold, click **Edit details** to view and change the following litigation hold settings:

  - **Hold date**: This read-only box indicates date and time when the mailbox was put on litigation hold.

  - **Put on hold by**: This read-only box indicates the user who put the mailbox on litigation hold.

  - **Note**: Use this box to notify the user about the litigation hold, explain why the mailbox is on litigation hold, or provide additional guidance to the user, such as informing them that the litigation hold won't affect their day-to-day use of email.

  - **URL**: Use this box to provide a URL to a website that provides information or guidance about the litigation hold on the mailbox.

> **NOTE**
>
> The text from these boxes appears in the user's mailbox only if they're using Outlook 2010 or later versions. It doesn't appear in Outlook on the web or other email clients. To view the text from the Note and URL boxes in Outlook, click the **File** tab and, on the **Info** page, under **Account Settings**, you'll see the litigation hold comment.

- **Archiving**: If an archive mailbox doesn't exist for the user, this feature is disabled. To enable an archive mailbox, click **Enable**. If the user has an archive mailbox, the size of the archive mailbox and usage statistics are displayed. Click **Edit details** to view and change the following archive mailbox settings:
  - **Status**: This read-only box indicates whether an archive mailbox exists.
  - **Database**: This read-only box shows the name of the mailbox database that hosts the archive mailbox.
  - **Name**: Type the name of the archive mailbox in this box. This name is displayed under the folder list in Outlook or Outlook on the web.
  - **Quota usage**: This read-only area shows the total size of the archive mailbox and the percentage of the total archive mailbox quota that has been used.
  - **Quota value (GB)**: This box shows the total size of the archive mailbox. To change the size, type a new value in the box or select a value from the drop-down list.
  - **Issue warning at (GB)**: This box shows the maximum storage limit for the archive mailbox before a warning is issued to the user. If the archive mailbox size reaches or exceeds the value specified, Exchange sends a warning message to the user. To change this limit, type a new value in the box or select a value from the drop-down list.
- **Delivery Options**: Use Delivery Options to forward email messages sent to the user to another recipient and to set the maximum number of recipients that the user can send a message to. Click **Edit details** to view and change these settings.
  - **Forwarding address**: Select the **Enable forwarding** check box and then click **Browse** to display the **Select Mail User and Mailbox** page. Use this page to select a recipient to whom you want to forward all email messages that are sent to this mailbox. Messages will be delivered to both the linked mailbox and the forwarding address.
  - **Recipient limit**: This setting controls the maximum number of recipients the user can send a message to. Select the **Maximum recipients** check box to limit the number of recipients allowed on the To:, Cc:, and Bcc: lines of an email message, and then specify the maximum number of recipients.

    > **NOTE**
    >
    > For on-premises Exchange organizations, the recipient limit is unlimited. For Exchange Online organizations, the limit is 500 recipients.

- **Message Size Restrictions**: These settings control the size of messages that the user can send and receive. Click **Edit details** to view and change the maximum size for sent and received messages.
  - **Sent messages**: To specify a maximum size for messages sent by this user, select the **Maximum message size (KB)** check box and type a value in the box. The message size must be between 0 and 2,097,151 KB. If the user sends a message larger than the specified size, the message will be returned to the user with a descriptive error message.

  - **Received messages**: To specify a maximum size for messages received by this user, select the **Maximum message size (KB)** check box and type a value in the box. The message size must be between 0 and 2,097,151 KB. If the user receives a message larger than the specified size, the message will be returned to the sender with a descriptive error message.

- **Message Delivery Restrictions**: These settings control who can send email messages to this user. Click **Edit details** to view and change these restrictions.

  - **Accept messages from**: Use this section to specify who can send messages to this user.

  - **All senders**: Select this option to specify that the user can accept messages from all senders. This includes both senders in your Exchange organization and external senders. This option is selected by default. This option includes external users only if you clear the **Require that all senders are authenticated** check box. If you select this check box, messages from external users will be rejected.

  - **Only senders in the following list**: Select this option to specify that the user can accept messages only from a specified set of senders in your Exchange organization. Click **Add** to display the **Select Recipients** page, which displays a list of all recipients in your Exchange organization. Select the recipients you want, add them to the list, and then click **OK**. You can also search for a specific recipient by typing the recipient's name in the search box and then clicking **Search**.

  - **Require that all senders are authenticated**: Select this option to prevent anonymous users from sending messages to the user.

  - **Reject messages from**: Use this section to block people from sending messages to this user.

  - **No senders**: Select this option to specify that the mailbox won't reject messages from any senders in the Exchange organization. This option is selected by default.

  - **Senders in the following list**: Select this option to specify that the mailbox will reject messages from a specified set of senders in your Exchange organization. Click **Add** to display the **Select Recipient** page, which displays a list of all recipients in your Exchange organization. Select the recipients you want to reject messages from, add them to the list, and then click **OK**. You can also search for a specific recipient by typing the recipient's name in the search box and then clicking **Search**.

Member Of

Use the **Member Of** section to view a list of the distribution groups or security groups to which this user belongs. You can't change membership information on this page. Note that the user may match the criteria for one or more dynamic distribution groups in your organization. However, dynamic distribution groups aren't displayed on this page because their membership is calculated each time they're used.

MailTip

Use the **MailTip** section to add a MailTip to alert users of potential issues if they send a message to this recipient. A MailTip is text that's displayed in the InfoBar when a recipient is added to the To, Cc, or Bcc lines of a new email message.

> **NOTE**
>
> MailTips can include HTML tags, but scripts aren't allowed. The length of a custom MailTip can't exceed 175 displayed characters. HTML tags aren't counted in the limit.

Mailbox Delegation

Use the **Mailbox Delegation** section to assign permissions to other users (also called *delegates*) to allow them to sign in to the user's mailbox or send messages on behalf of the user. You can assign the following permissions:

- **Send As**: This permission allows users other than the mailbox owner to use the mailbox to send messages.

After this permission is assigned to a delegate, any message that a delegate sends from this mailbox will appear as if it was sent by the mailbox owner. However, this permission doesn't allow a delegate to sign in to the user's mailbox.

- **Send on Behalf Of**: This permission also allows a delegate to use this mailbox to send messages. However, after this permission is assigned to a delegate, the **From:** address in any message sent by the delegate indicates that the message was sent by the delegate on behalf of the mailbox owner.

- **Full Access**: This permission allows a delegate to sign in to the user's mailbox and view the contents of the mailbox. However, after this permission is assigned to a delegate, the delegate can't send messages from the mailbox. To allow a delegate to send email from the user's mailbox, you still have to assign the delegate the Send As or the Send on Behalf Of permission.

To assign permissions to delegates, click **Add** under the appropriate permission to display the **Select Recipient** page, which displays a list of all recipients in your Exchange organization that can be assigned the permission. Select the recipients you want assign delegate permissions to, add them to the list, and then click **OK**. You can also search for a specific recipient by typing the recipient's name in the search box and then clicking **Search**.

Use the Exchange management Shell to change linked mailbox properties

Use the **Get-Mailbox** and **Set-Mailbox** cmdlets to view and change properties for linked mailboxes. One advantage of using the Exchange Management Shell is the ability to change the properties for multiple linked mailboxes. For information about what parameters correspond to mailbox properties, see the following topics:

- Get-Mailbox

- Set-Mailbox

Here are some examples of using the Exchange Management Shell to change linked mailbox properties.

This example uses the **Get-Mailbox** command to find all the linked mailboxes in the organization.

```
Get-Mailbox -ResultSize unlimited -Filter "RecipientTypeDetails -eq 'LinkedMailbox'"
```

This example uses the **Set-Mailbox** command to limit the number of recipients allowed on the To:, Cc:, and Bcc: lines of an email message to 500. This limit applies to all linked mailboxes in the organization.

```
Get-Mailbox -ResultSize unlimited -Filter "RecipientTypeDetails -eq 'LinkedMailbox'" | Set-Mailbox -
RecipientLimits 500
```

This example changes the linked master account in the fabrikam.com account forest that is associated with a linked mailbox in an Exchange forest.

```
Set-Mailbox -Identity "Ayla Kol" -LinkedDomainController DC1.fabrikam.com -LinkedMasterAccount
"fabrikam\robinw" -LinkedCredential:(Get-Credential fabrikam\administrator)
```

How do you know this worked?

To verify that you have successfully changed properties for a linked mailbox, do the following:

- In the EAC, select the linked mailbox and then click **Edit** to view the property or feature that you changed. Depending on the property that you changed, it might be displayed in the Details pane for the selected mailbox.

- In the Exchange Management Shell, use the **Get-Mailbox** cmdlet to verify the changes. One advantage of using the Exchange Management Shell is that you can view multiple properties for multiple linked mailboxes. In the example above where the recipient limit was changed, running the following command will verify the new value.

```
Get-Mailbox -ResultSize unlimited -Filter "RecipientTypeDetails -eq 'LinkedMailbox'" | Format-List
Name,RecipientLimits
```

For the example above where the linked master account was changed, run the following command to verify the new value.

```
Get-Mailbox "Ayla Kol" | Format-List LinkedMasterAccount
```

Manage distribution groups

8/3/2020 • 15 minutes to read • Edit Online

Use the Exchange admin center (EAC) or the Exchange Management Shell to create a new distribution group in your Exchange organization or to mail-enable an existing group in Active Directory.

There are two types of groups that can be used to distribute messages:

- *Mail-enabled universal distribution groups* (also called *distribution groups*) can be used only to distribute messages.

- *Mail-enabled universal security groups* (also called *security groups*) can be used to distribute messages as well as to grant access permissions to resources in Active Directory. For more information, see Manage mail-enabled security groups in Exchange Server.

It's important to note the terminology differences between Active Directory and Exchange. In Active Directory, a distribution group refers to any group that doesn't have a security context, whether it's mail-enabled or not. In contrast, in Exchange, all mail-enabled groups are referred to as distribution groups, whether they have a security context or not.

NOTE

You can create or mail-enable only universal distribution groups. To convert a domain-local or a global group to a universal group, you can use the Set-Group cmdlet using the Exchange Management Shell. You may have mail-enabled groups that were migrated from previous versions of Exchange that are not universal groups. You can use the EAC or the Exchange Management Shell to manage these groups

What do you need to know before you begin?

- Estimated time to complete: 2 to 5 minutes.

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Distribution groups" entry in the Recipients Permissions topic.

- If your organization has configured a group naming policy, it's applied only to groups created by users. When you or other administrators use the EAC to create distribution groups, the group naming policy is ignored and isn't applied to the group name. However, if you use the Exchange Management Shell to create or rename a distribution group, the policy is applied unless you use the *IgnoreNamingPolicy* parameter to override the group naming policy. For more information, see:

  - Create a Distribution Group Naming Policy

  - Override the Distribution Group Naming Policy

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Create a distribution group

Use the EAC to create a distribution group

1. In the EAC, navigate to **Recipients** > **Groups**.

2. Click **New** + > **Distribution group**.

3. On the **New distribution group** page, complete the following boxes:

   - *** Display name**: Use this box to type the display name. This name appears in your organization's address book, on the To: line when email is sent to this group, and in the Groups list in the EAC. The display name is required and should be user-friendly so people recognize what it is. It also must be unique in the forest.

   - *** Alias**: Use this box to type the name of the alias for the group. The alias can't exceed 64 characters and must be unique in the forest. When a user types the alias in the To: line of an email message, it resolves to the group's display name.

   - **Description**: Use this box to describe the group so people know what the purpose of the group is. This description appears in the address book.

   - **Organizational unit**: You can select an organizational unit (OU) other than the default (which is the recipient scope). If the recipient scope is set to the forest, the default value is set to the Users container in the Active Directory domain that contains the computer on which the EAC is running. If the recipient scope is set to a specific domain, the Users container in that domain is selected by default. If the recipient scope is set to a specific OU, that OU is selected by default.

     To select a different OU, click **Browse**. The dialog box displays all OUs in the forest that are within the specified scope. Select the OU you want, and then click **OK**.

   - *** Owners**: By default, the person who creates a group is the owner. All groups must have at least one owner. You can add owners by clicking **Add** +.

   - **Members**: Use this section to add members and to specify whether approval is required for people to join or leave the group.

     Group owners don't have to be members of the group. Use **Add group owners as members** to add or remove the owners as members.

     To add members to the group, click **Add** +. When you've finished adding members, click **OK** to return to the **New distribution group** page.

     Under **Choose whether owner approval is required to join the group**, specify whether approval is required for people to join the group. Select one of the following settings:

     - **Open: Anyone can join this group without being approved by the group owners**: This is the default setting.

     - **Closed: Members can be added only by the group owners. All requests to join will be rejected automatically**

   - **Owner Approval: All requests are manually approved or rejected by the group owners**: If you select this option, the group owner or owners will receive an email message requesting approval to join the group.

     Under **Choose whether the group is open to leave**, specify whether approval is required for people to leave the group. Select one of the following settings:

   - **Open: Anyone can leave this group without being approved by the group owners**: This is the default setting.

  - **Closed**: Members can be removed only by the group owners. All requests to leave will be rejected automatically

4. When you've finished, click **Save** to create the distribution group.

> **NOTE**
>
> By default, new distribution groups require that all senders be authenticated. This prevents external senders from sending messages to distribution groups. To configure a distribution group to accept messages from all senders, you must modify the message delivery restriction settings for that distribution group.

Use the Exchange Management Shell to create a distribution group

This example creates a distribution group with an alias **itadmin** and the name **IT Administrators**. The distribution group is created in the default OU, and anyone can join this group without approval by the group owners.

```
New-DistributionGroup -Name "IT Administrators" -Alias itadmin -MemberJoinRestriction open
```

For more information about using the Exchange Management Shell to create distribution groups, see New-DistributionGroup.

How do you know this worked?

To verify that you've successfully created a distribution group, do one of the following:

- In the EAC, navigate to **Recipients** > **Groups**. The new distribution group is displayed in the group list. Under **Group Type**, the type is **Distribution group**.

- In the Exchange Management Shell, run the following command to display information about the new distribution group.

  ```
  Get-DistributionGroup <Name> | Format-List Name,RecipientTypeDetails,PrimarySmtpAddress
  ```

Change distribution group properties

Use the EAC to change distribution group properties

1. In the EAC, navigate to **Recipients** > **Groups**.

2. In the list of groups, click the distribution group that you want to view or change, and then click **Edit** ✎.

3. On the group properties page, click one of the following sections to view or change properties.

   - General
   - Ownership
   - Membership
   - Membership approval
   - Delivery management
   - Message approval
   - Email options

- MailTip
- Group delegation

General

Use this section to view or change basic information about the group.

- *** Display name**: This name appears in the address book, on the To: line when email is sent to this group, and in the Groups list. The display name is required and should be user-friendly so people recognize what it is. It also has to be unique in your domain.

  If you've implemented a group naming policy, the display name has to conform to the naming format defined by the policy.

- *** Alias**: This is the portion of the email address that appears to the left of the at (@) symbol. If you change the alias, the primary SMTP address for the group will also be changed, and contain the new alias. Also, the email address with the previous alias will be kept as a proxy address for the group.
- **Description**: Use this box to describe the group so people know what the purpose of the group is. This description appears in the address book and in the Details pane in the EAC.
- **Hide this group from address lists**: Select this check box if you don't want users to see this group in the address book. To send email to this group, a sender has to type the group's alias or email address on the To: or Cc: lines.

  > **TIP**
  >
  > Consider hiding security groups because they're typically used to assign permissions to group members and not to send email.

- **Organizational unit**: This read-only box displays the organizational unit (OU) that contains the distribution group. You have to use Active Directory Users and Computers to move the group to a different OU.

Ownership

Use this section to assign group owners. The group owner can add members to the group, approve or reject requests to join or leave the group, and approve or reject messages sent to the group. By default, the person who creates a group is the owner. All groups must have at least one owner.

You can add owners by clicking **Add** +. You can remove an owner by selecting the owner and then clicking **Remove** −.

Membership

Use this section to add or remove members. Group owners don't have to be members of the group. Under **Members**, you can add members by clicking **Add** +. You can remove a member by selecting a user in the member list and then clicking **Remove** −.

Membership approval

Use this section to specify whether approval is required for users to join or leave the group.

- **Choose whether owner approval is required to join the group**: Select one of the following settings:
  - **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**: If you select this option, the group owner or owners receive an email requesting approval to join the group.

- **Choose whether the group is open to leave**: Select one of the following settings:
  - **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**

Delivery management

Use this section to manage who can send email to this group.

- **Only senders inside my organization**: Select this option to allow only senders in your organization to send messages to the group. This means that if someone outside of your organization sends an email message to this group, it will be rejected. This is the default setting.
- **Senders inside and outside of my organization**: Select this option to allow anyone to send messages to the group.

  You can further limit who can send messages to the group by allowing only specific senders to send messages to this group. Click **Add** + and then select one or more recipients. If you add senders to this list, they are the only ones who can send mail to the group. Mail sent by anyone not in the list will be rejected.

  To remove a person or a group from the list, select them in the list and then click **Remove** −.

> **IMPORTANT**
>
> If you've configured the group to allow only senders inside your organization to send messages to the group, email sent from a mail contact will be rejected, even if they are added to this list.

Message approval

Use this section to set options for moderating the group. Moderators approve or reject messages sent to the group before they reach the group members.

- **Messages sent to this group have to be approved by a moderator**: This check box isn't selected by default. If you select this check box, incoming messages are reviewed by the group moderators before delivery. Group moderators can approve or reject incoming messages.
- **Group moderators**: To add group moderators, click **Add** +. To remove a moderator, select the moderator, and then click **Remove** −. If you've selected "Messages sent to this group have to be approved by a moderator" and you don't select a moderator, messages to the group are sent to the group owners for approval.
- **Senders who don't require message approval**: To add people or groups that can bypass moderation for this group, click **Add** +. To remove a person or a group, select the item, and then click **Remove** −.
- **Select moderation notifications**: Use this section to set how users are notified about message approval.
  - **Notify all senders when their messages aren't approved**: This is the default setting. Notify all senders, inside and outside your organization, when their message isn't approved.
  - **Notify senders in your organization when their messages aren't approved**: When you select this option, only people or groups in your organization are notified when a message that they sent to the group isn't approved by a moderator.
  - **Don't notify anyone when a message isn't approved**: When you select this option, notifications aren't sent to message senders whose messages aren't approved by the group moderators.

Email options

Use this section to view or change the email addresses associated with the group. This includes the group's primary SMTP addresses and any associated proxy addresses. The primary SMTP address (also known as the *reply* *address*) is displayed in bold text in the address list, with the uppercase **SMTP** value in the **Type** column.

- **Add**: Click **Add** + to add a new email address for this mailbox. Select one of following address types:
  - **SMTP**: This is the default address type. Click this button and then type the new SMTP address in the *** Email address** box.
  - **Custom address type**: Click this button and type one of the supported non-SMTP email address types in the *** Email address** box.

    > **NOTE**
    >
    > With the exception of X.400 addresses, Exchange doesn't validate custom addresses for correct formatting. You must make sure that the custom address you specify complies with the format requirements for that address type.

- **Edit**: To change an email address associated with the group, select it in the list, and then click **Edit** ✎.
- **Remove**: To delete an email address associated with the group, select it in the list, and then click **Remove** −.
- **Automatically update email addresses based on the email address policy applied to this recipient**: Select this check box to have the recipient's email addresses automatically updated based on changes made to email address policies in your organization. This box is selected by default.

MailTip

Use this section to add a MailTip to alert users of potential issues if they send a message to this group. A MailTip is text that's displayed in the InfoBar when this group is added to the To, Cc, or Bcc lines of a new email message. For example, you could add a MailTip to large groups to warn potential senders that their message will be sent to lots of people.

> **NOTE**
>
> MailTips can include HTML tags, but scripts aren't allowed. The length of a custom MailTip can't exceed 175 displayed characters. HTML tags aren't counted in the limit.

Group delegation

Use this section to assign permissions to a user (called a *delegate*) to allow them to send messages as the group or send messages on behalf of the group. You can assign the following permissions:

- **Send As**: This permission allows the delegate to send messages as the group. After this permission is assigned, the delegate has the option to add the group to the **From** line to indicate that the message was sent by the group.
- **Send on Behalf Of**: This permission also allows a delegate to send messages on behalf of the group. After this permission is assigned, the delegate has the option to add the group in the **From** line. The message will appear to be sent by the group and will say that it was sent by the delegate on behalf of the group.

To assign permissions to delegates, click **Add** under the appropriate permission to display the **Select Recipient** page, which displays a list of all recipients in your Exchange organization that can be assigned the permission. Select the recipients you want, add them to the list, and then click **OK**. You can also search for a specific recipient by typing the recipient's name in the search box and then clicking **Search**.

Use the Exchange Management Shell to change distribution group properties

Use the **Get-DistributionGroup** and **Set-DistributionGroup** cmdlets to view and change properties for distribution groups. Advantages of using the Exchange Management Shell are the ability to change the properties that aren't available in the EAC and to change properties for multiple groups. For information about which parameters correspond to distribution group properties, see the following topics:

- Get-DistributionGroup
- Set-DistributionGroup

Here are some examples of using the Exchange Management Shell to change distribution group properties.

This example changes the primary SMTP address (also called the reply address) for the Seattle Employees distribution group from employees@contoso.com to sea.employees@contoso.com. Also, the previous reply address will be kept as a proxy address.

```
Set-DistributionGroup "Seattle Employees" -EmailAddresses
SMTP:sea.employees@contoso.com,smtp:employees@contoso.com
```

This example limits the maximum message size that can be sent to all distribution groups in the organization to 10 megabytes (MB).

```
Get-DistributionGroup -ResultSize unlimited -Filter "RecipientTypeDetails -eq
'MailUniversalDistributionGroup'" | Set-DistributionGroup -MaxReceiveSize 10MB
```

This example enables moderation for the distribution group Customer Support and sets the moderator to Amy. In addition, this moderated distribution group will notify senders who send mail from within the organization if their messages aren't approved.

```
Set-DistributionGroup -Identity "Customer Support" -ModeratedBy "Amy" -ModerationEnabled $true -
SendModerationNotifications 'Internal'
```

This example changes the user-created distribution group Dog Lovers to require the group manager to approve users' requests to join the group. In addition, by using the *BypassSecurityGroupManagerCheck* parameter, the group manager will not be notified that a change was made to the distribution group's settings.

```
Set-DistributionGroup -Identity "Dog Lovers" -MemberJoinRestriction 'ApprovalRequired' -
BypassSecurityGroupManagerCheck
```

How do you know this worked?

To verify that you've successfully changed properties for a distribution group, do the following:

- In the EAC, select the group and then click **Edit** ✎ to view the property or feature that you changed. Depending on the property that you changed, it might be displayed in the Details pane for the selected group.

- In the Exchange Management Shell, use the **Get-DistributionGroup** cmdlet to verify the changes. One advantage of using the Exchange Management Shell is that you can view multiple properties for multiple groups. In the example above where the recipient limit was changed, run the following command to verify the new value.

  ```
  Get-Mailbox -ResultSize unlimited -Filter "RecipientTypeDetails -eq 'UserMailbox'" | Format-List
  Name,RecipientLimits
  ```

  For the example above where the message limits were changed, run this command.

```
Get-Mailbox -OrganizationalUnit "Marketing" | Format-List
Name,IssueWarningQuota,ProhibitSendQuota,ProhibitSendReceiveQuota,UseDatabaseQuotaDefaults
```

Manage mail-enabled security groups in Exchange Server

8/3/2020 • 23 minutes to read • Edit Online

You can use mail-enabled security groups to distribute messages as well as grant access permissions to resources in Exchange and Active Directory. You can create, modify, and remove mail-enabled security groups in the Exchange admin center (EAC) or in the Exchange Management Shell. For more information about mail-enabled security groups, see Recipients.

What do you need to know before you begin?

- Estimated time to complete each procedure: 5 minutes.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Distribution groups" entry in the Recipients Permissions topic.
- For more information about accessing and using the EAC, see Exchange admin center in Exchange Server.
- To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- If you add users to or remove users from a mail-enabled security group, the users need to log out and log in for the permission changes to take effect.
- For mail-enabled security groups, users can't add or remove themselves from the group, nor can they send requests to the group owners to join or leave the group. A group owner needs to manually add and remove group members from a mail-enabled security group.
- If a mail-enabled security group contains members that aren't mail-enabled, a non-delivery report (also known as an NDR or bounce message) is returned for those non-mail-enabled members when you send a message to the group. To prevent NDRs, you can expand the group members in the **To** field of the message before you send the message (only the mail-enabled members of the group will appear).
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Create mail-enabled security groups

- When you create groups in the EAC, the value of the **Display name** property is used for the value of the unseen **Name** property (the unique identifier for the group object in the forest). Because the value of **Name** has a maximum length of 64 characters, the value of **Display name** also has a maximum length of 64 characters when you create groups in the EAC.
- When you create groups in the Exchange Management Shell, the *Name* parameter is required, the value must be unique, and the value has a maximum length of 64 characters. The *DisplayName* parameter is optional (the value of *Name* is used if you don't use it), the value isn't required to be unique, and the value has a maximum length of 256 characters.

- When you create groups in the EAC, the groups are automatically configured to only accept messages from authenticated (internal) senders. When you create groups in the Exchange Management Shell, you can use the *RequireSenderAuthenticationEnabled* parameter with the value `$false` so the group can accept messages from authenticated an unauthenticated (internal and external) senders. After you create the group, you can use the EAC or the Exchange Management Shell to change this setting.

Use the EAC to create a mail-enabled security group

1. In the EAC, go to **Recipients** > **Groups**.

2. Click **New** ➕ and then select **Security group** in the drop down list that appears.

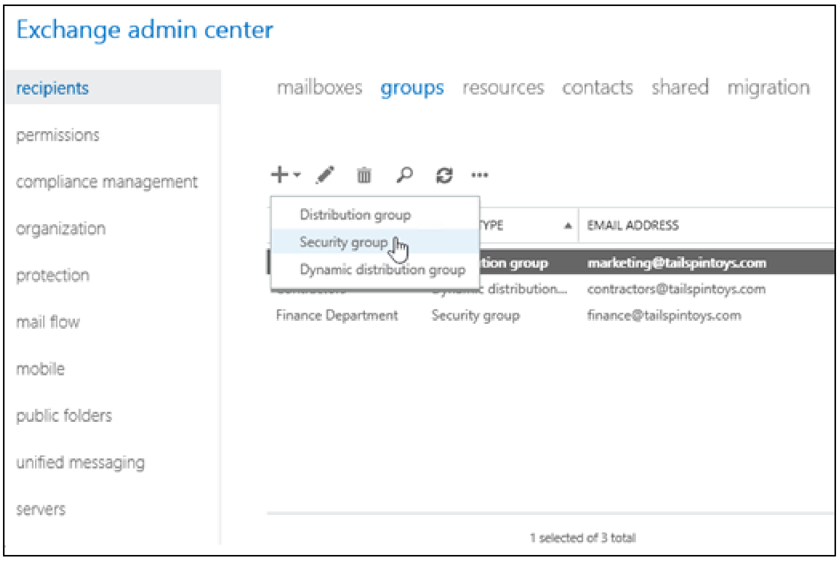

3. On the **New security group** page that opens, configure these settings (values marked with an * are required):

   - * **Display name**: This value should help users immediately recognize what the group is used for. This name appears in the global address list, on the To: line when email is sent to this group, and in the **Groups** list in the EAC. The maximum length in the EAC is 64 characters, and the value must be unique.

     > **NOTE**
     >
     > If a group naming policy is applied, you need to follow the naming constraints that are enforced for your organization. For more information, see Create a Distribution Group Naming Policy. If you want to override your organization's group naming policy, see Override a Distribution Group Naming Policy.

   - * **Alias**: This value is used to generate the primary email address (*<alias>*@ *<domain>*). This value can contain letters, numbers and the characters !, #, $, %, &, ', *, +, -, /, =, ?, ^, _, `, {, |, } and ~. Periods (.) are allowed, but each period must be surrounded by other valid characters (for example, help.desk). Unicode characters from U+00A1 to U+00FF are also allowed, but are mapped to best-fit US-ASCII text characters in the primary email address (for example, U+00F6 (ö) is changed to oe). The alias can't exceed 64 characters and must be unique in the forest. When a user types the alias on the To: line of an email message, it resolves to the group's display name.

   - **Notes**: Use this box to describe the purpose of the group. This description appears in the global address list and in the details pane in the EAC.

   - **Organizational unit**: The default location in Active Directory depends on the recipient scope that's configured:

     - If the recipient scope is the Active Directory forest, the default location is the Users container in the domain where the computer that's running the EAC is located.

  - If the recipient scope is a specific domain, the default location is the Users container in that domain.

  - If the recipient scope is a specific organizational unit (OU), the default location is that OU.

  To select a different OU, click **Browse**. The **Select an organizational unit** dialog box that opens shows all of the available OUs in the forest that are within the specified recipient scope. Select the desired OU, and then click **OK**.

- *** Owners**: By default, the person who creates the group is the owner. All groups must have at least one owner. Group owners can:

  - Modify the properties of the group

  - Add or remove group members

  - Delete the group

  - Approve messages sent to the group (if moderation is enabled)

  To add owners, click **Add** +. In the **Select Owners** dialog that appears, select one or more owners, click **Add**, and then click **OK**.

  To remove owners, select the owner in the list, and then click **Remove** −.

- **Members**

  To add members to the group, click **Add** +. In the **Select Members** dialog that appears, select one or more members, click **Add**, and then click **OK**.

  To remove members, select the member in the list, and then click **Remove** −.

- **Add group owners as members**: When this check box is selected, you don't need to manually include group owners in the list of members. If you don't want the group owners to be members of the group, clear this check box.

- **Owner approval is required**: For mail-enabled security groups, user requests to join the group aren't sent to the group owners, regardless of the state of this check box (selected or not selected). A group owner needs to manually add and remove group members from a mail-enabled security group.

When you've finished, click **Save**.

Use the Exchange Management Shell to create a mail-enabled security group

To create a mail-enabled security group, use this syntax:

```
New-DistributionGroup -Type Security -Name <UniqueName> [-IgnoreNamingPolicy] [-Alias <Alias>] [-DisplayName "
<DisplayName>"] [-Notes "<Description>"] [-OrganizationalUnit <OU>] [-ManagedBy "<owner1>","<owner2>"...] [-
Members "<member1>","<member2>"...] [-CopyOwnerToMember] [-MemberJoinRestriction <Closed | ApprovalRequired>]
[-RequireSenderAuthenticationEnabled <$true | $false>]
```

This example creates a security group with these settings:

- **Name**: File Server Managers. This value is also used for the display name because we aren't using the *DisplayName* parameter. If a group naming policy is applied, you can use the *IgnoreNamingPolicy* switch to override the policy.

- **Alias**: fsadmin. If we didn't use the *Alias* parameter, the value of the *Name* parameter would be used, with spaces removed (FileServerManagers in this example).

- **Description**: None, because we aren't using the *Notes* parameter.

- **Organizational Unit**: The default location that's specified by the recipient scope, because we aren't using the *OrganizationalUnit* parameter.

- **Owners**: The user account that's creating the group is the only owner, because we aren't using the *ManagedBy* parameter.

- **Members**: Bishamon Tamura and Valeria Barrios. Because we're using the *CopyOwnerToMember* switch, the group owner is also a member.

- **User requests to join the group**: For mail-enabled security groups, user requests to join the group aren't sent to the group owners, regardless of the *MemberJoinRestriction* parameter value (`ApprovalRequired` or `Closed`). A group owner needs to manually add and remove group members from a mail-enabled security group.

- **Accept messages from external senders**: No, because we're aren't using the *RequireSenderAuthenticationEnabled* parameter, and the default value is `$true`.

```
New-DistributionGroup -Type Security -Name "File Server Managers" -Alias fsadmin -Members "Bishamon
Tamura","Valeria Barrios" -CopyOwnerToMember
```

For detailed syntax and parameter information, see New-DistributionGroup.

How do you know this worked?

To verify that you've successfully created a mail-enabled security group, do any of these steps:

- In the EAC, go to **Recipients** > **Groups**. Verify that the group is listed, and the **Group Type** value is **Security group**.

- In the Exchange Management Shell, run this command and verify that the group is listed:

  ```
  Get-DistributionGroup -Filter "RecipientType -eq 'MailUniversalSecurityGroup'"
  ```

- In the Exchange Management Shell, replace *<GroupIdentity>* with the identity of the group (for example, name, alias, or email address), and run this command to verify the property values:

  ```
  Get-DistributionGroup -Identity <GroupIdentity> | Format-List
  ```

View or modify mail-enabled security groups

- When you modify groups in the EAC, the maximum length for the **Display name** property is now 256 characters, and the value doesn't need to be unique. This value no longer affects the value of the unseen **Name** property (the unique identifier for the group object in the forest). You can't use the EAC to modify the **Name** value of an existing group.

- When you modify groups in the Exchange Management Shell, the maximum length for the *Name* parameter value is still 64 characters, and the value must be unique. The maximum length for the *DisplayName* parameter value is still 256 characters, and the value doesn't need to be unique.

Use the EAC to view or modify a mail-enabled security group

1. In the EAC, go to **Recipients** > **Groups**.

2. In the list of groups, find the mail-enabled security group that you want to view or modify. You can:

   - Scroll through the list of groups.

   - Click **Search** 🔍 and enter part of the group's name, email address, or alias.

   - Click **More options** ••• > **Advanced search** to find the group.

   - Click the **Group Type** column header to sort the groups by **Security group**.

     Once you've found the mail-enabled security group that you want to modify, select it, and then click **Edit** ✎.

3. On the **Edit Security Group** page that opens, click one of the tabs to view or change the settings of the group:

   - General

   - Ownership

   - Membership

   - Membership approval

   - Delivery management

   - Message approval

   - Email options

   - MailTip

   - Group delegation

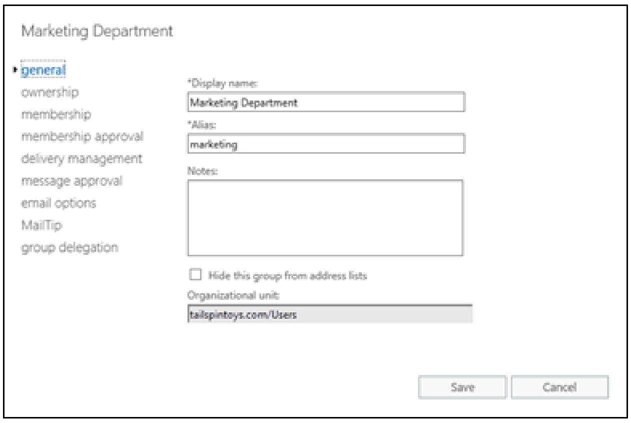

Marketing Department

*Display name:
Marketing Department

*Alias:
marketing

Notes:

□ Hide this group from address lists

Organizational unit:
tailspintoys.com/Users

   When you're finished, click **Save** or **Cancel**.

General

Use this tab to view or change basic information about the group.

- **Display name**: This value should help users immediately recognize what the group is used for. This name appears in the global address list, on the To: line when email is sent to this group, and in the **Groups** list in the EAC. The maximum length is 256 characters, and the value doesn't need to be unique.

- **Alias**: This value is used to generate the primary email address (*<alias>*@ *<domain>*). This value can contain letters, numbers and the characters !, #, $, %, &, ', *, +, -, /, =, ?, ^, _, `, {, |, } and ~. Periods (.) are allowed, but each period must be surrounded by other valid characters (for example, help.desk). Unicode characters from U+00A1 to U+00FF are also allowed, but are mapped to best-fit US-ASCII text characters in the primary email address (for example, U+00F6 (ö) is changed to oe). The alias can't exceed 64 characters and must be unique in the forest. When a user types the alias on the To: line of an email message, it resolves to the group's display name.

When you change the alias value, the previous primary email address is kept as a proxy address for the group.

- **Notes**: Use this box to describe the purpose of the group. This description appears in the global address list and in the details pane in the EAC.
- **Hide this group from address lists**: Select this check box if you don't want users to see the group in the global address list. If this check box is selected, a sender has to know and type the group's alias or email address to send messages to the group.
- **Organizational unit**: This read-only box displays the location of the group object in Active Directory. You need to use Active Directory Users and Computers to move the group to a different OU.

Ownership

Use this section to assign group owners. All groups must have at least one owner. Group owners can:

- Modify the properties of the group
- Add or remove group members
- Delete the group
- Approve member depart or join requests (if available)
- Approve messages sent to the group (if moderation is enabled)

To add owners, click **Add** +. In the **Select Owners** dialog that appears, select one or more owners, click **Add**, and then click **OK**.

To remove owners, select the owner in the list, and then click **Remove** −.

Membership

Use this tab to add or remove group members. Group owners aren't required to be members of the group.

To add members to the group, click **Add** +. In the **Select Members** dialog that opens, select one or more members, click **Add**, and then click **OK**.

To remove members, select the member in the list, and then click **Remove** −.

Membership approval

For mail-enabled security groups, user requests to join the group aren't sent to the group owners, regardless of the state of the **Owner approval is required** check box (selected or not selected). A group owner needs to manually add or remove group members from a mail-enabled security group.

Delivery management

Use this tab to control who's allowed to send messages to the group.

- **Only senders inside my organization**: The group only accepts messages from authenticated (internal) senders. This is the default setting.
- **Senders inside and outside of my organization**: The group accepts messages from authenticated and unauthenticated (internal and external) senders.
- Restrict the internal senders who can send messages to the group by clicking **Add** +. In the **Select Allowed Senders** dialog that appears, select one or more senders, click **Add**, and then click **OK**. Only the specified senders can send messages to the group.

  To remove internal senders that are allowed to send messages to the group, select the sender in the list, and then click **Remove** −.

IMPORTANT

Mail contacts are always considered unauthenticated (external) senders. If you select **Only senders inside my organization** and add the mail contact to the list of approved internal senders, messages sent to the group by the mail contact will be rejected.

Message approval

Use this tab to configure the moderation settings for messages that are sent to the group.

- **Messages sent to this group have to be approved by a moderator**: This check box isn't selected by default. If you select this check box, messages that are sent to the group must be approved by the specified moderators before they're delivered to the group members. When you select this option, you can configure these additional settings:

  - **Group moderators**: For mail-enabled security groups, the group owners aren't automatically used as moderators. You need to specify at least one moderator here when moderation is enabled.

    To add moderators, click **Add** +. In the **Select Group Moderators** dialog that appears, select one or more moderators (which can include any of the group owners), click **Add**, and then click **OK**.

    To remove moderators, select the moderator in the list, and then click **Remove** −.

  - **Senders who don't require message approval**: To configure senders who can bypass moderation for the group (send messages directly to the group members), click **Add** +. In the **Select Senders** dialog that appears, select one or more senders, click **Add**, and then click **OK**.

    To remove senders, select the sender in the list, and then click **Remove** −.

    You don't need to include moderators in the list of senders who bypass moderation. Messages that are sent to the group by a moderator aren't moderated.

  - **Select moderation notifications**: This setting configures how message senders are notified when their messages aren't approved by a moderator:

  - **Notify all senders when their messages aren't approved**: Authenticated and unauthenticated (internal and external) senders are notified when their messages aren't approved by a group moderator. This is the default value.

  - **Notify senders in your organization when their messages aren't approved**: Only authenticated (internal) senders are notified when their messages aren't approved by a group moderator.

  - **Don't notify anyone when a message isn't approved**: Senders aren't notified when their messages aren't approved by a group moderator.

Email options

Use this tab to view or change the email addresses that are configured for the group.

- **Email address**: By default, you use this setting to add additional email addresses for the group (also known as proxy addresses).

  By default, the primary email address (also known as the Reply To or reply address) is configured by the email address policy that's applied to the group. For more information about email address policies, see Email address policies in Exchange Server. The primary email address that's shown here is bold, and has the uppercase **SMTP** value in the **Type** column.

  To manually specify the group's primary email address here, you need to clear the check box **Automatically update email addresses based on the email address policy applied to this**

recipient. Note that clearing this check box prevents automatic updates to the email addresses of the group by email address policies.

  - To add a new email address for the group, click **Add** ＋. In the **New email address** page that opens, select one of these options:

  - **Email address type**: Select **SMTP**. In the **Email address** box, type the email address (for example, helpdesk@contoso.com). The domain must be an accepted domain that's configured for your organization. For more information, see Accepted domains in Exchange Server.

  - On the previous page, if you left **Automatically update email addresses based on the email address policy applied to this recipient** check box selected, the email address is added to the group as a proxy address (there's no **Make this the reply address** check box on this page).

  - On the previous page, if you cleared the **Automatically update email addresses based on the email address policy applied to this recipient** check box, you can select **Make this the reply address**. This setting adds the new email address as the primary email address, and the previous primary email address is kept as a proxy address. If you don't select **Make this the reply address**, the email address is added as a proxy address, and the primary email address is unaffected.

  - **Email address type**: Select **Enter a custom address type**. Type the custom email address type (for example, X400). In the **Email address** box, type the custom email address.

  Note: With the exception of X.400 addresses, Exchange doesn't validate custom email addresses for correct formatting. You need to make sure that the custom email address complies with the format requirements for that address type.

  When you're finished, click **OK**.

  - To modify an existing email address for the group, select it in the list, and then click **Edit** ✎. Note that you can't change the email address type, just the email address.

  - To remove an existing email address from the group, select it in the list, and then click **Remove** −. Note that you can't remove the primary email address.

MailTip

Use this tab to add a custom MailTip for the group. MailTips alert users to potential issues before they send a message to the group. For more information about MailTips, see Configure Custom MailTips for Recipients.

> **NOTE**
>
> MailTips can include HTML tags, but scripts aren't allowed. The length of a custom MailTip can't exceed 175 displayed characters. HTML tags aren't counted in the limit.

Group delegation

Use this tab to assign permissions to the group for a user (called a *delegate*).

- **Send As**: The specified users can send messages that appear to be sent by the group. The actual sender isn't revealed, and replies to these messages are delivered to the group.

- **Send on Behalf**: The specified users can send on behalf of the group. Although messages send on behalf of the group clearly show the sender in the From line (*<Sender>* on behalf of *<Group>*), replies to these messages are delivered to the group, not the sender.

To add delegates, click **Add** ＋ for the appropriate permission. In the resulting dialog that appears, select one or more delegates, click **Add**, and then click **OK**.

After you assign one of these permissions, the delegate can select the group for the **From** line in Outlook or

Outlook on the web (formerly known as Outlook Web App).

To remove delegates, select the delegate in the appropriate list, and then click **Remove** −.

Use the Exchange Management Shell to modify a mail-enabled security group

You use the **Set-DistributionGroup** cmdlet to modify mail-enabled security groups. Here are some interesting settings that you can configure using the **Set-DistributionGroup** cmdlet that aren't available in the EAC or on the **New-DistributionGroup** cmdlet:

- Configure values for the **CustomAttribute1** through **CustomAttribute15** properties (the *CustomAttribute1* through *CustomAttribute15* parameters).

- Configure MailTips in different languages (the *MailTipTranslations* parameter).

- Configure the maximum message size that can be sent to or sent from the group (the *MaxReceiveSize* and *MaxSendSize* parameters).

- Instead of specifying the internal recipients who *are* allowed to send messages to the group, you can specify the internal recipients who *aren't* allowed to send messages to the group (the *RejectMessagesFromSendersOrMembers* parameter).

For detailed syntax and parameter information, see Set-DistributionGroup.

This example configures the value DoNotMigrate for the **CustomAttribute5** property of the group named Experimental Project.

```
Set-DistributionGroup -Identity "Experimental Project" -CustomAttribute5 DoNotMigrate
```

This example adds the Spanish translation for the existing English MailTip, "Please allow 4 business days for a response to messages sent to this group" that's configured on the mail-enabled security group events@contoso.com.

```
Set-DistributionGroup -Identity events@contoso.com -MailTipTranslations @{Add="ES:Espere 4 días hábiles para
responder a los mensajes enviados a este grupo."}
```

How do you know this worked?

To verify that you've successfully modified a mail-enabled security group, do any of these steps:

- In the EAC, go to **Recipients** > **Groups** > select the mail-enabled security group (the **Group Type** value is **Security group**) > click **Edit** ✏ and verify the property values.

- In the Exchange Management Shell, replace *<GroupIdentity>* with the identity of the group (for example, name, alias, or email address), and run this command to verify the property values:

  ```
  Get-DistributionGroup -Identity <GroupIdentity> | Format-List
  ```

Use the Exchange Management Shell to view mail-enabled security groups

You use the **Get-DistributionGroup** cmdlet to view mail-enabled security groups.

This example returns a summary list of all security groups in the organization.

```
Get-DistributionGroup -ResultSize unlimited -Filter "RecipientTypeDetails -eq 'MailUniversalSecurityGroup'"
```

This example returns detailed information for the mail-enabled security group named Help Desk.

```
Get-DistributionGroup -Identity "Help Desk" | Format-List
```

For detailed syntax and parameter information, see Get-DistributionGroup.

Remove mail-enabled security groups

Use the EAC to remove a mail-enabled security group

1. In the EAC, go to **Recipients** > **Groups**.

2. In the list of groups, find the security group that you want to remove. You can:

   - Scroll through the list of groups.

   - Click **Search** 🔎 and enter part of the group's name, email address, or alias.

   - Click **More options** ••• > **Advanced search** to find the group.

   - Click the **Group Type** column header to sort the groups by **Security group**.

   Once you've found the security group that you want to remove, select it, click **Delete** 🗑, and then click **Yes** in the warning message that appears.

Use the Exchange Management Shell to remove a mail-enabled security group

To remove a mail-enabled security group, use this syntax:

```
Remove-DistributionGroup -Identity <GroupIdentity>
```

This example removes the mail-enabled security group that has the alias value contractors.

```
Remove-DistributionGroup -Identity contractors
```

How do you know this worked?

To verify that you've successfully removed a mail-enabled security group, do any of these steps:

- In the EAC, go to **Recipients** > **Groups**, and verify that the group isn't listed. Note that you might need to click **Refresh** 🔄.

- In the Exchange Management Shell, run this command and verify that the group isn't listed:

  ```
  Get-DistributionGroup -Filter "RecipientType -eq 'MailUniversalSecurityGroup'"
  ```

- In the Exchange Management Shell, replace *<GroupIdentity>* with the identity of the group (for example, name, alias, or email address), and run this command to verify that the group isn't returned:

  ```
  Get-DistributionGroup -Identity <GroupIdentity> | Format-List
  ```

- In the Exchange Management Shell, run this command and verify that the group is listed:

  ```
  Get-Group -Filter "RecipientTypeDetails -eq 'UniversalSecurityGroup'"
  ```

Mail-enable or mail-disable existing security groups

To mail-enable an existing universal security group that's not already mail-enabled, or to mail-disable an existing mail-enabled security group, you can't use the EAC. You can only use the Exchange Management Shell.

Use the Exchange Management Shell to mail-enable an existing security group

To mail-enable an existing universal security group, use this syntax:

```
Enable-DistributionGroup -Identity <GroupIdentity> [-Alias <Alias>] [-DisplayName <DisplayName>] [-
PrimarySMTPAddress <EmailAddress>]
```

This example mail-enables the existing universal security group named Help Desk with the following settings:

- **Alias**: hdesk. If we didn't use the *Alias* parameter, the value of the *Name* parameter would be used, with spaces removed (HelpDesk in this example).
- **Display name**: Because we aren't using the *DisplayName* parameter, the group's existing **Name** property value is used for the display name.
- **Primary email address**: Because we're using the *Alias* parameter, the group's primary email address is *<alias>*@ *<domain>*, where <domain> is specified by the email address policy that applies to the group. If we specified a value for the *PrimarySMTPAddress* parameter, the **EmailAddressPolicyEnabled** property would be set to the value `$false`, which means the email addresses of the group aren't automatically updated by email address policies.

```
Enable-DistributionGroup -Identity "Help Desk" -Alias hdesk
```

After you mail-enable the security group, the group will be visible to all other ***-DistributionGroup** cmdlets.

For detailed syntax and parameter information, see Enable-DistributionGroup.

How do you know this worked?

To verify that you've successfully mail-enabled an existing security group, do any of these steps:

- In the EAC, go to **Recipients** > **Groups**. Verify that the group is listed, and the **Group Type** value is **Security group**. Note that you might need to click **Refresh** ↻ if the EAC was already open.
- In the Exchange Management Shell, run this command and verify that the group is listed:

  ```
  Get-DistributionGroup -Filter "RecipientType -eq 'MailUniversalSecurityGroup'"
  ```

- In the Exchange Management Shell, replace *<GroupIdentity>* with the identity of the group (for example, name, alias, or email address), and run this command to verify the property values:

  ```
  Get-DistributionGroup -Identity <GroupIdentity> | Format-List
  ```

Use the Exchange Management Shell to mail-disable an existing mail-enabled security group

To mail-disable an existing mail-enabled universal security group, use this syntax:

```
Disable-DistributionGroup -Identity <GroupIdentity> [-IgnoreDefaultScope]
```

This example mail-disables the mail-enabled security group named Human Resources.

```
Disable-DistributionGroup -Identity "Human Resources"
```

Notes:

- If the distribution group isn't visible to you because of a restricted recipient scope, you'll need to use the *IgnoreDefaultScope* switch to see all groups in the Active Directory forest. But, when you use this switch, you'll need to identify the group by its distinguished name (DN). For example, `"CN=<Group Name>,CN=North America,CN=Users,DC=contoso,DC=com"`.
- After you mail-disable the security group, the group will be invisible to all ***-DistributionGroup** cmdlets except **Enable-DistributionGroup**.

For detailed syntax and parameter information, see Disable-DistributionGroup.

How do you know this worked?

To verify that you've successfully mail-disabled an existing mail-enabled universal security group, do any of these steps:

- In the EAC, go to **Recipients** > **Groups**, and verify that the group isn't listed. Note that you might need to click **Refresh** ⟳ if the EAC was already open.
- In the Exchange Management Shell, run this command and verify that the group isn't listed:

```
Get-DistributionGroup -Filter "RecipientType -eq 'MailUniversalSecurityGroup'"
```

- In the Exchange Management Shell, replace *<GroupIdentity>* with the name of the group, and run this command to verify that the group isn't returned:

```
Get-DistributionGroup -Identity <GroupIdentity> | Format-List
```

- In the Exchange Management Shell, run this command and verify that the group is listed:

```
Get-Group -Filter "RecipientTypeDetails -eq 'UniversalSecurityGroup'"
```

Manage dynamic distribution groups

8/3/2020 • 16 minutes to read • Edit Online

Dynamic distribution groups are mail-enabled Active Directory group objects that are created to expedite the mass sending of email messages and other information within a Microsoft Exchange organization.

Unlike regular distribution groups that contain a defined set of members, the membership list for dynamic distribution groups is calculated each time a message is sent to the group, based on the filters and conditions that you define. When an email message is sent to a dynamic distribution group, it's delivered to all recipients in the organization that match the criteria defined for that group.

> **IMPORTANT**
>
> A dynamic distribution group includes any recipient in Active Directory with attribute values that match its filter. If a recipient's properties are modified to match the filter, the recipient could inadvertently become a group member and start receiving messages that are sent to the group. Well-defined, consistent account provisioning processes will reduce the chances of this issue occurring.

What do you need to know before you begin?

- Estimated time to complete: 2 to 5 minutes.
- To open the EAC, see Exchange admin center in Exchange Server. To open the Exchange Management Shell, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Dynamic distribution groups" entry in the Recipients Permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Create a dynamic distribution group

Use the EAC to create a dynamic distribution group

1. In the EAC, navigate to **Recipients** > **Groups** > **New** > **Dynamic distribution group**.
2. On the **New dynamic distribution group** page, complete the following boxes:
   - *** Display name**: Use this box to type the display name. This name appears in the shared address book, on the To: line when email is sent to this group, and in the Groups list in the EAC. The display name is required and should be user-friendly so people recognize what it is. It also must be unique in the forest.

   > **NOTE**
   >
   > Group naming policy isn't applied to dynamic distribution groups.

- *** Alias**: Use this box to type the name of the alias for the group. The alias cannot exceed 64 characters and must be unique in the forest. When a user types the alias in the To: line of an email message, it resolves to the group's display name.

- **Description**: Use this box to describe the group so people know what the purpose of the group is. This description appears in the shared address book.

- **Organizational unit**: You can select an organizational unit (OU) other than the default (which is the recipient scope). If the recipient scope is set to the forest, the default value is set to the Users container in the Active Directory domain that contains the computer on which the EAC is running. If the recipient scope is set to a specific domain, the Users container in that domain is selected by default. If the recipient scope is set to a specific OU, that OU is selected by default.

  To select a different OU, click **Browse**. The dialog box displays all OUs in the forest that are within the specified scope. Select the OU you want, and then click **OK**.

- **Owner**: An owner for a dynamic distribution group is optional. You can add owners by clicking **Browse** and then selecting users from the list.

3. Use the **Members** section to specify the types of recipients for the group and set up rules that will determine membership. Select one of the following boxes:

   - **All recipient types**: Choose this option to send messages that meet the criteria defined for this group to all recipient types.

   - **Only the following recipient types**: Messages that meet the criteria defined for this group will be sent to one or more of the following recipient types:

   - **Users with Exchange mailboxes**: Select this check box if you want to include users that have Exchange mailboxes. Users that have Exchange mailboxes are those that have a user domain account and a mailbox in the Exchange organization.

   - **Users with external email addresses**: Select this check box if you want to include users that have external email addresses. Users that have external email accounts have user domain accounts in Active Directory, but use email accounts that are external to the organization. This enables them to be included in the global address list (GAL) and added to distribution lists.

   - **Resource mailboxes**: Select this check box if you want to include Exchange resource mailboxes. Resource mailboxes allow you to administer company resources through a mailbox, such as a conference room or a company vehicle.

   - **Contacts with external email addresses**: Select this check box if you want to include contacts that have external email addresses. Contacts that have external email addresses don't have user domain accounts in Active Directory, but the external email address is available in the GAL.

   - **Mail-enabled groups**: Select this check box if you want to include security groups or distribution groups that have been mail-enabled. Mail-enabled groups are similar to distribution groups. Email messages that are sent to a mail-enabled group account will be delivered to several recipients.

4. Click **Add a rule** to define the criteria for membership in this group.

5. Select one of the following recipient attributes from the drop-down list and provide a value. If the value for the selected attribute matches that value you define, the recipient receives a message sent to this group.

ATTRIBUTE	SEND MESSAGE TO A RECIPIENT IF...
Recipient container	The recipient object resides in the specified domain or OU.

ATTRIBUTE	SEND MESSAGE TO A RECIPIENT IF...
State or province	The specified value matches the recipient's State or province property.
Company	The specified value matches the recipient's Company property.
Department	The specified value matches the recipient's Department property.
Custom attributeN (where N is a number from 1 to 15)	The specified value matches the recipient's CustomAttributeN property.

> **IMPORTANT**
>
> The values that you enter for the selected attribute must exactly match those that appear in the recipient's properties. For example, if you enter **Washington** for **State or province**, but the value for the recipient's property is **WA**, the condition will not be met. Also, text-based values that you specify aren't case-sensitive. For example, if you specify **Contoso** for the **Company** attribute, messages will be sent to a recipient if this value is **contoso**.

6. In the **Specify words or phrases** window, type the value in the text box. Click **Add** and then click **OK**.

7. To add another rule to define the criteria for membership, click **Add a rule** under the previous rule that you created.

> **IMPORTANT**
>
> If you add multiple rules to define membership, a recipient must meet the criteria of each rule to receive a message sent to the group. In other words, each rule is connected with the Boolean operator **AND**.

8. When you've finished, click **Save** to create the dynamic distribution group.

> **NOTE**
>
> If you want to specify rules for attributes other than the ones available in the EAC, you must use the Exchange Management Shell to create a dynamic distribution group. Keep in mind that the filter and condition settings for dynamic distribution groups that have custom recipient filters can be managed only by using the Exchange Management Shell. For an example of how to create a dynamic distribution group with a custom query, see the next section on using the Exchange Management Shell to create a dynamic distribution group.

Use the Exchange Management Shell to create a dynamic distribution group

> **NOTE**
>
> If you do not specify an OU in your cmdlets, the default OU scope will be the local OU (the OU in which the dynamic distribution group is being created). With the `New-DynamicDistributionGroup` cmdlet, use the `RecipientContainer` parameter to specify an OU.

This example creates the dynamic distribution group "Mailbox Users DDG" that contains only mailbox users.

```
New-DynamicDistributionGroup -IncludedRecipients MailboxUsers -Name "Mailbox Users DDG" -RecipientContainer
Users
```

This example creates a dynamic distribution group with a custom recipient filter. The dynamic distribution group contains all mailbox users on a server called Server1.

```
New-DynamicDistributionGroup -Name "Mailbox Users on Server1" -RecipientContainer Users -RecipientFilter "
(RecipientTypeDetails -eq 'UserMailbox') -and (ServerName -eq 'Server1')"
```

This example creates a dynamic distribution group with a custom recipient filter. The dynamic distribution group contains all mailbox users that have a value of "FullTimeEmployee" in the **CustomAttribute10** property.

```
New-DynamicDistributionGroup -Name "Full Time Employees" -RecipientFilter "(RecipientTypeDetails -eq
'UserMailbox') -and (CustomAttribute10 -eq 'FullTimeEmployee')"
```

For detailed syntax and parameter information, see New-DynamicDistributionGroup.

How do you know this worked?

To verify that you've successfully created a dynamic distribution group, do one of the following:

- In the EAC, navigate to **Recipients** > **Groups**. The new dynamic distribution group is displayed in the group list. Under **Group Type**, the type is **Dynamic distribution group**.

- In the Exchange Management Shell, run the following command to display information about the new dynamic distribution group.

  ```
  Get-DynamicDistributionGroup | Format-List Name,RecipientTypeDetails,RecipientFilter,PrimarySmtpAddress
  ```

Change dynamic distribution group properties

Use the EAC to change dynamic distribution group properties

1. In the EAC, navigate to **Recipients** > **Groups**.

2. In the list of groups, click the dynamic distribution group that you want to view or change, and then click **Edit** ✏.

3. On the group's properties page, click one of the following sections to view or change properties.

   - General
   - Ownership
   - Membership
   - Delivery management
   - Message approval
   - Email options
   - MailTip
   - Group delegation

General

Use this section to view or change basic information about the group.

- *** Display name**: This name appears in the address book, on the To: line when email is sent to this group, and in the Groups list. The display name is required and should be user-friendly so people recognize what it is. It also has to be unique in your domain.

- *** Alias**: This is the portion of the email address that appears to the left of the at (@) symbol. If you change the alias, the primary SMTP address for the group will also be changed, and contain the new alias. Also, the email address with the previous alias will be kept as a proxy address for the group.

- **Description**: Use this box to describe the group so people know what the purpose of the group is. This description appears in the address book and in the Details pane in the EAC.

- **Hide this group from address lists**: Select this check box if you don't want users to see this group in the address book. To send email to this group, a sender has to type the group's alias or email address on the To: or Cc: lines.

- **Organizational unit**: This read-only box displays the organizational unit (OU) that contains the dynamic distribution group. You have to use Active Directory Users and Computers to move the group to a different OU.

Ownership

Use this section to assign a group owner. A dynamic distribution group can have only one owner. The group owner appears on the **Managed by** tab of the object in Active Directory Users and Computers.

You can add owners by clicking **Browse** and selecting the owner from the list. To remove the owner, click **Clear** (**X**) and then click **Save**.

Membership

Use this section to change the criteria used to determine membership of the group. You can delete or change existing membership rules and add new rules. For procedures that tell you how to do this, see Use the EAC to create a dynamic distribution group in the procedures for configuring membership when you use the EAC to create a new dynamic distribution group.

Delivery management

Use this section to manage who can send email to this group.

- **Only senders inside my organization**: Select this option to allow only senders in your organization to send messages to the group. This means that if someone outside your organization sends an email message to this group, it is rejected. This is the default setting.

- **Senders inside and outside of my organization**: Select this option to allow anyone to send messages to the group.

  You can further limit who can send messages to the group by allowing only specific senders to send messages to this group. Click **Add** + and then select one or more recipients. If you add senders to this list, they are the only ones who can send mail to the group. Mail sent by anyone not in the list will be rejected.

  To remove a person or a group from the list, select them in the list and then click **Remove** −.

> **IMPORTANT**
>
> If you've configured the group to allow only senders inside your organization to send messages to the group, email sent from a mail contact is rejected, even if they're added to this list.

Message approval

Use this section to set options for moderating the group. Moderators approve or reject messages sent to the group before they reach the group members.

- **Messages sent to this group have to be approved by a moderator**: This check box isn't selected by default. If you select this check box, incoming messages are reviewed by the group moderators before delivery. Group moderators can approve or reject incoming messages.

- **Group moderators**: To add group moderators, click **Add** +. To remove a moderator, select the moderator, and then click **Remove** −. If you've selected "Messages sent to this group have to be approved by a moderator" and you don't select a moderator, messages to the group are sent to the group owners for approval.
- **Senders who don't require message approval**: To add people or groups that can bypass moderation for this group, click **Add** +. To remove a person or a group, select the item, and then click **Remove** −.
- **Select moderation notifications**: Use this section to set how users are notified about message approval.
  - **Notify all senders when their messages aren't approved**: This is the default setting. Notify all senders, inside and outside your organization, when their message isn't approved.
  - **Notify senders in your organization only when their messages aren't approved**: When you select this option, only people or groups in your organization are notified when a message that they sent to the group isn't approved by a moderator.
  - **Don't notify anyone when a message isn't approved**: When you select this option, notifications aren't sent to message senders whose messages aren't approved by the group moderators.

Email options

Use this section to view or change the email addresses associated with the group. This includes the group's primary SMTP addresses and any associated proxy addresses. The primary SMTP address (also known as the *reply address*) is displayed in bold text in the address list, with the uppercase **SMTP** value in the **Type** column.

- **Add**: Click **Add** + to add a new email address for this mailbox. Select one of following address types:
  - **SMTP**: This is the default address type. Click this button and then type the new SMTP address in the *** Email address** box.

    > **NOTE**
    >
    > To make the new address the primary SMTP address for the group, select the **Make this the reply address** check box.

  - **Custom address type**: Click this button and type one of the supported non-SMTP email address types in the *** Email address** box.

    > **NOTE**
    >
    > With the exception of X.400 addresses, Exchange doesn't validate custom addresses for proper formatting. You must make sure that the custom address you specify complies with the format requirements for that address type.

- **Edit**: To change an email address associated with the group, select it from the list, and then click **Edit** ✏.

  > **NOTE**
  >
  > To make an existing address the primary SMTP address for the group, select the **Make this the reply address** check box.

- **Remove**: To delete an email address associated with the group, select it from the list, and then click

Remove −.

- **Automatically update email addresses based on the email address policy applied to this recipient**: Select this check box to have the recipient's email addresses automatically updated based on changes made to email address policies in your organization. This box is selected by default.

MailTip

Use this section to add a MailTip to alert users of potential issues before they send a message to this group. A MailTip is text that's displayed in the InfoBar when this group is added to the To, Cc, or Bcc lines of a new email message. For example, you could add a MailTip to large groups to warn potential senders that their message will be sent to lots of people.

> **NOTE**
>
> MailTips can include HTML tags, but scripts aren't allowed. The length of a custom MailTip can't exceed 175 displayed characters. HTML tags aren't counted in the limit.

Group delegation

Use this section to assign permissions to a user (called a *delegate*) to allow them to send messages as the group or send messages on behalf of the group. You can assign the following permissions:

- **Send As**: This permission allows the delegate to send messages as the group. After this permission is assigned, the delegate has the option to add the group to the **From** line to indicate that the message was sent by the group.

- **Send on Behalf Of**: This permission also allows a delegate to send messages on behalf of the group. After this permission is assigned, the delegate has the option to add the group on the **From** line. The message will appear to be sent by the group and will say that it was sent by the delegate on behalf of the group.

To assign permissions to delegates, click **Add** under the appropriate permission to display the **Select Recipient** page, which displays a list of all recipients in your Exchange organization that can be assigned the permission. Select the recipients you want, add them to the list, and then click **OK**. You can also search for a specific recipient by typing the recipient's name in the search box and then clicking **Search**.

Use the Exchange Management Shell to change dynamic distribution group properties

Use the **Get-DynamicDistributionGroup** and **Set-DynamicDistributionGroup** cmdlets to view and change properties for dynamic distribution groups. Advantages of using the Exchange Management Shell are the ability to change the properties that aren't available in the EAC and change properties for multiple groups. For information about what parameters correspond to distribution group properties, see the following topics:

- Get-DynamicDistributionGroup

- Set-DynamicDistributionGroup

Here are some examples of using the Exchange Management Shell to change dynamic distribution group properties.

This example changes the following parameters for all dynamic distribution groups in the organization:

- Hide all dynamic distribution groups from the address book

- Set the maximum message size that can be sent to the group to 5MB

- Enable moderation

- Assign the administrator as the group moderator

```
Get-DynamicDistributionGroup -ResultSize unlimited | Set-DynamicDistributionGroup -
HiddenFromAddressListsEnabled $true -MaxReceiveSize 5MB -ModerationEnabled $true -ModeratedBy administrator
```

This example adds the proxy SMTP email address, Seattle.Employees@contoso.com, to the All Employees group.

```
Set-DynamicDistributionGroup -Identity "All Employees" -EmailAddresses SMTP:All.Employees@contoso.com,
smtp:Seattle.Employees@contoso.com
```

How do you know this worked?

To verify that you've successfully changed properties for a dynamic distribution group, do the following:

- In the EAC, select the group and then click **Edit** ✏ to view the property or feature that you changed. Depending on the property that you changed, it might be displayed in the Details pane for the selected group.

- In the Exchange Management Shell, use the **Get-DynamicDistributionGroup** cmdlet to verify the changes. One advantage of using the Exchange Management Shell is that you can view multiple properties for multiple groups. In the first example, you would run the following command to verify the new values.

  ```
  Get-DynamicDistributionGroup -ResultSize unlimited | Format-List
  Name,HiddenFromAddressListsEnabled,MaxReceiveSize,ModerationEnabled,ModeratedBy
  ```

  For the example above where the message limits were changed, run this command.

  ```
  Get-Mailbox -OrganizationalUnit "Marketing" | Format-List
  Name,IssueWarningQuota,ProhibitSendQuota,ProhibitSendReceiveQuota,UseDatabaseQuotaDefaults
  ```

View members of a dynamic distribution group

8/3/2020 • 2 minutes to read • Edit Online

Dynamic distribution groups are distribution groups whose membership is based on specific recipient filters rather than a defined set of recipients. For more information, see Manage dynamic distribution groups.

You can't use the Exchange admin center (EAC) to view the members of a dynamic distribution group. You can only use the Exchange Management Shell.

What do you need to know before you begin?

- Estimated time to complete: 2 minutes.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Dynamic distribution groups" entry in the Recipients Permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to view the members of a dynamic distribution group

This example returns the list of members for the dynamic distribution group named Full Time Employees. The first command stores the dynamic distribution group object in the variable `$FTE`. The second command uses the **Get-Recipient** cmdlet to list the recipients that match the criteria defined for the dynamic distribution group.

```
$FTE = Get-DynamicDistributionGroup "Full Time Employees"
```

```
Get-Recipient -RecipientPreviewFilter $FTE.RecipientFilter -OrganizationalUnit $FTE.RecipientContainer
```

For detailed syntax and parameter information, see Get-DynamicDistributionGroup and Get-Recipient.

Manage mail contacts

8/3/2020 • 9 minutes to read • Edit Online

Mail contacts are essentially contacts for people outside your Exchange or organization. Each mail contact has an external email address. For more information about mail contacts, see Recipients.

What do you need to know before you begin?

- Estimated time to complete: 2 minutes.
- To open the EAC, see Exchange admin center in Exchange Server. To open the Exchange Management Shell, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Recipient Provisioning Permissions" section in the Recipients Permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Create a mail contact

Use the EAC to create a mail contact

1. In the EAC, navigate to **Recipients** > **Contacts**.
2. Click **New** ＋ > **Mail contact**.
3. Complete the following boxes on the **New mail contact** page:
   - **First name**: Use this box to type the contact's first name.
   - **Initials**: Use this box to type the contact's initials.
   - **Last name**: Use this box to type the contact's last name.
   - *** Display name**: Use this box to type a display name for the contact. This is the name that's listed in the contacts list in the EAC and in your organization's address book. By default, this box is populated with the names you enter in the **First name**, **Initials**, and **Last name** boxes. If you didn't use those boxes, you must still type a name in this box because it's required. The name can't exceed 64 characters.
   - *** Name**: Use this box to type a name for the contact. This is the name that's listed in the directory service. Like the display name, this box is populated by default with the names you enter in the **First name**, **Initials**, and **Last name** boxes. If you didn't use those boxes, you must still type a name in this box because it's required. The name can't exceed 64 characters.
   - *** Alias**: Use this box to type an alias (64 characters or less) for the contact. This box is required.

   - *** External email address**: Use this box to type the outside email account of the contact. This box is required. Email sent to this contact is forwarded to this email address.

   - **Organizational unit**: You can select an organizational unit (OU) other than the default, which is the recipient scope. If the recipient scope is set to the forest, the default value is set to the Users container in the domain that contains the computer on which the EAC is running. If the recipient scope is set to a specific domain, the Users container in that domain is selected by default. If the recipient scope is set to a specific OU, that OU is selected by default.

     To select a different OU, click **Browse**. The dialog box displays all OUs in the forest that are within the specified scope. Select the OU you want, and then click **OK**.

4. When you've finished, click **Save**.

Use the Exchange Management Shell to create a mail contact

This example creates a mail contact for Debra Garcia in Exchange Server 2016.

```
New-MailContact -Name "Debra Garcia" -ExternalEmailAddress dgarcia@tailspintoys.com -OrganizationalUnit Users
```

This example mail-enables an existing contact named Karen Toh in Exchange Server 2016.

```
Enable-MailContact -Identity "Karen Toh" -ExternalEmailAddress ktoh@tailspintoys.com
```

How do you know this worked?

To verify that you've successfully created a mail contact, do one of the following:

- In the EAC, navigate to **Recipients** > **Contacts**. The new mail contact is displayed in the contact list. Under **Contact Type**, the type is **Mail contact**.

- In the Exchange Management Shell, run the following command to display information about the new mail contact.

  ```
  Get-MailContact <Name> | Format-List Name,RecipientTypeDetails,ExternalEmailAddress
  ```

Change mail contact properties

Use the EAC to change mail contact properties

1. In the EAC, navigate to **Recipients** > **Contacts**.

2. In the list of mail contacts and mail users, click the mail contact that you want to change the properties for, and then click **Edit** ✎.

3. On the mail contact properties page, click one of the following sections to view or change properties.

   - General

   - Contact Information

   - Organization

   - Email Options

   - MailTip

General

Use the **General** section to view or change basic information about the mail contact.

- **First name**, **Initials**, **Last name**
- * **Name**: This is the name that's listed in Active Directory. If you change this name, it can't exceed 64 characters.
- * **Display name**: This name appears in your organization's address book, on the To and From lines in email, and in the Mailbox list. This name can't contain empty spaces before or after the display name.
- * **Alias**: This is the mail contact's alias. If you change it, it must be unique in the organization and must be 64 characters or less.
- * **External email address**: This is mail contact's primary SMTP address and their outside email account. Email sent to this contact is forwarded to this email address.
- Click **More options** to display the OU that contains the mail contact account. You have to use Active Directory Users and Computers to move the contact to a different OU.

Contact Information

Use the **Contact Information** section to view or change the recipient's contact information, such as mailing address and telephone numbers. This information is displayed in the address book.

Organization

Use the **Organization** section to record detailed information about the mail contact's role in the organization. This information is displayed in the address book. Also, you can create a virtual organization chart that's accessible from email clients such as Outlook.

- **Title**: Use this box to view or change the contact's title.
- **Department**: Use this box to view or change the department in which the contact works. You can use this box to create recipient conditions for dynamic distribution groups and address lists.
- **Company**: Use this box to view or change the company for which the contact works. You can also use this box to create recipient conditions for dynamic distribution groups.
- **Manager**: To add a manager, click **Browse**. In **Select Manager**, select a person, and then click **OK**.
- **Direct reports**: You can't modify this box. A *direct report* is a recipient who reports to a specific manager. If you've specified a manager for the recipient, that recipient appears as a direct report in the details of the manager's mailbox. For example, Toby manages Ann and Spencer, who are mail contacts, so Toby is specified in the **Manager** box in the organization properties for Ann and Spencer, and Ann and Spencer appear in the **Direct reports** box in the properties of Toby's mailbox.

Email Options

Use the **Email Options** section to add or remove proxy addresses for the mail contact or edit existing proxy addresses. The mail contact's primary SMTP address is also displayed in this section, but you can't change it. To change it, you have to change the contact's external email address in the **General** section.

MailTip

Use the **MailTip** section to add a MailTip to alert users of potential issues before they send a message to this recipient. A MailTip is text that's displayed in the InfoBar when this recipient is added to the To, Cc, or Bcc lines of a new email message.

> **NOTE**
>
> MailTips can include HTML tags, but scripts aren't allowed. The length of a custom MailTip can't exceed 175 displayed characters. HTML tags aren't counted in the limit.

Use the Exchange Management Shell to change mail contact properties

Properties for a mail contact are stored in both Active Directory and Exchange. In general, use the **Get-Contact** and **Set-Contact** cmdlets to view and change organization and contact information properties. Use the **Get-MailContact** and **Set-MailContact** cmdlets to view or change mail-related properties, such as email addresses, the MailTip, custom attributes, and whether the contact is hidden from address lists.

For more information, see the following topics:

- Get-Contact
- Set-Contact
- Get-MailContact
- Set-MailContact

Here are some examples of using the Exchange Management Shell to change mail contact properties.

This example configures the Title, Department, Company, and Manager properties for the mail contact Kai Axford.

```
Set-Contact "Kai Axford" -Title Consultant -Department "Public Relations" -Company Fabrikam -Manager "Karen Toh"
```

This example sets the CustomAttribute1 property to a value of PartTime for all mail contacts and hides them from the organization's address book.

```
Get-MailContact | Set-MailContact -CustomAttribute1 PartTime -HiddenFromAddressListsEnabled $true
```

This example sets the CustomAttribute15 property to a value of TemporaryEmployee for all mail contacts in the Public Relations department.

```
Get-Contact -Filter "Department -eq 'Public Relations'" | Set-MailContact -CustomAttribute15 TemporaryEmployee
```

How do you know this worked?

To verify that you've successfully changed properties for a mail contact, do the following:

- In the EAC, select the mail contact, and then click **Edit** ✎ to view the property that you changed.
- In the Exchange Management Shell, use the **Get-Contact** and **Get-MailContact** cmdlets to verify the changes. One advantage of using the Exchange Management Shell is that you can view multiple properties for multiple mail contacts. In the example above where all mail contacts had the CustomAttribute1 property set to PartTime and were hidden from the address book, run the following command to verify the changes.

  ```
  Get-MailContact | Format-List Name,CustomAttribute1,HiddenFromAddressListsEnabled
  ```

  In the example above where the CustomAttribute15 was set for all mail contacts in the Public Relations department, run the following command to verify the changes.

  ```
  Get-Contact -Filter "Department -eq 'Public Relations'" | Get-MailContact | Format-List Name,CustomAttribute15
  ```

Bulk edit mail contacts

You can use the EAC to change selected properties for multiple mail contacts. When you select two or more mail contacts from the contacts list in the EAC, the properties that can be bulk edited are displayed in the Details pane.

When you change one of these properties, the change is applied to all selected recipients.

When you bulk edit mail contacts, you can change the following property areas:

- **Contact Information**: Change shared properties such as street, postal code, and city name.
- **Organization**: Change shared properties such as department name, company name, and the manager that the selected mail contacts or mail users report to.

Use the EAC to bulk edit mail contacts

1. In the EAC, navigate to **Recipients** > **Contacts**.
2. In the list of contacts, select two or more mail contacts. You can't bulk edit a combination of mail contacts and mail users.

   > **TIP**
   >
   > You can select multiple adjacent mail contacts by holding down the Shift key and clicking the first mail contact, and then clicking the last mail contact you want to edit. You can also select multiple mail contacts by holding down the Ctrl key and clicking each one that you want to edit.

3. In the Details pane, under **Bulk Edit**, click **Update** under **Contact Information** or **Organization**.
4. Make the changes on the properties page and then save your changes.

How do you know this worked?

To verify that you've successfully bulk edited mail contacts, do one of the following:

- In the EAC, select each of the mail contacts that you bulk edited, and then click **Edit** ✎ to view the properties that you changed.
- In the Exchange Management Shell, use the **Get-Contact** cmdlet to verify the changes. For example, say you used the bulk edit feature in the EAC to change the manager and the office for all mail contacts from a vendor company named A. Datum Corporation. To verify these changes, you could run the following command in the Exchange Management Shell.

  ```
  Get-Contact -ResultSize unlimited -Filter "Company -eq 'Adatum'" | Format-List Name,Office,Manager
  ```

Manage mail users

8/3/2020 • 16 minutes to read • Edit Online

Mail users are similar to mail contacts. Both have external email addresses and both contain information about people outside your Exchange organization that can be displayed in the shared address book and other address lists. However, unlike a mail contact, a mail user has logon credentials in your Exchange organization and can access resources. For more information, see Recipients.

What do you need to know before you begin?

- Estimated time to complete: 2 minutes.
- To open the EAC, see Exchange admin center in Exchange Server. To open the Exchange Management Shell, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Recipient Provisioning Permissions" section in the Recipients Permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Create a mail user

Use the EAC to create a mail user

1. In the EAC, navigate to **Recipients** > **Contacts** > **New** > **Mail user**.
2. On the **New mail user** page, in the * **Alias** box, type the alias for the mail user. The alias can't exceed 64 characters and must be unique in the forest. This is required.
3. Do one of the following to specify the email address type for the mail user:
   - To specify an SMTP email address for the mail user's external email address, click **SMTP**.

   > **NOTE**
   >
   > Exchange validates SMTP addresses for correct formatting. If your entry is inconsistent with the SMTP format, an error message will be displayed when you click **Save** to create the mail user.

   - To specify a custom address type, click the option button and then type the custom address type. For example, you can specify an X.500, GroupWise, or Lotus Notes address.
4. In the * **External email address** box, type the mail user's external email address. Email sent to this mail user is forwarded to this email address. This is required.
5. Select one of the following options:
   - **Existing user**: Select to mail-enable an existing user.

  Click **Browse** to open the **Select User - Entire Forest** dialog box. This dialog box displays a list of user accounts in the organization that aren't mail-enabled or don't have mailboxes. Select the user account you want to mail-enable, and then click **OK**. If you select this option, you don't have to provide user account information because this information already exists in Active Directory.

- **New user**: Select to create a new user account in Active Directory and mail-enable the user. If you select this option, you'll have to provide the required user account information.

6. If you selected **New User** in Step 5, complete the following information on the **New mail user** page. Otherwise skip to Step 7.

   - **First name**: Type the first name of the mail user.

   - **Initials**: Type the initials of the mail user.

   - **Last name**: Type the last name of the mail user.

   - *** Display name**: Use this box to type a display name for the user. This is the name that's listed in the contacts list in the EAC and in your organization's address book. By default, this box is populated with the names you enter in the **First name**, **Initials**, and **Last name** boxes. If you didn't use those boxes, you must still type a name in this box because it's required. The name can't exceed 64 characters.

   - *** Name**: Use this box to type a name for the mail user. This is the name that's listed in the directory service. This box is also populated with the names you enter in the **First name**, **Initials**, and **Last name** boxes. If you didn't use those boxes, you must still type a name because this is required. This name also can't exceed 64 characters.

   - **Organizational unit**: You can select an organizational unit (OU) other than the default (which is the recipient scope). If the recipient scope is set to the forest, the default value is set to the Users container in the domain that contains the computer on which the EAC is running. If the recipient scope is set to a specific domain, the Users container in that domain is selected by default. If the recipient scope is set to a specific OU, that OU is selected by default.

     To select a different OU, click **Browse**. The dialog box displays all OUs in the forest that are within the specified scope. Select the OU you want, and then click **OK**.

   - *** User logon name**: Type the name that the mail user will use to log on to the domain. The user logon name consists of a username on the left side of the at (@) symbol and a suffix on the right side. Typically, the suffix is the domain name the user account resides in.

   - *** New Password**: Type the password that the mail user must use to log on to the domain. Make sure that the password you supply complies with the password length, complexity, and history requirements of the domain you're creating the user account in.

   - *** Confirm password**: Use this box to confirm the password that you typed in the **Password** box.

   - **Require password change on next logon**: Select this check box if you want mail users to reset the password when they first log on to the domain.

     If you select this check box, at first logon, the new mail user will be prompted with a dialog box in which to change the password. The mail user won't be allowed to perform any tasks until the password is changed successfully.

7. When you've finished, click **Save** to create the mail user.

Use the Exchange Management Shell to create a mail user

This example creates a mail-enabled user account for Jeffrey Zeng in Exchange Server 2016 with the following details:

- The name and display name is Jeffrey Zeng.
- The alias is jeffreyz.
- The external email address is jzeng@tailspintoys.com.
- The first name is Jeffrey and the last name is Zeng.
- The logon name is jeffreyz@contoso.com.
- The password is Pa$$word1.
- The mail user will be created in the default OU. To specify a different OU, you can use the *OrganizationalUnit* parameter.

```
New-MailUser -Name "Jeffrey Zeng" -Alias jeffreyz -ExternalEmailAddress jzeng@tailspintoys.com -FirstName
Jeffrey -LastName Zeng -UserPrincipalName jeffreyz@contoso.com -Password (ConvertTo-SecureString -String
'Pa$$word1' -AsPlainText -Force)
```

For detailed syntax and parameter information, see New-MailUser.

How do you know this worked?

To verify that you've successfully created a mail user, do one of the following:

- In the EAC, navigate to **Recipients** > **Contacts**. The new mail user is displayed in the list of contacts. Under **Contact Type**, the type is **Mail user**.
- In the Exchange Management Shell, run the following command to display information about the new mail user.

  ```
  Get-MailUser <Name> | Format-List Name,RecipientTypeDetails,ExternalEmailAddress
  ```

Change mail user properties

After you create a mail user, you can make changes and set additional properties by using the EAC or the Exchange Management Shell.

You can also change properties for multiple user mailboxes at the same time. For more information, see Use the EAC to bulk edit mail users.

The estimated time to complete this task will vary based on the number of properties you want to view or change.

Use the EAC to change user mailbox properties

1. In the EAC, navigate to **Recipients** > **Contacts**.
2. In the list of contacts, click the mail user that you want to change the properties for, and then click **Edit** ✎.
3. On the mail user properties page, click one of the following sections to view or change properties.
   - General
   - Contact Information
   - Organization
   - Email Addresses
   - Mail Flow Settings

  - Member Of

  - MailTip

General

Use the **General** section to view or change basic information about the mail user.

- **First name**, **Initials**, **Last name**

- * **Name**: This is the name that's listed in Active Directory. If you change this name, it can't exceed 64 characters.

- * **Display name**: This name appears in your organization's address book, on the To: and From: lines in email, and in the list of contacts in the EAC. This name can't contain empty spaces before or after the display name.

- * **User logon name**: This is the name that the user uses to log on to the domain.

- **Hide from address lists**: Select this check box to prevent the mail user from appearing in the address book and other address lists that are defined in your Exchange organization. After you select this check box, users can still send messages to the recipient by using the email address.

- **Require password change on next logon**: Select this check box if you want the user to reset their password the next time they log on to the domain.

Click **More options** to view or change these additional properties:

- **Organizational unit**: This read-only box displays the organizational unit (OU) that contains the mail user account. You have to use Active Directory Users and Computers to move the account to a different OU.

- **Custom attributes**: This section displays the custom attributes defined for the mail user. To specify custom attribute values, click **Edit** ✎. You can specify up to 15 custom attributes for the recipient.

Contact Information

Use the **Contact Information** section to view or change the user's contact information. The information on this page is displayed in the address book. Click **More options** to display additional boxes.

> **TIP**
>
> You can use the **State/Province** box to create recipient conditions for dynamic distribution groups, email address policies, or address lists.

Organization

Use the **Organization** section to record detailed information about the user's role in the organization. This information is displayed in the address book. Also, you can create a virtual organization chart that's accessible from email clients such as Outlook.

- **Title**: Use this box to view or change the recipient's title.

- **Department**: Use this box to view or change the department in which the user works. You can use this box to create recipient conditions for dynamic distribution groups, email address policies, or address lists.

- **Company**: Use this box to view or change the company for which the user works. You can use this box to create recipient conditions for dynamic distribution groups, email address policies, or address lists.

- **Manager**: To add a manager, click **Browse**. In **Select Manager**, select a person, and then click **OK**.

- **Direct reports**: You can't modify this box. A *direct report* is a user who reports to a specific manager. If you've specified a manager for the user, that user appears as a direct report in the details of the manager's mailbox. For example, Kari manages Chris and Kate, so Kari is specified in the **Manager** box for Chris and Kate, and Chris and Kate appear in the **Direct reports** box in the properties of Kari's account.

Email Addresses

Use the **Email Addresses** section to view or change the email addresses associated with the mail user. This includes the mail user's primary SMTP address, their external email address, and any associated proxy addresses. The primary SMTP address (also known as the *default reply address*) is displayed in bold text in the address list, with the uppercase **SMTP** value in the **Type** column. By default, after the mail user is created, the primary SMTP address and the external email address are the same.

- **Add**: Click **Add** ＋ to add a new email address for this mailbox. Select one of following address types:
  - **SMTP**: This is the default address type. Click this button and then type the new SMTP address in the *** Email address** box.
  - **Custom address type**: Click this button and type one of the supported non-SMTP email address types in the *** Email address** box.

    Note: With the exception of X.400 addresses, Exchange doesn't validate custom addresses for correct formatting. You must make sure that the custom address you specify complies with the format requirements for that address type.
- **Set the external email address**: Use this box to change the mail user's external address. Email sent to this mail user is forwarded to this email address.
- **Automatically update email addresses based on the email address policy applied to this recipient**: Select this check box to have the recipient's email addresses automatically updated based on changes made to email address policies in your organization. This box is selected by default.

Mail Flow Settings

Use the **Mail Flow Settings** section to view or change the following settings:

- **Message Size Restrictions**: These settings control the size of messages that the mail user can send and receive. Click **View details** to view and change maximum size for sent and received messages.
  - **Sent messages**: To specify a maximum size for messages sent by this user, select the **Maximum message size (KB)** check box and type a value in the box. The message size must be between 0 and 2,097,151 KB. If the user sends a message larger than the specified size, the message will be returned to the user with a descriptive error message.
  - **Received messages**: To specify a maximum size for messages received by this user, select the **Maximum message size (KB)** check box and type a value in the box. The message size must be between 0 and 2,097,151 KB. If the user receives a message larger than the specified size, the message will be returned to the sender with a descriptive error message.
- **Message Delivery Restrictions**: These settings control who can send email messages to this mail user. Click **View details** to view and change these restrictions.
  - **Accept messages from**: Use this section to specify who can send messages to this user.
  - **All senders**: Select this option to specify that the user can accept messages from all senders. This includes both senders in your Exchange organization and external senders. This option is selected by default. This option includes external users only if you clear the **Require that all senders are authenticated** check box. If you select this check box, messages from external users will be rejected.
  - **Only senders in the following list**: Select this option to specify that the user can accept messages only from a specified set of senders in your Exchange organization. Click **Add** ＋ to display the **Select Recipients** page, which displays a list of all recipients in your Exchange organization. Select the recipients you want, add them to the list, and then click **OK**. You can also search for a specific

recipient by typing the recipient's name in the search box and then clicking **Search** 🔎.

  - **Require that all senders are authenticated**: Select this option to prevent anonymous users from sending messages to the user.

  - **Reject messages from**: Use this section to block people from sending messages to this user.

  - **No senders**: Select this option to specify that the mailbox won't reject messages from any senders in the Exchange organization. This option is selected by default.

  - **Senders in the following list**: Select this option to specify that the mailbox will reject messages from a specified set of senders in your Exchange organization. Click **Add** + to display the **Select Recipients** page, which displays a list of all recipients in your Exchange organization. Select the recipients you want, add them to the list, and then click **OK**. You can also search for a specific recipient by typing the recipient's name in the search box and then clicking **Search** 🔎.

Member Of

Use the **Member Of** section to view a list of the distribution groups or security groups to which this user belongs. You can't change membership information on this page. Note that the user may match the criteria for one or more dynamic distribution groups in your organization. However, dynamic distribution groups aren't displayed on this page because their membership is calculated each time they're used.

MailTip

Use the **MailTip** section to add a MailTip to alert users of potential issues before they send a message to this recipient. A MailTip is text that's displayed in the InfoBar when this recipient is added to the To, Cc, or Bcc lines of a new email message.

> **NOTE**
>
> MailTips can include HTML tags, but scripts aren't allowed. The length of a custom MailTip can't exceed 175 displayed characters. HTML tags aren't counted in the limit.

Use the Exchange Management Shell to change mail user properties

Properties for a mail user are stored in both Active Directory and Exchange. In general, use the **Get-User** and **Set-User** cmdlets to view and change organization and contact information properties. Use the **Get-MailUser** and **Set-MailUser** cmdlets to view or change mail-related properties, such email addresses, the MailTip, custom attributes, and whether the mail user is hidden from address lists.

Use the **Get-MailUser** and **Set-MailUser** cmdlets to view and change properties for mail users. For information, see the following topics:

- Get-User

- Set-User

- Get-MailUser

- Set-MailUser

Here are some examples of using the Exchange Management Shell to change mail user properties.

This example sets the external email address for Pilar Pinilla.

```
Set-MailUser "Pilar Pinilla" -ExternalEmailAddress pilarp@tailspintoys.com
```

This example hides all mail users from the organization's address book.

```
Get-MailUser | Set-MailUser -HiddenFromAddressListsEnabled $true
```

This example sets the Company property for all mail users to Contoso.

```
Get-User -ResultSize unlimited -Filter "RecipientTypeDetails -eq 'mailuser'" | Set-User -Company Contoso
```

This example sets the CustomAttribute1 property to a value of ContosoEmployee for all mail users that have a value of Contoso in the Company property.

```
Get-User -ResultSize unlimited -Filter "(RecipientTypeDetails -eq 'mailuser') -and (Company -eq 'Contoso')" |
Set-MailUser -CustomAttribute1 ContosoEmployee
```

For detailed syntax and parameter information, see Set-MailUser.

How do you know this worked?

To verify that you've successfully changed properties for mail users, do the following:

- In the EAC, select the mail user and then click **Edit** ✎ to view the property that you changed.

- In the Exchange Management Shell, use the **Get-User** and **Get-MailUser** cmdlets to verify the changes. One advantage of using the Exchange Management Shell is that you can view multiple properties for multiple mail contacts.

  ```
  Get-MailUser | Format-List Name,CustomAttribute1
  ```

  In the example above where the Company property was set to Contoso for all mail contacts, run the following command to verify the changes:

  ```
  Get-User -ResultSize unlimited -Filter "RecipientTypeDetails -eq 'mailuser'" | Format-List Name,Company
  ```

  In the example above where all mail users had the CustomAttribute1 property set to ContosoEmployee, run the following command to verify the changes.

  ```
  Get-MailUser | Format-List Name,CustomAttribute1
  ```

Bulk edit mail users

You can also use the EAC to change selected properties for multiple mail users. When you select two or more mail users from the contacts list in the EAC, the properties that can be bulk edited are displayed in the Details pane. When you change one of these properties, the change is applied to all selected recipients.

When you bulk edit mail users, you can change the following property areas:

- **Contact Information**: Change shared properties such as street, postal code, and city name.

- **Organization**: Change shared properties such as department name, company name, and the manager that the selected mail contacts or mail users report to.

Use the EAC to bulk edit mail users

1. In the EAC, navigate to **Recipients** > **Contacts**.

2. In the list of contacts, select two or more mail users. You can't bulk edit a combination of mail contacts and

mail users.

Note: You can select multiple adjacent mail users by holding down the Shift key and clicking the first mail user, and then clicking the last mail user you want to edit. You can also select multiple mail users by holding down the Ctrl key and clicking each one that you want to edit.

3. In the Details pane, under **Bulk Edit**, click **Update** under **Contact Information** or **Organization**.

4. Make the changes on the properties page and then save your changes.

How do you know this worked?

To verify that you've successfully bulk edited mail users, do one of the following:

- In the EAC, select each of the mail users that you bulk edited and then click **Edit** ✏ to view the properties that you changed.

- In the Exchange Management Shell, use the **Get-User** cmdlet to verify the changes. For example, say you used the bulk edit feature in the EAC to change the manager and the office for all mail users from a vendor company named A. Datum Corporation. To verify these changes, you could run the following command in the Exchange Management Shell:

```
Get-User -ResultSize unlimited -Filter "(RecipientTypeDetails -eq 'mailuser') -and (Company -eq
'Adatum')" | Format-List Name,Office,Manager
```

Create and manage room mailboxes

8/3/2020 • 18 minutes to read • Edit Online

A room mailbox is a resource mailbox that's assigned to a physical location, such as a conference room, an auditorium, or a training room. With room mailboxes, users can easily reserve these rooms by including room mailboxes in their meeting requests. When they do this, the room mailbox uses options you can configure to decide whether the invite should be accepted or denied.

To create a room mailbox, you need to be an administrator who's a member of either the Organization Management or Recipient Management role groups.

If you want to grant someone access to a room mailbox so they can directly manage its calendar (for example, an assistant who needs to make room for an executive meeting), you can do so using the instructions in Manage permissions for recipients. After a user's been granted permissions to access a room mailbox, they can open the mailbox using the instructions in Open and use a shared mailbox in Outlook for Windows.

IMPORTANT Room mailboxes should never be set as the organizer of a meeting, nor should room mailboxes be accessed directly by users in order to make changes to a meeting. Rooms should only be added to meetings in the Attendee or Location fields. Otherwise you will override the Resource Booking Assistant (RBA), which manages and processes all calendar items sent to the room mailbox, and unexpected errors may occur. If your organization has one or more users who need to manage a room and its mailbox, then assign users to the room as resource delegates for the room mailbox, as described later in this article. When a delegate is assigned, all items sent to the room's mailbox will be directed to the booking delegate, who can then accept or decline from their own Inbox. If your organization wants to use a room mailbox like a team calendar, consider using Exchange's shared calendar features.

If you want to learn about the types of recipients that are available in Exchange Server, check out Recipients. For info about another type of resource mailbox, check out Manage equipment mailboxes.

What do you need to know before you begin?

- Estimated time to complete: 5 minutes.
- To open the Exchange admin center (EAC), see Exchange admin center in Exchange Server. To open the Exchange Management Shell, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the"Recipient Provisioning Permissions" section in the Recipients Permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.
- If you're using room or equipment mailboxes in Microsoft 365 or Office 365, see [Room and equipment mailboxes(https://docs.microsoft.com/microsoft-365/admin/manage/room-and-equipment-mailboxes) for more information.

IMPORTANT

If you're running Exchange 2013 in a hybrid scenario, make sure you create the room mailboxes in the appropriate place. Create your room mailboxes for your on-premises organization on-premises, and room mailboxes for Exchange Online should be created in the cloud. > Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Create a room mailbox

1. In the Exchange admin center, navigate to **Recipients** > **Resources**.

2. To create a room mailbox, click **New** + > **Room mailbox**.

3. Use the options on the page to specify the settings for the new resource mailbox.

   - **Room name**: Use this box to type a name for the room mailbox. This is the name that's listed in the resource mailbox list in the Exchange admin center and in your organization's address book. This name is required and it can't exceed 64 characters.

     TIP

     Although there are other fields that describe the details of the room (for example, Location and Capacity) consider summarizing the most important details in the room name using a consistent naming convention. Why? So users can easily see the details when they select the room from the address book in the meeting request.

   - **Alias**: A room mailbox has an email address so it can receive booking requests. The email address consists of an alias on the left side of the @ symbol, which must be unique in the forest, and your domain name on the right. The alias is required.

   - **Location**, **Phone**, **Capacity**: You can use these fields to enter details about the room. However, as explained earlier, you can include some or all of this information in the room name so users can see it.

4. When you're finished, click **Save** to create the room mailbox.

After you've created a room mailbox, you can Change how a room mailbox handles meeting requests (including whether it responds automatically or someone needs to decide what to do). By default, it'll automatically accept or decline requests depending on whether the requests conflict with any existing meetings on its calendar. It'll also allow meetings that repeat, and allow meetings up to 180 days from the current date (and decline any requests beyond that) that are up to 24 hours in duration. If you want to change other options, head down to Change other room mailbox properties.

For information on how to create a room mailbox using the Exchange Management shell, see Examples 2 and 3 in New-Mailbox.

Change how a room mailbox handles meeting requests

1. In the Exchange admin center, navigate to **Recipients** > **Resources**.

2. In the list of resource mailboxes, click the room mailbox that you want to change the properties for, and then click **Edit** ✎.

3. On the room mailbox properties page, click **Booking Delegates** (allow automatic responses or not) or **Booking Options** (allow repeating meetings, decline meetings that are scheduled too far out, etc).

Use the **Booking Delegates** section to view or change how the room mailbox handles meeting requests and to define who can accept or decline booking requests if it isn't done automatically.

- Select one of the following options to handle meeting requests.
  - **Accept or decline booking requests automatically**: If this is selected, the meeting request will automatically be declined if there's a scheduling conflict with an existing reservation, or if the booking request violates the scheduling limits of the resource, for example, the reservation duration is too long.
  - **Select delegates who can accept or decline booking requests**: If this is selected, one of the people you added to the **Delegates** list below will be responsible for accepting or declining meeting requests that are sent to the room mailbox. If you assign more than one resource delegate, only one of them has to act on a specific meeting request.

Use the **Booking Options** section to view or change the settings for the booking policy that defines when the room can be scheduled, how long it can be reserved, and how far in advance it can be reserved.

- **Allow repeating meetings**: This setting allows or prevents repeating meetings for the room.
- **Allow scheduling only during working hours**: This setting accepts or declines meeting requests that aren't during the working hours defined for the room, which are, by default, 8:00 A.M. to 5:00 P.M. Monday through Friday. You can configure the working hours of the room mailbox either by logging into the mailbox using Outlook on the web and going to the **Options** > **Calendar** > **Calendar appearance** page, or by using Set-MailboxCalendarConfiguration.
- **Always decline if the end date is beyond this limit**: This setting controls the behavior of repeating meetings that extend beyond the date specified by the maximum booking lead time setting.
  - If you enable this setting, a repeating booking request is automatically declined if the bookings start on or before the date specified by the value in the **Maximum booking lead time** box, and they extend beyond the specified date. This is the default setting.
  - If you disable this setting, a repeating booking request is automatically accepted if booking requests start on or before the date specified by the value in the **Maximum booking lead time** box, and they extend beyond the specified date. However, the number of bookings is reduced so bookings won't occur after the specified date.
- **Maximum booking lead time (days)**: This setting specifies the maximum number of days in advance that the room can be booked. Valid input is an integer between 0 and 1080. The default value is 180 days.
- **Maximum duration (hours)**: This setting specifies the maximum duration that the room can be reserved in a booking request. A valid value is an integer from 0 through 35791394. The default value is 24 hours. When the value is set to 0, the maximum duration of a meeting is unlimited.

  For repeating booking requests, the maximum booking duration applies to the length of each instance of the repeating booking request.

There's also a box on this page that you can use to write a message that will be sent to users who send booking requests to reserve the room.

Change resource scheduling settings

An Admin or user with full access to the resource mailbox can make changes to the resource scheduling settings.

1. Log in to Outlook Web App and click on **Your name** in the top right corner.
2. Click **Open another mailbox**. Locate the meeting room resource you want and click **Open**.
3. Go to **settings** and click **Calendar**.

4. Navigate to **Resource scheduling**.
5. Configure the Scheduling Options and Scheduling Permissions as needed. (See the descriptions of all options in the following two sections for details.)
6. Click **Save** after you have finished making your changes.

Scheduling Options

- **Automatically process meeting requests and cancellations**

Enables or disables all options below as well as the options under Scheduling Permissions. If not checked, the owner must manage every request manually. By default, this is not checked.

- **Disable reminders**

Enables or disables reminders for events in this calendar. This setting applies only to the resource; the organizer and attendees will still receive reminders if they have elected to do so.

- **Maximum number of days in advance resources can be booked**

Limits how far in advance an event can be scheduled. The default is 180 days.

- **Always decline if the end date is beyond this limit**

Requests beyond the maximum number of days specified will be automatically declined. Valid values are between 0 (today) and 1080 (about three years in the future).

- **Limit meeting duration and Maximum allowed minutes**

Limits the amount of time for which a room can be scheduled within a single day. Unchecking the box will mean a meeting has no limit. Checking the box allows for a limit between 0 to 1440 minutes.

- **Allow scheduling only during working hours**

If checked, an event can only be scheduled during the hours specified under Calendar Work Week in the Calendar tab. Events outside of working hours will be automatically declined.

- **Allow repeating meetings**

Allows booking of the resource room at a regular interval. The event can be set to repeat over a specified duration of time (also called recurring).

- **Allow conflicts**

Allow or prevent conflicting meeting requests (double booking). If repeating meetings are allowed as well, this setting will only apply to repeating meetings. When the resource is invited, it will need to be entered into the Attendees field as opposed to being chosen with the Add Rooms button.

- **Allow up to this number of individual conflicts**

This setting specifies the maximum number of conflicts that are allowed for new repeating meeting requests. When set to 0, a recurring event will fail to schedule if one or more conflicting appointments already appear. When set to a number greater than 0, a recurring event is allowed the specified number of conflicts before being denied.

- **Allow up to this percentage of individual conflicts**

This setting specifies the maximum percentage of meeting conflicts that are allowed for new repeating meeting requests. This is similar to specifying a number of individual conflicts (explained above), but in this case, a recurring event is allowed the specified percentage of conflicts before being denied.

Scheduling Permissions

- **These users can schedule automatically if the resource is available**

By default, everyone can schedule this resource without the manual approval of the resource owner. If **Select users and groups** is selected, only the users and groups specified can schedule automatically. All other users or groups will receive a decline message. If **Select users and groups** is selected but no users or groups are specified, this option will be ignored.

- **These users can submit a request for owner approval if the resource is available**

If everyone is selected, then all requests must receive manual approval by the resource owner. If **Select users and groups** is selected, only the specified users and groups require manual approval by the resource owner. **Select users and groups** is selected and left blank by default so that all requests are approved automatically.

- **These users can schedule automatically if the resource is available and can submit a request for owner approval if the resource is unavailable**

When everyone is selected (the default setting), any request during an open time frame will be automatically approved. If the room is booked at the requested time, a form is submitted to the resource owner for manual approval. If **Select users and groups** is selected, only those specified will have the option to have the request manually approved; all others will have a conflicting request denied without the option of manual approval by the resource owner.

Change other room mailbox properties

After you create a room mailbox, you can make changes and set additional properties by using the Exchange admin center or the Exchange Management Shell.

Use the Exchange admin center to change room mailbox properties

1. In the Exchange admin center, navigate to **Recipients** > **Resources**.

2. In the list of resource mailboxes, click the room mailbox that you want to change the properties for, and then click **Edit** ✎.

3. On the room mailbox properties page, click one of the following sections to view or change properties (for booking options, see Change how a room mailbox handles meeting requests.

   - General

   - Contact Information

   - Email Address

   - MailTip

General

Use the **General** section to view or change basic information about the resource.

- **Room name**: This name appears in the resource mailbox list in the Exchange admin center and in your organization's address book. It can't exceed 64 characters if you change it.

- **Email address**: This read-only box displays the email address for the room mailbox. You can change it in the Email Address section.

- **Capacity**: Use this box to enter the maximum number of people who can safely occupy the room.

Click **More options** to view or change these additional properties:

- **Organizational unit**: This read-only box displays the organizational unit (OU) that contains the account for the room mailbox. You have to use Active Directory Users and Computers to move the account to a different OU.

- **Mailbox database**: This read-only box displays the name of the mailbox database that hosts the room mailbox. Use the **Migration** page in the Exchange admin center to move the mailbox to a different database.
- **Alias** Use this box to change the alias for the room mailbox.
- **Hide from address lists**: Select this check box to prevent the room mailbox from appearing in the address book and other address lists that are defined in your Exchange organization. After you select this check box, users can still send booking messages to the room mailbox by using the email address.
- **Department**: Use this box to specify a department name that the room is associated with. You can use this property to create recipient conditions for dynamic distribution groups and address lists.
- **Company**: Use this box to specify a company that the room is associated with, if applicable. Like the Department property, you can use this property to create recipient conditions for dynamic distribution groups and address lists.
- **Address book policy**: Use this option to specify an address book policy (ABP) for the room mailbox. ABPs contain a global address list (GAL), an offline address book (OAB), a room list, and a set of address lists. To learn more, see Address book policies in Exchange Server.

  In the drop-down list, select the policy that you want associated with this mailbox.
- **Custom attributes**: This section displays the custom attributes defined for the room mailbox. To specify custom attribute values, click **Edit** ✎. You can specify up to 15 custom attributes for the recipient.

Contact Information

Use the **Contact Information** section to view or change the contact information for the room. The information on this page is displayed in the address book.

> **TIP**
>
> You can use the **State/Province** box to create recipient conditions for dynamic distribution groups, email address policies, or address lists.

Email Address

Use the **Email Address** section to view or change the email addresses associated with the room mailbox. This includes the mailbox's primary SMTP address and any associated proxy addresses. The primary SMTP address (also known as the *reply address*) is displayed in bold text in the address list, with the uppercase **SMTP** value in the **Type** column.

- **Add**: Click **Add** + to add a new email address for this mailbox. Select one of following address types:
  - **SMTP**: This is the default address type. Click this button and then type the new SMTP address in the **Email address** box.
  - **EUM**: An EUM (Exchange Unified Messaging) address is used by the Microsoft Exchange Unified Messaging service in Exchange 2016 to locate UM-enabled recipients within an Exchange organization. EUM addresses consist of the extension number and the UM dial plan for the UM-enabled user. Click this button and type the extension number in the **Address/Extension** box. Then click **Browse** and select a dial plan for the mailbox. (**Note**: Unified Messaging is not available in Exchange 2019.)
  - **Custom address type**: Click this button and type one of the supported non-SMTP email address types in the **Email address** box.

Notes:

- With the exception of X.400 addresses, Exchange doesn't validate custom addresses for correct formatting. You must make sure that the custom address you specify complies with the format requirements for that address type.

- When you add a new email address, you have the option to make it the primary SMTP address.

- **Automatically update email addresses based on the email address policy applied to this recipient**: Select this check box to have the recipient's email addresses automatically updated based on changes made to email address policies in your organization.

MailTip

Use the **MailTip** section to add a MailTip to alert users of potential issues before they send a booking request to the room mailbox. A MailTip is text that's displayed in the InfoBar when this recipient is added to the To, Cc, or Bcc lines of a new email message.

> **NOTE**
>
> MailTips can include HTML tags, but scripts aren't allowed. The length of a custom MailTip can't exceed 175 displayed characters. HTML tags aren't counted in the limit.

Use the Exchange Management Shell to change room mailbox properties

Use the following sets of cmdlets to view and change room mailbox properties: **Get-Mailbox** and **Set-Mailbox** cmdlets to view and change general properties and email addresses for room mailboxes. Use the **Get-CalendarProcessing** and **Set-CalendarProcessing** cmdlets to view and change delegates and booking options.

- **Get-User** and **Set-User**: Use these cmdlets to view and set general properties such as location, department, and company names.

- **Get-Mailbox** and **Set-Mailbox**: Use these cmdlets to view and set mailbox properties, such as email addresses and the mailbox database.

- **Get-CalendarProcessing** and **Set-CalendarProcessing**: Use these cmdlets to view and set booking options and delegates.

For information about these cmdlets, see the following topics:

- Get-User

- Set-User

- Get-Mailbox

- Set-Mailbox

- Get-CalendarProcessing

- Set-CalendarProcessing

Here are some examples of using the Exchange Management Shell to change room mailbox properties.

This example changes the display name, the primary SMTP address (called the default reply address), and the room capacity. Also, the previous reply address is kept as a proxy address.

```
Set-Mailbox "Conf Room 123" -DisplayName "Conf Room 31/123 (12)" -EmailAddresses
SMTP:Rm33.123@contoso.com,smtp:rm123@contoso.com -ResourceCapacity 12
```

This example configures room mailboxes to allow booking requests to be scheduled only during working hours and sets a maximum duration of 9 hours.

```
Get-Mailbox -ResultSize unlimited -Filter "RecipientTypeDetails -eq 'RoomMailbox'" | Set-CalendarProcessing -
ScheduleOnlyDuringWorkHours $true -MaximumDurationInMinutes 540
```

This example uses the **Get-User** cmdlet to find all room mailboxes that correspond to private conference rooms, and then uses the **Set-CalendarProcessing** cmdlet to send booking requests to a delegate named Robin Wood to accept or decline.

```
Get-User -ResultSize unlimited -Filter "(RecipientTypeDetails -eq 'RoomMailbox') -and (DisplayName -like
'Private*')" | Set-CalendarProcessing -AllBookInPolicy $false -AllRequestInPolicy $true -ResourceDelegates
"Robin Wood"
```

Create a room list

If you're planning to have more to have hundreds of rooms, use multiple room lists to help you organize your rooms. If your company has several buildings with rooms that can be booked for meetings, it might help to create room lists for each building. Room lists are specially marked distribution groups that you can use the same way you use distribution groups. However, you can only create room lists using the Exchange Management Shell.

> **NOTE**
>
> Although there is no hard limit to the number of rooms you can have in a Room List, the maximum number of rooms that can be returned in request for a Room List is 100. A possible workaround would be to further break down your rooms into smaller lists.

Use the Exchange Management Shell to create a room list

This example creates a room list for building 32.

```
New-DistributionGroup -Name "Building 32 Conference Rooms" -OrganizationalUnit "contoso.com/rooms" -RoomList
```

Use the Exchange Management Shell to add a room to a room list

This example adds confroom3223 to the building 32 room list.

```
Add-DistributionGroupMember -Identity "Building 32 Conference Rooms" -Member confroom3223@contoso.com
```

Use the Exchange Management Shell to convert a distribution group to a room list

You may already have created distribution groups in the past that contain your conference rooms. You don't need to recreate them; we can convert them quickly into a room list.

This example converts the distribution group, building 34 conference rooms, to a room list.

```
Set-DistributionGroup -Identity "Building 34 Conference Rooms" -RoomList
```

How do you know this worked?

To verify that you've successfully changed properties for a room mailbox, do the following:

- In the Exchange admin center, select the mailbox and then click **Edit** ✎ to view the property or feature that you changed. Depending on the property that you changed, it might be displayed in the Details pane for the selected mailbox.

- In the Exchange Management Shell, use the **Get-Mailbox** cmdlet to verify the changes. One advantage of using the Exchange Management Shell is that you can view multiple properties for multiple mailboxes. In the example above where booking requests could be scheduled only during working hours and have a maximum duration of 9 hours, run the following command to verify the new values.

  ```
  Get-Mailbox -ResultSize unlimited -Filter "RecipientTypeDetails -eq 'RoomMailbox'" | Get-
  CalendarProcessing | Format-List Identity,ScheduleOnlyDuringWorkHours,MaximumDurationInMinutes
  ```

Manage equipment mailboxes

8/3/2020 • 12 minutes to read • Edit Online

In Exchange Server, an *equipment mailbox* is a resource mailbox assigned to a resource that's not location specific, such as a portable computer, projector, microphone, or a company car. After an administrator creates an equipment mailbox, users can easily reserve the piece of equipment by including the corresponding equipment mailbox in a meeting request. You can use the Exchange admin center (EAC) and the Exchange Management Shell to create an equipment mailbox or change equipment mailbox properties. For more information, see Recipients.

For information about another type of resource mailbox, a room mailbox, see Create and manage room mailboxes.

What do you need to know before you begin?

- Estimated time to complete: 2 to 5 minutes.
- To open the EAC, see Exchange admin center in Exchange Server. To open the Exchange Management Shell, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Recipient Provisioning Permissions" section in the Recipients Permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Create an equipment mailbox

Use the EAC to create an equipment mailbox

1. In the EAC, navigate to **Recipients** > **Resources**.
2. To create an equipment mailbox, click **New** > **Equipment mailbox**. To create a room mailbox, click **New** > **Room mailbox**.
3. Use the options on the page to specify the settings for the new resource mailbox.

   - *** Equipment name**: Use this box to type a name for the equipment mailbox. This is the name that's listed in the resource mailbox list in the EAC and in your organization's address book. This name is required and it can't exceed 64 characters.

   > **TIP**
   >
   > Although there are other fields that describe the details of the room, for example, Capacity, consider summarizing the most important details in the equipment name using a consistent naming convention. Why? So users can easily see the details when they select the equipment from the address book in a meeting request.

   - *** Email address**: An equipment mailbox has an email address to receive booking requests. The email address consists of an alias on the left side of the @ symbol, which must be unique in the forest, and your domain name on the right. The email address is required.

4. When you're finished, click **Save** to create the equipment mailbox.

Once you've created your equipment mailbox, you can edit your equipment mailbox to update info about booking options, MailTips and delegates. Check out the Change equipment mailbox properties section below to change room mailbox properties.

Use the Exchange Management Shell to create an equipment mailbox

This example creates an equipment mailbox with the following configuration:

- The equipment mailbox resides on Mailbox Database 1.
- The equipment's name is MotorVehicle2 and the name will display in the GAL as Motor Vehicle 2.
- The email address is MotorVehicle2@contoso.com.
- The mailbox is in the Equipment organizational unit.
- The *Equipment* parameter specifies that this mailbox will be created as an equipment mailbox.

```
New-Mailbox -Database "Mailbox Database 1" -Name MotorVehicle2 -OrganizationalUnit Equipment -DisplayName
"Motor Vehicle 2" -Equipment
```

For detailed syntax and parameter information, see New-Mailbox.

How do you know this worked?

To verify that you've successfully created a user mailbox, do one of the following:

- In the EAC, navigate to **Recipients** > **Resources**. The new user mailbox is displayed in the mailbox list. Under **Mailbox Type**, the type is **Equipment**.
- In the Exchange Management Shell, run the following command to display information about the new equipment mailbox.

  ```
  Get-Mailbox <Name> | Format-List Name,RecipientTypeDetails,PrimarySmtpAddress
  ```

Change equipment mailbox properties

After you create an equipment mailbox, you can make changes and set additional properties by using the EAC or the Exchange Management Shell.

Use the EAC to change equipment mailbox properties

1. In the EAC, navigate to **Recipients** > **Resources**.
2. In the list of resource mailboxes, click the equipment mailbox that you want to change the properties for, and then click **Edit** ✎.
3. On the equipment mailbox properties page, click one of the following sections to view or change properties.
   - General
   - Delegates
   - Booking Options
   - Contact Information

- Email Address
- MailTip

General

Use the **General** section to view or change basic information about the resource.

- *** Equipment name**: This name appears in the resource mailbox list in the EAC and in your organization's address book. It can't exceed 64 characters if you change it.
- *** Email address**: This read-only box displays the email address for the equipment mailbox. You can change it in the Email Address section.
- **Capacity**: Use this box to enter the maximum number of people who can use this resource, if applicable, For example, if the equipment mailbox corresponds to a compact car, you could enter **4**.

Click **More options** to view or change these additional properties:

- **Organizational unit**: This read-only box displays the organizational unit (OU) that contains the account for the equipment mailbox. You have to use Active Directory Users and Computers to move the account to a different OU.
- **Mailbox database**: This read-only box displays the name of the mailbox database that hosts the equipment mailbox. Use the **Migration** page in the EAC to move the mailbox to a different database.
- *** Alias** Use this box to change the alias for the equipment mailbox.
- **Hide from address lists**: Select this check box to prevent equipment mailbox from appearing in the address book and other address lists that are defined in your Exchange organization. After you select this check box, users can still send booking messages to the equipment mailbox by using the email address.
- **Department**: Use this box to specify a department name that the resource is associated with. You can use this property to create recipient conditions for dynamic distribution groups and address lists.
- **Company**: Use this box to specify a company that the resource is associated with. Like the Department property, you can use this property to create recipient conditions for dynamic distribution groups and address lists.
- **Address book policy**: Use this option to specify an address book policy (ABP) for the resource. ABPs contain a global address list (GAL), an offline address book (OAB), a room list, and a set of address lists. To learn more, see Address book policies in Exchange Server.

  In the drop-down list, select the policy that you want associated with this mailbox.
- **Custom attributes**: This section displays the custom attributes defined for the equipment mailbox. To specify custom attribute values, click **Edit** ✎. You can specify up to 15 custom attributes for the recipient.

Delegates

Use this section to view or change how the equipment mailbox handles reservation requests and to define who can accept or decline booking requests if it isn't done automatically.

- **Booking requests**: Select one of the following options to handle booking requests.
  - **Accept or decline booking requests automatically**: A valid meeting request automatically reserves the resource. If there's a scheduling conflict with an existing reservation, or if the booking request violates the scheduling limits of the resource, for example, the reservation duration is too long, the meeting request is automatically declined.
  - **Select delegates who can accept or decline booking requests**: Resource delegates are responsible for accepting or declining meeting requests that are sent to the equipment mailbox. If

you assign more than one resource delegate, only one of them has to act on a specific meeting request.

- **Delegates**: If you selected the option requiring that booking requests be sent to delegates, the specified delegates are listed. Click **Add** + or **Remove** − to add or remove delegates from this list.

Booking Options

Use the **Booking Options** section to view or change the settings for the booking policy that defines when the resource can be scheduled, how long it can be reserved, and how far in advance it can be reserved.

- **Allow repeating meetings**: This setting allows or prevents repeating meetings for the resource. By default, this setting is enabled, so repeating meetings are allowed.

- **Allow scheduling only during working hours**: This setting accepts or declines meeting requests that aren't during the working hours defined for the resource. By default, this setting is disabled, so meeting requests are allowed outside the working hours.By default, working hours are 8:00 AM to 5:00 PM Monday through Friday. You can configure the working hours of the equipment mailbox in the Appearance section on the Calendar page.

- **Always decline if the end date is beyond this limit**: This setting controls the behavior of repeating meetings that extend beyond the date specified by the maximum booking lead time setting.

  - If you enable this setting, a repeating booking request is automatically declined if the bookings start on or before the date specified by the value in the **Maximum booking lead time** box, and they extend beyond the specified date. This is the default setting.

  - If you disable this setting, a repeating booking request is automatically accepted if the booking requests start on or before the date specified by the value in the **Maximum booking lead time** box, and they extend beyond the specified date. However, the number of bookings is reduced so bookings won't occur after the specified date.

- **Maximum booking lead time (days)**: This setting specifies the maximum number of days in advance that the resource can be booked. Valid input is an integer between 0 and 1080. The default value is 180 days.

- **Maximum duration (hours)**: This setting specifies the maximum duration that the resource can be reserved in a booking request. The default value is 24 hours.

  For repeating booking requests, the maximum booking duration applies to the length of each instance of the repeating booking request.

There is also a box on this page that you can use to write a message that will be sent to users who send meeting requests to reserve the resource.

Contact Information

Use the **Contact Information** section to view or change the contact information for the resource. The information on this page is displayed in the address book.

> **TIP**
>
> You can use the **State/Province** box to create recipient conditions for dynamic distribution groups, email address policies, or address lists.

Email Address

Use the **Email Address** section to view or change the email addresses associated with the equipment mailbox. This includes the mailbox's primary SMTP address and any associated proxy addresses. The primary SMTP address (also known as the *reply address*) is displayed in bold text in the address list, with the uppercase **SMTP** value in the **Type** column.

- **Add**: Click **Add** + to add a new email address for this mailbox. Select one of following address types:
  - **SMTP**: This is the default address type. Click this button and then type the new SMTP address in the * **Email address** box.
  - **EUM**: An Exchange Unified Messaging (EUM) address is used by the Exchange Unified Messaging service in Exchange 2016 to locate UM-enabled recipients in an Exchange organization. EUM addresses consist of the extension number and the UM dial plan for the UM-enabled user. Click this button and type the extension number in the **Address/Extension** box. Then click **Browse** and select a dial plan for the mailbox.(**Note**: Unified Messaging is not available in Exchange 2019.)
  - **Custom address type**: Click this button and type one of the supported non-SMTP email address types in the * **Email address** box.

    Notes:

    - With the exception of X.400 addresses, Exchange doesn't validate custom addresses for correct formatting. You must make sure that the custom address you specify complies with the format requirements for that address type.
    - When you add a new email address, you have the option to make it the primary SMTP address.
- **Automatically update email addresses based on the email address policy applied to this recipient**: Select this check box to have the recipient's email addresses automatically updated based on changes made to email address policies in your organization.

MailTip

Use the **MailTip** section to add a MailTip to alert users of potential issues before they send a booking request to the equipment mailbox. A MailTip is text that's displayed in the InfoBar when this recipient is added to the To, Cc, or Bcc lines of a new email message.

> **NOTE**
>
> MailTips can include HTML tags, but scripts aren't allowed. The length of a custom MailTip can't exceed 175 displayed characters. HTML tags aren't counted in the limit.

Use the Exchange Management Shell to change equipment mailbox properties

Use the following sets of cmdlets to view and change equipment mailbox properties: **Get-Mailbox** and **Set-Mailbox** cmdlets to view and change general properties and email addresses for equipment mailboxes. Use the **Get-CalendarProcessing** and **Set-CalendarProcessing** cmdlets to view and change delegates and booking options.

- **Get-User** and **Set-User**: Use these cmdlets to view and set general properties such as department and company names.
- **Get-Mailbox** and **Set-Mailbox**: Use these cmdlets to view and set mailbox properties, such as email addresses and the mailbox database.
- **Get-CalendarProcessing** and **Set-CalendarProcessing**: Use these cmdlets to view and set booking options and delegates.

For information about these cmdlets, see the following topics:

- Get-User

- Set-User
- Get-Mailbox
- Set-Mailbox
- Get-CalendarProcessing
- Set-CalendarProcessing

Here are some examples of using the Exchange Management Shell to change equipment mailbox properties.

This example changes the display name and primary SMTP address (called the default reply address) for the MotorPool 1 equipment mailbox. The previous reply address is kept as a proxy address.

```
Set-Mailbox "MotorPool 1" -DisplayName "Motor Pool 1 - Compact" -EmailAddresses
SMTP:MP1.compact@contoso.com,smtp:MP.1@contoso.com
```

This example configures equipment mailboxes to allow booking requests to be scheduled only during working hours.

```
Get-Mailbox -ResultSize unlimited -Filter "RecipientTypeDetails -eq 'EquipmentMailbox'" | Set-
CalendarProcessing -ScheduleOnlyDuringWorkHours $true
```

This example uses the **Get-User** cmdlet to find all equipment mailboxes in the Audio Visual department, and then uses the **Set-CalendarProcessing** cmdlet to send booking requests to a delegate named Ann Beebe to accept or decline.

```
Get-User -ResultSize unlimited -Filter "(RecipientTypeDetails -eq 'EquipmentMailbox') -and (Department -eq
'Audio Visual')" | Set-CalendarProcessing -AllBookInPolicy $false -AllRequestInPolicy $true -ResourceDelegates
"Ann Beebe"
```

How do you know this worked?

To verify that you've successfully changed properties for an equipment mailbox, do the following:

- In the EAC, select the mailbox and then click **Edit** ✎ to view the property or feature that you changed. Depending on the property that you changed, it might be displayed in the Details pane for the selected mailbox.

- In the Exchange Management Shell, use the **Get-Mailbox** cmdlet to verify the changes. One advantage of using the Exchange Management Shell is that you can view multiple properties for multiple mailboxes. In the example above where booking requests could be scheduled only during working hours, run the following command to verify the new value.

  ```
  Get-Mailbox -ResultSize unlimited -Filter "RecipientTypeDetails -eq 'EquipmentMailbox'" | Get-
  CalendarProcessing | Format-List Identity,ScheduleOnlyDuringWorkHours
  ```

Disconnected mailboxes

8/3/2020 • 8 minutes to read • Edit Online

Each Microsoft Exchange mailbox consists of an Active Directory user account and the mailbox data stored in the Exchange mailbox database. All configuration data for a mailbox is stored in the Exchange attributes of the Active Directory user object. The mailbox database contains the mail data that's in the mailbox associated with the user account. The following figure shows the components of a mailbox.

Mailbox components

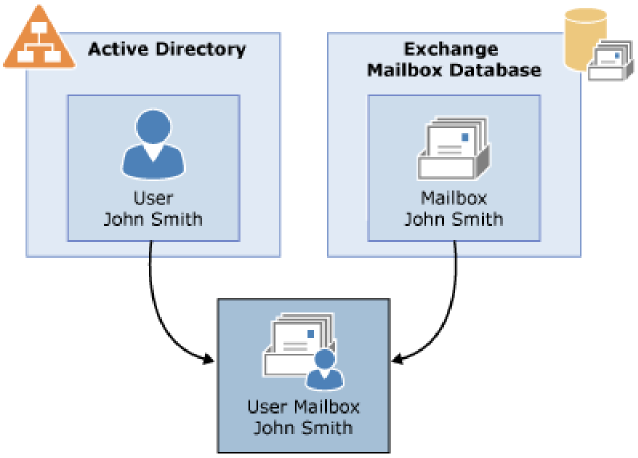

A *disconnected mailbox* is a mailbox object in the mailbox database that isn't associated with an Active Directory user account. There are two types of disconnected mailboxes:

- **Disabled mailboxes**: When a mailbox is disabled or deleted in the Exchange admin center (EAC) or using the **Disable-Mailbox** or **Remove-Mailbox** cmdlet in the Exchange Management Shell, Exchange retains the deleted mailbox in the mailbox database, and switches the mailbox to a disabled state. This is why mailboxes that are either disabled or deleted are referred to as *disabled mailboxes*. The difference is that when you disable a mailbox, the Exchange attributes are removed from the corresponding Active Directory user account, but the user account is retained. When you delete a mailbox, both the Exchange attributes and the Active Directory user account are deleted.

  Disabled and deleted mailboxes are retained in the mailbox database until the deleted mailbox retention period expires, which is 30 days by default. After the retention period expires, the mailbox is permanently deleted (also called *purged*). If a mailbox is deleted using the **Remove-Mailbox** cmdlet, it's also retained for the duration of the retention period.

  > **IMPORTANT**
  >
  > If a mailbox is deleted using the **Remove-Mailbox** cmdlet and either the *Permanent* or *StoreMailboxIdentity* parameter, it will be immediately deleted from the mailbox database.

  To identify the disabled mailboxes in your organization, run the following commands in the Exchange Management Shell:

  ```
  $dbs = Get-MailboxDatabase
  $dbs | foreach {Get-MailboxStatistics -Database $_.DistinguishedName} | where {$_.DisconnectReason -eq
  "Disabled"} | Format-Table DisplayName,Database,DisconnectDate
  ```

- **Soft-deleted mailboxes**: When a mailbox is moved to a different mailbox database, Exchange doesn't fully delete the mailbox from the source mailbox database when the move is complete. Instead, the mailbox in the source mailbox database is switched to a *soft-deleted* state. Like disabled mailboxes, soft-deleted mailboxes are retained in the source database either until the deleted mailbox retention period expires or until the **Remove-StoreMailbox** cmdlet is used to purge the mailbox.

  Run the following commands to identify soft-deleted mailboxes in your organization.

  ```
  $dbs = Get-MailboxDatabase
  $dbs | foreach {Get-MailboxStatistics -Database $_.DistinguishedName} | where {$_.DisconnectReason -eq
  "SoftDeleted"} | Format-Table DisplayName,Database,DisconnectDate
  ```

Working with disabled mailboxes

You can perform several operations on a disabled mailbox before it's purged from the mailbox database:

- Reconnect it to the same user account.

- Connect it to a different user account that isn't mail-enabled, which means the user account doesn't have a mailbox.

- Restore it to a user account that has an existing mailbox. For example, if a user whose mailbox was deleted has a new mailbox, you can restore the user's disabled mailbox to their new mailbox.

- Permanently delete it from the Exchange mailbox database.

Connecting or restoring a disabled mailbox

Here are scenarios in which you may want to connect or restore a disabled mailbox before the mailbox retention period expires or before it's permanently deleted:

- You disabled a mailbox and now want to reconnect the mailbox to the same Active Directory user account.

- You deleted a mailbox by using the EAC or the Remove-Mailbox cmdlet and now want to reconnect the mailbox to a different Active Directory user account.

- You deleted a mailbox and now want to restore the mailbox to an existing mailbox. For example, if a user whose mailbox was deleted has a new mailbox, you can restore the user's disabled mailbox to their new mailbox.

- You want to convert a user mailbox to a linked mailbox associated with a user account that's external to the forest in which your Exchange organization exists. The resource forest scenario is an example of when you would want to associate a mailbox with an external account. In this scenario, user objects in the Exchange forest have mailboxes, but the user objects are disabled for logon. You must associate a mailbox in the Exchange forest with a user account in the external account forest.

There are two ways you can reconnect or restore a disabled mailbox. The first method is to use the EAC or the **Connect-Mailbox** cmdlet to connect a disabled mailbox to a user account. For procedures to reconnect disabled mailboxes, see Connect a disabled mailbox.

The second method uses the **New-MailboxRestoreRequest** cmdlet to merge the contents of the disabled mailbox with an existing mailbox. This cmdlet uses the Mailbox Replication Service (MRS) to restore the mailbox. For procedures to restore disabled mailboxes, see Connect or restore a deleted mailbox.

Permanently deleting a disabled mailbox

As stated previously, Exchange retains disabled mailboxes in the mailbox database based on the deleted mailbox retention settings configured for that mailbox database. After the specified retention period, a disabled mailbox is purged from the Exchange mailbox database. You can also permanently delete a disabled mailbox and all its message content from the mailbox database by using the **Remove-StoreMailbox** cmdlet. After a disabled mailbox is automatically purged or permanently deleted by an administrator, the data loss is permanent and the mailbox can't be recovered.

For more information, see Permanently delete a mailbox.

Working with disabled archive mailboxes

Archive mailboxes become disconnected when they're disabled. Similar to a disabled primary mailbox, a disconnected archive mailbox can be connected by using the **Connect-Mailbox** cmdlet with the *Archive* parameter.

The primary mailbox and the archive mailbox share the same legacy distinguished name (DN), so you must connect the archive mailbox to the same user mailbox that it was previously connected to. You can't connect the archive mailbox to a different user mailbox.

You can perform two operations on a disconnected archive mailbox:

- **Connect it to an existing primary mailbox**: Like a disconnected primary mailbox, a disconnected archive mailbox is retained in the mailbox database until the deleted mailbox retention period expires, which is 30 days by default. During this time, you can recover the archive mailbox by reconnecting it to the same user account that it was connected to before it was disabled.

  > **NOTE**
  >
  > If you disable an archive mailbox for a user mailbox and then enable an archive mailbox for that same user, that user mailbox will get a new archive mailbox. While you can use the **Connect-Mailbox** cmdlet to connect a primary mailbox to a user, you must use the **Enable-Mailbox** cmdlet to connect a disabled archive mailbox to an existing mailbox.

  For more information, see Manage In-Place Archives in Exchange Server.

- **Permanently delete it from the Exchange mailbox database**: Exchange retains disconnected archive mailboxes based on the deleted mailbox retention settings configured for the mailbox database. The default retention period is 30 days. After the specified mailbox retention period, a disconnected archive mailbox is purged from the Exchange mailbox database.

  Like a disabled primary mailbox, you can permanently delete a disabled archive mailbox at any time by using the **Remove-StoreMailbox** cmdlet. For more information, see Permanently delete a mailbox.

Working with soft-deleted mailboxes

A soft-deleted mailbox is created when a mailbox is moved from one Exchange mailbox database to any other mailbox database. Exchange doesn't fully delete the mailbox from the source database after a move in case an error occurs during the move that causes the mailbox on the destination database to fail. You can always restore the source mailbox and try again. Exchange will retain the soft-deleted mailbox for the duration of the mailbox retention period.

You can perform two operations on a soft-deleted mailbox:

- Restore it to an existing mailbox.

- Permanently delete it from the Exchange mailbox database.

The procedures for restoring and permanently deleting a soft-deleted mailbox are similar to those for a disabled mailbox. For more information, see the following topics:

- Connect or restore a deleted mailbox
- Permanently delete a mailbox

Summary of working with disconnected mailboxes

The following table summarizes the information about disconnected mailboxes, including how the mailbox was disconnected, what happens to the corresponding Active Directory user account when a mailbox is disconnected, and the options and tools you have to connect or restore disconnected mailboxes.

HOW MAILBOX WAS DISABLED	VALUE OF *DISCONNECTREASON* PROPERTY	IS ACTIVE DIRECTORY USER ACCOUNT RETAINED?	CONNECT OR RESTORE OPTIONS	TOOLS
The EAC: **Recipients** > **Mailboxes** > **Disable** The Exchange Management Shell: **Disable-Mailbox** cmdlet	Disabled	Yes	Connect to same user account	The EAC: **Recipients** > **Mailboxes** > **Connect a Mailbox** The Exchange Management Shell: **Connect-Mailbox** cmdlet
The EAC: **Recipients** > **Mailboxes** > **Delete** The Exchange Management Shell: **Remove-Mailbox** cmdlet	Disabled	No	Connect to a different user account Restore to a different mailbox	The EAC: **Recipients** > **Mailboxes** > **Connect a Mailbox** The Exchange Management Shell: • **Connect-Mailbox** cmdlet • **Enable-Mailbox** cmdlet • **New-MailboxRestore** cmdlet
Moved to a different mailbox database	SoftDeleted	Yes	Connect to a different user account Restore to a different mailbox	The EAC: **Recipients** > **Mailboxes** > **Connect a Mailbox** The Exchange Management Shell: • **Connect-Mailbox** cmdlet • **Enable-Mailbox** cmdlet • **New-MailboxRestore** cmdlet

Disconnected mailbox documentation

The following table contains links to topics that will help you manage disconnected mailboxes. This includes managing disconnected user mailboxes, linked mailboxes, resource mailboxes, and shared mailboxes.

TOPIC	DESCRIPTION
Disable or delete a mailbox in Exchange Server	Learn how to disable or delete mailboxes.

TOPIC	DESCRIPTION
Connect a disabled mailbox	Learn how to connect a disabled mailbox to an existing user account.
Connect or restore a deleted mailbox	Learn how to connect a deleted mailbox to a user account or restore the contents of a deleted mailbox to an existing mailbox.
Manage Mailbox Restore Requests	Learn how to manage mailbox restore requests using the Exchange Management Shell.
Permanently delete a mailbox	Learn how to permanently delete a mailbox.

Disable or delete a mailbox in Exchange Server

8/3/2020 • 7 minutes to read • Edit Online

In Exchange Server, you can use the Exchange admin center (EAC) or the Exchange Management Shell to disable or delete mailboxes. Disabled or deleted mailboxes are also known as *disconnected mailboxes*. For more information about disconnected mailboxes, see Disconnected mailboxes.

Note: If you need to delete a mailbox in Microsoft 365 or Office 365, see Delete or Restore User Mailboxes in Exchange Online.

What do you need to know before you begin?

- Estimated time to complete each procedure: 2 minutes.
- For more information about accessing and using the EAC, see Exchange admin center in Exchange Server. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Recipient Provisioning Permissions" section in the Recipients Permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Disable mailboxes

When you disable a mailbox, all Exchange attributes are removed from the associated user account in Active Directory. The disconnected mailbox is hidden and marked for removal. The disconnected mailbox is permanently deleted (purged) based on the **MailboxRetention** property value for the mailbox database (the default value is 30 days). Before the mailbox is purged, you can reconnect it to a new or existing user account that doesn't already have an associated mailbox. For more information, see Connect a disabled mailbox.

Note: Disabling a mailbox that has an associated archive marks both the primary and archive mailboxes for removal. To only mark the archive mailbox for removal without affecting the primary mailbox, see Disable an archive mailbox.

Use the EAC to disable a mailbox

1. In the EAC, go to **Recipients**, and click the tab for the type of mailbox that you want to disable:
   - **Mailboxes** for user mailboxes and linked mailboxes.
   - **Shared** for shared mailboxes.
2. Find and select the mailbox that you want to disable. For example:
   - Scroll through the list. You can also click the column headers to sort the mailboxes.

  - Click **Search** and enter the text to filter the list of mailboxes.

  - Select multiple mailboxes by selecting a mailbox, holding the Shift key, and selecting a mailbox farther down in the list, or by holding down the CTRL key as you select each mailbox.

3. After you've selected the mailbox or mailboxes that you want to disable, click **More** •••, select **Disable**, and then click **Yes** in the warning message that appears.

Use the Exchange Management Shell to disable a mailbox

To disable a mailbox, use this syntax:

```
Disable-Mailbox <MailboxIdentity> [-Arbitration] [-Archive] [-PublicFolder] [-RemoteArchive]
```

This example disables the user mailbox that has the alias value danj.

```
Disable-Mailbox danj
```

This example disables the room mailbox named Conf Room 31/1234 (12).

```
Disable-Mailbox "Conf Room 31/1234 (12)"
```

This example disables the shared mailbox that has the email address sharedmbx@contoso.com.

```
Disable-Mailbox sharedmbx@contoso.com
```

For detailed syntax and parameter information, see Disable-Mailbox.

How do you know this worked?

To verify that you've successfully disabled a mailbox, do any of these steps:

- In the EAC, click **Recipients**, go to the appropriate tab for the type of mailbox that you disabled, and verify that the mailbox is no longer listed. Note that you might need to click **Refresh** 🔄.

- In Active Directory Users and Computers, right-click the user account whose mailbox you disabled, and then click **Properties**. On the **General** tab, verify that the **E-mail** field is blank.

- In the Exchange Management Shell, replace *<DisplayName>* with the user's display name, and run the following commands to verify the **DisconnectReason** property value is `Disabled` (which indicates the mailbox has been marked for removal):

  ```
  $dbs = Get-MailboxDatabase
  $dbs | foreach {Get-MailboxStatistics -Database $_.DistinguishedName} | where {$_.DisplayName -eq "
  <DisplayName>"} | Format-List DisconnectReason,DisconnectDate
  ```

  Notes:

  - The **DisconnectReason** property doesn't distinguish between disabled and deleted mailboxes (the value for both is `Disabled`). The presence of the associated user account indicates whether the mailbox was disabled.

    When you delete a mailbox, the value of the **DisconnectReason** property is also `Disabled`, but the corresponding Active Directory user account is also deleted.

  - If the command returns no results, replace *<DatabaseName>* with the name of the mailbox database where the disconnected mailbox resides, and run this command to synchronize the mailbox state for all disconnected mailboxes on the database:

```
Get-MailboxStatistics -Database "<DatabaseName>" | foreach {Update-StoreMailboxState -Database
$_.Database -Identity $_.MailboxGuid -Confirm:$false}
```

Then, run the previous command, which should now return results.

- In the Exchange Management Shell, replace *<UserIdentity>* with the name or user principal name of the user (for example, user@contoso.com), and run this command to verify that the **RecipientType** property value is `User`, not `UserMailbox`.

```
Get-User -Identity <UserIdentity>
```

Delete mailboxes

When you delete a mailbox, the mailbox is disconnected from the associated user account, and the account is removed from Active Directory. The disconnected mailbox is hidden and marked for removal. The disconnected mailbox is permanently deleted (purged) based on the **MailboxRetention** property value for the mailbox database (the default value is 30 days). Before the mailbox is purged, you can reconnect it to a new or existing user account that doesn't already have an associated mailbox. For more information, see Connect or restore a deleted mailbox.

Note: Deleting a mailbox that has an associated archive marks both the primary and archive mailboxes for removal. To only mark the archive mailbox for removal without affecting the primary mailbox, see Disable an archive mailbox.

Use the EAC to delete a mailbox

1. In the EAC, go to the location for the type of mailbox that you want to delete:
   - **Recipients** > **Mailboxes** for user mailboxes and linked mailboxes.
   - **Recipients** > **Resources** for room and equipment mailboxes.
   - **Recipients** > **Shared** for shared mailboxes.
   - **Public folders** > **Public folder mailboxes** for public folder mailboxes.
2. Find and select the mailbox that you want to disable. For example:
   - Scroll through the list. You can also click the column headers to sort the mailboxes.
   - Click **Search** and enter the text to filter the list of mailboxes.
   - Select multiple mailboxes by selecting a mailbox, holding the Shift key, and selecting a mailbox farther down in the list, or by holding down the CTRL key as you select each mailbox.
3. After you've selected the mailbox or mailboxes that you want to delete, click **Delete** 🗑, and then click **Yes** in the warning message that appears.

Use the Exchange Management Shell to delete a mailbox

To delete a mailbox, use this syntax:

```
Remove-Mailbox <MailboxIdentity> [-Arbitration] [-PublicFolder]
```

This example deletes the mailbox that has the email address pilarp@contoso.com.

```
Remove-Mailbox pilarp@contoso.com
```

This example deletes the equipment mailbox named Fleet Van (16).

```
Remove-Mailbox "Fleet Van (16)"
```

This example deletes the mailbox that has the alias value corpprint.

```
Remove-Mailbox corpprint
```

For detailed syntax and parameter information, see Remove-Mailbox.

Note: If you use the **Remove-Mailbox** cmdlet with the *Purge* switch, the mailbox is immediately purged and isn't recoverable. For more information, see Permanently delete a mailbox.

How do you know this worked?

To verify that you've successfully deleted a mailbox, do any of these steps:

- In the EAC, click **Recipients**, go to the appropriate tab for the type of mailbox that you deleted, and verify that the mailbox is no longer listed. Note that you might need to click **Refresh** 🔄.

- In Active Directory Users and Computers, verify that the associated account is no longer listed. Note that mailbox types other than user and linked mailboxes also have associated user accounts that are disabled (for example, room, equipment, arbitration, shared, and public folder mailboxes).

- In the Exchange Management Shell replace *<DisplayName>* with the user's display name, and run the following commands to verify the **DisconnectReason** property value is `Disabled` (which indicates the mailbox has been marked for removal):

  ```
  $dbs = Get-MailboxDatabase
  $dbs | foreach {Get-MailboxStatistics -Database $_.DistinguishedName} | where {$_.DisplayName -eq "
  <DisplayName>"} | Format-List DisconnectReason,DisconnectDate
  ```

  Notes:

  - The **DisconnectReason** property doesn't distinguish between disabled and deleted mailboxes (the value for both is `Disabled`). The absence of the associated user account indicates whether the mailbox was deleted.

  - If the command returns no results, replace *<DatabaseName>* with the name of the mailbox database where the disconnected mailbox resides, and run the following command to synchronize the mailbox state for all disconnected mailboxes on the database:

    ```
    Get-MailboxStatistics -Database "<DatabaseName>" | foreach {Update-StoreMailboxState -Database
    $_.Database -Identity $_.MailboxGuid -Confirm:$false}
    ```

    Then, run the previous command, which should now return results.

- In the Exchange Management Shell, replace *<UserIdentity>* with the name or user principal name of the user (for example, user@contoso.com), and run this command to verify that the user can't be found.

  ```
  Get-User <UserIdentity>
  ```

More information

When you delete the Active Directory user account that's associated with a mailbox, Exchange will detect that the mailbox is no longer connected to a user account, and will mark the mailbox for removal, even if the mailbox has been placed on Litigation Hold or In-Place Hold. To retain the mailbox, do these steps:

- Instead of deleting the user account, disable the user account.
- Change the properties of the mailbox to restrict its use and who has access to the mailbox. For example, set send and receive quotas equal to 1, block who can send messages to the mailbox, and restrict who has access the mailbox.
- Retain the mailbox until all data has been expunged, or until preserving the data is no longer required.

For more information, see In-Place Hold and Litigation Hold in Exchange Server.

Connect a disabled mailbox

8/3/2020 • 5 minutes to read • Edit Online

When you disable a mailbox, Exchange retains the mailbox in the mailbox database and switches the mailbox to a disabled state. The Exchange attributes are also removed from the corresponding Active Directory user account, but the user account is retained. The mailbox is retained until the deleted mailbox retention period expires, which is 30 days by default, before it's then deleted permanently (or *purged*) from the mailbox database.

Until a disabled mailbox is permanently deleted from the Exchange mailbox database, you can use the EAC or the Exchange Management Shell to reconnect it to the original Active Directory user account.

To learn more about disconnected mailboxes and perform other related management tasks, see the following topics:

- Disconnected mailboxes
- Disable or delete a mailbox in Exchange Server
- Connect or restore a deleted mailbox
- Permanently delete a mailbox

What do you need to know before you begin?

- Estimated time to complete: 2 minutes.
- To open the EAC, see Exchange admin center in Exchange Server. To open the Exchange Management Shell, see Open the Exchange Management Shell.
- Run the **Get-User** cmdlet in theExchange Management Shell to verify that the Active Directory user account that you want to connect the disabled mailbox to exists and that it isn't already associated with another mailbox. To connect a disabled mailbox to a user account, the account must exist and the value for the *RecipientType* property has to be `User`, which indicates that the account isn't already mailbox-enabled.

  You can also verify this information in Active Directory Users and Computers.

- Replace *<DisplayName>* with the display name of the mailbox, and run the following commands in the Exchange Management Shell to verify that the disabled mailbox that you want to connect to a user account exists and isn't a soft-deleted mailbox.

  ```
  $dbs = Get-MailboxDatabase
  $dbs | foreach {Get-MailboxStatistics -Database $_.DistinguishedName} | where {$_.DisplayName -eq "
  <DisplayName>"} | Format-List DisplayName,Database,DisconnectReason
  ```

  To be able to connect a disabled mailbox, the mailbox has to exist in the mailbox database and the value for the *DisconnectReason* property has to be `Disabled`. If the mailbox has been purged from the database, the command won't return any results.

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Recipient Provisioning Permissions" section in the Recipients Permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to connect a disabled mailbox

The following procedure shows how to connect a disabled user mailbox. You can also reconnect disabled linked mailboxes and disabled shared mailboxes to the corresponding user account.

1. In the EAC, navigate to **Recipients** > **Mailboxes**.
2. Click **More** •••, and then click **Connect a mailbox**.

   A list of mailboxes that are disconnected on the selected Exchange server in your Exchange organization will be displayed.

   > **NOTE**
   >
   > This list of disconnected mailboxes includes disabled mailboxes, deleted mailboxes, and soft-deleted mailboxes.

3. Click the disabled mailbox that you want to reconnect, and then click **Connect**.
4. In the window that asks if you're sure that you want to reconnect the mailbox, click **Yes**.

   Exchange will reconnect the disabled mailbox to the corresponding user account.

Use the Exchange Management Shell to connect a disabled mailbox or personal archive

Use the **Connect-Mailbox** cmdlet in the Exchange Management Shell to connect a user account to a disabled mailbox. You have to specify the type of mailbox that you're connecting. The following examples show the syntax for reconnecting user, linked, shared, and archive mailboxes.

This example connects a user mailbox. The *Identity* parameter specifies the disconnected mailbox in the Exchange database. The *User* parameter specifies the Active Directory user account to reconnect the mailbox to.

```
Connect-Mailbox -Identity "Jeffrey Zeng" -Database MBXDB01 -User "Jeffrey Zeng"
```

This example connects a linked mailbox. The *Identity* parameter specifies the disconnected mailbox in the Exchange database. The *LinkedMasterAccount* parameter specifies the Active Directory user account in the account forest that you want to reconnect the mailbox to. The *Alias* parameter specifies the alias, which is the portion of the email address on the left side of the at (@) symbol, for the reconnected mailbox.

```
Connect-Mailbox -Identity "Kai Axford" -Database MBXDB02 -LinkedDomainController FabrikamDC01 -
LinkedMasterAccount kai.axford@fabrikam.com -Alias kaia
```

This example connects a shared mailbox.

```
Connect-Mailbox -Identity "Corporate Shared Mailbox" -Database "Mailbox Database 03" -User "Corporate Shared
Mailbox" -Alias corpshared -Shared
```

> **NOTE**
>
> If you don't include the *Alias* parameter when you run the **Connect-Mailbox** cmdlet, the value specified in the *User* or *LinkedMasterAccount* parameter is used to create the email address alias for the reconnected mailbox.

This example connects a personal archive to the primary mailbox using the mailbox GUID stored in mailbox database DB01.

```
Connect-Mailbox -Identity "95352f8b-e5aa-496f-ac7f-ce93357d7b0c" -Archive -User "Megan Bown" -Database "DB01"
```

If you do not know the name of the personal archive, you can view it in the Exchange Management Shell by running the following command. This example returns all personal archive mailboxes in mailbox database DB01.

```
Get-MailboxDatabase "DB01" | Get-MailboxStatistics | Where {($_.DisconnectDate -ne $null) -and
($_.IsArchiveMailbox -eq $true)} | Format-Table DisplayName,MailboxGuid -AutoSize
```

> **NOTE**
>
> You can connect a personal archive mailbox to any primary mailbox you wish, even if it is not the original owner's mailbox. Use the *AllowLegacyDNMismatch* parameter to allow the connection of the archive mailbox to a different primary mailbox.

For detailed syntax and parameter information, see Connect-Mailbox.

How do you know this worked?

To verify that you've successfully connected a disabled mailbox to a user account, do one of the following:

- In the EAC, click **Recipients**, navigate to the appropriate page for the mailbox type that you reconnected, click **Refresh** ⟳, and verify that the mailbox is listed.

- In Active Directory Users and Computers, right-click the user account whose mailbox you disabled, and then click **Properties**. On the **General** tab, notice that the **E-mail** box is populated with the email address for the reconnected mailbox.

- In theExchange Management Shell,replace <Identity> with the name of the user account and run the following command:

  ```
  Get-User "<Identity>"
  ```

  The **UserMailbox** value for the *RecipientType* property indicates that the user account and the mailbox are connected. You can also run the **Get-Mailbox** cmdlet to verify that the mailbox exists.

Connect or restore a deleted mailbox

8/3/2020 • 7 minutes to read • Edit Online

When you delete a mailbox, Exchange retains the mailbox in the mailbox database and switches the mailbox to a disabled state. The associated Active Directory user account is also deleted. The mailbox is retained until the deleted mailbox retention period expires, which is 30 days by default, and then it's permanently deleted (or *purged*) from the mailbox database.

Until a deleted mailbox is permanently deleted from the Exchange mailbox database, you can use the EAC or the Exchange Management Shell to connect it to an Active Directory user account. You can also use the Exchange Management Shell to restore the contents of the deleted mailbox to an existing mailbox.

To learn more about disconnected mailboxes and perform other related management tasks, see the following topics:

- Disconnected mailboxes
- Disable or delete a mailbox in Exchange Server
- Connect a disabled mailbox
- Permanently delete a mailbox

What do you need to know before you begin?

- Estimated time to complete: 2 minutes.
- To open the EAC, see Exchange admin center in Exchange Server. To open the Exchange Management Shell, see Open the Exchange Management Shell.
- Create a new user account in Active Directory to connect the deleted mailbox to. Or use the **Get-User** cmdlet in the Exchange Management Shell to verify that the Active Directory user account that you want to connect the deleted mailbox to exists and that it isn't already associated with another mailbox. To connect a deleted mailbox to a user account, the account must exist and the value for the *RecipientType* property has to be `User`, which indicates that the account isn't already mailbox-enabled.

  For on-premises Exchange organizations, you can also verify this information in Active Directory Users and Computers.

  > **IMPORTANT**
  > When you connect deleted linked mailboxes, resource mailboxes, or shared mailboxes, the Active Directory user account that you're connecting the mailbox to must be disabled.

- To verify that the deleted mailbox that you want to connect a user account to exists in the mailbox database and isn't a soft-deleted mailbox, run the following command:

  ```
  Get-MailboxDatabase | foreach {Get-MailboxStatistics -Database $_.name} | where {$_.DisplayName -eq "
  <display name>"} | Format-List DisplayName,Database,DisconnectReason
  ```

  The deleted mailbox has to exist in the mailbox database and the value for the *DisconnectReason* property has to be `Disabled`. If the mailbox has been purged from the database, the command won't return any results.

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Recipient Provisioning Permissions" section in the Recipients Permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.
- Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Connect a deleted mailbox

When you connect a deleted mailbox, you associate the mailbox with a user account that isn't mail-enabled, which means that it doesn't have an existing mailbox. To connect a deleted mailbox to a user account that has a mailbox, you have to restore the deleted mailbox. For more information, see Restore a deleted mailbox later in this topic.

Use the EAC to connect a deleted mailbox

The following procedure shows how to connect a deleted user mailbox to a user account. You can also use this procedure to connect linked mailboxes, resource mailboxes, and shared mailboxes that have been deleted to a user account.

1. In the EAC, go to **Recipients** > **Mailboxes**.
2. Click **More** •••, and then click **Connect a mailbox**.

   A list of mailboxes that are disconnected on the selected Exchange server in your Exchange organization will be displayed.

   > **NOTE**
   >
   > This list of disconnected mailboxes includes disabled mailboxes, deleted mailboxes, and soft-deleted mailboxes.

3. Click the deleted mailbox that you want to connect a user to, and then click **Connect**.
4. In the window that asks if you're sure that you want to connect the mailbox, click **Yes**.

   A list of user accounts that aren't mail-enabled is displayed.

5. Click the user that you want to connect the deleted mailbox to, and then click **OK**.

   Exchange will connect the deleted mailbox to the user account that you selected.

Use the Exchange Management Shell to connect a deleted mailbox

Use the **Connect-Mailbox** cmdlet in the Exchange Management Shell to connect a deleted mailbox to a user account that isn't mail enabled. You have to specify the type of mailbox that you're connecting. The following examples show the syntax for reconnecting user, linked, room, equipment, and shared mailboxes. In all examples, the optional *Alias* parameter is used to specify the email alias, which is the portion of the email address on the left side of the at (@) symbol. If you don't include the *Alias* parameter, the value specified in the *User* or *LinkedMasterAccount* parameter is used to create the alias for the email address for the reconnected mailbox.

> **NOTE**
>
> As previously stated, when you connect linked, resource, or shared mailboxes, the Active Directory user account that you're linking the mailbox to must be disabled.

This example connects a deleted user mailbox to a user account that isn't mail enabled. The *Identity* parameter specifies the display name of the deleted mailbox retained in the mailbox database named MBXDB01. The *User* parameter specifies the Active Directory user account to connect the mailbox to.

```
Connect-Mailbox -Identity "Paul Cannon" -Database MBXDB01 -User "Robin Wood" -Alias robinw
```

> **NOTE**
>
> You can also use the values for the `LegacyDN` or `MailboxGuid` properties to identify the deleted mailbox.

This example connects a linked mailbox. The *Identity* parameter specifies the deleted mailbox on the mailbox database named MBXDB02. The *LinkedMasterAccount* parameter specifies the Active Directory user account in the account forest that you want to connect the mailbox to. The *LinkedDomainController* parameter specifies a domain controller in the account forest.

```
Connect-Mailbox -Identity "Temp User" -Database MBXDB02 -LinkedDomainController FabrikamDC01 -
LinkedMasterAccount danpark@fabrikam.com -Alias dpark
```

This example connects a room mailbox.

```
Connect-Mailbox -Identity "rm2121" -Database "MBXResourceDB" -User "Conference Room 2121" -Alias ConfRm2121 -
Room
```

This example connects an equipment mailbox.

```
Connect-Mailbox -Identity "MotorPool01" -Database "MBXResourceDB" -User "Van01 (12 passengers)" -Alias van01
-Equipment
```

This example connects a shared mailbox.

```
Connect-Mailbox -Identity "Printer Support" -Database MBXDB01 -User "Corp Printer Support" -Alias corpprint -
Shared
```

> **NOTE**
>
> You can also use the `LegacyDN` or `MailboxGuid` values to identify the deleted mailbox.

For detailed syntax and parameter information, see Connect-Mailbox.

How do you know this worked?

To verify that you've successfully connected a deleted mailbox to a user account, do one of the following steps:

- In the EAC, click **Recipients**, go to the appropriate page for the mailbox type that you connected, click **Refresh** ↻, and verify that the mailbox is listed.

- In Active Directory Users and Computers, right-click the user account that you connected to the mailbox, and then click **Properties**. On the **General** tab, notice that the **E-mail** box is populated with the email address for the connected mailbox.

- In the Exchange Management Shell, run the following command.

```
Get-User <identity>
```

The **UserMailbox** value for the *RecipientType* property indicates that the user account and the mailbox are connected. You can also run the **Get-Mailbox <identity>** command to verify that the mailbox was connected.

Restore a deleted mailbox

You can use the Exchange Management Shell to restore a deleted mailbox to an existing mailbox using the **New-MailboxRestoreRequest** cmdlet. When you restore a deleted mailbox, its contents are copied to an existing mailbox, which is referred to as the *target mailbox*. After a deleted mailbox is restored, it's still retained in the mailbox database until it's permanently deleted by an administrator or purged after the deleted mailbox retention period expires.

After a mailbox restore request is successfully completed, it's retained for 30 days, by default, before it's removed. You can remove the mailbox sooner by using the **Remove-StoreMailbox** cmdlet.

> **NOTE**
>
> You can't use the EAC to restore a deleted mailbox.

Use the Exchange Management Shell to restore a deleted mailbox

To create a mailbox restore request, you have to use the display name, legacy distinguished name (DN), or mailbox GUID of the deleted mailbox. Use the **Get-MailboxStatistics** cmdlet to display the values of the `DisplayName`, `MailboxGuid`, and `LegacyDN` properties for the deleted mailbox that you want to restore. For example, run the following commands to return this information for all disabled and deleted mailboxes in your organization.

```
$dbs = Get-MailboxDatabase
$dbs | foreach {Get-MailboxStatistics -Database $_.DistinguishedName} | where {$_.DisconnectReason -eq
"Disabled"} | Format-Table DisplayName,MailboxGuid,Database,DisconnectDate
```

This example restores the deleted mailbox, which is identified by the *SourceStoreMailbox* parameter and is located on the MBXDB01 mailbox database, to the target mailbox Debra Garcia. The *AllowLegacyDNMismatch* parameter is used so the source mailbox can be restored to a different mailbox, one that doesn't have the same legacy DN value.

```
New-MailboxRestoreRequest -SourceStoreMailbox e4890ee7-79a2-4f94-9569-91e61eac372b -SourceDatabase MBXDB01 -
TargetMailbox "Debra Garcia" -AllowLegacyDNMismatch
```

This example restores Pilar Pinilla's deleted archive mailbox to her current archive mailbox. The *AllowLegacyDNMismatch* parameter isn't necessary because a primary mailbox and its corresponding archive mailbox have the same legacy DN.

```
New-MailboxRestoreRequest -SourceStoreMailbox "Personal Archive - Pilar Pinilla" -SourceDatabase "MDB01" -
TargetMailbox pilarp@contoso.com -TargetIsArchive
```

For detailed syntax and parameter information, see New-MailboxRestoreRequest.

How do you know this worked?

To verify that you've successfully restored a deleted mailbox to the target mailbox, run the **Get-**

MailboxRestoreRequest cmdlet to display information about the restore request. If the restore request was successfully created, the *Status* property will have a value of `Queued`, `InProgress`, or `Completed`. After the restore request is completed, the contents from the deleted mailbox will appear in the target mailbox.

For more information, see:

- Manage Mailbox Restore Requests
- Get-MailboxRestoreRequest
- Get-MailboxRestoreRequestStatistics

Permanently delete a mailbox

8/3/2020 • 3 minutes to read • Edit Online

When you permanently delete active mailboxes and disconnected mailboxes, all mailbox contents are purged from the Exchange mailbox database, and the data loss is permanent. When you permanently delete an active mailbox, the associated Active Directory user account is also deleted.

An alternative to permanently deleting a mailbox is to disconnect it. After you disconnect a mailbox, by default, Exchange retains the data in the mailbox database for 30 days. This gives you the opportunity to reconnect or restore a mailbox before it's purged from the database.

To learn more about disconnected mailboxes and perform other related management tasks in Exchange, see the following topics:

- Disconnected mailboxes
- Disable or delete a mailbox in Exchange Server
- Connect a disabled mailbox
- Connect or restore a deleted mailbox

> **NOTE**
>
> You can't use the Exchange admin center (EAC) to permanently delete an active mailbox or a disconnected mailbox.

What do you need to know before you begin?

- Estimated time to complete: 2 minutes.
- The procedures in this topic require the Exchange Management Shell. For more information, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Recipient Provisioning Permissions" section in the Recipients Permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to permanently delete an active mailbox

If you don't include the *Permanent* parameter when you delete a mailbox, the deleted mailbox is retained in the mailbox database for 30 days (by default) before it's permanently deleted.

Run the following command to permanently delete an active mailbox and the associated Active Directory user account:

```
Remove-Mailbox -Identity <Identity> -Permanent $true
```

For detailed syntax and parameter information, see Remove-Mailbox.

How do you know this worked?

To verify that you've permanently deleted an active mailbox, do the following:

1. Verify that the mailbox is no longer listed in the Exchange admin center (EAC).

2. Verify that the associated user account is no longer listed in Active Directory Users and Computers.

3. Replace *<DisplayName>* with the display name of the mailbox and run the following commands in the Exchange Management Shell to verify that the mailbox was successfully purged from the Exchange mailbox database:

   ```
   $dbs = Get-MailboxDatabase
   $dbs | foreach {Get-MailboxStatistics -Database $_.DistinguishedName} | where {$_.DisplayName -eq "
   <DisplayName>"}
   ```

   If you successfully purged the mailbox, the command won't return any results. If the mailbox wasn't purged, the command will return information about the mailbox.

Use the Exchange Management Shell to find the disconnected mailbox type

A disconnected mailbox can be either disabled or soft-deleted. You need to specify the correct type to permanently delete a disconnected mailbox. If you don't, the command will fail.

Replace *<DisplayName>* with the display name of the mailbox and run the following command to determine whether a disconnected mailbox is disabled or soft-deleted:

```
$dbs = Get-MailboxDatabase
$dbs | foreach {Get-MailboxStatistics -Database $_.DistinguishedName} | where {$_.DisplayName -eq "
<DisplayName>"} | Format-List DisplayName,MailboxGuid,Database,DisconnectReason
```

The value for the *DisconnectReason* property will be either `Disabled` or `SoftDeleted`.

You can run the following commands to display the type for all disconnected mailboxes in your organization:

```
$dbs = Get-MailboxDatabase
$dbs | foreach {Get-MailboxStatistics -Database $_.DistinguishedName} | where {$_.DisconnectReason -ne $null}
| Format-List DisplayName,MailboxGuid,Database,DisconnectReason
```

Use the Exchange Management Shell to permanently delete a disconnected mailbox

Caution

When you use the **Remove-StoreMailbox** cmdlet to permanently delete a disconnected mailbox, all its contents are purged from the mailbox database and the data loss is permanent.

This example permanently deletes the disabled mailbox with the GUID 2ab32ce3-fae1-4402-9489-c67e3ae173d3 from mailbox database named MBD01.

```
Remove-StoreMailbox -Database MBD01 -Identity "2ab32ce3-fae1-4402-9489-c67e3ae173d3" -MailboxState Disabled
```

This example permanently deletes the soft-deleted mailbox for Dan Jump from mailbox database named MBD01.

```
Remove-StoreMailbox -Database MBD01 -Identity "Dan Jump" -MailboxState SoftDeleted
```

This example permanently deletes all soft-deleted mailboxes from mailbox database named MBD01.

```
Get-MailboxStatistics -Database MBD01 | where {$_.DisconnectReason -eq "SoftDeleted"} | ForEach {Remove-
StoreMailbox -Database $_.Database -Identity $_.MailboxGuid -MailboxState SoftDeleted}
```

For detailed syntax and parameter information, see Remove-StoreMailbox and Get-MailboxStatistics.

How do you know this worked?

To verify that you've permanently deleted a disconnected mailbox and that it was successfully purged from the mailbox database, replace *<DisplayName>* with the display name of the mailbox and run the following command:

```
$dbs = Get-MailboxDatabase
$dbs | foreach {Get-MailboxStatistics -Database $_.DistinguishedName} | where {$_.DisplayName -eq "
<DisplayName>"}
```

If you successfully purged the mailbox, the command won't return any results. If the mailbox wasn't purged, the command will return information about the mailbox.

Custom attributes

8/3/2020 • 3 minutes to read • Edit Online

Exchange Server includes 15 extension attributes that you can use to add information about a recipient, such as an employee ID, organizational unit (OU), or some other custom value for which there isn't an existing attribute.

In earlier versions of Exchange, if you wanted to store this information in Active Directory, you had to create an attribute by extending the Active Directory schema. Schema extension requires planning, procuring object identifiers (OIDs) for new attributes, and testing the extension process in a test environment before you implement it in a production environment. Exchange Server doesn't let you use schema extensions in recipient filters that are used by address lists, e-mail address policies, and dynamic distribution groups.

The custom attributes available to Exchange Server are labeled in Active Directory as **ms-Exch-Extension-Attribute1** through **ms-Exch-Extension-Attribute15**. In the Exchange Management Shell, the corresponding parameters are *CustomAttribute1* through *CustomAttribute15*. These attributes aren't used by any Exchange components. They can be used to store Active Directory data without having to extend the Active Directory schema.

> **NOTE**
>
> **ms-Exch-Extension-Attribute-16** to **ms-Exch-Extension-Attribute-45** are present in Active Directory, but aren't available in the Exchange admin center (EAC) or the Exchange Management Shell. Don't use non-Exchange tools to edit these attributes because they might be used for future Exchange features.

Advantages of custom attributes

There are several advantages to using custom attributes:

- You avoid extending the Active Directory schema.
- You don't have to do the work, because the attributes are created by Exchange Setup.
- You can use the EAC or the Exchange Management Shell to manage the attributes. You don't need to build custom controls or write scripts to populate and display these attributes.
- You can filter and reuse the attributes, as attributes are filterable properties that can be used in the *Filter* parameter with recipient cmdlets such as **Get-Mailbox**. They can also be used in the EAC and the Exchange Management Shell to create filters for e-mail address policies, address lists, and dynamic distribution groups.

Multivalued custom attributes

Starting with Exchange 2010 Service Pack 2 (SP2), five multivalued custom attributes were added to Exchange to allow you to store additional information for mail recipients if the traditional custom attributes didn't meet your needs. The *ExtensionCustomAttribute1* to *ExtensionCustomAttribute5* parameters can hold up to 1,300 values each. You can specify multiple values as a comma-delimited list. The following cmdlets support these new parameters:

- Set-DistributionGroup
- Set-DynamicDistributionGroup
- Set-Mailbox
- Set-MailContact

- Set-MailPublicFolder
- Set-RemoteMailbox

For more information about multivalued properties, see Modifying multivalued properties.

Custom attribute examples

A common scenario in many Exchange deployments is that of creating an e-mail address policy for all recipients in an OU. The OU isn't a filterable property that can be used in the *RecipientFilter* parameter of an e-mail address policy or an address list.

> **NOTE**
>
> Dynamic distribution groups have an additional parameter that you can use to restrict it to recipients in a particular OU or container.

If the recipients in a particular OU don't share any common properties that you can filter by, such as department or location, you can populate one of the custom attributes with a common value, as shown in this example.

```
Get-Mailbox -OrganizationalUnit Sales | Set-Mailbox CustomAttribute1 "SalesOU"
```

With that done, now you can create an e-mail address policy for all recipients that have the *CustomAttribute1* property that equals SalesOU, as shown in this example.

```
New-EmailAddressPolicy -Name "Sales" -RecipientFilter "CustomAttribute1 -eq 'SalesOU'" -
EnabledEmailAddressTemplates "SMTP:%s%2g@sales.contoso.com"
```

Custom attribute example using the ConditionalCustomAttributes parameter

When creating dynamic distribution groups, email address policies, or address lists, you don't need to use the *RecipeintFilter* parameter to specify custom attributes. You can use the *ConditionalCustomAttribute1* to *ConditionalCustomAttribute15* parameters instead.

This example creates a dynamic distribution group based on the recipients whose *CustomAttribute1* is set to SalesOU.

```
New-DynamicDistributionGroup -Name "Sales Users and Contacts" -IncludedRecipients "MailboxUsers,MailContacts"
-ConditionalCustomAttribute1 "SalesOU"
```

> **NOTE**
>
> You need to use the *IncludedRecipients* parameter if you use a *Conditional* parameter. In addition, you can't use *Conditional* parameters if you use the *RecipientFilter* parameter. If you want to include additional filters to create your dynamic distribution group, email address policies, or address lists, you should use the *RecipientFilter* parameter.

Custom attribute example using ExtensionCustomAttributes parameter

In this example, the mailbox for Kweku will have *ExtensionCustomAttribute1* updated to reflect that he's enrolled in the following educational classes: MATH307, ECON202, and ENGL300.

```
Set-Mailbox -Identity Kweku -ExtensionCustomAttribute1 MATH307,ECON202,ENGL300
```

Next, a dynamic distribution group for all students enrolled MATH307 is created by using the *RecipientFilter* parameter where *ExtensionCustomAttribute1* is equal to MATH307. When using the *ExtentionCustomAttributes* parameters, you can use the `-eq` operator instead of the `-like` operator.

```
New-DynamicDistributionGroup -Name Students_MATH307 -RecipientFilter "ExtensionCustomAttribute1 -eq 'MATH307'"
```

In this example, Kweku's *ExtensionCustomAttribute1* values are updated to reflect that he's added the class ENGL210 and removed the class ECON202.

```
Set-Mailbox -Identity Kweku -ExtensionCustomAttribute1 @{Add="ENGL210"; Remove="ECON202"}
```

Manage permissions for recipients

8/3/2020 • 11 minutes to read • Edit Online

In Exchange Server, you can use the Exchange admin center (EAC) or the Exchange Management Shell to assign permissions to a mailbox or group so that other users can access the mailbox (the Full Access permission), or send email messages that appear to come from the mailbox or group (the Send As or Send on Behalf permissions). The users that are assigned these permissions on other mailboxes or groups are called *delegates*.

The permissions that you can assign to delegates for mailboxes and groups in Exchange Server are described in the following table:

Note: Although you can use the Exchange Management Shell to assign some or all of these permissions to other delegate types on other kinds of recipient objects, this topic focuses on the delegate and recipient object types that produce useful results.

PERMISSION	DESCRIPTION	RECIPIENT TYPES IN THE EAC	ADDITIONAL RECIPIENT TYPES IN POWERSHELL	DELEGATE TYPES IN THE EAC	ADDITIONAL DELEGATE TYPES IN THE POWERSHELL
Full Access	Allows the delegate to open the mailbox, and view, add and remove the contents of the mailbox. Doesn't allow the delegate to send messages from the mailbox. If you assign the Full Access permission to a mailbox that's hidden from address lists, the delegate won't be able to open the mailbox. By default, arbitration and discovery mailboxes are hidden from address lists. By default, the mailbox auto-mapping feature uses Autodiscover to automatically open the mailbox in the delegate's Outlook profile (in addition to their own	User mailboxes Linked mailboxes Resource mailboxes Shared mailboxes	Arbitration mailboxes Discovery mailboxes	Mailboxes with user accounts Mail users with accounts Mail-enabled security groups	User accounts that aren't mail-enabled. Universal, global, and domain local security groups that aren't mail-enabled.

PERMISSION	DESCRIPTION	RECIPIENT TYPES IN THE EAC	ADDITIONAL RECIPIENT TYPES IN POWERSHELL	DELEGATE TYPES IN THE EAC	ADDITIONAL DELEGATE TYPES IN THE POWERSHELL
	their own mailbox). Note that auto-mapping will only work for individual users granted the proper permissions and will not work for any kind of group. If you don't want mailboxes to be auto-mapped, you need to take one of the following actions: • Use the **Add-MailboxPermission** cmdlet in the Exchange Management Shell to assign the Full Access permission with the `-AutoMapping $false` setting. For more information, see the Use the Exchange Management Shell to assign the Full Access permission to mailboxes section in this topic. • Assign the Full Access permission to a (mail-enabled) security group. The mailbox won't open in the Outlook profile of each member.				

PERMISSION	DESCRIPTION	RECIPIENT TYPES IN THE EAC	ADDITIONAL RECIPIENT TYPES IN POWERSHELL	DELEGATE TYPES IN THE EAC	ADDITIONAL DELEGATE TYPES IN THE POWERSHELL
Send As	Allows the delegate to send messages as if they came directly from the mailbox or group. There's no indication that the message was sent by the delegate. Doesn't allow the delegate to read the contents of the mailbox. If you assign the Send As permission to a mailbox that's hidden from address lists, the delegate won't be able to send messages from the mailbox.	User mailboxes Linked mailboxes Resource mailboxes Shared mailboxes Distribution groups Dynamic distribution groups Mail-enabled security groups	n/a	Mailboxes with user accounts Mail users with accounts Mail-enabled security groups	n/a

PERMISSION	DESCRIPTION	RECIPIENT TYPES IN THE EAC	ADDITIONAL RECIPIENT TYPES IN POWERSHELL	DELEGATE TYPES IN THE EAC	ADDITIONAL DELEGATE TYPES IN THE POWERSHELL
Send on Behalf	Allows the delegate to send messages from the mailbox or group. The From address of these messages clearly shows that the message was sent by the delegate ("*<Delegate>* on behalf of *<MailboxOrGroup>*"). However, replies to these messages are sent to the mailbox or group, not to the delegate. Doesn't allow the delegate to read the contents of the mailbox. If you assign the Send on Behalf permission to a mailbox that's hidden from address lists, the delegate won't be able to send messages from the mailbox.	User mailboxes Linked mailboxes Resource mailboxes Distribution groups Dynamic distribution groups Mail-enabled security groups	Shared mailboxes	Mailboxes with user accounts Mail users with accounts Mail-enabled security groups Distribution groups	n/a

NOTE

If a user has both Send As and Send on Behalf permissions to a mailbox or group, the Send As permission is always used.|User mailboxes

What do you need to know before you begin?

- Estimated time to complete each procedure: 2 minutes.
- You need to be assigned permissions before you can perform the procedures in this topic. To see what permissions you need, see the "Recipient provisioning permissions" entry in the Recipients Permissions topic.
- To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- Procedures in this topic require specific permissions. See each procedure for its permissions information.

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Use the EAC to assign permissions to individual mailboxes

1. In the EAC, click **Recipients** in the feature pane. Depending on the type of mailbox that you want to assign permissions for, click on one of the following tabs:

   - **Mailboxes**: User or linked mailboxes.

   - **Resources**: Room or equipment mailboxes.

   - **Shared**: Shared mailboxes.

2. In the list of mailboxes, select the mailbox that you want to assign permissions for, and then click **Edit** ✎.

3. On the mailbox properties page that opens, click **Mailbox delegation** and configure one or more of the following permissions:

   - **Send As**: Messages sent by a delegate appear to come from the mailbox.

   - **Send on Behalf**: Messages sent by a delegate have " *<Delegate>* on behalf of *<Mailbox>*" in the From address. Note that this permission isn't available in the EAC for shared mailboxes.

   - **Full Access**: The delegate can open the mailbox and do anything except send messages.

   To assign permissions to delegates, click **Add** + under the appropriate permission. A dialog box appears that lists the users or groups that can have the permission assigned to them. Select the user or group from the list, and then click **Add**. Repeat this process as many times as necessary. You can also search for users or groups in the search box by typing all or part of the name, and then clicking **Search** 🔍. When you're finished selecting delegates, click **OK**.

   To remove a permission from a delegate, select the delegate in the list under the appropriate permission, and then click **Remove** −.

4. When you're finished, click **Save**.

Use the EAC to assign permissions to multiple mailboxes at the same time

1. In the EAC, navigate to **Recipients** > **Mailboxes**.

2. Select the mailboxes that you want to assign permissions for. Use click + Shift key + click to select a range of mailboxes, or Ctrl key + click to select multiple individual mailboxes. The title of the details pane changes to **Bulk Edit** as shown in the following diagram.

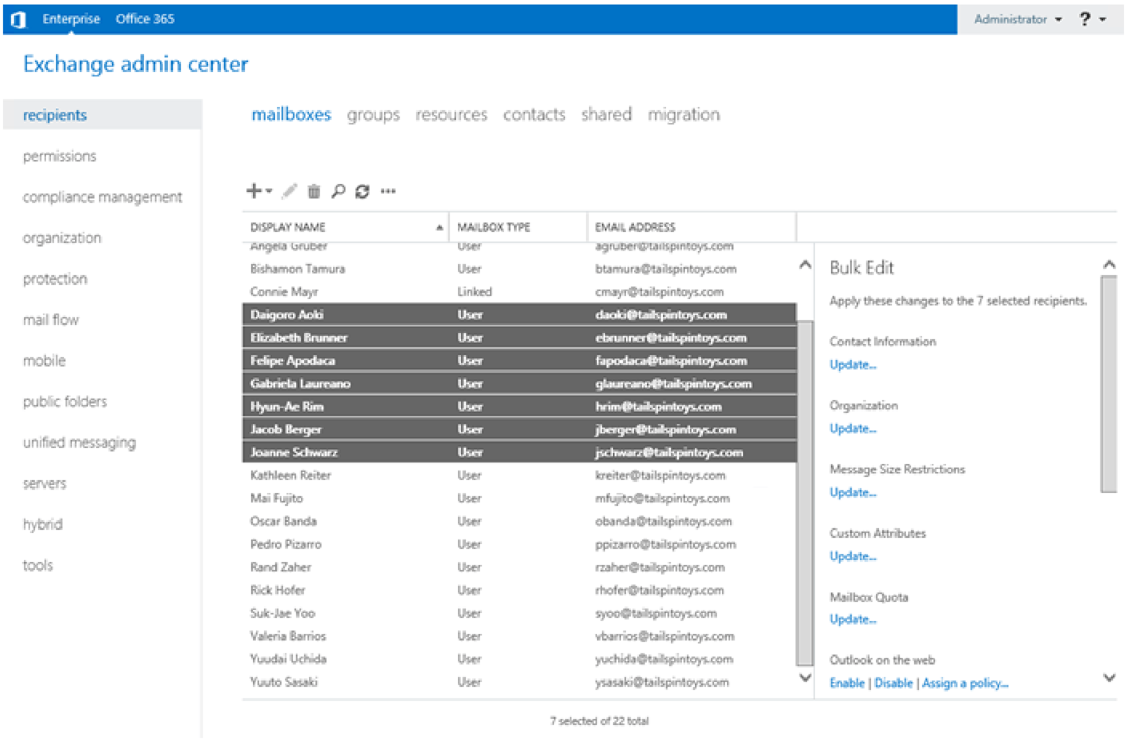

DISPLAY NAME	MAILBOX TYPE	EMAIL ADDRESS
Angela Gruber	User	agruber@tailspintoys.com
Bishamon Tamura	User	btamura@tailspintoys.com
Connie Mayr	Linked	cmayr@tailspintoys.com
Daigoro Aoki	User	daoki@tailspintoys.com
Elizabeth Brunner	User	ebrunner@tailspintoys.com
Felipe Apodaca	User	fapodaca@tailspintoys.com
Gabriela Laureano	User	glaureano@tailspintoys.com
Hyun-Ae Rim	User	hrim@tailspintoys.com
Jacob Berger	User	jberger@tailspintoys.com
Joanne Schwarz	User	jschwarz@tailspintoys.com
Kathleen Reiter	User	kreiter@tailspintoys.com
Mai Fujito	User	mfujito@tailspintoys.com
Oscar Banda	User	obanda@tailspintoys.com
Pedro Pizarro	User	ppizarro@tailspintoys.com
Rand Zaher	User	rzaher@tailspintoys.com
Rick Hofer	User	rhofer@tailspintoys.com
Suk-Jae Yoo	User	syoo@tailspintoys.com
Valeria Barrios	User	vbarrios@tailspintoys.com
Yuudai Uchida	User	yuchida@tailspintoys.com
Yuuto Sasaki	User	ysasaki@tailspintoys.com

7 selected of 22 total

Bulk Edit

Apply these changes to the 7 selected recipients.

Contact Information
Update...

Organization
Update...

Message Size Restrictions
Update...

Custom Attributes
Update...

Mailbox Quota
Update...

Outlook on the web
Enable | Disable | Assign a policy...

Note that the mailboxes that you select need to be the same type. For example, if you select both user mailboxes and linked mailboxes, you'll get a warning in the details pane that says bulk edit won't work.

3. At the bottom of the details pane, click **More options**. Under the **Mailbox Delegation** option that appears, choose **Add** or **Remove**. Depending on your selection, do one of the following steps:

   - **Add**: In the **Bulk Add Delegation** dialog box that appears, click **Add** ✚ under the appropriate permission (**Send As**, **Send on Behalf**, or **Full Access**). When you're finished selecting users or groups to add as delegates, click **Save**.

   - **Remove**: In the **Bulk Remove Delegation** dialog box that appears, click **Add** ✚ under the appropriate permission (**Send As**, **Send on Behalf**, or **Full Access**). When you're finished selecting users or groups to remove from the existing delegates, click **Save**.

Use the EAC to assign permissions to groups

1. In the EAC, navigate to **Recipients** > **Groups**.

2. In the list of groups, select the group that you want to assign permissions for, and then click **Edit** ✎.

3. On the group properties page that opens, click **Group delegation** and configure one of the following permissions:

   - **Send As**: Messages sent by a delegate appear to come from the group.

   - **Send on Behalf**: Messages sent by a delegate have " *<Delegate>* on behalf of *<Group>*" in the From address.

4. To assign permissions to delegates, click **Add** ✚ under the appropriate permission. A dialog box appears that lists the users or groups that can have the permission assigned to them. Select the user or group from the list, and then click **Add**. Repeat this process as many times as necessary. You can also search for users or groups in the search box by typing all or part of the name, and then clicking **Search** 🔍. When you're finished selecting delegates, click **OK**.

   To remove a permission from a delegate, select the delegate in the list under the appropriate permission, and then click **Remove** ━.

5. When you're finished, click **Save**.

Use the Exchange Management Shell to assign the Full Access permission to mailboxes

You use the **Add-MailboxPermission** and **Remove-MailboxPermission** cmdlets to manage the Full Access permission for mailboxes. These cmdlets use the same basic syntax:

```
Add-MailboxPermission -Identity <MailboxIdentity> -User <DelegateIdentity> -AccessRights FullAccess -
InheritanceType All [-AutoMapping $false]
```

For more information, see Add-MailboxPermission.

```
Remove-MailboxPermission -Identity <MailboxIdentity> -User <DelegateIdentity> -AccessRights FullAccess -
InheritanceType All
```

For more information, see Remove-MailboxPermission.

This example assigns the delegate Raymond Sam the Full Access permission to the mailbox of Terry Adams.

```
Add-MailboxPermission -Identity "Terry Adams" -User raymonds -AccessRights FullAccess -InheritanceType All
```

This example assigns Esther Valle the Full Access permission to the organization's default discovery search mailbox, and prevents the mailbox from automatically opening in Esther Valle's Outlook.

```
Add-MailboxPermission -Identity "DiscoverySearchMailbox{D919BA05-46A6-415f-80AD-7E09334BB852}" -User estherv -
AccessRights FullAccess -InheritanceType All -AutoMapping $false
```

This example assigns members of the Helpdesk mail-enabled security group the Full Access permission to the shared mailbox named Helpdesk Tickets.

```
Add-MailboxPermission -Identity "Helpdesk Tickets" -User Helpdesk -AccessRights FullAccess -InheritanceType
All
```

This example removes Full Access permission for Jim Hance from Ayla Kol's mailbox.

```
Remove-MailboxPermission -Identity ayla -User "Jim Hance" -AccessRights FullAccess -InheritanceType All
```

How do you know this worked?

To verify that you've successfully assigned or removed the Full Access permission for a delegate on a mailbox, use either of the following procedures:

- In the properties of the mailbox in the EAC, verify the delegate is or isn't listed in **Mailbox delegation** > **Full Access**.

- Replace *<MailboxIdentity>* with the identity of the mailbox and run the following command in the Exchange Management Shell to verify that the delegate is or isn't listed..

  ```
  Get-MailboxPermission <MailboxIdentity> | where {$_.AccessRights -like 'Full*'} | Format-Table -Auto
  User,Deny,IsInherited,AccessRights
  ```

For more information, see Get-MailboxPermission.

Use the Exchange Management Shell to assign the Send As permission to mailboxes and groups

You use the **Add-AdPermission** and **Remove-AdPermission** cmdlets to manage the Send As permission for mailboxes. These cmdlets use the same basic syntax:

```
<Add-AdPermission | Remove-AdPermission> -Identity <MailboxOrGroupNameOrDN> -User <DelegateIdentity> [-
AccessRights ExtendedRight] -ExtendedRights "Send As"
```

For more information, see Add-AdPermission and Remove-AdPermission.

Notes:

- The *Identity* parameter requires you to use the **Name** or **DistinguishedName** (DN) value of the mailbox or group.
  - **Name**: This value may or may not be the same as the display name. For example, `Felipe Apodaca`.
  - **DistinguishedName**: This value always contains the **Name** value and uses Active Directory LDAP syntax. For example, `CN=Felipe Apodaca,CN=Users,DC=contoso,DC=com`.

  To find these values for a mailbox or group, you can use the **Get-Recipient** cmdlet, which accepts many different values for the *Identity* parameter. For example:

  ```
  Get-Recipient -Identity helpdesk@contoso.com | Format-List Name,DistinguishedName
  ```

- The commands work with or without `-AccessRights ExtendedRight`, which is why it's shown as optional in the syntax.

This example assigns the Send As permission to the Helpdesk mail-enabled security group on the shared mailbox named Helpdesk Support Team.

```
Add-ADPermission -Identity "Helpdesk Support Team" -User Helpdesk -ExtendedRights "Send As"
```

This example removes the Send As permission for the user Pilar Pinilla on the mailbox of James Alvord.

```
Remove-ADPermission -Identity "James Alvord" -User pilarp -ExtendedRights "Send As"
```

How do you know this worked?

To verify that you've successfully assigned or removed the Send As permission for a delegate on a mailbox or group, use either of the following procedures:

- In the properties of the mailbox or group in the EAC, verify the delegate is or isn't listed in **Mailbox delegation** > **Send As** or **Group delegation** > **Send As**.
- Replace *<MailboxOrGroupNameOrDN>* with the name or distinguished name of the mailbox or group and run the following command in the Exchange Management Shell to verify that the delegate is or isn't listed.

  ```
  Get-ADPermission -Identity <MailboxOrGroupNameOrDN> | where {$_.ExtendedRights -like 'Send*'} | Format-
  Table -Auto User,Deny,ExtendedRights
  ```

For more information, see Get-AdPermission.

Use the Exchange Management Shell to assign the Send on Behalf permission to mailboxes and groups

You use the *GrantSendOnBehalfTo* parameter on the various mailbox and group **Set-** cmdlets to manage the Send on Behalf permission for mailboxes and groups:

- **Set-Mailbox**
- **Set-DistributionGroup**: Distribution groups and mail-enabled security groups.
- **Set-DynamicDistributionGroup**

The basic syntax for these cmdlets is:

```
<Cmdlet> -Identity <MailboxOrGroupIdentity> -GrantSendOnBehalfTo <Delegates>
```

The *GrantSendOnBehalfTo* parameter has the following options for delegate values:

- **Replace existing delegates**: `<DelegateIdentity>` or `"<DelegateIdentity1>","<DelegateIdentity2>",...`
- **Add or remove delegates without affecting other delegates**: `@{Add="\<value1\>","\<value2\>"...; Remove="\<value1\>","\<value2\>"...}`
- **Remove all delegates**: Use the value `$null`.

This example assigns the delegate Holly Holt the Send on Behalf permission to the mailbox of Sean Chai.

```
Set-Mailbox -Identity seanc@contoso.com -GrantSendOnBehalfTo hollyh
```

This example adds the group tempassistants@contoso.com to the list of delegates that have Send on Behalf permission to the Contoso Executives shared mailbox.

```
Set-Mailbox "Contoso Executives" -GrantSendOnBehalfTo @{Add="tempassistants@contoso.com"}
```

This example assigns the delegate Sara Davis the Send on Behalf permission to the Printer Support distribution group.

```
Set-DistributionGroup -Identity printersupport@contoso.com -GrantSendOnBehalfTo sarad
```

This example removes the Send on Behalf permission that was assigned to the administrator on the All Employees dynamic distribution group.

```
Set-DynamicDistributionGroup "All Employees" -GrantSendOnBehalfTo @{Remove="Administrator"}
```

How do you know this worked?

To verify that you've successfully assigned or removed the Send on Behalf permission for a delegate on a mailbox or group, use either of the following procedures:

- In the properties of the mailbox or group in the EAC, verify the delegate is or isn't listed in **Mailbox delegation** > **Send As** or **Group delegation** > **Send As**.
- Replace *<MailboxIdentity>* or *<GroupIdentity>* with the identity of the mailbox or group and run the one of

the following commands in the Exchange Management Shell to verify that the delegate is or isn't listed.

- Mailbox:

  ```
  Get-Mailbox -Identity <MailboxIdentity> | Format-List GrantSendOnBehalfTo
  ```

- Group:

  ```
  Get-DistributionGroup -Identity <GroupIdentity> | Format-List GrantSendOnBehalfTo
  ```

- Dynamic distribution group:

  ```
  Get-DynamicDistributionGroup -Identity <GroupIdentity> | Format-List GrantSendOnBehalfTo
  ```

Next steps

For more information about how delegates can use the permissions that are assigned to them on mailboxes and groups, see the following topics:

- Access another person's mailbox
- Open and use a shared mailbox in Outlook for Windows
- Open and use a shared mailbox in Outlook on the web

Mailbox moves in Exchange Server

8/3/2020 • 4 minutes to read • Edit Online

You use mailbox moves to move mailboxes to, from, and within your Exchange organization. These are the basic types of mailbox moves that are available:

- **Local mailbox moves**: You move mailboxes from one mailbox database to another on Exchange servers within a single Active Directory forest. For instructions, see Manage on-premises mailbox moves in Exchange Server.
- **Cross-forest mailbox moves**: You move mailboxes to Exchange servers in a different Active Directory forest. You can initiate the move from the target forest where you want to move the mailboxes (known as a *pull* move type), or from the source forest that currently hosts the mailboxes (known as a *push* move type). For more information, see Prepare mailboxes for cross-forest move requests.
- **Remote mailbox moves in hybrid deployments**: In hybrid deployments between on-premises Exchange and Microsoft Office 365, you can move mailboxes from Exchange to Microsoft 365 or Office 365 (known as *onboarding remote move migrations*) and from Microsoft 365 or Office 365 to Exchange (know as *offboarding remote move migrations*). For more information, see Move mailboxes between on-premises and Exchange Online organizations in hybrid deployments.

NOTE

For more information about migrating on-premises Exchange organizations to Microsoft 365 or Office 365, see Ways to migrate multiple email accounts to Microsoft 365 or Office 365.

Mailbox moves in Exchange 2016 and Exchange 2019 use the batch move architecture that was introduced in Exchange 2013. The batch move architecture gives you the ability to move mailboxes in large batches. The enhanced management capabilities in the batch move architecture includes:

- Email notification during move with reporting.
- Automatic retry and automatic prioritization of moves.
- Move primary and personal archive mailboxes together or separately.
- Option for manual move request finalization to let you review your move before completion.
- Periodic incremental syncs to update migration changes.

You can move mailboxes in the Exchange admin center (EAC), or by using the New-MoveRequest or New-MigrationBatch cmdlets in the Exchange Management Shell.

Scenarios for local and cross-forest mailbox moves

These are some scenarios for local mailbox moves:

- **Upgrade**: When you upgrade from an earlier version of Exchange, you move mailboxes from the existing Exchange servers to an Exchange Mailbox server.
- **Realignment**: For example, you might want to move a mailbox to a database that has a larger mailbox size limit.
- **Investigate an issue**: If you need to investigate an issue with a mailbox, you can move that mailbox to a

different server. For example, you can move all mailboxes that have high activity to another server.

- **Corrupted mailboxes**: If you encounter corrupted mailboxes, you can move the mailboxes to a different server or database. The corrupted messages won't be moved.
- **Physical location changes**: You can move mailboxes to a server in a different Active Directory site. For example, if a user moves to a different physical location, you can move that user's mailbox to a server closer to the new location.

These are some scenarios for cross-forest mailbox moves:

- **Separation of administrative roles**: You might want to separate Exchange administration from Active Directory user account administration. To do this, you can move mailboxes from a single forest into a resource forest scenario. In this scenario, the Exchange mailboxes reside in one forest and their associated Active Directory user accounts reside in a different forest.
- **Outsourced email administration**: You might want to outsource the administration of email and retain the administration of Active Directory user accounts. To do this, you can move mailboxes from a single forest into a resource forest scenario.
- **Integrate email and user account administration**: You might want to change from a separated or outsourced email administration model to a model where email and user accounts can be managed from within the same forest. To do this, you can move mailboxes from a resource forest scenario to a single forest. In this scenario, the Exchange mailboxes and Active Directory user accounts reside in the same forest.

CSV files for mailbox moves

One of the major benefits of the batch move architecture is the ability to use a comma-separated value (CSV) to specify the mailboxes to move. The information that's required in the CSV file depends on the type of move. For more information, see CSV Files for Mailbox Migration.

Migration endpoints for cross-forest and remote mailbox moves

You use migration endpoints for cross-forest mailbox moves, and remote mailbox moves between Exchange and Microsoft 365 or Office 365 in hybrid deployments. You don't use migration endpoints for local mailbox moves.

Migration endpoints specify the remote server information, source throttling settings, and the required credentials for migrating the mailboxes.

- Cross-forest mailbox moves require an ExchangeRemoteMove migration endpoint.
- Onboarding mailbox move migrations in hybrid organizations (from Exchange to Microsoft 365 or Office 365) require an ExchangeRemoteMove migration endpoint as the *source* of the migration batch.
- Offboarding mailbox move migrations in hybrid organizations (from Microsoft 365 or Office 365 to Exchange) require an ExchangeRemoteMove migration endpoint as the *target* of the migration batch.

You can create migration endpoints in the EAC or by using the New-MigrationEndpoint cmdlet in the Exchange Management Shell.

MRS Proxy endpoints for cross-forest and remote mailbox moves

The Mailbox Replication Service Proxy (MRS Proxy) facilitates cross-forest mailbox moves and remote move migrations. By default, the Client Access services on Mailbox servers aren't configured to accept incoming move requests, so you'll need to enable the MRS Proxy endpoint.

- For cross-forest moves from the target forest (pull moves), you need to enable the MRS Proxy endpoint in the Client Access services on Mailbox servers in the source forest.

- For cross-forest moves from the source forest (push moves), you need to enable the MRS Proxy endpoint in the Client Access services on Mailbox servers in the target forest.

- For both onboarding and offboarding remote move migrations in hybrid deployments, you need to enable the MRS Proxy endpoint in the Client Access services on Mailbox servers in the on-premises Exchange organization.

For more information, see Enable the MRS Proxy endpoint for remote moves.

Mailbox imports and exports in Exchange Server

8/3/2020 • 5 minutes to read • Edit Online

Exchange Server uses the Microsoft Exchange Mailbox Replication service (MRS) to import .pst files to mailboxes, and export mailboxes to .pst files. The advantages of using MRS instead of Outlook to import and export mailboxes are:

- Import and export requests are asynchronous (you can import and export multiple .pst files at the same time).
- Imports and exports take advantage of the queuing and throttling that's provided by the MRS.
- You can import a .pst file directly to a user's archive mailbox.
- The source or destination .pst files can reside on any network share that's accessible by your Exchange servers.

This feature was introduced in Exchange 2010 Service Pack 1 (SP1). In Exchange 2010, the MRS runs on Client Access servers. In Exchange 2013 or later, the MRS runs in the backend services on Mailbox servers (not in the frontend Client Access services).

> **NOTE**
> Mailbox imports and exports are available only in the Mailbox Import Export role, and by default, that role isn't assigned to a role group. To use these features, you need to add the Mailbox Import Export role to a role group that you belong to (for example, the Organization Management role group). For more information, see Add a role to a role group.

Reasons to import or export mailboxes

As an administrator, you might need to import .pst files to mailboxes or export mailboxes to .pst files. For example:

- **Compliance requirements**: You can export a mailbox to a .pst file for legal discovery purposes. After the export is complete, you can import the .pst file to a mailbox that's specifically used for compliance purposes.
- **Create a point-in-time snapshot of a mailbox**: Suppose you're keeping a backup of an entire mailbox database for just a few mailboxes. By exporting those mailboxes to .pst files, you can eliminate the mailbox database backup.
- **Get content out of .pst files and into mailboxes**: Typically, Outlook users can save their email messages locally in .pst files. You can import users' .pst files to their primary mailboxes or archive mailboxes. This is an easy method for transferring email from a user's local computer to an Exchange server.

Considerations

Before you import .pst files to mailboxes, or export mailboxes to .pst files, consider these issues:

- You need to use a UNC network share (\ *<Server>*\ *<Share>*\ or \ *<LocalServerName>*\c$). The Exchange Trusted Subsystem security group requires permissions to the network share (Read for imports, Read/Write for exports). If the share doesn't have these permissions, you'll get errors when you try to import or export .pst files.
- We recommend that you don't try to import or export .pst files that are larger than 50 gigabytes (GB), because 50 GB is the maximum .pst file size that's supported by current versions of Outlook. You can export

mailboxes that are larger than 50 GB to .pst files by using multiple export requests that include or exclude specific folders, or by using a content filter.

- The operations may take several hours depending on the size of the .pst files or mailboxes, the available network bandwidth, and MRS throttling.
- You can't import .pst files to public folders.

Import .pst files to mailboxes

Here are some things to consider when you import .pst files to mailboxes:

- You can create new mailbox import requests in the EAC or the Exchange Management Shell. To view, modify, suspend, resume, or remove mailbox import requests, you need to use the Exchange Management Shell.
- You can import the .pst file to a different mailbox. For example, you can export data from john@contoso.com and import it to legaldiscovery@contoso.com.
- You can import the .pst file directly to the user's personal archive instead of their primary mailbox.
- By default, associated messages are imported if they exist in the .pst file. Associated messages are special messages that contain hidden data with information about rules, views, and forms. You can change this setting in the Exchange Management Shell (the *AssociatedMessagesCopyOption* parameter).
- By default, the Recoverable Items folder is imported if it exists in the .pst file. You can change this setting in the Exchange Management Shell (the *ExcludeDumpster* switch).
- In the Exchange Management Shell, you can include or exclude specific folders to import (the *IncludeFolders*, *ExcludeFolders*, or *SourceRootFolder* parameters).
- In the Exchange Management Shell, you can specify the destination folder for imported items in the target mailbox (the *TargetRootFolder* parameter).
- In the Exchange Management Shell, you can increase or decrease the priority value for mailbox import requests (the *Priority* parameter).

For mailbox import procedures, see Procedures for mailbox imports from .pst files in Exchange Server.

Export mailboxes to .pst files

Here are some things to consider when you export mailboxes to .pst files:

- You can create new mailbox export requests in the EAC or the Exchange Management Shell. To view, modify, suspend, resume, or remove mailbox export requests, you need to use the Exchange Management Shell.
- You can export a mailbox or a user's archive mailbox to a .pst file.
- By default, associated messages are exported from the mailbox. Associated messages are special messages that contain hidden data with information about rules, views, and forms. You can change this setting in the Exchange Management Shell (the *AssociatedMessagesCopyOption* parameter).
- By default, the Recoverable Items folder is exported from the mailbox. You can change this setting in the Exchange Management Shell (the *ExcludeDumpster* switch).
- In the Exchange Management Shell, you can filter the messages to export from the mailbox (the *ContentFilter* parameter). You can filter by message content, attachment, senders, recipients, Inbox category, importance, message type, message size, and when the message was sent, received, or expired.
- In the Exchange Management Shell, you can include or exclude specific folders to export (the *IncludeFolders*, *ExcludeFolders*, or *SourceRootFolder* parameters).

- In the Exchange Management Shell, you can specify the destination folder for exported items in the target .pst file (the *TargetRootFolder* parameter).

- In the Exchange Management Shell, you can increase or decrease the priority value for mailbox export requests (the *Priority* parameter).

For mailbox export procedures, see Procedures for mailbox exports to .pst files in Exchange Server.

Procedures for mailbox imports from .pst files in Exchange Server

8/3/2020 • 9 minutes to read • Edit Online

Mailbox import requests use the Microsoft Exchange Mailbox Replication service (MRS) to import the contents of .pst files into mailboxes. For more information, see Mailbox imports and exports in Exchange Server.

This topic shows you how to:

- Create mailbox import requests
- View mailbox import requests.
- Modify mailbox import requests that haven't completed.
- Suspend mailbox import requests that haven't completed or failed.
- Resume suspended or failed mailbox import requests
- Remove mailbox import requests.

What do you need to know before you begin?

IMPORTANT

The procedures in this topic require the Mailbox Import Export role, which isn't assigned to any role groups by default. To assign the role to a role group that you belong to, see Add a role to a role group. Note that changes in permission require you to log off and log on for the changes to take effect.

- Estimated time to complete each procedure: 5 minutes
- You need to import the .pst files from a UNC network share (\ *<Server>*\ *<Share>*\ or \ *<LocalServerName>*\c$). The Exchange Trusted Subsystem security group requires the Read permission to the network share. If the share doesn't have this permission, you'll get errors when you try to import .pst files to mailboxes.
- You can create mailbox import requests in the Exchange admin center (EAC) or in the Exchange Management Shell. All other procedures can only be done in the Exchange Management Shell. For more information about accessing and using the EAC, see Exchange admin center in Exchange Server. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Create mailbox import requests

Use the EAC to create a mailbox import request

1. In the EAC, go to **Recipients** > **Mailboxes** > click **More options** •••, and select **Import PST**.

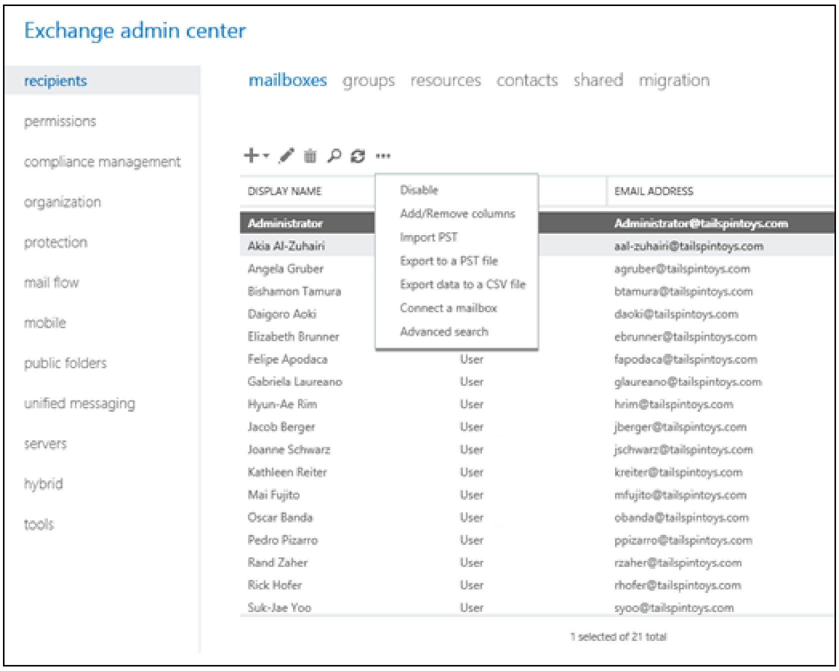

DISPLAY NAME		EMAIL ADDRESS
Administrator		Administrator@tailspintoys.com
Akia Al-Zuhairi		aal-zuhairi@tailspintoys.com
Angela Gruber		agruber@tailspintoys.com
Bishamon Tamura		btamura@tailspintoys.com
Daigoro Aoki		daoki@tailspintoys.com
Elizabeth Brunner		ebrunner@tailspintoys.com
Felipe Apodaca	User	fapodaca@tailspintoys.com
Gabriela Laureano	User	glaureano@tailspintoys.com
Hyun-Ae Rim	User	hrim@tailspintoys.com
Jacob Berger	User	jberger@tailspintoys.com
Joanne Schwarz	User	jschwarz@tailspintoys.com
Kathleen Reiter	User	kreiter@tailspintoys.com
Mai Fujito	User	mfujito@tailspintoys.com
Oscar Banda	User	obanda@tailspintoys.com
Pedro Pizarro	User	ppizarro@tailspintoys.com
Rand Zaher	User	rzaher@tailspintoys.com
Rick Hofer	User	rhofer@tailspintoys.com
Suk-Jae Yoo	User	syoo@tailspintoys.com

1 selected of 21 total

2. The **Import from a .pst** wizard opens. On the first page, enter the UNC path and filename of the source .pst file.

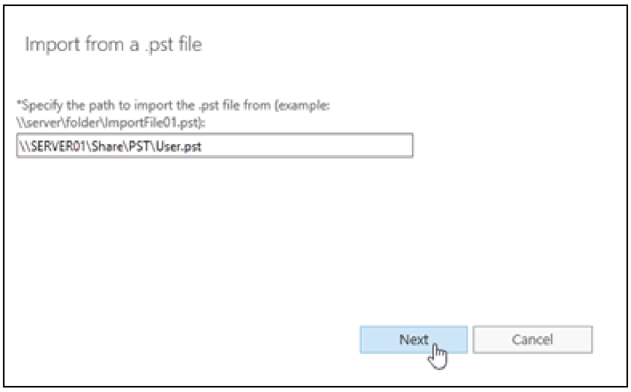
Import from a .pst file

*Specify the path to import the .pst file from (example: \\server\folder\ImportFile01.pst):

\\SERVER01\Share\PST\User.pst

When you're finished, click **Next**.

3. On the next page, select the target mailbox, and then select one of these options:

- **Import to this mailbox**
- **Import to this mailbox's archive**

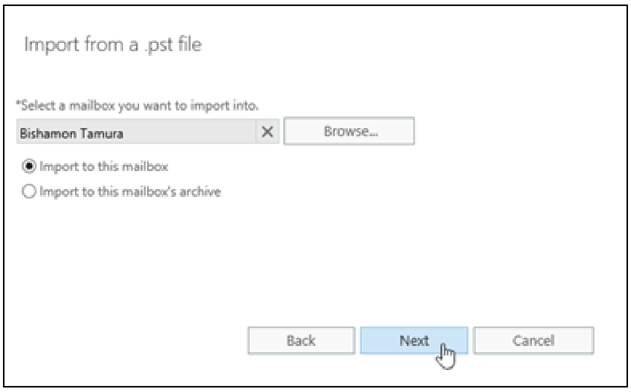
Import from a .pst file

*Select a mailbox you want to import into.

Bishamon Tamura

◉ Import to this mailbox

○ Import to this mailbox's archive

When you're finished, click **Next**.

4. On the last page, configure one of these settings:

- Leave the **Send email to the mailbox below when the .pst file has been exported** check box selected. Click **Browse** to add or remove notification recipients.

- Clear the **Send email to the mailbox below when the .pst file has been exported** check box.

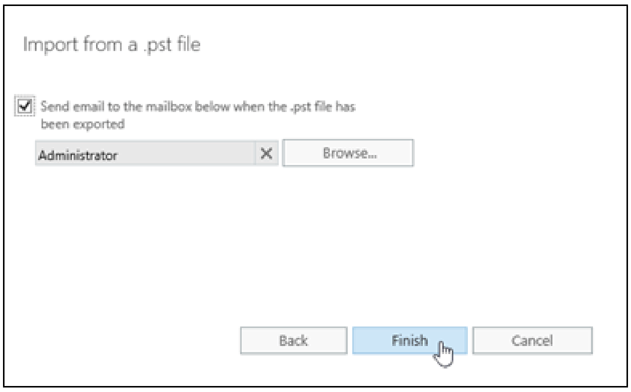

Import from a .pst file

☑ Send email to the mailbox below when the .pst file has been exported

Administrator

When you're finished, click **Finish**.

Use the Exchange Management Shell to create a mailbox import request

To create a mailbox import request, use this syntax:

```
New-MailboxImportRequest  [-Name <UniqueName>] -FilePath <UNCPathToPST> -Mailbox <TargetMailboxIdentity> [-IsArchive] [-SourceRootFolder <PSTFolder>] [-TargetRootFolder <MailboxFolder>] [-IncludeFolders <MailboxFolder1>,<MailboxFolder2>...] [-ExcludeFolders <MailboxFolder1>,<MailboxFolder2>...] [-Priority <PriorityValue>]
```

This example creates a new mailbox import request with these settings:

- **Mailbox import request name**: The default value `MailboxImport` is used, because we aren't using the *Name* parameter. The unique identity of the mailbox import request is `<MailboxIdentity>\MailboxImportX` (*X* is either not present, or has the value 0 to 9).

- **Source .pst file**: \\SERVER01\PSTFiles\Archives\Vbarrios.pst

- **Target mailbox**: Valeria Barrios

- **Content and folders**: Content in all folder paths in the .pst file is replicated in the target mailbox. Content is merged under existing folders and new folders are created if they don't already exist.

- **Priority**: `Normal`, because we aren't using the *Priority* parameter.

```
New-MailboxImportRequest -FilePath \\SERVER01\PSTFiles\Archives\Vbarrios.pst -Mailbox "Valeria Barrios"
```

This example creates a new mailbox import request with these settings:

- **Mailbox import request name**: The custom name Kathleen Reiter Import is specified by the *Name* parameter. Specifying a custom name allows more than 10 mailbox import requests for the mailbox. The unique identity value of the mailbox import request is `<MailboxIdentity>\<MailboxImportRequestName>` (for example, `kreiter\Kathleen Reiter Import`).

- **Source .pst file**: \\SERVER01\PSTFiles\Archives\Recovered.pst

- **Target mailbox**: The archive mailbox for Kathleen Reiter (Kathleen's primary mailbox alias is kreiter).

- **Content and folders**: Only content in the Inbox folder of the .pst file is imported (regardless of the localized name of the folder), and it's imported to the Recovered Files folder in the target mailbox.

- **Priority**: `High`

```
New-MailboxImportRequest -Name "Kathleen Reiter Import" -FilePath \\SERVER01\PSTFiles\Recovered.pst -Mailbox
kreiter -IsArchive -IncludeFolders "#Inbox#" -TargetRootFolder "Recovered Files" -Priority High
```

For detailed syntax and parameter information, see New-MailboxImportRequest.

How do you know this worked?

To verify that you've successfully created a mailbox import request, do any of these steps:

- In the EAC, click the notification viewer 🔔 to view the status of the request.
- If you created the mailbox import request in the EAC, and selected the option to send notification email messages, check the notification messages. The sender is Microsoft Exchange. The first message has the subject `Your Import PST request has been received`. If the import request completed successfully, you'll receive another message with the subject `Import PST has finished`.
- Replace *<MailboxIdentity>* with the name, email address, or alias of the target mailbox, and run this command in the Exchange Management Shell to verify the basic property values:

  ```
  Get-MailboxImportRequest -Mailbox "<MailboxIdentity>" | Format-List Name,FilePath,Mailbox,Status
  ```

- Replace *<MailboxIdentity>* and *<MailboxImportRequestName>* with the appropriate values, and run this command in the Exchange Management Shell to verify the details:

  ```
  Get-MailboxImportRequestStatistics -Identity "<MailboxIdentity>\<MailboxImportRequestName>"
  ```

Use the Exchange Management Shell to view mailbox import requests

By default, the **Get-MailboxImportRequest** cmdlet returns the name, target mailbox, and status of mailbox import requests. If you pipeline the command to the **Format-List** cmdlet, you'll only get a limited number of additional useful details:

- **FilePath**: The source .pst file.
- **RequestGUID**: The unique GUID value of the mailbox import request.
- **RequestQueue**: The mailbox database that the import request is being run on.
- **BatchName**: The optional batch name for the mailbox import request.
- **Identity**: The unique identity value of the mailbox import request (*<MailboxIdentity>*\ *<MailboxImportRequestName>*).

By default, the **Get-MailboxImportRequestStatistics** cmdlet returns the name, status, alias of the target mailbox, and the completion percentage of mailbox import requests. If you pipeline the command to the **Format-List** cmdlet, you'll see detailed information about the mailbox import request.

This example returns the summary list of all mailbox import requests.

```
Get-MailboxImportRequest
```

This example returns additional information for mailbox import requests to the mailbox Akia Al-Zuhairi.

```
Get-MailboxImportRequest -Mailbox "Akia Al-Zuhairi" | Format-List
```

This example returns the summary list of in-progress mailbox import requests for mailboxes that reside on the mailbox database named DB01.

```
Get-MailboxImportRequest -Status InProgress -Database DB01
```

This example returns the summary list of completed mailbox import requests in the batch named Import DB01 PSTs.

```
Get-MailboxImportRequest -Status Completed -BatchName "Import DB01 PSTs"
```

For detailed syntax and parameter information, see Get-MailboxImportRequest.

To view detailed information about a mailbox import request, use this syntax:

```
Get-MailboxImportRequestStatistics -Identity <MailboxImportRequestIdentity> [-IncludeReport] | Format-List
```

Where *<MailboxImportRequestIdentity>* is the identity value of the mailbox import request (*<MailboxIdentity>*\ *<MailboxImportRequestName>* or *<RequestGUID>*).

This example returns detailed information for the mailbox import request named MailboxImport for Akia Al-Zuhairi's mailbox, including the log of actions in the **Report** property.

```
Get-MailboxImportRequestStatistics -Identity "aal-zuhairi\MailboxImport" -IncludeReport | Format-List
```

For detailed syntax and parameter information, see Get-MailboxImportRequestStatistics.

Use the Exchange Management Shell to modify mailbox import requests

You can modify mailbox import requests that haven't completed. You can't modify the fundamental settings of an existing request (for example, the source .pst file, target mailbox, the source content in the .pst file, or the destination in the target mailbox).

To modify a mailbox import request, use this syntax:

```
Set-MailboxImportRequest -Identity <MailboxIdentity>\<MailboxImportRequestName> [-BadItemLimit <value>] [-LargeItemLimit <value>] [-AcceptLargeDataLoss]
```

This example modifies the failed mailbox import request for the mailbox of Valeria Barrios to accept up to five corrupted mailbox items.

```
Set-MailboxImportRequest -Identity "Valeria Barrios\MailboxImport" -BadItemLimit 5
```

For detailed syntax and parameter information, see Set-MailboxImportRequest.

Note: After you modify a suspended or failed mailbox import request, you need to resume it by using the **Resume-MailboxImportRequest** cmdlet.

How do you know this worked?

To verify that you've successfully modified a mailbox import request, replace *<MailboxIdentity>* and *<MailboxImportRequestName>* with the appropriate values, and run this command in the Exchange Management Shell to verify the details:

```
Get-MailboxImportRequestStatistics -Identity "<MailboxIdentity>\<MailboxImportRequestName>" | Format-List
```

Use the Exchange Management Shell to suspend mailbox import requests

You can suspend mailbox import requests that are in progress. You can't suspend completed or failed mailbox import requests.

To suspend a mailbox import request, use this syntax:

```
Suspend-MailboxImportRequest -Identity <MailboxIdentity>\<MailboxImportRequestName> [-SuspendComment "
<Descriptive Comment>"]
```

This example suspends the mailbox import request to Kathleen Reiter's mailbox that's named Kathleen Reiter Import.

```
Suspend-MailboxImportRequest -Identity "kreiter@contoso.com\Kathleen Reiter Import"
```

This example suspends all in-progress mailbox import requests with the comment "OK to resume after 10 P.M. on Monday 6/19"

```
Get-MailboxImportRequest -Status InProgress | Suspend-MailboxImportRequest -SuspendComment "OK to resume after
10 P.M. on Monday 6/19"
```

For detailed syntax and parameter information, see Suspend-MailboxImportRequest.

Notes:

- You can also use the **New-MailboxImportRequest** cmdlet with the *Suspend* switch to create a suspended mailbox import request.

- You use the **Resume-MailboxImportRequest** parameter to resume suspended mailbox import requests.

How do you know this worked?

To verify that you've successfully suspended a mailbox import request, do any of these steps:

- Replace *<MailboxIdentity>* with the name, email address, or alias of the target mailbox, run this command in the Exchange Management Shell, and verify that the **Status** property has the value `Suspended`:

  ```
  Get-MailboxImportRequest -Mailbox "<MailboxIdentity>" | Format-List Name,FilePath,Mailbox,Status
  ```

- Run this command in the Exchange Management Shell, and verify that the suspended mailbox import request is listed:

  ```
  Get-MailboxImportRequest -Status Suspended
  ```

Use the Exchange Management Shell to resume mailbox import requests

You can resume suspended or failed mailbox import requests.

To resume a mailbox import request, use this syntax:

```
Resume-MailboxImportRequest -Identity <MailboxIdentity>\<MailboxImportRequestName>
```

This example resumes the failed mailbox import request for Valeria Barrios' mailbox.

```
Resume-MailboxImportRequest -Identity vbarrios\MailboxImport
```

This example resumes all suspended mailbox import requests.

```
Get-MailboxImportRequest -Status Suspended | Resume-MailboxImportRequest
```

For detailed syntax and parameter information, see Resume-MailboxImportRequest.

How do you know this worked?

To verify that you've successfully resumed a mailbox import request, replace *<MailboxIdentity>* with the name, email address, or alias of the target mailbox, run this command in the Exchange Management Shell, and verify that the **Status** property doesn't have the value `Suspended`:

```
Get-MailboxImportRequest -Mailbox <MailboxIdentity> | Format-List Name,FilePath,Mailbox,Status
```

Use the Exchange Management Shell to remove mailbox import requests

You can remove fully or partially completed mailbox import requests.

- If you remove a partially completed mailbox import request, the request is removed from the MRS job queue. Any content that's already been imported from the .pst file isn't removed from the target mailbox.
- By default, completed mailbox import request are removed after 30 days (you can override this value with the *CompletedRequestAgeLimit* parameter), and failed requests aren't automatically removed. But, if you use the *RequestExpiryInterval* parameter when you create or modify a mailbox import request, these results are available:
  - **RequestExpiryInterval with a timespan value**: Completed and failed requests are automatically removed after the specified timespan.
  - **RequestExpiryInterval with the value unlimited**: Completed and failed requests aren't automatically removed.

This example removes the mailbox import request named MailboxImport for Akia Al-Zuhairi's mailbox.

```
Remove-MailboxImportRequest -Identity "aal-zuhairi\MailboxImport"
```

This example removes all completed mailbox import requests.

```
Get-MailboxImportRequest -Status Completed | Remove-MailboxImportRequest
```

For detailed syntax and parameter information, see Remove-MailboxImportRequest.

How do you know this worked?

To verify that you've successfully removed a mailbox import request, replace *<MailboxIdentity>* with the name, email address, or alias of the target mailbox, run this command in the Exchange Management Shell, and verify that the mailbox import request isn't listed:

```
Get-MailboxImportRequest -Mailbox <MailboxIdentity> | Format-List Name,FilePath,Mailbox,Status
```

Procedures for mailbox exports to .pst files in Exchange Server

8/3/2020 • 9 minutes to read • Edit Online

Mailbox export requests use the Microsoft Exchange Mailbox Replication service (MRS) to export the contents of mailboxes to .pst files. For more information, see Mailbox imports and exports in Exchange Server.

This topic shows you how to:

- Create mailbox export requests.
- View mailbox export requests.
- Modify mailbox export requests that haven't completed.
- Suspend mailbox export requests that haven't completed or failed.
- Resume suspended or failed mailbox export requests
- Remove mailbox export requests.

What do you need to know before you begin?

IMPORTANT

The procedures in this topic require the Mailbox Import Export role, which isn't assigned to any role groups by default. To assign the role to a role group that you belong to, see Add a role to a role group. Note that changes in permission require you to log off and log on for the changes to take effect.

- Estimated time to complete each procedure: 5 minutes
- You need to export mailboxes to .pst files on a UNC network share (\ *<Server>*\ *<Share>*\ or \ *<LocalServerName>*\c$). The Exchange Trusted Subsystem security group requires the Read/Write permission to the network share. If the share doesn't have this permission, you'll get errors when you try to export mailboxes to .pst files.
- You can create mailbox export requests in the Exchange admin center (EAC) or in the Exchange Management Shell. All other procedures can only be done in the Exchange Management Shell. For more information about accessing and using the EAC, see Exchange admin center in Exchange Server. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Create mailbox export requests

Use the EAC to create a mailbox export request

1. In the EAC, go to **Recipients** > **Mailboxes** > click **More options** ***, and select **Export to a PST file**.

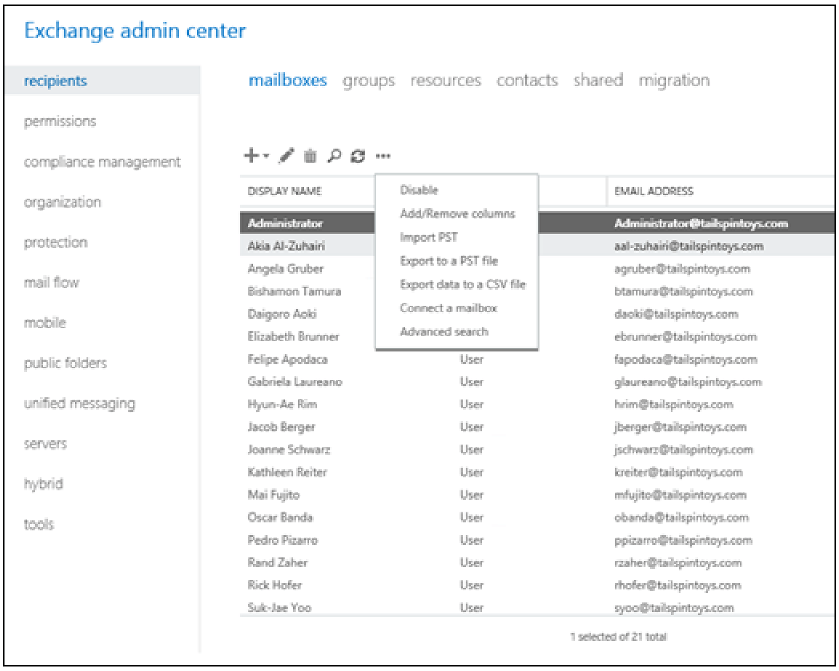

DISPLAY NAME		EMAIL ADDRESS
Administrator		Administrator@tailspintoys.com
Akia Al-Zuhairi		aal-zuhairi@tailspintoys.com
Angela Gruber		agruber@tailspintoys.com
Bishamon Tamura		btamura@tailspintoys.com
Daigoro Aoki		daoki@tailspintoys.com
Elizabeth Brunner		ebrunner@tailspintoys.com
Felipe Apodaca	User	fapodaca@tailspintoys.com
Gabriela Laureano	User	glaureano@tailspintoys.com
Hyun-Ae Rim	User	hrim@tailspintoys.com
Jacob Berger	User	jberger@tailspintoys.com
Joanne Schwarz	User	jschwarz@tailspintoys.com
Kathleen Reiter	User	kreiter@tailspintoys.com
Mai Fujito	User	mfujito@tailspintoys.com
Oscar Banda	User	obanda@tailspintoys.com
Pedro Pizarro	User	ppizarro@tailspintoys.com
Rand Zaher	User	rzaher@tailspintoys.com
Rick Hofer	User	rhofer@tailspintoys.com
Suk-Jae Yoo	User	syoo@tailspintoys.com

1 selected of 21 total

2. The **Export to a .pst file** wizard opens. On the first page, select the source mailbox, and then select one of these options:

   - **Export only the contents of this mailbox**

   - **Export only the contents of this mailbox's archive**

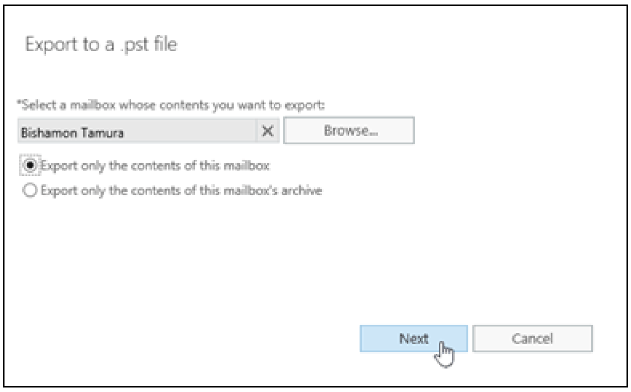

Export to a .pst file

*Select a mailbox whose contents you want to export:

Bishamon Tamura

● Export only the contents of this mailbox

○ Export only the contents of this mailbox's archive

When you're finished, click **Next**.

3. On the next page, enter the UNC path and filename of the target .pst file.

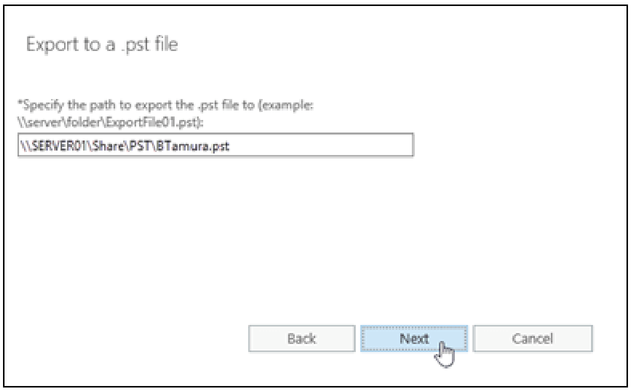

Export to a .pst file

*Specify the path to export the .pst file to (example: \\server\folder\ExportFile01.pst):

\\SERVER01\Share\PST\BTamura.pst

When you're finished, click **Next**.

4. On the last page, configure one of these settings:

- Leave the **Send email to the mailbox below when the .pst file has been exported** check box selected. Click **Browse** to add or remove notification recipients.
- Clear the **Send email to the mailbox below when the .pst file has been exported** check box.

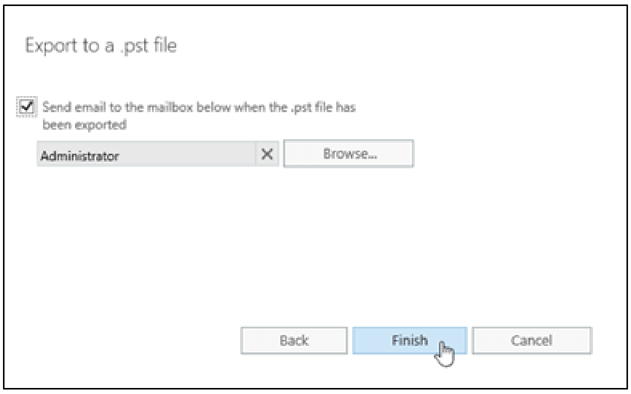

Export to a .pst file

☑ Send email to the mailbox below when the .pst file has been exported

Administrator

When you're finished, click **Finish**.

Use the Exchange Management Shell to create a mailbox export request

To create a mailbox export request, use this syntax:

```
New-MailboxExportRequest  [-Name <UniqueName>] -Mailbox <TargetMailboxIdentity> -FilePath <UNCPathToPST> [-IsArchive] [-SourceRootFolder <MailboxFolder>] [-TargetRootFolder <PSTFolder>] [-IncludeFolders <MailboxFolder1>,<MailboxFolder2>...] [-ExcludeFolders <MailboxFolder1>,<MailboxFolder2>...] [-ContentFilter <Filter>] [-Priority <PriorityValue>]
```

This example creates a new mailbox export request with these settings:

- **Mailbox export request name**: The default value `MailboxExport` is used, because we aren't using the *Name* parameter. The unique identity of the mailbox export request is `<MailboxIdentity>\MailboxExportX` (*X* is either not present, or has the value 0 to 9).
- **Source mailbox**: Valeria Barrios
- **Target .pst file**: \SERVER01\PSTFiles\Vbarrios.pst
- **Content and folders**: Content in all folder paths in the source mailbox is replicated in the target .pst file.
- **Priority**: `Normal`, because we aren't using the *Priority* parameter.

```
New-MailboxExportRequest -Mailbox "Valeria Barrios" -FilePath \\SERVER01\PSTFiles\Vbarrios.pst
```

This example creates a new mailbox export request with these settings:

- **Mailbox export request name**: The custom name Kathleen Reiter Export is specified by the *Name* parameter. Specifying a custom name allows more than 10 mailbox export requests for the mailbox. The unique identity value of the mailbox export request is `<MailboxIdentity>\<MailboxExportRequestName>` (for example, `kreiter\Kathleen Reiter Export`).
- **Source mailbox**: The archive mailbox for Kathleen Reiter (Kathleen's primary mailbox alias is kreiter).
- **Target .pst file**: \SERVER01\PSTFiles\Archives\Kathleen Reiter.pst
- **Content and folders**: Only content in the Inbox folder of the mailbox is exported (regardless of the localized name of the folder).
- **Priority**: `High`

```
New-MailboxExportRequest -Name "Kathleen Reiter Export" -Mailbox kreiter -FilePath
"\\SERVER01\PSTFiles\Kathleen Reiter.pst" -IsArchive -IncludeFolders "#Inbox#" -Priority Hight
```

For detailed syntax and parameter information, see New-MailboxExportRequest.

How do you know this worked?

To verify that you've successfully created a mailbox export request, do any of these steps:

- In the EAC, click the notification viewer 🔔 to view the status of the request.
- If you created the mailbox export request in the EAC, and selected the option to send notification email messages, check the notification messages. The sender is Microsoft Exchange. The first message has the subject `Your Export PST request has been received`. If the export request completed successfully, you'll receive another message with the subject `Export PST has finished`.
- Replace *<MailboxIdentity>* with the name, email address, or alias of the source mailbox, and run this command in the Exchange Management Shell to verify the basic property values:

  ```
  Get-MailboxExportRequest -Mailbox "<MailboxIdentity>" | Format-List Name,FilePath,Mailbox,Status
  ```

- Replace *<MailboxIdentity>* and *<MailboxExportRequestName>* with the appropriate values, and run this command in the Exchange Management Shell to verify the details:

  ```
  Get-MailboxExportRequestStatistics -Identity "<MailboxIdentity>\<MailboxExportRequestName>"
  ```

Use the Exchange Management Shell to view mailbox export requests

By default, the **Get-MailboxExportRequest** cmdlet returns the name, source mailbox, and status of mailbox export requests. If you pipeline the command to the **Format-List** cmdlet, you'll only get a limited number of additional useful details:

- **FilePath**: The target .pst file.
- **RequestGUID**: The unique GUID value of the mailbox export request.
- **RequestQueue**: The mailbox database that the export request is being run on.
- **BatchName**: The optional batch name for the mailbox export request.
- **Identity**: The unique identity value of the mailbox export request (*<MailboxIdentity>*\ *<MailboxExportRequestName>*).

By default, the **Get-MailboxExportRequestStatistics** cmdlet returns the name, status, alias of the source mailbox, and the completion percentage of mailbox export requests. If you pipeline the command to the **Format-List** cmdlet, you'll see detailed information about the mailbox export request.

This example returns the summary list of all mailbox export requests.

```
Get-MailboxExportRequest
```

This example returns additional information for mailbox export requests from the mailbox Akia Al-Zuhairi.

```
Get-MailboxExportRequest -Mailbox "Akia Al-Zuhairi" | Format-List
```

This example returns the summary list of in-progress mailbox export requests for mailboxes that reside on the mailbox database named DB01.

```
Get-MailboxExportRequest -Status InProgress -Database DB01
```

This example returns the summary list of completed mailbox export requests in the batch named Export DB01 PSTs.

```
Get-MailboxExportRequest -Status Completed -BatchName "Export DB01 PSTs"
```

For detailed syntax and parameter information, see Get-MailboxExportRequest.

To view detailed information about a mailbox export request, use this syntax:

```
Get-MailboxExportRequestStatistics -Identity <MailboxExportRequestIdentity> [-IncludeReport] | Format-List
```

Where *<MailboxExportRequestIdentity>* is the identity value of the mailbox export request (*<MailboxIdentity>\<MailboxExportRequestName>* or *<RequestGUID>*).

This example returns detailed information for the mailbox export request named MailboxExport for Akia Al-Zuhairi's mailbox, including the log of actions in the **Report** property.

```
Get-MailboxExportRequestStatistics -Identity "aal-zuhairi\MailboxExport" -IncludeReport | Format-List
```

For detailed syntax and parameter information, see Get-MailboxExportRequestStatistics.

Use the Exchange Management Shell to modify mailbox export requests

You can modify mailbox export requests that haven't completed. You can't modify the fundamental settings of an existing request (for example, the source mailbox, target .pst file, the source content in the mailbox, or the destination in the target .pst file).

To modify a mailbox export request, use this syntax:

```
Set-MailboxExportRequest -Identity <MailboxIdentity>\<MailboxExportRequestName> [-BadItemLimit <value>] [-
LargeItemLimit <value>] [-AcceptLargeDataLoss]
```

This example modifies the failed mailbox export request for the mailbox of Valeria Barrios to accept up to five corrupted mailbox items.

```
Set-MailboxExportRequest -Identity "Valeria Barrios\MailboxExport" -BadItemLimit 5
```

For detailed syntax and parameter information, see Set-MailboxExportRequest.

Note: After you modify a suspended or failed mailbox export request, you need to resume it by using the **Resume-MailboxExportRequest** cmdlet.

How do you know this worked?

To verify that you've successfully modified a mailbox export request, replace *<MailboxIdentity>* and *<MailboxExportRequestName>* with the appropriate values, and run this command in the Exchange Management Shell to verify the details:

```
Get-MailboxExportRequestStatistics -Identity "<MailboxIdentity>\<MailboxExportRequestName>" | Format-List
```

Use theExchange Management Shell to suspend mailbox export requests

You can suspend mailbox export requests that are in progress. You can't suspend completed or failed mailbox export requests.

To suspend a mailbox export request, use this syntax:

```
Suspend-MailboxExportRequest -Identity <MailboxIdentity>\<MailboxExportRequestName> [-SuspendComment "
<Descriptive Comment>"]
```

This example suspends the mailbox export request from Kathleen Reiter's mailbox that's named Kathleen Reiter Export.

```
Suspend-MailboxExportRequest -Identity "kreiter@contoso.com\Kathleen Reiter Export"
```

This example suspends all in-progress mailbox export requests with the comment "OK to resume after 10 P.M. on Monday 6/19"

```
Get-MailboxExportRequest -Status InProgress | Suspend-MailboxExportRequest -SuspendComment "OK to resume after
10 P.M. on Monday 6/19"
```

For detailed syntax and parameter information, see Suspend-MailboxExportRequest.

Notes:

- You can also use the **New-MailboxExportRequest** cmdlet with the *Suspend* switch to create a suspended mailbox export request.
- You use the **Resume-MailboxExportRequest** parameter to resume suspended mailbox export requests.

How do you know this worked?

To verify that you've successfully suspended a mailbox export request, do any of these steps:

- Replace *<MailboxIdentity>* with the name, email address, or alias of the source mailbox, run this command in the Exchange Management Shell, and verify that the **Status** property has the value `Suspended`:

  ```
  Get-MailboxExportRequest -Mailbox "<MailboxIdentity>" | Format-List Name,FilePath,Mailbox,Status
  ```

- Run this command in the Exchange Management Shell, and verify that the suspended mailbox export request is listed:

  ```
  Get-MailboxExportRequest -Status Suspended
  ```

Use the Exchange Management Shell to resume mailbox export requests

You can resume suspended or failed mailbox export requests.

To resume a mailbox export request, use this syntax:

```
Resume-MailboxExportRequest -Identity <MailboxIdentity>\<MailboxExportRequestName>
```

This example resumes the failed mailbox export request for Valeria Barrios' mailbox.

```
Resume-MailboxExportRequest -Identity vbarrios\MailboxExport
```

This example resumes all suspended mailbox export requests.

```
Get-MailboxExportRequest -Status Suspended | Resume-MailboxExportRequest
```

For detailed syntax and parameter information, see Resume-MailboxExportRequest.

How do you know this worked?

To verify that you've successfully resumed a mailbox export request, replace *<MailboxIdentity>* with the name, email address, or alias of the source mailbox, run this command in the Exchange Management Shell, and verify that the **Status** property doesn't have the value `Suspended`:

```
Get-MailboxExportRequest -Mailbox <MailboxIdentity> | Format-List Name,FilePath,Mailbox,Status
```

Use the Exchange Management Shell to remove mailbox export requests

You can remove fully or partially completed mailbox export requests.

- If you remove a partially completed mailbox export request, the request is removed from the MRS job queue. Any content that's already been exported from the source mailbox isn't removed from the target .pst file.
- By default, completed mailbox export request are removed after 30 days (you can override this value with the *CompletedRequestAgeLimit* parameter), and failed requests aren't automatically removed. But, if you use the *RequestExpiryInterval* parameter when you create or modify a mailbox export request, these results are available:
    - **RequestExpiryInterval with a timespan value**: Completed and failed requests are automatically removed after the specified timespan.
    - **RequestExpiryInterval with the value unlimited**: Completed and failed requests aren't automatically removed.

This example removes the mailbox export request named MailboxExport for Akia Al-Zuhairi's mailbox.

```
Remove-MailboxExportRequest -Identity "aal-zuhairi\MailboxExport"
```

This example removes all completed mailbox export requests.

```
Get-MailboxExportRequest -Status Completed | Remove-MailboxExportRequest
```

For detailed syntax and parameter information, see Remove-MailboxExportRequest.

How do you know this worked?

To verify that you've successfully removed a mailbox export request, replace *<MailboxIdentity>* with the name, email address, or alias of the source mailbox, run this command in the Exchange Management Shell, and verify that the mailbox export request isn't listed:

```
Get-MailboxExportRequest -Mailbox <MailboxIdentity> | Format-List Name,FilePath,Mailbox,Status
```

Clients and mobile in Exchange Server

8/3/2020 • 2 minutes to read • Edit Online

There are many different clients that you can use to access information in an Exchange mailbox. These clients include desktop programs such as Outlook, Outlook on the web (formerly known as Outlook Web App), and mobile clients such as mobile phones, tablets, and other mobile devices. Each of these clients offers a variety of features.

The following table contains links to topics that will help you learn about and manage some of the clients and client access methods that you can use to access your Exchange mailbox.

TOPIC	DESCRIPTION
MAPI over HTTP in Exchange Server	Learn about the latest client access method that provides connectivity to Outlook.
Outlook Anywhere	Learn about the earlier client access method that provides connectivity to Outlook. (This feature was formerly known as RPC/HTTP.)
Exchange ActiveSync	Learn about the protocol that provides connectivity to a wide variety of mobile phones and tablets. Using Exchange ActiveSync, users can access email, calendar, contact, and task information.
POP3 and IMAP4 in Exchange Server	Learn about how users can access their Exchange mailbox by using email programs that use POP3 or IMAP4.
Outlook for iOS and Android	Learn about the Outlook for iOS and Android app and how it allows your users to securely access their mailbox data remotely with their iOS and Android devices.
Install Office Online Server in an Exchange organization	Learn about how the integration of Office Online Server helps provide rich attachment preview functionality in Outlook on the web.
Outlook on the web in Exchange Server	Learn about Outlook on the web, which provides users access to their Exchange mailbox through a web browser.
MailTips in Exchange	Learn about MailTips, the informative messages displayed to users while they're composing a message.

MAPI over HTTP in Exchange Server

8/3/2020 • 3 minutes to read • Edit Online

Messaging Application Programming Interface (MAPI) over HTTP is a transport protocol that improves the reliability and stability of the Outlook and Exchange connections by moving the transport layer to the industry-standard HTTP model. This allows a higher level of visibility of transport errors and enhanced recoverability. Additional functionality includes support for an explicit pause-and-resume function. This enables supported clients to change networks or resume from hibernation while maintaining the same server context.

Implementing MAPI over HTTP does not mean that it is the only protocol that can be used for Outlook to access Exchange. Outlook clients that are not MAPI over HTTP capable can still use Outlook Anywhere (RPC over HTTP) to access Exchange through a MAPI-enabled Client Access server.

In Exchange 2016 and Exchange 2019, MAPI over HTTP can be applied across your entire organization, or at the individual mailbox level.

Benefits of MAPI over HTTP

MAPI over HTTP offers the following benefits to the clients that support it:

- Enables future innovation in authentication by using an HTTP based protocol.
- Provides faster reconnection times after a communications break because only TCP connections (not RPC connections) need to be rebuilt. Examples of a communication break include:
  - Device hibernation
  - Changing from a wired network to a wireless or cellular network
- Offers a session context that is not dependent on the connection. The server maintains the session context for a configurable period of time, even if the user changes networks.

MAPI over HTTP when upgrading Exchange

In Exchange 2016 or later, MAPI over HTTP is enabled by default at the organization level, although you still need to configure the virtual directories as described in Configure MAPI over HTTP for users to take advantage of it.

The scenarios where MAPI over HTTP is enabled or disabled by default at the organization level are described in the following table:

	EXCHANGE 2019	EXCHANGE 2016
Upgrading from an Exchange 2016 environment	MAPI over HTTP is enabled by default	n/a
Upgrading from an environment that contains any Exchange 2013 servers	MAPI over HTTP is disabled by default	MAPI over HTTP is disabled by default
Upgrading from an Exchange 2010 environment	n/a	MAPI over HTTP is enabled by default

During the upgrade from an organization that contains Exchange 2013 servers, administrators will receive the

MAPI over HTTP isn't enabled [WarnMapiHttpNotEnabled] readiness check warning, and enabling MAPI over HTTP post-installation is recommended. In any organization that contains Exchange 2013 servers, MAPI over HTTP won't be enabled by default, and administrators will need to follow the steps in Configure MAPI over HTTP to enable it.

Supportability and Prerequisites

Consider the following requirements to enable MAPI over HTTP.

Supportability

Use the following matrix to verify that your clients and servers support MAPI over HTTP.

PRODUCT	EXCHANGE 2019	EXCHANGE 2016	EXCHANGE 2013 SP1	EXCHANGE 2013 RTM	EXCHANGE 2010 SP3
Outlook 2013 SP1 and all later versions of Outlook	MAPI over HTTP Outlook Anywhere	MAPI over HTTP Outlook Anywhere	MAPI over HTTP Outlook Anywhere	Outlook Anywhere	RPC Outlook Anywhere
Outlook 2010 SP2 with updates KB2956191 and KB2965295 (April 14, 2015)	MAPI over HTTP Outlook Anywhere	MAPI over HTTP Outlook Anywhere	MAPI over HTTP Outlook Anywhere	Outlook Anywhere	RPC Outlook Anywhere
Outlook 2013 RTM	Outlook Anywhere	Outlook Anywhere	Outlook Anywhere	Outlook Anywhere	RPC Outlook Anywhere
All earlier versions of Outlook	Outlook Anywhere	Outlook Anywhere	Outlook Anywhere	Outlook Anywhere	RPC Outlook Anywhere

Prerequisites

The following conditions are required for clients and servers to support MAPI over HTTP with Exchange Server. Once the following prerequisites are in place, see Configure MAPI over HTTP to enable it in your organization.

- Supported Outlook clients (see the table in the previous section).
- .NET Framework 4.5.2 or later. Note that this is no longer an issue for Exchange 2016 CU5 or later. For more information about the .NET Framework requirements for Exchange 2016, see Supported .NET Framework versions for Exchange 2016.

Configure MAPI over HTTP in Exchange Server

8/3/2020 • 5 minutes to read • Edit Online

In Exchange 2016 and Exchange 2019, you can configure MAPI over HTTP at the organization level or at the individual mailbox level. Mailbox-level settings always take precedence over organization-wide settings.

The scenarios where MAPI over HTTP is enabled or disabled by default at the organization level are described in the following table:

	EXCHANGE 2019	EXCHANGE 2016
Upgrading from an Exchange 2016 environment	MAPI over HTTP is enabled by default	n/a
Upgrading from an environment that contains any Exchange 2013 servers	MAPI over HTTP is disabled by default	MAPI over HTTP is disabled by default
Upgrading from an Exchange 2010 environment	n/a	MAPI over HTTP is enabled by default

> **NOTE**
> When MAPI over HTTP is enabled at the organization level, the *MapiHttpEnabled* property value that's returned by the **Get-OrganizationConfig** cmdlet is `True`.

This topic describes how to configure and then enable MAPI over HTTP for Exchange organizations that contain Exchange 2013 servers, or for any topology where MAPI over HTTP has been previously disabled. You can also use the procedures in this article to disable MAPI over HTTP at the organization level.

This topic also describes how to enable or disable MAPI over HTTP for an individual mailbox. At the mailbox level, you have the ability to allow or block MAPI over HTTP connections internally, externally, or both. In all cases, when MAPI over HTTP is disabled, connections will be made with Outlook Anywhere.

Configure MAPI over HTTP

Complete the following steps to configure MAPI over HTTP for your organization. These steps assume you have already configured the prerequisites described in MAPI over HTTP in Exchange Server. Once configured (steps 1-3), use step 4 to enable or disable specific permission scenarios at the organization level, at the mailbox level, or both.

1. **Virtual directory configuration**: By default, Exchange creates a virtual directory for MAPI over HTTP. You use the **Set-MapiVirtualDirectory** cmdlet to configure the virtual directory. You need to configure an internal URL, an external URL, or both. For more information see, Set-MapiVirtualDirectory.

   For example, to configure the default MAPI virtual directory on the local Exchange server by setting the internal URL value to https://contoso.com/mapi, and the authentication method to `Negotiate`, run the following command:

```
Set-MapiVirtualDirectory -Identity "Contoso\mapi (Default Web Site)" -InternalUrl
https://Contoso.com/mapi -IISAuthenticationMethods Negotiate
```

2. **Certificate configuration**: The digital certificate used by your Exchange environment must include the same *InternalURL* and *ExternalURL* values that are defined on the MAPI virtual directory. For more information on Exchange certificate management, see Digital certificates and encryption in Exchange Server. Make sure the Exchange certificate is trusted on the Outlook client workstation and that there are no certificate errors, especially when you access the URLs configured on the MAPI virtual directory.

3. **Update server rules**: Verify that your load balancers, reverse proxies, and firewalls are configured to allow access to the MAPI over HTTP virtual directory.

4. Use the following steps to enable MAPI over HTTP in your entire Exchange organization, or enable MAPI over HTTP for one or more individual mailboxes.

   > **NOTE**
   >
   > After you run the commands below, Outlook clients with MAPI over HTTP enabled will see a message to restart Outlook to use MAPI over HTTP.

   #### Enable MAPI over HTTP in your Exchange organization

   To enable or disable MAPI over HTTP at the organizational level, use the **Set-OrganizationConfig** cmdlet with the *MapiHttpEnabled* parameter. Valid values are:

   - **$true**: MAPI over HTTP connections are allowed for all mailboxes in the organization (unless MAPI over HTTP is disabled on a specific mailbox).

   - **$false**: MAPI over HTTP connections aren't allowed for all mailboxes in the organization (unless MAPI over HTTP is enabled on a specific mailbox).

   The following example enables MAPI over HTTP connections for the entire organization:

   ```
   Set-OrganizationConfig -MapiHttpEnabled $true
   ```

   #### Enable MAPI over HTTP for an individual mailbox

   To enable or disable MAPI over HTTP at the mailbox level, use the **Set-CasMailbox** cmdlet with the *MapiHttpEnabled* parameter. Valid values are:

   - **$null**: The mailbox follows organization-level settings. This is the default value.

   - **$true**: Enable MAPI over HTTP for the mailbox. If MAPI over HTTP is disabled at the organizational level, it's enabled for the mailbox.

   - **$false**: Disable MAPI over HTTP for the mailbox. If MAPI over HTTP is enabled at the organizational level, it's disabled for the mailbox, so the mailbox will use Outlook Anywhere connections.

   The following example enables MAPI over HTTP connections for a single mailbox:

   ```
   Set-CasMailbox <user or mailbox ID> -MapiHttpEnabled $true
   ```

Test MAPI over HTTP connections

You can test the end-to-end MAPI over HTTP connection by using the **Test-OutlookConnectivity** cmdlet. To use the **Test-OutlookConnectivity** cmdlet, the Microsoft Exchange Health Manager (MSExchangeHM) service must be started.

The following example tests the MAPI over HTTP connection from the Exchange server named ContosoMail.

```
Test-OutlookConnectivity -RunFromServerId ContosoMail -ProbeIdentity OutlookMapiHttpSelfTestProbe
```

A successful test returns output that's similar to the following example:

```
MonitorIdentity                                          StartTime              EndTime                Result
Error     Exception
---------------                                          ---------              -------                ------
-----     ---------
OutlookMapiHttp.Protocol\OutlookMapiHttpSelfTestProbe    2/14/2018 7:15:00 AM   2/14/2018 7:15:10 AM
Succeeded
```

For more information, see Test-OutlookConnectivity.

Logs for MAPI over HTTP activity are at the following locations:

- %ExchangeInstallPath%Logging\MAPI Address Book Service\
- %ExchangeInstallPath%Logging\MAPI Client Access\
- %ExchangeInstallPath%Logging\HttpProxy\Mapi\

Combining MAPI over HTTP configurations and internal or external connections

In addition to the organization and mailbox settings described earlier in this topic, you can use the *MapiBlockOutlookExternalConnectivity* parameter on the **Set-CasMailbox** cmdlet to allow or deny external Outlook Anywhere or MAPI over HTTP connections to a specific mailbox. Valid values are:

- **True**: Only internal connections are allowed to the mailbox.
- **False**: Internal and external connections are allowed to the mailbox. This is the default value.

The following table summarizes the results of the different setting combinations at the organization level and on individual mailboxes.

MAPIHTTPENABLED VALUE ON SET-ORGANIZATIONCONFIG	MAPIHTTPENABLED VALUE ON SET-CASMAILBOX	MAPIBLOCKOUTLOOKEXTERNALCONNECTIVITY VALUE ON SET-CASMAILBOX	AUTODISCOVER RESULT
$true	$null	$false	MAPI over HTTP, internal and external
$true	$null	$true	MAPI over HTTP, internal only
$true	$true	$false	MAPI over HTTP, internal and external
$true	$true	$true	MAPI over HTTP, internal only
$true	$false	$false	Outlook Anywhere, internal and external

MAPIHTTPENABLED VALUE ON SET-ORGANIZATIONCONFIG	MAPIHTTPENABLED VALUE ON SET-CASMAILBOX	MAPIBLOCKOUTLOOKEXTERNALCONNECTIVITY VALUE ON SET-CASMAILBOX	AUTODISCOVER RESULT
$true	$false	$true	Outlook Anywhere, internal only
$false	$null	$false	Outlook Anywhere, internal and external
$false	$null	$true	Outlook Anywhere, internal only
$false	$true	$false	MAPI over HTTP, internal and external
$false	$true	$true	MAPI over HTTP, internal only
$false	$false	$false	Outlook Anywhere, internal and external
$false	$false	$true	Outlook Anywhere, internal only

Manage MAPI over HTTP

You can manage the configuration of MAPI over HTTP by using the following cmdlets:

- Set-MapiVirtualDirectory
- Get-MapiVirtualDirectory
- New-MapiVirtualDirectory
- Remove-MapiVirtualDirectory

Enable or disable MAPI access to mailboxes in Exchange Server

8/3/2020 • 4 minutes to read • Edit Online

MAPI is a client protocol that lets users access their mailbox by using Outlook or other MAPI email clients. By default, MAPI access to a user mailbox is enabled. Disabling MAPI access to a mailbox prevents the user from using Outlook to access their mailbox in Exchange mode. It doesn't prevent the user from using Outlook on the web or Outlook using other protocols (for example, POP3, IMAP4, or Exchange ActiveSync) to access their mailbox.

Administrators can use the Exchange admin center (EAC) or the Exchange Management Shell to enable or disable MAPI access to user mailbox.

For additional management tasks related to user access to mailboxes, see these topics:

- Enable or disable Outlook on the web access to mailboxes in Exchange Server
- Enable or disable Exchange ActiveSync access to mailboxes in Exchange Server
- Enable or disable POP3 or IMAP4 access to mailboxes in Exchange Server

What do you need to know before you begin?

- Estimated time to complete: 5 minutes.
- For more information about accessing and using the EAC, see Exchange admin center in Exchange Server.
- To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Client Access user settings" entry in the Clients and mobile devices permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Enable or disable MAPI access to a single mailbox

Use the EAC to Enable or disable MAPI access to a single mailbox

1. In the EAC, go to **Recipients** > **Mailboxes**.
2. In the list of mailboxes, find the mailbox that you want to modify. You can:
   - Scroll through the list of mailboxes.
   - Click **Search** 🔍 and enter part of the user's name, email address, or alias.
   - Click **More options** ••• > **Advanced search** to find the mailbox.

Once you've found the mailbox that you want to modify, select it, and then click **Edit** ✎.

3. On the mailbox properties page that opens, click **Mailbox features**.

4. In the **Email Connectivity** section, configure one of these settings:

   - If you see **MAPI: Enabled**, click **Disable** to disable it, and then click **Yes** in the warning message that appears.

   - If you see **MAPI: Disabled**, click **Enable** to enable it.

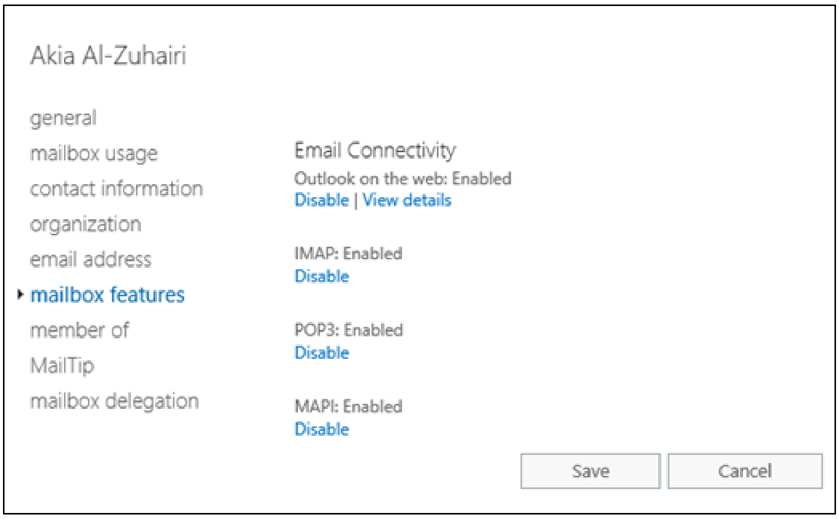

Akia Al-Zuhairi

Email Connectivity

Outlook on the web: Enabled
Disable | View details

IMAP: Enabled
Disable

POP3: Enabled
Disable

MAPI: Enabled
Disable

When you're finished, click **Save**.

Use the Exchange Management Shell to enable or disable MAPI access to a mailbox

To enable or disable MAPI access to a single mailbox, use this syntax:

```
Set-CasMailbox -Identity <MailboxIdentity> -MAPIEnabled <$true | $false>
```

This example disables MAPI access to the mailbox named Ken Sanchez.

```
Set-CasMailbox -Identity "Ken Sanchez" -MAPIEnabled $false
```

This example enables MAPI access to the mailbox named Esther Valle.

```
Set-CasMailbox -Identity "Esther Valle" -MAPIEnabled $true
```

For detailed syntax and parameter information, see Set-CASMailbox.

Enable or disable MAPI access to multiple mailboxes

Use the EAC to enable or disable MAPI access to multiple mailboxes

1. In the EAC, go to **Recipients** > **Mailboxes**.

2. In the list of mailboxes, find the mailboxes that you want to modify. You can:

   - Scroll through the list of mailboxes.

   - Click **Search** 🔎 and enter part of the user's name, email address, or alias.

   - Click **More options** ••• > **Advanced search** to find the mailbox.

3. In the list of mailboxes, select multiple mailboxes of the same type (for example, **User**) from the list. For

example:

- Select a mailbox, hold down the Shift key, and select another mailbox that's farther down in the list.
- Hold down the CTRL key as you select each mailbox.

After you select multiple mailboxes of the same type, the title of the details pane changes to **Bulk Edit**.

4. In the details pane, scroll down to **MAPI**, click **Enable** or **Disable**, and then click **OK** in the warning message that appears.

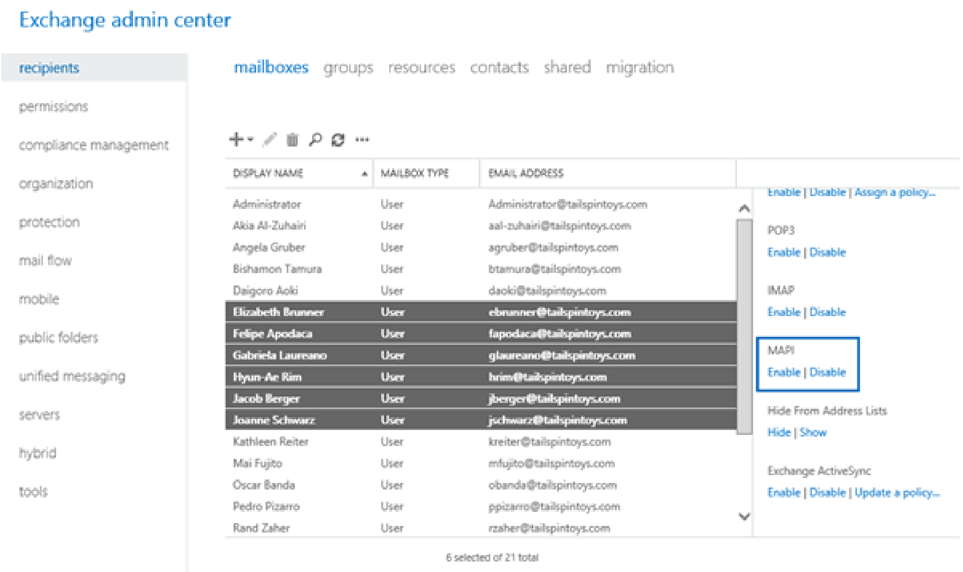

DISPLAY NAME	MAILBOX TYPE	EMAIL ADDRESS
Administrator	User	Administrator@tailspintoys.com
Akia Al-Zuhairi	User	aal-zuhairi@tailspintoys.com
Angela Gruber	User	agruber@tailspintoys.com
Bishamon Tamura	User	btamura@tailspintoys.com
Daigoro Aoki	User	daoki@tailspintoys.com
Elizabeth Brunner	User	ebrunner@tailspintoys.com
Felipe Apodaca	User	fapodaca@tailspintoys.com
Gabriela Laureano	User	glaureano@tailspintoys.com
Hyun-Ae Rim	User	hrim@tailspintoys.com
Jacob Berger	User	jberger@tailspintoys.com
Joanne Schwarz	User	jschwarz@tailspintoys.com
Kathleen Reiter	User	kreiter@tailspintoys.com
Mai Fujito	User	mfujito@tailspintoys.com
Oscar Banda	User	obanda@tailspintoys.com
Pedro Pizarro	User	ppizarro@tailspintoys.com
Rand Zaher	User	rzaher@tailspintoys.com

6 selected of 21 total

Use the Exchange Management Shell to enable or disable MAPI access to multiple mailboxes

You can use the **Get-Mailbox**, **Get-User** or **Get-Content** cmdlets to identify the mailboxes that you want to modify. For example:

- Use the *OrganizationalUnit* parameter to filter the mailboxes by organizational unit (OU).
- Use the *Filter* parameter to create OPATH filters that identify the mailboxes. For more information, see Filterable Properties for the -Filter Parameter.
- Use a text file to specify the mailboxes. The text file contains one mailbox (email address, name, or other unique identifier) on each line like this:

  > ebrunner@tailspintoys.com
  > fapodaca@tailspintoys.com
  > glaureano@tailspintoys.com
  > hrim@tailspintoys.com

This example disables MAPI access to all user mailboxes in the North America\Finance OU.

```
$NAFinance = Get-Mailbox -OrganizationalUnit "OU=Marketing,OU=North America,DC=contoso,DC=com" -Filter
"RecipientTypeDetails -eq 'UserMailbox'" -ResultSize Unlimited; $NAFinance | foreach {Set-CasMailbox
$_.Identity -MAPIEnabled $false}
```

This example disables MAPI access to all user mailboxes in the Engineering department in Washington state.

```
Get-User -Filter "RecipientType -eq 'UserMailbox' -and Department -like 'Engineering*' -and StateOrProvince -
eq 'WA'" | Set-CasMailbox -MAPIEnabled $false
```

This example uses the text file C:\My Documents\Accounts.txt to disable MAPI access to the specified mailboxes.

```
Get-Content "C:\My Documents\Accounts.txt" | foreach {Set-CasMailbox $_ -MAPIEnabled $false}
```

For detailed syntax and parameter information, see Get-Mailbox and Get-User.

How do you know this worked?

To verify that you've successfully enabled or disabled MAPI access to a mailbox, do any of these steps:

- In the EAC, go to **Recipients** > **Mailboxes** > select the mailbox > click **Edit** ✎ > **Mailbox features** and verify the **MAPI** value in the **Email Connectivity** section.

  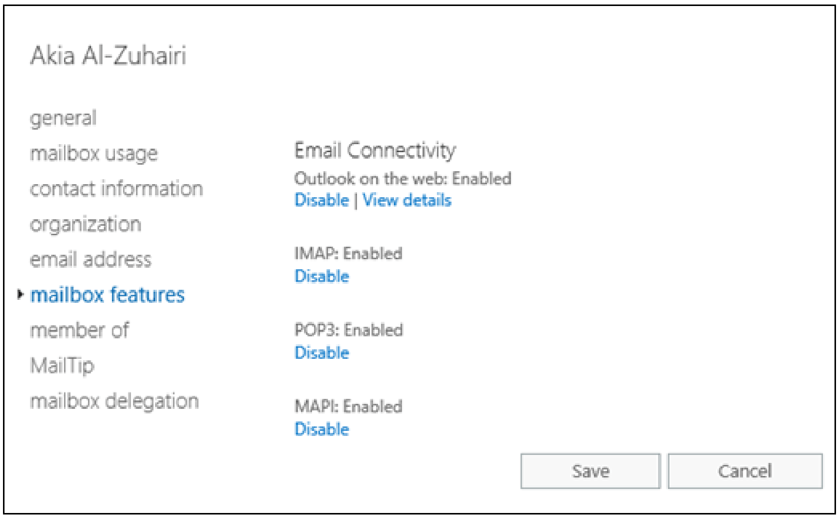

  Email Connectivity
  Outlook on the web: Enabled
  IMAP: Enabled
  POP3: Enabled
  MAPI: Enabled

- In the Exchange Management Shell, replace *<MailboxIdentity>* with the identity of the mailbox (for example, name, alias, or email address), and run this command:

  ```
  Get-CasMailbox -Identity "<MailboxIdentity>"
  ```

- Use the same filter that you used to identify the mailboxes, but use the **Get-CasMailbox** cmdlet instead of **Set-CasMailbox**. For example:

  ```
  Get-User -Filter "RecipientType -eq 'UserMailbox' -and Department -like 'Engineering*' -and
  StateOrProvince -eq 'WA'" | Get-CasMailbox
  ```

- In the Exchange Management Shell, run this command to show all mailboxes where Outlook on the web access is disabled:

  ```
  Get-CasMailbox -ResultSize unlimited -Filter "MAPIEnabled -eq `$false"
  ```

Exchange ActiveSync

8/3/2020 • 3 minutes to read • Edit Online

Exchange ActiveSync is an Exchange synchronization protocol that's optimized to work together with high-latency and low-bandwidth networks. The protocol, based on HTTP and XML, lets mobile phones access an organization's information on a server that's running Microsoft Exchange.

Overview of Exchange ActiveSync

Exchange ActiveSync lets mobile phone users access their email, calendar, contacts, and tasks, and lets them continue to access this information when they're working offline.

Standard encryption services add security to mobile communication with the server. You can configure Exchange ActiveSync to use Secure Sockets Layer (SSL) encryption for communications between the Exchange server and the mobile device.

Features in Exchange ActiveSync

Exchange ActiveSync provides the following:

- Support for HTML messages
- Support for follow-up flags
- Conversation grouping of email messages
- Ability to synchronize or not synchronize an entire conversation
- Synchronization of Short Message Service (SMS) messages with a user's Exchange mailbox
- Support for viewing message reply status
- Support for fast message retrieval
- Meeting attendee information
- Enhanced Exchange Search
- PIN reset
- Enhanced device security through password policies
- Autodiscover for over-the-air provisioning
- Support for setting automatic replies when users are away, on vacation, or out of the office
- Support for task synchronization
- Direct Push
- Support for availability information for contacts

Managing Exchange ActiveSync

By default, Exchange ActiveSync is enabled. All users who have an Exchange mailbox can synchronize their mobile device with the Microsoft Exchange server.

You can perform the following Exchange ActiveSync tasks:

- Enable and disable Exchange ActiveSync for users
- Set policies such as minimum password length, device locking, and maximum failed password attempts
- Initiate a remote wipe to clear all data from a lost or stolen mobile phone
- Run a variety of reports for viewing or exporting into a variety of formats
- Control which types of mobile devices can synchronize with your organization through device access rules

Managing mobile device access in Exchange ActiveSync

You can control which mobile devices can synchronize with Exchange Server. You do this by monitoring new mobile devices as they connect to your organization or by setting up rules that determine which types of mobile devices are allowed to connect. Regardless of the method you choose to specify which mobile devices can synchronize, you can approve or deny access for any specific mobile device for a specific user at any time.

Device security features in Exchange ActiveSync

In addition to the ability to configure security options for communications between the Exchange server and your mobile devices, Exchange ActiveSync offers the following features to enhance the security of mobile devices:

- **Remote wipe**: If a mobile device is lost, stolen, or otherwise compromised, you can issue a remote wipe command from the Exchange Server computer or from any Web browser by using Outlook Web App. This command erases all data from the mobile device.
- **Device password policies**: Exchange ActiveSync lets you configure several options for device passwords. The device password options include the following:
  - **Minimum password length (characters)**: This option specifies the length of the password for the mobile device. The default length is 4 characters, but as many as 18 can be included.
  - **Minimum number of character sets**: Use this text box to specify the complexity of the alphanumeric password and force users to use a number of different sets of characters from among the following: lowercase letters, uppercase letters, symbols, and numbers.
  - **Require alphanumeric password**: This option determines password strength. You can enforce the usage of a character or symbol in the password in addition to numbers.
  - **Inactivity time (seconds)**: This option determines how long the mobile device must be inactive before the user is prompted for a password to unlock the mobile device.
  - **Enforce password history**: Select this check box to force the mobile phone to prevent the user from reusing their previous passwords. The number that you set determines the number of past passwords that the user won't be allowed to reuse.
  - **Enable password recovery**: Select this check box to enable password recovery for the mobile device. Administrators can use the **Get-ActiveSyncDeviceStatistics** cmdlet to look up the user's recovery password.
  - **Wipe device after failed (attempts)**: This option lets you specify whether you want the phone's memory to be wiped after multiple failed password attempts.
- **Device encryption policies**: There are a number of mobile device encryption policies that you can enforce for a group of users. These policies include the following:
  - **Require encryption on device**: Select this check box to require encryption on the mobile device. This increases security by encrypting all information on the mobile device.
  - **Require encryption on storage cards**: Select this check box to require encryption on the mobile

device's removable storage card. This increases security by encrypting all information on the storage cards for the mobile device.

Enable or disable Exchange ActiveSync access to mailboxes in Exchange Server

8/3/2020 • 5 minutes to read • Edit Online

ActiveSync is a client protocol that lets users synchronize their Exchange mailbox with a mobile device. By default, ActiveSync is enabled on new user mailboxes. Disabling ActiveSync on a mailbox prevents the user from synchronizing their mailbox with a mobile device (by using ActiveSync).

Administrators can use the Exchange admin center (EAC) or the Exchange Management Shell to enable or disable Exchange ActiveSync access to a mailbox.

For more information about ActiveSync, see Exchange ActiveSync.

For information about setting up email on your mobile device, see these topics:

- Set up Office apps and email on iOS devices
- Set up Office apps and email on Android

For additional management tasks related to user access to mailboxes, see these topics:

- Enable or disable Outlook on the web access to mailboxes in Exchange Server
- Enable or disable POP3 or IMAP4 access to mailboxes in Exchange Server
- Enable or disable MAPI access to mailboxes in Exchange Server

What do you need to know before you begin?

- Estimated time to complete: 5 minutes.
- For more information about accessing and using the EAC, see Exchange admin center in Exchange Server.
- To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Exchange ActiveSync settings" entry in the Clients and mobile devices permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Enable or disable Exchange ActiveSync access to a single mailbox

Use the EAC to enable or disable Exchange ActiveSync access to a mailbox

1. In the EAC, go to **Recipients** > **Mailboxes**.

2. In the list of mailboxes, find the mailbox that you want to modify. You can:

   - Scroll through the list of mailboxes.
   - Click **Search** and enter part of the user's name, email address, or alias.
   - Click **More options** ••• > **Advanced search** to find the mailbox.

   Once you've found the mailbox that you want to modify, select it, and then click **Edit** .

3. On the mailbox properties page that opens, click **Mailbox features**.

4. In the **Mobile Devices** section, configure one of these settings:

   - If ActiveSync is enabled on the mailbox, you'll see a **Disable Exchange ActiveSync** link. Click the link to disable ActiveSync, and then click **Yes** in the warning message that appears.
   - If ActiveSync is disabled on the mailbox, you'll see a **Enable Exchange ActiveSync** link. Click the link to enable ActiveSync.

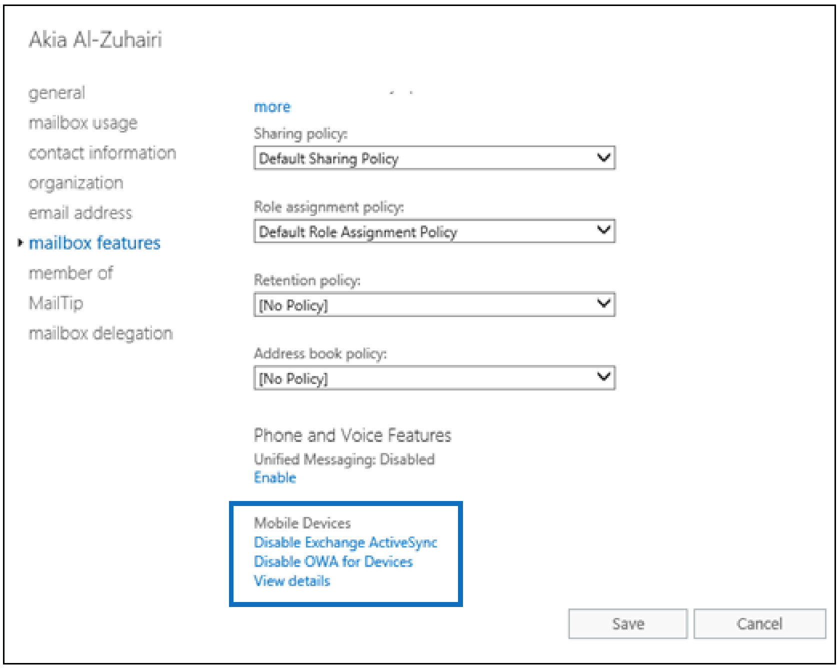

Akia Al-Zuhairi

Sharing policy: Default Sharing Policy

Role assignment policy: Default Role Assignment Policy

Retention policy: [No Policy]

Address book policy: [No Policy]

Phone and Voice Features
Unified Messaging: Disabled
Enable

Mobile Devices
Disable Exchange ActiveSync
Disable OWA for Devices
View details

   When you're finished, click **Save**.

Use the Exchange Management Shell to enable or disable Exchange ActiveSync access to a mailbox

To enable or disable ActiveSync access to a single mailbox, use this syntax:

```
Set-CasMailbox -Identity <MailboxIdentity> -ActiveSyncEnabled <$true | $false>
```

This example disables ActiveSync access to the mailbox named Yan Li.

```
Set-CasMailbox -Identity "Yan Li" -ActiveSyncEnabled $false
```

This example enables ActiveSync access to the mailbox named Elly Nkya.

```
Set-CasMailbox -Identity "Elly Nkya" -ActiveSyncEnabled $true
```

For detailed syntax and parameter information, see Set-CASMailbox.

Enable or disable Exchange ActiveSync access to multiple mailboxes

Use the EAC to enable or disable Exchange ActiveSync access to multiple mailboxes

1. In the EAC, go to **Recipients** > **Mailboxes**.

2. In the list of mailboxes, find the mailboxes that you want to modify. You can:

   - Scroll through the list of mailboxes.

   - Click **Search** 🔎 and enter part of the user's name, email address, or alias.

   - Click **More options** ••• > **Advanced search** to find the mailbox.

3. In the list of mailboxes, select multiple mailboxes of the same type (for example, **User**) from the list. For example:

   - Select a mailbox, hold down the Shift key, and select another mailbox that's farther down in the list.

   - Hold down the CTRL key as you select each mailbox.

   After you select multiple mailboxes of the same type, the title of the details pane changes to **Bulk Edit**.

4. In the details pane, scroll down to **Exchange ActiveSync**, click **Enable** or **Disable**, and then click **OK** in the warning message that appears.

   Bulk select mailboxes in the EAC to enable or disable Exchange ActiveSync

Use the Exchange Management Shell to enable or disable Exchange ActiveSync access to multiple mailboxes

You can use the **Get-Mailbox**, **Get-User** or **Get-Content** cmdlets to identify the mailboxes that you want to modify. For example:

- Use the *OrganizationalUnit* parameter to filter the mailboxes by organizational unit (OU).

- Use the *Filter* parameter to create OPATH filters that identify the mailboxes. For more information, see Filterable Properties for the -Filter Parameter.

- Use a text file to specify the mailboxes. The text file contains one mailbox (email address, name, or other unique identifier) on each line like this:

  ```
  ebrunner@tailspintoys.com
  fapodaca@tailspintoys.com
  glaureano@tailspintoys.com
  hrim@tailspintoys.com
  ```

This example disables ActiveSync access to all user mailboxes in the North America\Finance OU.

```
$NAFinance = Get-Mailbox -OrganizationalUnit "OU=Marketing,OU=North America,DC=contoso,DC=com" -Filter
"RecipientTypeDetails -eq 'UserMailbox'" -ResultSize Unlimited; $NAFinance | foreach {Set-CasMailbox
$_.Identity -ActiveSyncEnabled $false}
```

This example disables ActiveSync access to all user mailboxes in the Engineering department in Washington state.

```
Get-User -Filter "RecipientType -eq 'UserMailbox' -and Department -like 'Engineering*' -and StateOrProvince -
eq 'WA'" | Set-CasMailbox -ActiveSyncEnabled $false
```

This example uses the text file C:\My Documents\Accounts.txt to disable ActiveSync access to the specified mailboxes.

```
Get-Content "C:\My Documents\Accounts.txt" | foreach {Set-CasMailbox $_ -ActiveSyncEnabled $false}
```

For detailed syntax and parameter information, see Get-Mailbox and Get-User.

How do you know this worked?

To verify that you've successfully enabled or disabled Exchange ActiveSync access to a mailbox, do any of these steps:

- In the EAC, go to **Recipients** > **Mailboxes** > select the mailbox > click **Edit** ✎ > **Mailbox features** > and verify the Exchange ActiveSync value in the **Mobile Devices** section.
    - If ActiveSync access is enabled for the mailbox, you'll see **Disable Exchange ActiveSync**.
    - If ActiveSync access is disabled for the mailbox, you'll see **Enable Exchange ActiveSync**.

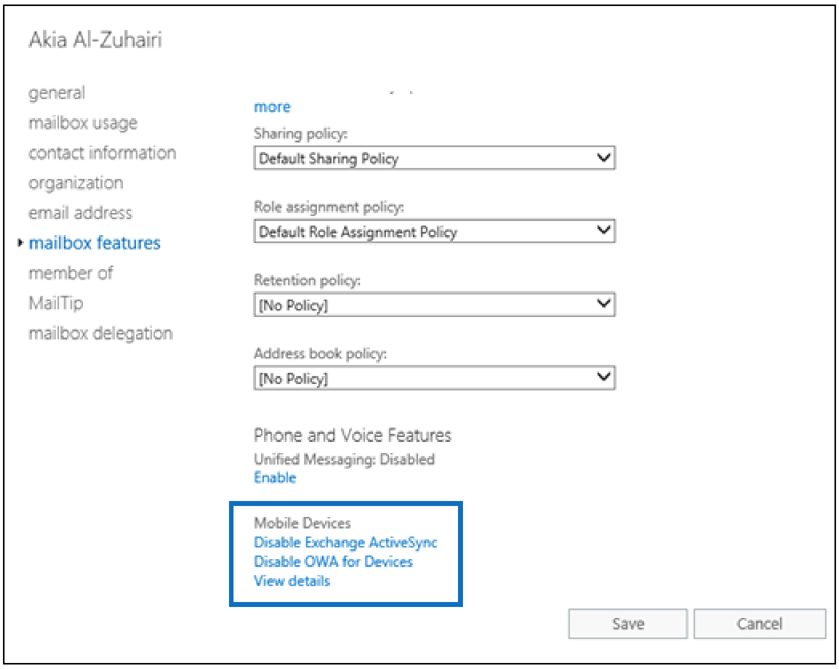

Akia Al-Zuhairi

Sharing policy: Default Sharing Policy

Role assignment policy: Default Role Assignment Policy

Retention policy: [No Policy]

Address book policy: [No Policy]

Phone and Voice Features
Unified Messaging: Disabled
Enable

Mobile Devices
Disable Exchange ActiveSync
Disable OWA for Devices
View details

- In the Exchange Management Shell, replace *<MailboxIdentity>* with the identity of the mailbox (for example, name, alias, or email address), and run this command:

```
Get-CasMailbox -Identity "<MailboxIdentity>"
```

- Use the same filter that you used to identify the mailboxes, but use the **Get-CasMailbox** cmdlet instead of **Set-CasMailbox**. For example:

```
Get-User -Filter "RecipientType -eq 'UserMailbox' -and Department -like 'Engineering*' -and
StateOrProvince -eq 'WA'" | Get-CasMailbox
```

- In the Exchange Management Shell, run this command to show all mailboxes where ActiveSync access is disabled:

```
Get-CasMailbox -ResultSize unlimited -Filter "ActiveSyncEnabled -eq `$false"
```

Mobile device mailbox policies

8/3/2020 • 9 minutes to read • Edit Online

In Exchange Server, you can create mobile device mailbox policies to apply a common set of policies or security settings to a collection of users. After you deploy Exchange ActiveSync in your Exchange Server organization, you can create new mobile device mailbox policies or modify existing policies. When you install Exchange Server, a default mobile device mailbox policy is created. All users are automatically assigned this default mobile device mailbox policy.

Overview of mobile device mailbox policies

You can use mobile device mailbox policies to manage many different settings. These include the following:

- Require a password
- Specify the minimum password length
- Require a number or special character in the password
- Designate how long a device can be inactive before requiring the user to re-enter a password
- Wipe a device after a specific number of failed password attempts

For more information about all the settings you can configure, see Mobile device policy settings.

Exchange mobile device mailbox policies

Exchange ActiveSync is a client protocol that lets you synchronize a mobile device with your Exchange mailbox. Exchange ActiveSync is enabled by default when you install Exchange Server.

You can create mobile device mailbox policies in the Exchange admin center (EAC) or the Exchange Management Shell. If you create a policy in the EAC, you can configure only a subset of the available settings. You can configure the rest of the settings using the Exchange Management Shell.

Mobile device password settings and biometrics

Many mobile devices support biometrics such as Apple Touch ID or Face ID. Exchange mobile device mailbox policies do not control whether biometrics can be used instead of typing the device PIN. Mobile device mailbox policies can be configured to require a device PIN, but then the users control whether they use biometrics after complying with the device PIN requirement.

Mobile device password settings and Android

Android 9.0 and earlier versions utilize Android's device admin functionality to manage device password settings defined in a mobile device mailbox policy.

With Android 10.0 and later, Android has removed device admin functionality. Instead, apps that require a screen lock query the device's (or the work profile's) screen lock complexity. Apps that require a stronger screen lock direct the user to the system screen lock settings, allowing the user to update the security settings to become compliant. At no time is the app aware of the user's password; the app is only aware of the password complexity level. Android supports the following four password complexity levels:

PASSWORD COMPLEXITY LEVEL	PASSWORD REQUIREMENTS
None	No password requirements are configured
Low	Password can be a pattern or a PIN with either repeating (4444) or ordered (1234, 4321, 2468) sequences
Medium	Passwords that meet one of the following criteria: - PIN with no repeating (4444) or ordered (1234, 4321, 2468) sequences with a minimum length of 4 characters - Alphabetic passwords with a minimum length of 4 characters - Alphanumeric passwords with a minimum length of 4 characters
High	Passwords that meet one of the following criteria: - PIN with no repeating (4444) or ordered (1234, 4321, 2468) sequences with a minimum length of 8 characters - Alphabetic passwords with a minimum length of 6 characters - Alphanumeric passwords with a minimum length of 6 characters

From the perspective of an Exchange mobile device mailbox policy, Android's password complexity levels are mapped to the following policy settings:

MOBILE DEVICE MAILBOX POLICY SETTING	ANDROID PASSWORD COMPLEXITY LEVEL
Password enabled = false	None
Allow simple password = true Min password length < 4	Low
Allow simple password = true Min password length < 6	Medium
Allow simple password = false Alphanumeric password required = true Min password length < 6	Medium
Allow simple password = true Min password length > 6	High
Allow simple password = false Alphanumeric password required = true Min password length >= 6	High

Mobile device mailbox policy settings

The following table summarizes the settings you can specify using mobile device mailbox policies.

SETTING	DESCRIPTION
Allow Bluetooth	This setting specifies whether a mobile device allows Bluetooth connections. The available options are Disable, HandsFree Only, and Allow. The default value is Allow.

SETTING	DESCRIPTION
Allow Browser	This setting specifies whether Pocket Internet Explorer is allowed on the mobile device. This setting doesn't affect third-party browsers installed on the mobile device. The default value is `$true`.
Allow Camera	This setting specifies whether the mobile device camera can be used. The default value is `$true`.
Allow Consumer EMail	This setting specifies whether the mobile device user can configure a personal email account (either POP3 or IMAP4) on the mobile device. The default value is `$true`. This setting doesn't control access to email accounts that are using third-party mobile device email programs.
Allow Desktop Sync	This setting specifies whether the mobile device can synchronize with a computer through a cable, Bluetooth, or IrDA connection. The default value is `$true`.
Allow External Device Management	This setting specifies whether an external device management program is allowed to manage the mobile device.
Allow HTML Email	This setting specifies whether email synchronized to the mobile device can be in HTML format. If this setting is set to `$false`, all email is converted to plain text.
Allow Internet Sharing	This setting specifies whether the mobile device can be used as a modem for a desktop or a portable computer. The default value is `$true`.
AllowIrDA	This setting specifies whether infrared connections are allowed to and from the mobile device.
Allow Mobile OTA Update	This setting specifies whether the mobile device mailbox policy settings can be sent to the mobile device over a cellular data connection. The default value is `$true`.
Allow non-provisionable devices	This setting specifies whether mobile devices that may not support application of all policy settings are allowed to connect to Exchange Server by using Exchange ActiveSync. Allowing non-provisionable mobile devices has security implications. For example, some non-provisionable devices may not be able to implement an organization's password requirements.
Allow POPIMAPEmail	This setting specifies whether the user can configure a POP3 or an IMAP4 email account on the mobile device. The default value is `$true`. This setting doesn't control access by third-party email programs.
Allow Remote Desktop	This setting specifies whether the mobile device can initiate a remote desktop connection. The default value is `true`.
Allow simple password	This setting enables or disables the ability to use a simple password such as 1111 or 1234. The default value is `$true`.

SETTING	DESCRIPTION
Allow S/MIME encryption algorithm negotiation	This setting specifies whether the messaging application on the mobile device can negotiate the encryption algorithm if a recipient's certificate doesn't support the specified encryption algorithm.
Allow S/MIME software certificates	This setting specifies whether S/MIME software certificates are allowed on the mobile device.
Allow storage card	This setting specifies whether the mobile device can access information that's stored on a storage card.
Allow text messaging	This setting specifies whether text messaging is allowed from the mobile device. The default value is `$true`.
Allow unsigned applications	This setting specifies whether unsigned applications can be installed on the mobile device. The default value is `$true`.
Allow unsigned installation packages	This setting specifies whether an unsigned installation package can be run on the mobile device. The default value is `$true`.
Allow Wi-Fi	This setting specifies whether wireless Internet access is allowed on the mobile device. The default value is `$true`.
Alphanumeric password required	This setting requires that a password contains numeric and non-numeric characters. The default value is `$true`.
Approved Application List	This setting stores a list of approved applications that can be run on the mobile device.
Attachments enabled	This setting enables attachments to be downloaded to the mobile device. The default value is `$true`.
Device encryption enabled	This setting enables encryption on the mobile device. Not all mobile devices can enforce encryption. For more information, see the device and mobile operating system documentation.
Device policy refresh interval	This setting specifies how often the mobile device mailbox policy is sent from the server to the mobile device.
IRM enabled	This setting specifies whether Information Rights Management (IRM) is enabled on the mobile device.
Max attachment size	This setting controls the maximum size of attachments that can be downloaded to the mobile device. The default value is Unlimited.

SETTING	DESCRIPTION
Max calendar age filter	This setting specifies the maximum range of calendar days that can be synchronized to the mobile device. The following values are accepted: 1: All 2: OneDay 3: ThreeDays 4: OneWeek 5: TwoWeeks 6: OneMonth
Max email age filter	This setting specifies the maximum number of days of email items to synchronize to the mobile device. The following values are accepted: 1: All 2: OneDay 3: ThreeDays 4: OneWeek 5: TwoWeeks 6: OneMonth
Max email body truncation size	This setting specifies the maximum size at which email messages are truncated when synchronized to the mobile device. The value is in kilobytes (KB).
Max email HTML body truncation size	This setting specifies the maximum size at which HTML email messages are truncated when synchronized to the mobile device. The value is in kilobytes (KB).
Max inactivity time lock	This value specifies the length of time that the mobile device can be inactive before a password is required to reactivate it. You can enter any interval between 30 seconds and 1 hour. The default value is 15 minutes.
Max password failed attempts	This setting specifies the number of attempts a user can make to enter the correct password for the mobile device. You can enter any number from 4 through 16. The default value is 8.
Min password complex characters	This setting specifies the number of character sets that are required in the password of the mobile device. The character sets are: * Lower case letters. * Upper case letters. * Digits 0 through 9. * Special characters (for example, exclamation marks). You can enter any number from 1 through 4. The default value is 1.
Min password length	This setting specifies the minimum number of characters in the mobile device password. You can enter any number from 1 through 16. The default value is 4.
Password enabled	This setting enables the mobile device password.
Password expiration	This setting enables the administrator to configure a length of time after which a mobile device password must be changed.

SETTING	DESCRIPTION
Password history	This setting specifies the number of past passwords that can be stored in a user's mailbox. A user can't reuse a stored password.
Password recovery enabled	When this setting is enabled, the mobile device generates a recovery password that's sent to the server. If the user forgets their mobile device password, the recovery password can be used to unlock the mobile device and enable the user to create a new mobile device password.
Require device encryption	This setting specifies whether device encryption is required. If set to `$true`, the mobile device must be able to support and implement encryption to synchronize with the server.
Require encrypted S/MIME messages	This setting specifies whether S/MIME messages must be encrypted. The default value is `$false`.
Require encryption S/MIME algorithm	This setting specifies what required algorithm must be used when encrypting S/MIME messages.
Require manual synchronization while roaming	This setting specifies whether the mobile device must synchronize manually while roaming. Allowing automatic synchronization while roaming will frequently lead to larger-than-expected data costs for the mobile device data plan.
Require signed S/MIME algorithm	This setting specifies what required algorithm must be used when signing a message.
Require signed S/MIME messages	This setting specifies whether the mobile device must send signed S/MIME messages.
Require storage card encryption	This setting specifies whether the storage card must be encrypted. Not all mobile device operating systems support storage card encryption. For more information, see your mobile device and mobile operating system documentation.
Unapproved InROM application list	This setting specifies a list of applications that cannot be run in ROM.

Mobile devices

8/3/2020 • 2 minutes to read • Edit Online

There are many different mobile phones and devices enabled for Exchange ActiveSync. These include Android phones and tablets, as well as the Apple iPhone, iPod, and iPad.

Both phone and non-phone mobile devices support Exchange ActiveSync, and in most Exchange ActiveSync documentation, we use the term *mobile device*. Unless the feature or features we're discussing require a cellular telephone signal, such as SMS message notification, the term mobile device applies to both mobile phones and other mobile devices such as tablets.

What Exchange ActiveSync does

Exchange ActiveSync is a communications protocol that enables over-the-air mobile access to email messages, scheduling data, contacts, and tasks. Exchange ActiveSync is available on third-party phones that are enabled for Exchange ActiveSync.

Exchange ActiveSync offers Direct Push technology. Direct Push uses an encrypted HTTPS connection that's established and maintained between the mobile device and the server to push new email messages and other Exchange data to the phone.

To use Direct Push with Microsoft Exchange Server 2013, your users must have a mobile device that's designed to support Direct Push.

Exchange ActiveSync features

Exchange ActiveSync provides access to many different features that enable you to enforce security policies on mobile devices. By using Exchange 2013, you can configure multiple mobile device mailbox policies and control which mobile devices can synchronize with your Exchange server. Exchange ActiveSync enables you to send a remote device wipe command that wipes all data from a mobile device in case that mobile device is lost or stolen. Users can also initiate a remote device wipe from Outlook Web App.

Exchange ActiveSync lets users generate a recovery password. This recovery password is saved on the mobile device and is used when a user forgets their password. The user generates the recovery password at the same time that they generate the device password or PIN. This recovery password can be used to unlock the mobile device. Immediately after this recovery password is used, the user will be required to create a new PIN.

POP3 and IMAP4 limitations

If your mobile device doesn't support Exchange ActiveSync, or you don't need the rich feature set that Exchange ActiveSync provides, you can use POP3 or IMAP4 to access your email on your mobile device. For more information about POP3 and IMAP4 access to your mailbox, see POP3 and IMAP4 in Exchange Server.

Configure mobile phones to access email

8/3/2020 • 2 minutes to read • Edit Online

This article is about enabling users in your organization to access their Exchange 2016 or Exchange 2019 mailboxes with their mobile devices using Exchange ActiveSync.

Prerequisites

- You've reviewed the manufacturer's documentation for the mobile phone you want to configure.
- Exchange ActiveSync is enabled in your organization.

> **NOTE**
>
> For device-specific information about setting up Exchange-based email on a phone or tablet, see Set up Office apps and email on a mobile device.

Configure a mobile phone or device to use Exchange ActiveSync

Most mobile phones and devices are capable of using Autodiscover in Exchange to configure the mobile email client to use Exchange ActiveSync. To configure an email account on most mobile devices, you'll need two pieces of information.

- The user's email address
- The user's password

If the mobile phone is unable to contact the Exchange server automatically through the Autodiscover service, you'll need to set up the mobile phone manually. Manual setup requires the user's email address and password, as well as the Exchange ActiveSync server name. In most organizations, the Exchange ActiveSync server name is the same as the Outlook on the web server name without the /owa, for example, mail.contoso.com.

Perform a remote wipe on a mobile phone

8/3/2020 • 5 minutes to read • Edit Online

Your users carry sensitive corporate information in their pockets every day. If one of them loses their mobile phone, your data can end up in the hands of another person. If one of your users loses their mobile phone, you can use the Exchange admin center (EAC) or the Exchange Management Shell to wipe their phone clean of all corporate and user information.

> **NOTE**
>
> This topic also provides instructions for how to use Outlook on the web to perform a remote wipe on a phone. The user must be signed in to Outlook on the web to perform a remote wipe.

What do you need to know before you begin?

- Estimated time to complete: 5 minutes.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Mobile devices" entry in the Clients and mobile devices permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **NOTE**
>
> Prior to EAS v16.1, remote wipe would perform a device-level wipe, restoring the device to factory conditions. With EAS v16.1 and later, EAS also supports account-only remote wipe. In order for this to work, the client must support the EAS v16.1 protocol. If the client doesn’t support v16.1, the wipe will fail and an error will be given.

Caution

Exchange ActiveSync v16.1 supports two different remote wipe processes: A **Wipe Data** remote wipe and also an **Account Only Remote Wipe Device** remote wipe. There are important differences between how Outlook responds and how native mail apps on iOS and Android respond to these different wipe commands.

Outlook for iOS and Outlook for Android support only the **Wipe Data** command, which wipes only data within Outlook. The Outlook app will reset and all Outlook email, calendar, contacts, and file data will be removed, but no other data is wiped from the device. The **Account Only Remote Wipe Device** command is therefore redundant and is not supported by Outlook for iOS or Android.

However, if a native iOS or Android mail app is connected to Exchange and receives a **Wipe Data** command from Exchange ActiveSync, all data on the device will be wiped, including photos, personal files, and so on.

If a native iOS or Android mail app is connected to Exchange and receives an **Account Only Remote Wipe Device** command from Exchange ActiveSync, only the native mail app's Exchange ActiveSync mail, calendar, and account data are wiped.

These commands are designed to destroy data. Exercise caution when using them.

After the remote wipe command is requested by the administrator, the wipe happens within seconds of the Outlook app's next connection to Exchange.

Since Outlook for iOS and Android appears as a single mobile device association under a user's mobile devices in

Exchange, a remote wipe command will remove data and delete sync relationships from all devices running Outlook (iPhone, iPad, Android) associated with that user.

A remote wipe action deletes the synchronization profile, so when the user adds his or her account to Outlook for iOS and Android, a new Device ID is generated and reported to Exchange on-premises.

> **NOTE**
>
> If you are using Intune, you should be using Intune to trigger data removal, not Exchange. Depending on the scenario, it could be accomplished via App Protection Policy selective wipe, or Device enrollment retire/wipe commands.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to wipe a user's phone

You can use the EAC to wipe a user's phone or cancel a remote wipe that has not yet completed.

1. In the EAC, navigate to **Recipients** > **Mailboxes**.

2. Select the user, and under **Mobile Devices**, choose **View details**.

3. On the **Mobile Device Details** page, select the lost mobile device, and then select **Wipe Data** (or **Account Only Remote Wipe Device** if desired).

4. Select **Save**.

Use the Exchange Management Shell to wipe a user's phone

You can use the **Clear-MobileDevice** cmdlet in the Exchange Management Shell to wipe a user's phone.

The following command wipes the device named WM_TonySmith and sends a confirmation message to admin@contoso.com.

```
Clear-MobileDevice -Identity WM_TonySmith -NotificationEmailAddresses "admin@contoso.com"
```

If the device connects to Exchange using a mail app other than Outlook, you can use the following command to wipe only the mail app's Exchange ActiveSync mail, calendar, and account data and leave all other data on the device intact:

```
Clear-MobileDevice -AccountOnly -Identity WM_TonySmith -NotificationEmailAddresses "admin@contoso.com"
```

The **-AccountOnly** switch has no effect on Outlook devices because an account-only remote wipe is the only type of wipe that is supported by Outlook. See Clear-MobileDevice for more information.

Use Outlook on the web to wipe a user's phone

Your users can use Outlook on the web to wipe their own phones.

1. In Outlook on the web, select the **Settings** icon.

2. Under **Your app settings**, select **Mail**.

3. Under **Options**, click to expand **General** if necessary, and then select **Mobile devices**.

4. Select the mobile phone.

5. Click or tap the **Wipe Device** icon (or the **Account Only Remote Wipe Device** icon if desired).

Use the New Outlook on the web to wipe a user's phone

1. In Outlook on the web, select the **Settings** icon.

2. Click on **View All Outlook settings**.

3. Click **General**, and then select **Mobile devices**.

4. Select the mobile phone.

5. Click or tap the **Wipe Device** icon (or the **Account Only Remote Wipe Device** icon if desired).

How do you know this worked?

There are several ways to verify that the remote wipe completed.

- Run the **Clear-MobileDevice** cmdlet with the *-NotificationEmailAddresses* parameter configured. A message will be sent to the supplied email address when the remote wipe has completed.

- In the EAC, check the status of the mobile device. The status will change from **Wipe Pending** to **Wipe Successful**.

- In Outlook on the web, check the status of the mobile device. The status will change from **Wipe Pending** to **Wipe Successful**.

> **NOTE**
>
> In a Microsoft 365 or Office 365-based environment, the result of a remote device wipe is not reported back to Exchange. Even when the wipe is successful, the status will display as **Pending**.

POP3 and IMAP4 in Exchange Server

8/3/2020 • 5 minutes to read • Edit Online

Although users typically access their Exchange mailboxes by using Outlook (MAPI), Outlook on the web (formerly known as Outlook Web App), and Exchange ActiveSync, POP3 and IMAP4 are available in Exchange Server 2016 and Exchange Server 2019. To support clients that still rely on these protocols, you need to start the services, and configure the settings for POP3 and IMAP4. For detailed instructions, see the following topics:

- Enable and configure POP3 on an Exchange server
- Enable and configure IMAP4 on an Exchange server
- Configure authenticated SMTP settings for POP3 and IMAP4 clients in Exchange Server

After you enable and configure POP3 or IMAP4 on the Exchange server, you can enable or disable POP3 or IMAP4 access to specific mailboxes. For more information, see Enable or disable POP3 or IMAP4 access to mailboxes in Exchange Server.

Note: Clients connect to the POP3 and IMAP4 services in the Client Access (frontend) services on the Mailbox server. They never connect directly to the POP3 and IMAP4 backend services. For more information, see Client access protocol architecture.

POP3 and IMAP4 improvements in Exchange Server

POP3 and IMAP4 functionality in Exchange 2016 and Exchange 2019 is basically unchanged from Exchange 2013. These are the improvements in POP3 and IMAP4 as compared to Exchange 2010:

- By default, the Client Access (frontend) services in Exchange 2016 and 2019 automatically proxy POP3 and IMAP4 client connections from one Active Directory site to the correct Mailbox server in a different Active Directory site. In previous versions of Exchange, you had to perform a manual configuration step to allow POP3 and IMAP4 clients to connect to their mailboxes from one site to another.
- You can't use the Anonymous or Guest accounts to access an Exchange 2016 or Exchange 2019 mailbox by using POP3 or IMAP4. Access is blocked to prevent security vulnerabilities when you use non-standard accounts for POP3 and IMAP4 access.
- You can't connect to the Administrator mailbox by using POP3 or IMAP4 (you can use Outlook or Outlook Web App). This limitation was intentionally included in Exchange 2016 to enhance security for the Administrator mailbox.

Overview of POP3 and IMAP4 functionality

The POP3 and IMAP4 protocols have the following benefits and limitations:

- **POP3**
  - Designed for offline message processing.
  - Can only download messages from a single folder (usually the Inbox) in the mailbox to a single folder in the POP3 application on the client computer or device.
  - By default, downloaded messages are removed from the email server, and are stored only on the local computer or device. Therefore, users can't access the same email messages from multiple computers or devices (although many POP3 applications can be configured to keep copies of

downloaded messages in the mailbox on the email server).

  - Doesn't offer advanced collaboration features such as calendaring, contacts, and tasks.

- IMAP4
  - Offers offline and online message processing.
  - Can synchronize messages from multiple folders in the mailbox with the client computer or device. For example, most IMAP4 applications can be configured to keep a copy of sent messages in the mailbox on the email server.
  - By default, copies of downloaded messages remain on the email server. Therefore, users can access the same messages from multiple computers.
  - Supports additional features. For example, you can download the message headers (the message's sender and subject) before you decide to download the complete message.
  - Doesn't offer advanced collaboration features such as calendaring, contacts, and tasks.

Note: POP3 and IMAP4 clients have limited access to Exchange calendar information. For more information, see Configure Calendar Options for POP3 and Configure Calendar Options for IMAP4.

POP3 and IMAP4 applications and settings

After you've enabled and configured the required services, users can connect to their Exchange mailboxes by using any application that support POP3 and IMAP4. For example, Outlook, Windows Mail, and Mozilla Thunderbird. POP3 and IMAP4 feature support varies by application, so check the application's documentation.

Verify the POP3 or IMAP4 email program is configured to keep a copy of all messages on the server. This allows the users to access their messages from different computers or applications.

Another important setting is how frequently the email program contacts the server to send and receive mail. There are three basic settings:

- **Send and receive messages when the application is started**
- **Send and receive messages manually**: Messages are only sent and received when the user clicks a "send and receive" option in the application. This is a good setting for computers that aren't always connected to the Internet (for example, dial-up or metered Internet connections).
- **Send and receive messages every set number of minutes**: The application connects to the email server periodically to send messages and to download any new messages. This is a good setting for computers that are always connected to the Internet, because the application is kept up-to-date with the most current messages from the mailbox.

Note: If the application and server both support the IMAP4 **IDLE** command, users can send and receive messages in near real time (Exchange supports the IMAP4 **IDLE** command). In most cases, users don't need to configure any settings in their IMAP4 application to use this connection method.

To configure a POP3 or IMAP4 client to connect to a mailbox, users need specific information about the POP3 or IMAP4 settings. By default, Exchange uses the following settings for **internal** POP3 connections:

- **POP3 server FQDN**: `<ServerFQDN>`. For example, `mailbox01.contoso.com`.
- **TCP port and encryption method**: 995 for always SSL/TLS encrypted connections, and 110 for unencrypted connections, or for opportunistic TLS (**STARTTLS**) that results in an encrypted connection after the initial plain text protocol handshake.

To allow **external** POP3 clients to connect to mailboxes, you need to configure these settings for external connections. For more information, see Enable and configure POP3 on an Exchange server.

By default, Exchange uses the following settings for **internal** IMAP4 connections:

- **IMAP4 server FQDN**: `<ServerFQDN>`. For example, `mailbox01.contoso.com`.
- **TCP port and encryption method**: 993 for always SSL/TLS encrypted connections, and 143 for unencrypted connections, or for opportunistic TLS (**STARTTLS**) that results in an encrypted connection after the initial plain text protocol handshake.

To allow **external** IMAP4 clients to connect to mailboxes, you need to configure these settings for external connections. For more information, see Enable and configure IMAP4 on an Exchange server.

Enable and configure POP3 on an Exchange server

8/3/2020 • 8 minutes to read • Edit Online

By default, POP3 client connectivity isn't enabled in Exchange. To enable POP3 client connectivity, you need to perform the following steps:

1. Start the POP3 services, and configure the services to start automatically:

   - **Microsoft Exchange POP3**: This is the Client Access (frontend) service that POP3 clients connect to.

   - **Microsoft Exchange POP3 Backend**: POP3 client connections from the Client Access service are proxied to the backend service on the server that hold the active copy of the user's mailbox. For more information, see Client Access protocol architecture.

2. Configure the POP3 settings for external clients.

   By default, Exchange uses the following settings for **internal** POP3 connections:

   - **POP3 server FQDN**: `<ServerFQDN>`. For example, `mailbox01.contoso.com`.

   - **TCP port and encryption method**: 995 for always TLS encrypted connections, and 110 for unencrypted connections, or for opportunistic TLS (**STARTTLS**) that results in an encrypted connection after the initial plain text protocol handshake.

   To allow **external** POP3 clients to connect to mailboxes, you need to configure the POP3 server FQDN, TCP port, and encryption method for external connections. This step causes the external POP3 settings to be displayed in Outlook on the web (formerly known as Outlook Web App) at **Settings** > **Options** > **Mail** > **Accounts** > **POP and IMAP**.

   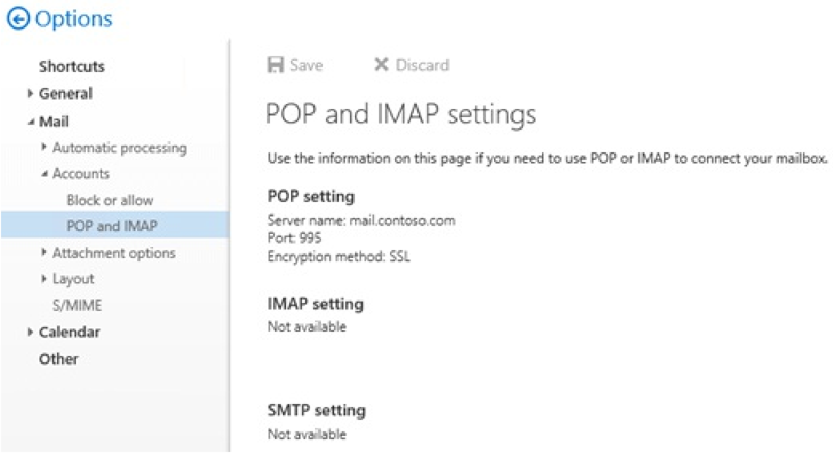

   POP and IMAP settings

   Use the information on this page if you need to use POP or IMAP to connect your mailbox.

   POP setting
   Server name: mail.contoso.com
   Port: 995
   Encryption method: SSL

   IMAP setting
   Not available

   SMTP setting
   Not available

3. Restart the POP3 services to save the changes.

4. Configure the authenticated SMTP settings for internal and external clients. For more information, see Configure authenticated SMTP settings for POP3 and IMAP4 clients in Exchange Server.

For more information about POP3, see POP3 and IMAP4 in Exchange Server.

What do you need to know before you begin?

- Estimated time to complete each procedure: 5 minutes.

- Secure Sockets Layer (SSL) is being replaced by Transport Layer Security (TLS) as the protocol that's used to encrypt data sent between computer systems. They're so closely related that the terms "SSL" and "TLS" (without versions) are often used interchangeably. Because of this similarity, references to "SSL" in Exchange topics, the Exchange admin center, and the Exchange Management Shell have often been used to encompass both the SSL and TLS protocols. Typically, "SSL" refers to the actual SSL protocol only when a version is also provided (for example, SSL 3.0). To find out why you should disable the SSL protocol and switch to TLS, check out Protecting you against the SSL 3.0 vulnerability.

- To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "POP3 and IMAP4 Permissions" section in the Clients and mobile devices permissions topic.

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Step 1: Start the POP3 services, and configure the services to start automatically

You can perform this step by using the Windows Services console, or the Exchange Management Shell.

Use the Windows Services console to start the POP3 services, and configure the services to start automatically

1. On the Exchange server, open the Windows Services console. For example:

   - Run the command `services.msc` from the **Run** dialog, a Command Prompt window, or the Exchange Management Shell.

   - Open Server Manager, and then click **Tools** > **Services**.

2. In the list of services, select **Microsoft Exchange POP3**, and then click **Action** > **Properties**.

3. The **Microsoft Exchange POP3 Properties** window opens. On the **General** tab, configure the following settings:

   - **Startup type**: Select **Automatic**.

   - **Service status**: Click **Start**.

   When you're finished, click **OK**.

4. In the list of services, select **Microsoft Exchange POP3 Backend**, and then click **Action** > **Properties**.

5. The **Microsoft Exchange POP3 Backend Properties** window opens. On the **General** tab, configure the following settings:

   - **Startup type**: Select **Automatic**.

   - **Service status**: Click **Start**.

   When you're finished, click **OK**.

Use the Exchange Management Shell to start the POP3 services, and configure the services to start automatically

1. Run the following command to start the POP3 services:

```
Start-Service MSExchangePOP3; Start-Service MSExchangePOP3BE
```

2. Run the following command to configure the POP3 services to start automatically:

```
Set-Service MSExchangePOP3 -StartupType Automatic; Set-Service MSExchangePOP3BE -StartupType Automatic
```

For more information about these cmdlets, see Start-Service and Set-Service.

How do you know this step worked?

To verify that you've successfully started the POP3 services, use either of the following procedures:

- On the Exchange server, open Windows Task Manager. On the **Services** tab, verify that the **Status** value for the **MSExchangePOP3** and **MSExchangePOP3BE** services is **Running**.

- In the Exchange Management Shell, run the following command to verify that the POP3 services are running:

```
Get-Service MSExchangePOP3; Get-Service MSExchangePOP3BE
```

Step 2: Use the Exchange Management Shell to configure the POP3 settings for external clients

To configure the POP3 settings for external clients, use the following syntax:

```
Set-PopSettings -ExternalConnectionSettings "<FQDN1>:<TCPPort1>:<SSL | TLS | blank>", "<FQDN2>:<TCPPort2>:<SSL
| TLS | blank>"...  -X509CertificateName <FQDN> [-SSLBindings "<IPv4Orv6Address1>:<TCPPort1>","
<IPv4Orv6Address2>:<TCPPort2>"...] [-UnencryptedOrTLSBindings "<IPv4Orv6Address1>:<TCPPort1>","
<IPv4Orv6Address2>:<TCPPort2>"...]
```

This example allows configures the following settings for external POP3 connections:

- **POP3 server FQDN**: mail.contoso.com

- **TCP port**: 995 for always TLS encrypted connections, and 110 for unencrypted connections or opportunistic TLS (STARTTLS) encrypted connections.

- **Internal Exchange server IP address and TCP port for always TLS encrypted connections**: All available IPv4 and IPv6 addresses on the server on port 995 (we aren't using the *SSLBindings* parameter, and the default value is `[::]:995,0.0.0.0:995`).

- **Internal Exchange server IP address and TCP port for unencrypted or opportunistic TLS (STARTTLS) encrypted connections**: All available IPv4 and IPv6 addresses on the server on port 110 (we aren't using the *UnencryptedOrTLSBindings* parameter, and the default value is `[::]:110,0.0.0.0:110`).

- **FQDN used for encryption**: mail.contoso.com. This value identifies the certificate that matches or contains the POP3 server FQDN.

```
Set-PopSettings -ExternalConnectionSettings "mail.contoso.com:995:SSL","mail.contoso.com:110:TLS" -
X509CertificateName mail.contoso.com
```

Notes:

- For detailed syntax and parameter information, see Set-PopSettings.

- The external POP3 server FQDN that you configure needs to have a corresponding record in your public DNS, and the TCP port (110 or 995) needs to be allowed through your firewall to the Exchange server.

- The combination of encryption methods and TCP ports that you use for the *ExternalConnectionSettings* parameter need to match the corresponding TCP ports and encryption methods that you use for the *SSLBindings* or *UnencryptedOrTLSBindings* parameters.

- Although you can use a separate certificate for POP3, we recommend that you use the same certificate as the other Exchange IIS (HTTP) services, which is likely a wildcard certificate or a subject alternative name (SAN) certificate from a commercial certification authority that's automatically trusted by all clients. For more information, see Certificate requirements for Exchange services.

- If you use a single subject certificate, or a SAN certificate, you also need to assign the certificate to the Exchange POP service. You don't need to assign a wildcard certificate to the Exchange POP service. For more information, see Assign certificates to Exchange Server services.

How you do know this step worked?

To verify that you've successfully configured the POP3 settings for external clients, run the following command in the Exchange Management Shell and verify the settings:

```
Get-PopSettings | Format-List *ConnectionSettings,*Bindings,X509CertificateName
```

For more information, see Get-POPSettings.

Step 3: Restart the POP3 services

After you enable and configure POP3, you need to restart the POP3 services on the server by using the Windows Services console, or the Exchange Management Shell.

Use the Windows Services console to restart the POP3 services

1. On the Exchange server, open the Windows Services console.

2. In the list of services, select **Microsoft Exchange POP3**, and then click **Action** > **Restart**.

3. In the list of services, select **Microsoft Exchange POP3 Backend**, and then click **Action** > **Restart**.

Use the Exchange Management Shell to restart the POP3 services

Run the following command to restart the POP3 services.

```
Restart-Service MSExchangePOP3; Restart-Service MSExchangePOP3BE
```

For more information about this cmdlet, see Restart-Service.

To verify that you've successfully restarted the POP3 services, run the following command:

```
Get-Service MSExchangePOP3; Get-Service MSExchangePOP3BE
```

Step 4: Configure the authenticated SMTP settings for POP3 clients

Because POP3 isn't used to send email messages, you need to configure the authenticated SMTP settings that are used by internal and external POP3 clients. For more information, see POP3 and IMAP4 in Exchange Server.

How do you know this task worked?

To verify that you have enabled and configured POP3 on the Exchange server, perform the following procedures:

1. Open a mailbox in Outlook on the web, and then click **Settings** > **Options**.

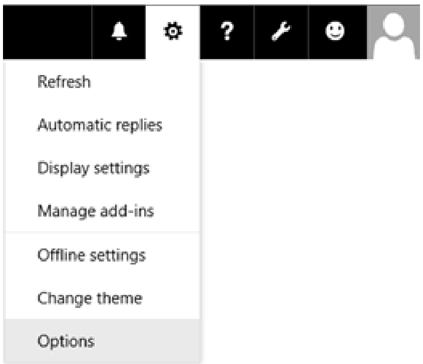

2. Click **Mail** > **Accounts** > **POP and IMAP** and verify the correct POP3 settings are displayed.

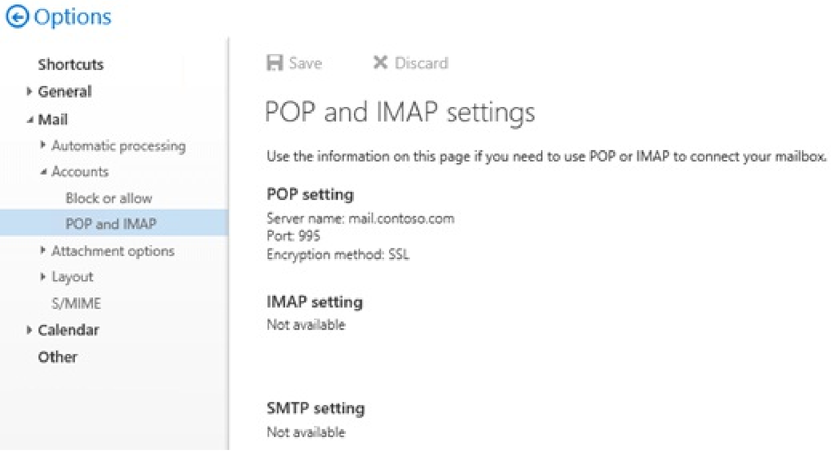

POP and IMAP settings

Use the information on this page if you need to use POP or IMAP to connect your mailbox.

POP setting
Server name: mail.contoso.com
Port: 995
Encryption method: SSL

IMAP setting
Not available

SMTP setting
Not available

Note: If you configured 995/SSL **and** 110/TLS values for the *ExternalConnectionSettings* parameter on the **Set-PopSettings** cmdlet, only the 995/SSL value is displayed in Outlook on the web. Also, if the external POP3 settings that you configured don't appear as expected in Outlook on the web after you restart the POP3 services, run the commands `net stop w3svc /y` and `net start w3svc` to restart Internet Information Services (IIS).

3. You can test POP3 client connectivity to the Exchange server by using the following methods:

- **Internal clients**: Use the **Test-PopConnectivity** cmdlet. For example,

```
Test-PopConnectivity -ClientAccessServer <ServerName> -Lightmode -MailboxCredential (Get-Credential)
```

. For more information, see Test-PopConnectivity.

Note: The *Lightmode* switch tells the command test POP3 logons to the server. To test sending (SMTP) and receiving (POP3) a message, you need to configure the authenticated SMTP settings as described in POP3 and IMAP4 in Exchange Server.

- **External clients**: Use the **POP Email** test in the Microsoft Remote Connectivity Analyzer.

Note: You can't use POP3 to connect to the Administrator mailbox. This limitation was intentionally included in Exchange 2016 and Exchange 2019 to enhance the security of the Administrator mailbox.

Next steps

To enabled or disable POP3 access to individual mailboxes, see Enable or disable POP3 or IMAP4 access to mailboxes in Exchange Server.

Enable and configure IMAP4 on an Exchange server

8/3/2020 • 8 minutes to read • Edit Online

By default, IMAP4 client connectivity isn't enabled in Exchange. To enable IMAP4 client connectivity, you need to perform the following steps:

1. Start the IMAP4 services, and configure the services to start automatically:

   - **Microsoft Exchange IMAP4**: This is the Client Access (frontend) service that IMAP4 clients connect to.

   - **Microsoft Exchange IMAP4 Backend**: IMAP4 client connections from the Client Access service are proxied to the backend service on the server that hold the active copy of the user's mailbox. For more information, see Exchange architecture.

2. Configure the IMAP4 settings for external clients.

   By default, Exchange uses the following settings for **internal** IMAP4 connections:

   - **IMAP4 server FQDN**: `<ServerFQDN>`. For example, `mailbox01.contoso.com`.

   - **TCP port and encryption method**: 993 for always TLS encrypted connections, and 143 for unencrypted connections, or for opportunistic TLS (**STARTTLS**) that results in an encrypted connection after the initial plain text protocol handshake.

   To allow **external** IMAP4 clients to connect to mailboxes, you need to configure the IMAP4 server FQDN, TCP port, and encryption method for external connections. This step causes the external IMAP4 settings to be displayed in Outlook on the web (formerly known as Outlook Web App) at **Settings** > **Options** > **Mail** > **Accounts** > **POP and IMAP**.

   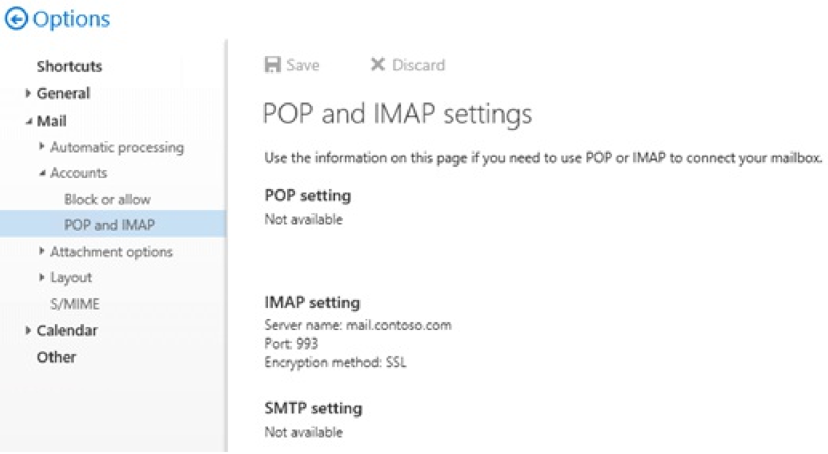

   POP and IMAP settings

   Use the information on this page if you need to use POP or IMAP to connect your mailbox.

   POP setting
   Not available

   IMAP setting
   Server name: mail.contoso.com
   Port: 993
   Encryption method: SSL

   SMTP setting
   Not available

3. Restart the IMAP4 services to save the changes.

4. Configure the authenticated SMTP settings for internal and external clients. For more information, see Configure authenticated SMTP settings for POP3 and IMAP4 clients in Exchange.

For more information about IMAP4, see POP3 and IMAP4 in Exchange Server.

What do you need to know before you begin?

- Estimated time to complete each procedure: 5 minutes.

- Secure Sockets Layer (SSL) is being replaced by Transport Layer Security (TLS) as the protocol that's used to encrypt data sent between computer systems. They're so closely related that the terms "SSL" and "TLS" (without versions) are often used interchangeably. Because of this similarity, references to "SSL" in Exchange topics, the Exchange admin center, and the Exchange Management Shell have often been used to encompass both the SSL and TLS protocols. Typically, "SSL" refers to the actual SSL protocol only when a version is also provided (for example, SSL 3.0). To find out why you should disable the SSL protocol and switch to TLS, check out Protecting you against the SSL 3.0 vulnerability.

- To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "POP3 and IMAP4 Permissions" section in the Clients and mobile devices permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Step 1: Start the IMAP4 services, and configure the services to start automatically

You can perform this step by using the Windows Services console, or the Exchange Management Shell.

Use the Windows Services console to start the IMAP4 services, and configure the services to start automatically

1. On the Exchange server, open the Windows Services console. For example:
   - Run the command `services.msc` from the **Run** dialog, a Command Prompt window, or the Exchange Management Shell.
   - Open Server Manager, and then click **Tools** > **Services**.
2. In the list of services, select **Microsoft Exchange IMAP4**, and then click **Action** > **Properties**.
3. The **Microsoft Exchange IMAP4 Properties** window opens. On the **General** tab, configure the following settings:
   - **Startup type**: Select **Automatic**.
   - **Service status**: Click **Start**.

   When you're finished, click **OK**.
4. In the list of services, select **Microsoft Exchange IMAP4 Backend**, and then click **Action** > **Properties**.
5. The **Microsoft Exchange IMAP4 Backend Properties** window opens. On the **General** tab, configure the following settings:
   - **Startup type**: Select **Automatic**.
   - **Service status**: Click **Start**.

   When you're finished, click **OK**.

Use the Exchange Management Shell to start the IMAP4 services, and configure the services to start

automatically

1. Run the following command to start the IMAP4 services:

   ```
   Start-Service MSExchangeIMAP4; Start-Service MSExchangeIMAP4BE
   ```

2. Run the following command to configure the IMAP4 services to start automatically:

   ```
   Set-Service MSExchangeIMAP4 -StartupType Automatic; Set-Service MSExchangeIMAP4BE -StartupType
   Automatic
   ```

For more information about these cmdlets, see Start-Service and Set-Service.

How do you know this step worked?

To verify that you've successfully started the IMAP4 services, use either of the following procedures:

- On the Exchange server, open Windows Task Manager. On the **Services** tab, verify that the **Status** value for the **MSExchangeIMAP4** and **MSExchangeIMAP4BE** services is **Running**.

- In the Exchange Management Shell, run the following command to verify that the IMAP4 services are running:

  ```
  Get-Service MSExchangeIMAP4; Get-Service MSExchangeIMAP4BE
  ```

Step 2: Use the Exchange Management Shell to configure the IMAP4 settings for external clients

To configure the IMAP4 settings for external clients, use the following syntax:

```
Set-ImapSettings -ExternalConnectionSettings "<FQDN1>:<TCPPort1>:<SSL | TLS | blank>", "<FQDN2>:<TCPPort2>:
<SSL | TLS | blank>"...  -X509CertificateName <FQDN> [-SSLBindings "<IPv4Orv6Address1>:<TCPPort1>","
<IPv4Orv6Address2>:<TCPPort2>"...] [-UnencryptedOrTLSBindings "<IPv4Orv6Address1>:<TCPPort1>","
<IPv4Orv6Address2>:<TCPPort2>"...]
```

This example allows configures the following settings for external IMAP4 connections:

- **IMAP4 server FQDN**: mail.contoso.com

- **TCP port**: 993 for always TLS encrypted connections, and 143 for unencrypted connections or opportunistic TLS (STARTTLS) encrypted connections.

- **Internal Exchange server IP address and TCP port for always TLS encrypted connections**: All available IPv4 and IPv6 addresses on the server on port 993 (we aren't using the *SSLBindings* parameter, and the default value is `[::]:993,0.0.0.0:993`).

- **Internal Exchange server IP address and TCP port for unencrypted or opportunistic TLS (STARTTLS) encrypted connections**: All available IPv4 and IPv6 addresses on the server on port 143 (we aren't using the *UnencryptedOrTLSBindings* parameter, and the default value is `[::]:143,0.0.0.0:143`).

- **FQDN used for encryption**: mail.contoso.com. This value identifies the certificate that matches or contains the IMAP4 server FQDN.

```
Set-ImapSettings -ExternalConnectionSettings "mail.contoso.com:993:SSL","mail.contoso.com:143:TLS" -
X509CertificateName mail.contoso.com
```

Notes:

- For detailed syntax and parameter information, see Set-IMAPSettings.
- The external IMAP4 server FQDN that you configure needs to have a corresponding record in your public DNS, and the TCP port (143 or 993) needs to be allowed through your firewall to the Exchange server.
- The combination of encryption methods and TCP ports that you use for the *ExternalConnectionSettings* parameter need to match the corresponding TCP ports and encryption methods that you use for the *SSLBindings* or *UnencryptedOrTLSBindings* parameters.
- Although you can use a separate certificate for IMAP4, we recommend that you use the same certificate as the other Exchange IIS (HTTP) services, which is likely a wildcard certificate or a subject alternative name (SAN) certificate from a commercial certification authority that's automatically trusted by all clients. For more information, see Certificate requirements for Exchange services.
- If you use a single subject certificate, or a SAN certificate, you also need to assign the certificate to the Exchange IMAP service. You don't need to assign a wildcard certificate to the Exchange IMAP service. For more information, see Assign certificates to Exchange Server services.

How you do know this step worked?

To verify that you've successfully configured the IMAP4 settings for external clients, run the following command in the Exchange Management Shell and verify the settings:

```
Get-ImapSettings | Format-List *ConnectionSettings,*Bindings,X509CertificateName
```

For more information, see Get-IMAPSettings.

Step 3: Restart the IMAP4 services

After you enable and configure IMAP4, you need to restart the IMAP4 services on the server by using the Windows Services console, or the Exchange Management Shell.

Use the Windows Services console to restart the IMAP4 services

1. On the Exchange server, open the Windows Services console.
2. In the list of services, select **Microsoft Exchange IMAP4**, and then click **Action** > **Restart**.
3. In the list of services, select **Microsoft Exchange IMAP4 Backend**, and then click **Action** > **Restart**.

Use the Exchange Management Shell to restart the IMAP4 services

Run the following command to restart the IMAP4 services.

```
Restart-Service MSExchangeIMAP4; Restart-Service MSExchangeIMAP4BE
```

For more information about this cmdlet, see Restart-Service.

To verify that you've successfully restarted the IMAP4 services, run the following command:

```
Get-Service MSExchangeIMAP4; Get-Service MSExchangeIMAP4BE
```

Step 4: Configure the authenticated SMTP settings for IMAP4 clients

Because IMAP4 isn't used to send email messages, you need to configure the authenticated SMTP settings that are used by internal and external IMAP4 clients. For more information, see Configure authenticated SMTP settings for

POP3 and IMAP4 clients in Exchange Server.

How do you know this task worked?

To verify that you have enabled and configured IMAP4 on the Exchange server, perform the following procedures:

1. Open a mailbox in Outlook on the web, and then click **Settings** > **Options**.

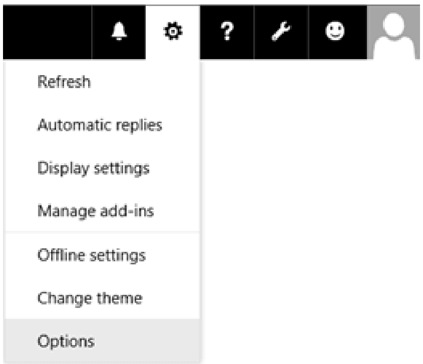

2. Click **Mail** > **Accounts** > **POP and IMAP** and verify the correct IMAP4 settings are displayed.

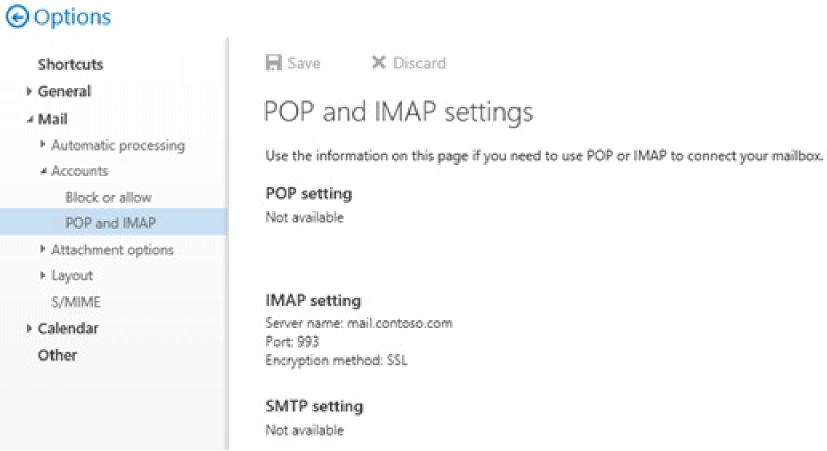

POP and IMAP settings

Use the information on this page if you need to use POP or IMAP to connect your mailbox.

POP setting
Not available

IMAP setting
Server name: mail.contoso.com
Port: 993
Encryption method: SSL

SMTP setting
Not available

Note: If you configured 993/SSL **and** 143/TLS values for the *ExternalConnectionSettings* parameter on the **Set-ImapSettings** cmdlet, only the 993/SSL value is displayed in Outlook on the web. Also, if the external IMAP4 settings that you configured don't appear as expected in Outlook on the web after you restart the IMAP4 services, run the commands `net stop w3svc /y` and `net start w3svc` to restart Internet Information Services (IIS).

3. You can test IMAP4 client connectivity to the Exchange server by using the following methods:

   - **Internal clients**: Use the **Test-ImapConnectivity** cmdlet. For example, `Test-ImapConnectivity -ClientAccessServer <ServerName> -Lightmode -MailboxCredential (Get-Credential)`. For more information, see Test-ImapConnectivity.

   Note: The *Lightmode* switch tells the command test IMAP4 logons to the server. To test sending (SMTP) and receiving (IMAP4) a message, you need to configure the authenticated SMTP settings as described in Configure authenticated SMTP settings for POP3 and IMAP4 clients in Exchange Server.

   - **External clients**: Use the **Imap Email** test in the Microsoft Remote Connectivity Analyzer.

   Note: You can't use IMAP4 to connect to the Administrator mailbox. This limitation was intentionally included in Exchange 2016 and Exchange 2019 to enhance the security of the Administrator mailbox.

Next steps

To enabled or disable IMAP4 access to individual mailboxes, see Enable or disable POP3 or IMAP4 access to mailboxes in Exchange Server.

Enable or disable POP3 or IMAP4 access to mailboxes in Exchange Server

8/3/2020 • 6 minutes to read • Edit Online

After you enable and configure POP3 or IMAP4 on an Exchange server as described in Enable and configure POP3 on an Exchange server and Enable and configure IMAP4 on an Exchange server, all user mailboxes (with the exception of the Administrator mailbox) can be accessed by using POP3 or IMAP4. You can use the procedures in this topic to disable POP3 and IMAP4 access to specific mailboxes.

For more information about POP3 and IMAP4, see POP3 and IMAP4 in Exchange Server.

What do you need to know before you begin?

- Estimated time to complete: 5 minutes.

- To open the Exchange admin center (EAC), see Exchange admin center in Exchange Server. To open the Exchange Management Shell, see Open the Exchange Management Shell.

- The procedures in this topic don't apply to the Administrator mailbox, because you can't use POP3 or IMAP4 to connect to the Administrator mailbox. This limitation was intentionally included in Exchange 2016 and Exchange 2019 to enhance the security of the Administrator mailbox.

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Recipient provisioning permissions" section in the Recipients Permissions topic.

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at Exchange Server.

Enable or disable POP3 or IMAP4 access to a single mailbox

Use the EAC to enable or disable POP3 or IMAP4 access to a mailbox

1. In the EAC, go to **Recipients** > **Mailboxes**.

2. In the list of mailboxes, find the mailbox that you want to modify. You can:

   - Scroll through the list of mailboxes.

   - Click **Search** (🔎) and enter part of the user's name, email address, or alias.

   - Click **More options** (•••) > **Advanced search** to find the mailbox.

   Once you've found the mailbox that you want to modify, select it, and then click **Edit** (✎).

3. In the mailbox properties window that opens, click **Mailbox Features**.

   In the **Email connectivity** section, configure one or more of the following settings:

   - **POP3**: To disable POP3 access to the mailbox, click **Disable**, and then click **Yes** in the warning message that appears. If POP3 is already disabled, click **Enable** to enable it.

- **IMAP**: To disable IMAP4 access to the mailbox, click **Disable**, and then click **Yes** in the warning message that appears. If IMAP4 is already disabled, click **Enable** to enable it.

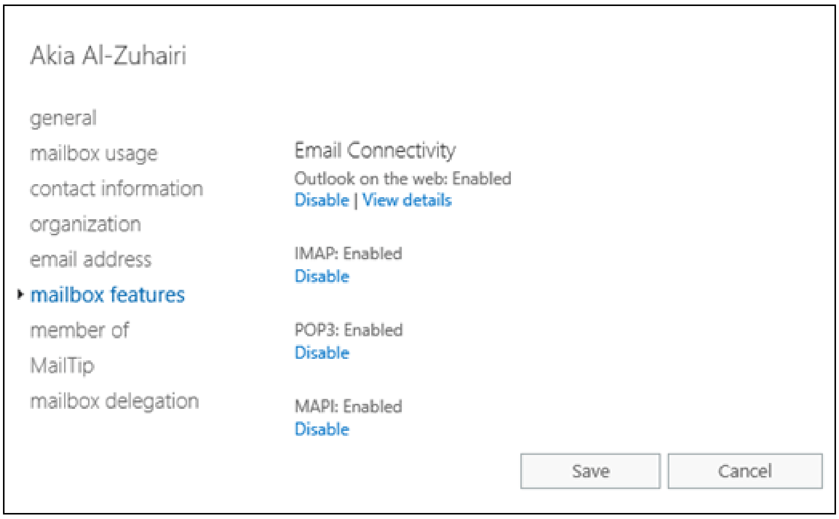

Akia Al-Zuhairi

Email Connectivity
Outlook on the web: Enabled

IMAP: Enabled

POP3: Enabled

MAPI: Enabled

When you're finished, click **Save**.

Use the Exchange Management Shell to enable or disable POP3 or IMAP4 access to a mailbox

To enable or disable POP3 or IMAP4 access to a single mailbox use the following syntax:

```
Set-CasMailbox -Identity <MailboxIdentity> -PopEnabled <$true | $false> -ImapEnabled <$true | $false>
```

This example disables POP3 and IMAP4 access to the mailbox named Rand Zaher.

```
Set-CasMailbox -Identity "Rand Zaher" -PopEnabled $false -ImapEnabled $false
```

This example enables POP3 and IMAP4 access to the mailbox named Rand Zaher.

```
Set-CasMailbox -Identity "Rand Zaher" -POPEnabled $true -ImapEnabled $true
```

For more information, see Set-CASMailbox.

Enable or disable POP3 or IMAP4 access to multiple mailboxes

Use the EAC to enable or disable POP3 or IMAP4 access to multiple mailboxes

1. In the EAC, go to **Recipients** > **Mailboxes**.

2. In the list of mailboxes, find the mailboxes that you want to modify. You can:

   - Scroll through the list of mailboxes.

   - Click **Search** (🔎) and enter part of the user's name, email address, or alias.

   - Click **More options** (•••) > **Advanced search** to find the mailbox.

3. In the list of mailboxes, select multiple mailboxes of the same type (for example, **User**) from the list. For example:

   - Select a mailbox, hold down the Shift key, and select another mailbox that's farther down in the list.

   - Hold down the CTRL key as you select each mailbox.

After you select multiple mailboxes of the same type, the title of the details pane changes to **Bulk Edit**.

4. In the details pane, go to **POP3** or **IMAP**, click **Enable** or **Disable**, and then click **OK** in the warning message that appears.

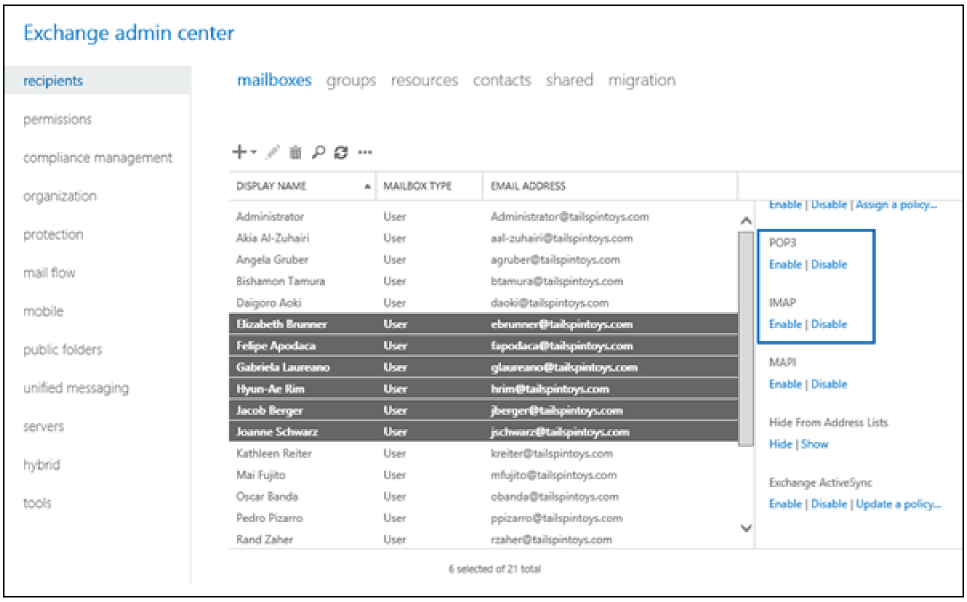

DISPLAY NAME	MAILBOX TYPE	EMAIL ADDRESS
Administrator	User	Administrator@tailspintoys.com
Akia Al-Zuhairi	User	aal-zuhairi@tailspintoys.com
Angela Gruber	User	agruber@tailspintoys.com
Bishamon Tamura	User	btamura@tailspintoys.com
Daigoro Aoki	User	daoki@tailspintoys.com
Elizabeth Brunner	User	ebrunner@tailspintoys.com
Felipe Apodaca	User	fapodaca@tailspintoys.com
Gabriela Laureano	User	glaureano@tailspintoys.com
Hyun-Ae Rim	User	hrim@tailspintoys.com
Jacob Berger	User	jberger@tailspintoys.com
Joanne Schwarz	User	jschwarz@tailspintoys.com
Kathleen Reiter	User	kreiter@tailspintoys.com
Mai Fujito	User	mfujito@tailspintoys.com
Oscar Banda	User	obanda@tailspintoys.com
Pedro Pizarro	User	ppizarro@tailspintoys.com
Rand Zaher	User	rzaher@tailspintoys.com

6 selected of 21 total

Use the Exchange Management Shell to enable or disable POP3 or IMAP4 access to multiple mailboxes

You can use the **Get-Mailbox**, **Get-User**, or **Get-Content** cmdlets to identify the mailboxes that you want to modify. For example:

- Use the *OrganizationalUnit* parameter to filter the mailboxes by organizational unit (OU).

- Use the *Filter* parameter to create OPATH filters that identify the mailboxes. For more information, see Filterable Properties for the -Filter Parameter.

- Use a text file to specify the mailboxes. The text file contains one mailbox (email address, name, or other unique identifier) on each line like this:

> ebrunner@tailspintoys.com
> fapodaca@tailspintoys.com
> glaureano@tailspintoys.com
> hrim@tailspintoys.com

This example disables POP3 and IMAP4 access to all user mailboxes in the North America\Finance OU.

```
$NAFinance = Get-Mailbox -OrganizationalUnit "OU=Marketing,OU=North America,DC=contoso,DC=com" -Filter
"RecipientTypeDetails -eq 'UserMailbox'" -ResultSize Unlimited; $NAFinance | foreach {Set-CasMailbox
$_.Identity -PopEnabled $false -ImapEnabled $false}
```

This example disables POP3 and IMAP4 access to all mailboxes in the Engineering department in Washington state.

```
Get-User -Filter "RecipientType -eq 'UserMailbox' -and Department -like 'Engineering*' -and StateOrProvince -
eq 'WA'" | Set-CasMailbox -PopEnabled $false -ImapEnabled $false
```

This example uses the text file C:\My Documents\Accounts.txt to disable POP3 or IMAP4 access to the specified mailboxes.

```
Get-Content "C:\My Documents\Accounts.txt" | foreach {Set-CASMailbox $_ -PopEnabled $false -ImapEnabled $false}
```

For more information, see Get-Mailbox and Get-User.

Restart the POP3 or IMAP4 services

After you change the POP3 or IMAP4 access settings on a mailbox, you need to restart the POP3 and IMAP4 services on the server. You can do this by using the Windows Services console, or the Exchange Management Shell.

Use the Windows Services console to restart the POP3 or IMAP4 services

1. On the Exchange server, open the Windows Services console. For example:

   - Run the command `services.msc` from the **Run** dialog, a Command Prompt window, or the Exchange Management Shell.

   - Open Server Manager, and then click **Tools** > **Services**.

2. In the list of services, do one or both of the following actions:

   - **POP3**:

     a. Select **Microsoft Exchange POP3**, and then click **Action** > **Restart**.

     b. Select **Microsoft Exchange POP3 Backend**, and then click **Action** > **Restart**.

   - **IMAP4**:

     a. Select **Microsoft Exchange IMAP4**, and then click **Action** > **Restart**.

     b. Select **Microsoft Exchange IMAP4 Backend**, and then click **Action** > **Restart**.

Use the Exchange Management Shell to restart the POP3 or IMAP4 services

To restart the POP3 services, run the following command:

```
Restart-Service MSExchangePOP3; Restart-Service MSExchangePOP3BE
```

To restart the IMAP4 services, run the following command:

```
Restart-Service MSExchangeIMAP4; Restart-Service MSExchangeIMAP4BE
```

For more information about this cmdlet, see Restart-Service.

To verify that you've successfully restarted the POP3 or IMAP4 services, run the following command:

```
Get-Service MSExchangePOP3; Get-Service MSExchangePOP3BE; Get-Service MSExchangeIMAP4; Get-Service MSExchangeIMAP4BE
```

How do you know this worked?

To verify that you've enabled or disabled POP3 or IMAP4 access to a mailbox, use any of the following procedures:

- In the EAC, go to **Recipients** > **Mailboxes** > select the mailbox > click **Edit** ✎ > **Mailbox features** > **Email connectivity**.

- If POP3 access is enabled for the mailbox, you'll see **POP3: Enabled** and the **Disable** link. If POP3 access is disabled, you'll see **POP3: Disabled** and the **Enable** link.

- If IMAP4 access is enabled for the mailbox, you'll see **IMAP4: Enabled** and a **Disable** link. If IMAP4 access is disabled, you'll see **IMAP4: Disabled** and the **Enable** link.

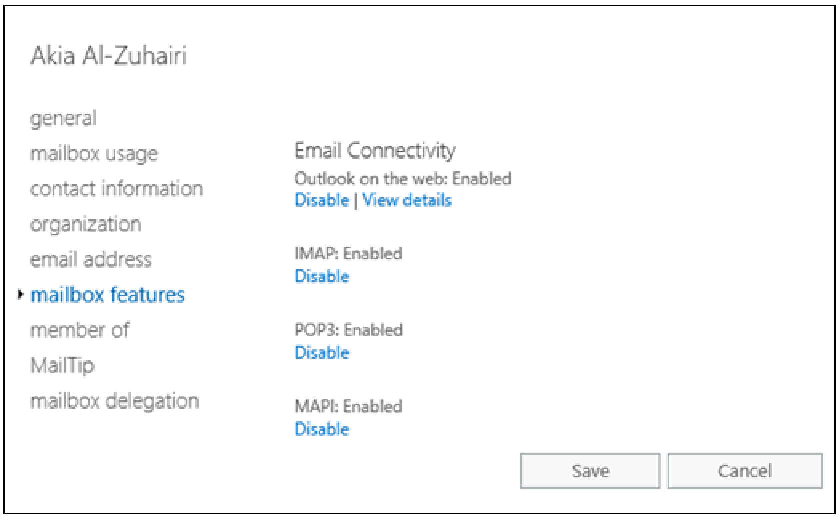
Akia Al-Zuhairi

Email Connectivity

Outlook on the web: Enabled
Disable | View details

IMAP: Enabled
Disable

POP3: Enabled
Disable

MAPI: Enabled
Disable

- In the Exchange Management Shell, replace *<MailboxIdentity>* with the identity of the mailbox (for example, name, alias, or email address), and run the following command:

```
Get-CasMailbox - Identity <MailboxIdentity>
```

- Use the same filter that you used to identify the mailboxes, but use the **Get-CasMailbox** cmdlet instead of **Set-CasMailbox**. For example:

```
Get-User -Filter "RecipientType -eq 'UserMailbox' -and Department -like 'Engineering*' -and
StateOrProvince -eq 'WA'" | Get-CasMailbox
```

- In the Exchange Management Shell, run this command to show all mailboxes where POP3 and IMAP4 access is disabled:

```
Get-CasMailbox -ResultSize unlimited -Filter "PopEnabled -eq `$false -and ImapEnabled -eq `$false"
```

Configure authenticated SMTP settings for POP3 and IMAP4 clients in Exchange Server

8/3/2020 • 6 minutes to read • Edit Online

After you enable and configure POP3 or IMAP4 on an Exchange server as described in Enable and configure POP3 on an Exchange server and Enable and configure IMAP4 on an Exchange server, you need to configure the authenticated SMTP settings for POP3 and IMAP4 clients so they can send email messages.

The default Receive connector named "Client Frontend *<Server name>*" in the Client Access services on the Mailbox server listens for authenticated SMTP client submissions on port 587. By default, this connector uses the following settings for **internal and external** client (authenticated) SMTP connections:

- **SMTP server**: `<ServerFQDN>`. For example, `mailbox01.contoso.com`.
- **TCP port**: 587
- **Encryption method**: TLS. Note that this is opportunistic TLS (**STARTTLS**) that results in an encrypted connection after the initial plain text protocol handshake.

For more information, see Default Receive connectors created during setup and Client access protocol architecture.

To configure the authenticated SMTP settings that are used by POP3 and IMAP4 clients, perform the following steps:

1. Configure the FQDN on the "Client Frontend *<Server name>*" Receive connector.
2. Specify the certificate that's used to encrypt authenticated SMTP client connections.
3. Configure Outlook on the web (formerly known as Outlook Web App) to display the SMTP settings for authenticated SMTP clients at **Settings** > **Options** > **Mail** > **Accounts** > **POP and IMAP**.

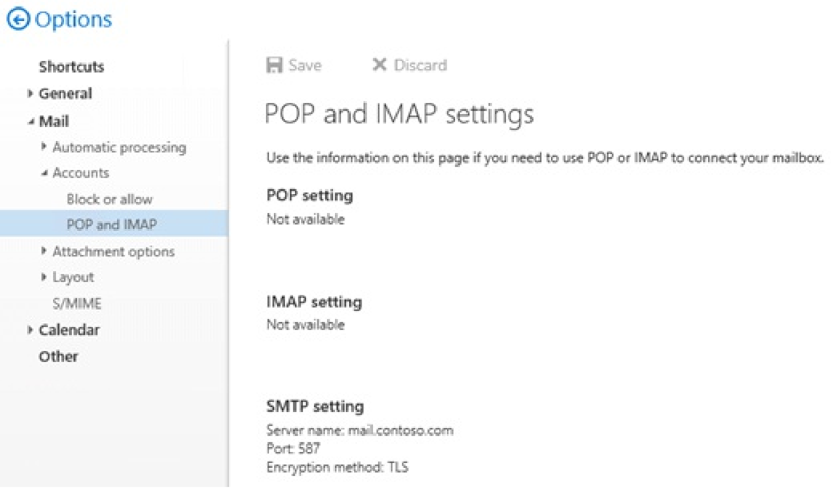

POP and IMAP settings

Use the information on this page if you need to use POP or IMAP to connect your mailbox.

POP setting
Not available

IMAP setting
Not available

SMTP setting
Server name: mail.contoso.com
Port: 587
Encryption method: TLS

For more information about POP3 and IMAP4, see POP3 and IMAP4 in Exchange Server.

What do you need to know before you begin?

- Estimated time to complete: 5 minutes.
- Secure Sockets Layer (SSL) is being replaced by Transport Layer Security (TLS) as the protocol that's used

to encrypt data sent between computer systems. They're so closely related that the terms "SSL" and "TLS" (without versions) are often used interchangeably. Because of this similarity, references to "SSL" in Exchange topics, the Exchange admin center, and the Exchange Management Shell have often been used to encompass both the SSL and TLS protocols. Typically, "SSL" refers to the actual SSL protocol only when a version is also provided (for example, SSL 3.0). To find out why you should disable the SSL protocol and switch to TLS, check out Protecting you against the SSL 3.0 vulnerability.

- If you have POP3 or IMAP4 clients that can only send SMTP email on port 25, you can configure port 25 on the "Client Frontend *<Server name>*" Receive connector to allow clients to send authenticated SMTP email. However, because port 25 is also configured on the "Client Frontend *<Server name>*" Receive connector for email from external SMTP servers, you'll need to modify the local IP addresses that are used to listen on port 25 on one or both of the connectors. For more information, see Receive connector local address bindings.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Receive connectors" entry in the Mail flow permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Step 1: Configure the FQDN on the "Client Frontend <Server name>" Receive connector

You can skip this step if you want to keep the default server FQDN value (for example, mailbox01.contoso.com). Or, you can specify an FQDN value that's more compatible with your Internet naming convention or a TLS certificate that you want to use.

If you change the FQDN value, and you want internal POP3 or IMAP4 clients to use this connector to send email, the new FQDN needs to have a corresponding record in your internal DNS.

Regardless of the FQDN value, if you want external POP3 or IMAP4 clients to use this connector to send email, the FQDN needs to have a corresponding record in your public DNS, and the TCP port (587) needs to be allowed through your firewall to the Exchange server.

Use the EAC to configure the FQDN for authenticated SMTP clients

1. In the EAC, go to **Mail flow** > **Receive connectors**.
2. In the list of Receive connectors, select **Client Frontend <Server name>**, and then click **Edit** (✏).
3. In the **Exchange Receive Connector** page that opens, click **Scoping**.
4. In the **FQDN** field, enter the SMTP server FQDN that you want to use for authenticated SMTP client connections (for example, mail.contoso.com) and then click **Save**.

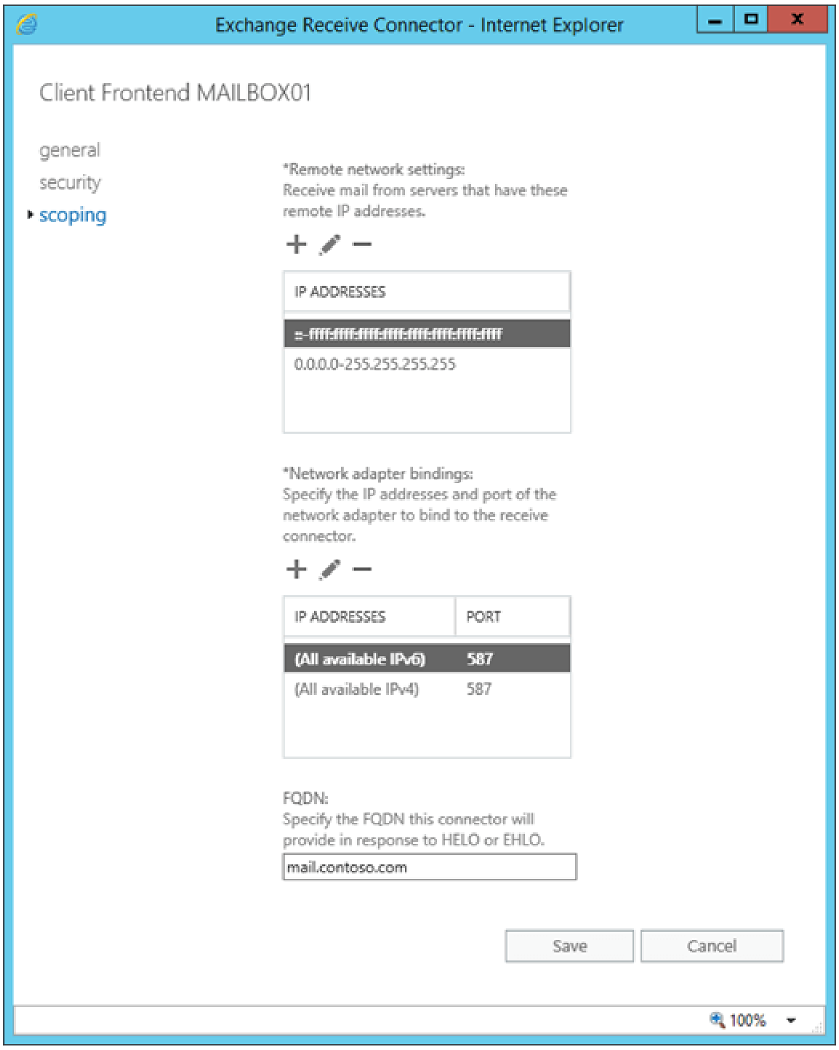

Client Frontend MAILBOX01

*Remote network settings:
Receive mail from servers that have these remote IP addresses.

IP ADDRESSES
::-ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff
0.0.0.0-255.255.255.255

*Network adapter bindings:
Specify the IP addresses and port of the network adapter to bind to the receive connector.

IP ADDRESSES	PORT
(All available IPv6)	587
(All available IPv4)	587

FQDN:
Specify the FQDN this connector will provide in response to HELO or EHLO.
mail.contoso.com

Use the Exchange Management Shell to configure the FQDN for authenticated SMTP clients

To configure the FQDN for authenticated SMTP clients, use the following syntax:

```
Get-ReceiveConnector -Identity "Client Frontend*" | Set-ReceiveConnector -Fqdn <FQDN>
```

This example configures the FQDN value mail.contoso.com.

```
Get-ReceiveConnector -Identity "Client Frontend*" | Set-ReceiveConnector -Fqdn mail.contoso.com
```

How do you know this step worked?

To verify that you've successfully the FQDN on the "Client Frontend *<Server name>* " Receive connector, use either of the following procedures:

- the EAC, go to **Mail flow** > **Receive connectors** > select **Client Frontend <Server name>**, click **Edit** (✎) > **Scoping**, and verify the value in the **FQDN** field.

- In the Exchange Management Shell, run the following command:

```
Get-ReceiveConnector -Identity "Client Frontend*" | Format-List Name,Fqdn
```

Step 2: Use the Exchange Management Shell to specify the certificate that's used to encrypt authenticated SMTP client connections

The certificate needs to match or contain the FQDN value that you specified in the previous step, and the POP3 and SMTP clients need to trust the certificate, which likely means a certificate from a commercial certification authority. For more information, see Certificate requirements for Exchange services.

Also, you need to assign the certificate to the Exchange SMTP service. For more information, see Assign certificates to Exchange Server services.

To specify the certificate that's used for authenticated SMTP client connections, use the following syntax:

```
$TLSCert = Get-ExchangeCertificate -Thumbprint <ThumbprintValue>
```

```
$TLSCertName = "<I>$($TLSCert.Issuer)<S>$($TLSCert.Subject)"
```

```
Get-ReceiveConnector -Identity "Client Frontend*" | Set-ReceiveConnector -TlsCertificateName $TLSCertName
```

This example uses the certificate that has the thumbprint value 434AC224C8459924B26521298CE8834C514856AB.

```
$TLSCert = Get-ExchangeCertificate -Thumbprint 434AC224C8459924B26521298CE8834C514856AB
```

```
$TLSCertName = "<I>$($TLSCert.Issuer)<S>$($TLSCert.Subject)"
```

```
Get-ReceiveConnector -Identity "Client Frontend*" | Set-ReceiveConnector -TlsCertificateName $TLSCertName
```

How do you know this step worked?

To verify that you've specified the certificate that's used to encrypt authenticated SMTP client connections, perform the following steps:

1. Run the following command in the Exchange Management Shell:

```
Get-ReceiveConnector -Identity "Client Frontend*" | Format-List Name,Fqdn,TlsCertificateName
```

2. Run the following command in the Exchange Management Shell:

```
Get-ExchangeCertificate | Format-List Thumbprint,Issuer,Subject,CertificateDomains,Services
```

3. Verify the **Subject** or **CertificateDomains** field of the certificate that you specified on the Receive connector contains the **Fqdn** value of the Receive connector (exact match or wildcard match).

Step 3: Use the Exchange Management Shell to configure Outlook on the web to display the SMTP settings for authenticated SMTP clients

To configure Outlook on the web to display the SMTP settings server for authenticated SMTP clients, run the following command:

```
Get-ReceiveConnector -Identity "Client Frontend*" | Set-ReceiveConnector -AdvertiseClientSettings $true
```

Note: To prevent the SMTP settings from being displayed in Outlook on the web, change the value from `$true` to

`$false`.

How do you know this step worked?

To verify that you've configured Outlook on the web to display the SMTP settings for authenticated SMTP clients, perform the following steps:

1. Open a mailbox in Outlook on the web, and then click **Settings** > **Options**.

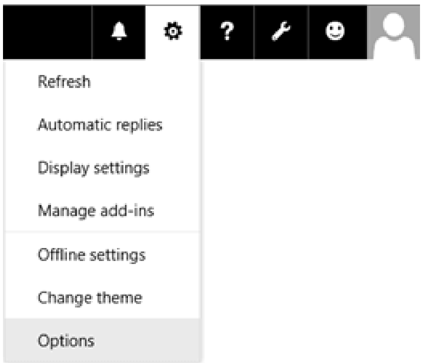

2. Click **Mail** > **Accounts** > **POP and IMAP** and verify the correct SMTP settings are displayed.

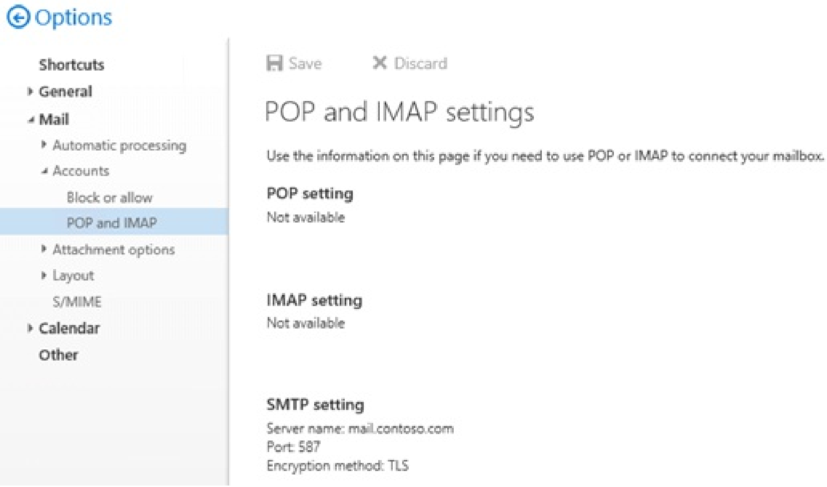

POP and IMAP settings

Use the information on this page if you need to use POP or IMAP to connect your mailbox.

POP setting
Not available

IMAP setting
Not available

SMTP setting
Server name: mail.contoso.com
Port: 587
Encryption method: TLS

Note: If the SMTP settings that you configured don't appear as expected in Outlook on the web, run the commands `net stop w3svc /y` and `net start w3svc` to restart Internet Information Services (IIS).

How do you know this task worked?

To verify that you've configured the authenticated SMTP settings on the Exchange server, perform one or more following procedures:

- Use the **Test-PopConnectivity** or **Test-ImapConnectivity** cmdlets, which use authenticated SMTP to send test messages. For more information, see Test-PopConnectivity and Test-ImapConnectivity.

- Enable protocol logging on the "Client Frontend *<Server name>*" Receive connector, configure a POP3 or IMAP4 client to connect to a mailbox, send a test message from an internal network connection and/or an external Internet connection, and view the results in the protocol log. For more information, see Protocol logging.

  Note: You can't use POP3 or IMAP4 to connect to the Administrator mailbox. This limitation was intentionally included in Exchange 2016 and Exchange 2019 to enhance the security of the Administrator mailbox.

Outlook for iOS and Android

8/3/2020 • 2 minutes to read • Edit Online

Outlook for iOS and Android supports two authentication types in Exchange on-premises environments: *Basic authentication* and *hybrid Modern Authentication*.

Outlook for iOS and Android uses Basic authentication with Exchange ActiveSync in the following environments:

- In Exchange Server 2010 environments
- When a hybrid relationship with Microsoft 365 or Office 365 has not been configured
- When hybrid Modern Authentication has not been enabled

For more information see Using Basic authentication with Outlook for iOS and Android.

For customers running Exchange Server 2013, Exchange Server 2016, or Exchange Server 2019 in a hybrid relationship with Microsoft 365 or Office 365, Outlook for iOS and Android can be configured to leverage hybrid Modern Authentication. For more information, see Using hybrid Modern Authentication with Outlook for iOS and Android.

> **NOTE**
>
> The Outlook for iOS and Android Help Center is available for users, including help for using the app on specific devices and troubleshooting information.

Using hybrid Modern Authentication with Outlook for iOS and Android

8/3/2020 • 25 minutes to read • Edit Online

The Outlook app for iOS and Android is designed as the best way to experience Microsoft 365 or Office 365 on your mobile device by leveraging Microsoft services to help find, plan, and prioritize your daily life and work. Outlook provides the security, privacy, and support you need while protecting corporate data via capabilities such as Azure Active Directory conditional access and Intune app protection policies. The following sections provide an overview of the hybrid Modern Authentication architecture, the required pre-requisites for its deployment, and how to securely deploy Outlook for iOS and Android for Exchange on-premises mailboxes.

Microsoft Cloud architecture for hybrid Exchange Server customers

Outlook for iOS and Android is a cloud-backed application. This means your experience consists of a locally installed app powered by a secure and scalable service running in the Microsoft Cloud.

For Exchange Server mailboxes, Outlook for iOS and Android's architecture is built directly into the Microsoft Cloud, providing customers the additional benefits of security, privacy, built-in compliance, and transparent operations that Microsoft commits to in the Microsoft Trust Center and Azure Trust Center.

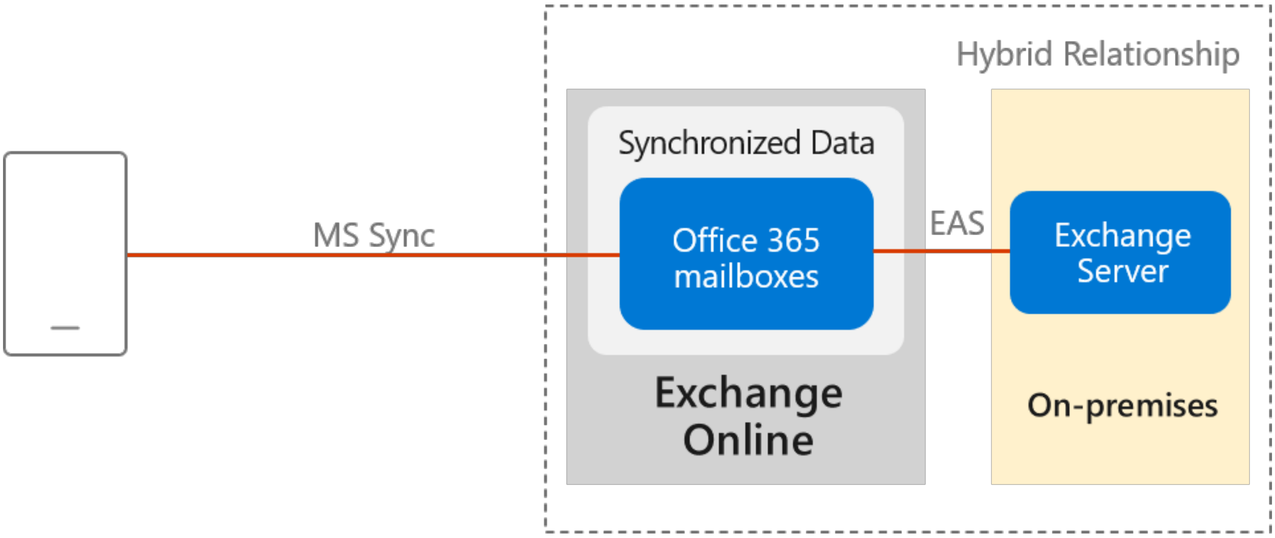

Within the Microsoft 365 or Office 365-based architecture, Outlook for iOS and Android utilizes the native Microsoft sync technology for data synchronization which is protected by a TLS-secured connection end-to-end, between Microsoft 365 or Office 365 and the app.

The Exchange ActiveSync (EAS) connection between Exchange Online and the on-premises environment enables synchronization of the users' on-premises data and includes four weeks of email, all calendar data, all contact data, and out-of-office status in your Exchange Online tenant. This data will be removed automatically from Exchange Online after 30 days when the account is deleted in Azure Active Directory.

Data synchronization between the on-premises environment and Exchange Online happens independent of user behavior. This ensures that we can send new messages to the devices very quickly.

Processing information in the Microsoft Cloud enables advanced features and capabilities, such as the categorization of email for the Focused Inbox, customized experience for travel and calendar, and improved search speed. Relying on the cloud for intensive processing and minimizing the resources required from users' devices enhances the app's performance and stability. Lastly, it allows Outlook to build features that work across all email accounts, regardless of the technological capabilities of the underlying servers (such as different versions of Exchange Server, Microsoft 365, or Office 365).

Specifically, this new architecture has the following improvements:

1. **Enterprise Mobility + Security support**: Customers can take advantage of Microsoft Enterprise Mobility + Security (EMS) including Microsoft Intune and Azure Active Directory Premium, to enable conditional access and Intune app protection policies, which control and secure corporate messaging data on the mobile device.

2. **Fully powered by Microsoft Cloud**: The on-premises mailbox data is synchronized into Exchange Online, which provides the benefits of security, privacy, compliance and transparent operations that Microsoft commits to in the Microsoft Trust Center.

3. **OAuth protects users' passwords**: Outlook leverages hybrid Modern Authentication (OAuth) to protect users' credentials. Hybrid Modern Authentication provides Outlook with a secure mechanism to access the Exchange data without ever touching or storing a user's credentials. At sign in, the user authenticates directly against an identity platform (either Azure Active Directory or an on-premises identity provider like ADFS) and receives an access token in return, which grants Outlook access to the user's mailbox or files. At no time does the service have access to the user's password.

4. **Provides Unique Device IDs**: Each Outlook connection is uniquely registered in Microsoft Intune and can therefore be managed as a unique connection.

5. **Unlocks new features on iOS and Android**: This update enables the Outlook app to take advantage of native Microsoft 365 or Office 365 features that are not supported in Exchange on-premises today, such as leveraging full Exchange Online search and Focused Inbox. These features will only be available when using Outlook for iOS and Android.

> **NOTE**
>
> Device management through the on-premises Exchange admin center (EAC) is not possible. Intune is required to manage mobile devices.

Data security, access, and auditing controls

With on-premises data being synchronized with Exchange Online, customers have questions about how the data is protected in Exchange Online. Encryption in the Microsoft Cloud discusses how BitLocker is used for volume-level encryption. Service Encryption with Customer Key is supported in the Outlook for iOS and Android architecture, but note that the user must have an Office 365 Enterprise E5 license (or the corresponding versions of those plans for Government or Education) to have an encryption policy assigned using the set-mailuser cmdlet.

By default, Microsoft engineers have zero standing administrative privileges and zero standing access to customer content in Microsoft 365 or Office 365. Administrative Access Controls discusses personnel screening, background checks, Lockbox and Customer Lockbox, and more.

ISO Audited Controls on Service Assurance documentation provides the status of audited controls from global information security standards and regulations that Microsoft 365 and Office 365 have implemented.

Connection flow

When Outlook for iOS and Android is enabled with hybrid Modern Authentication, the connection flow is as follows.

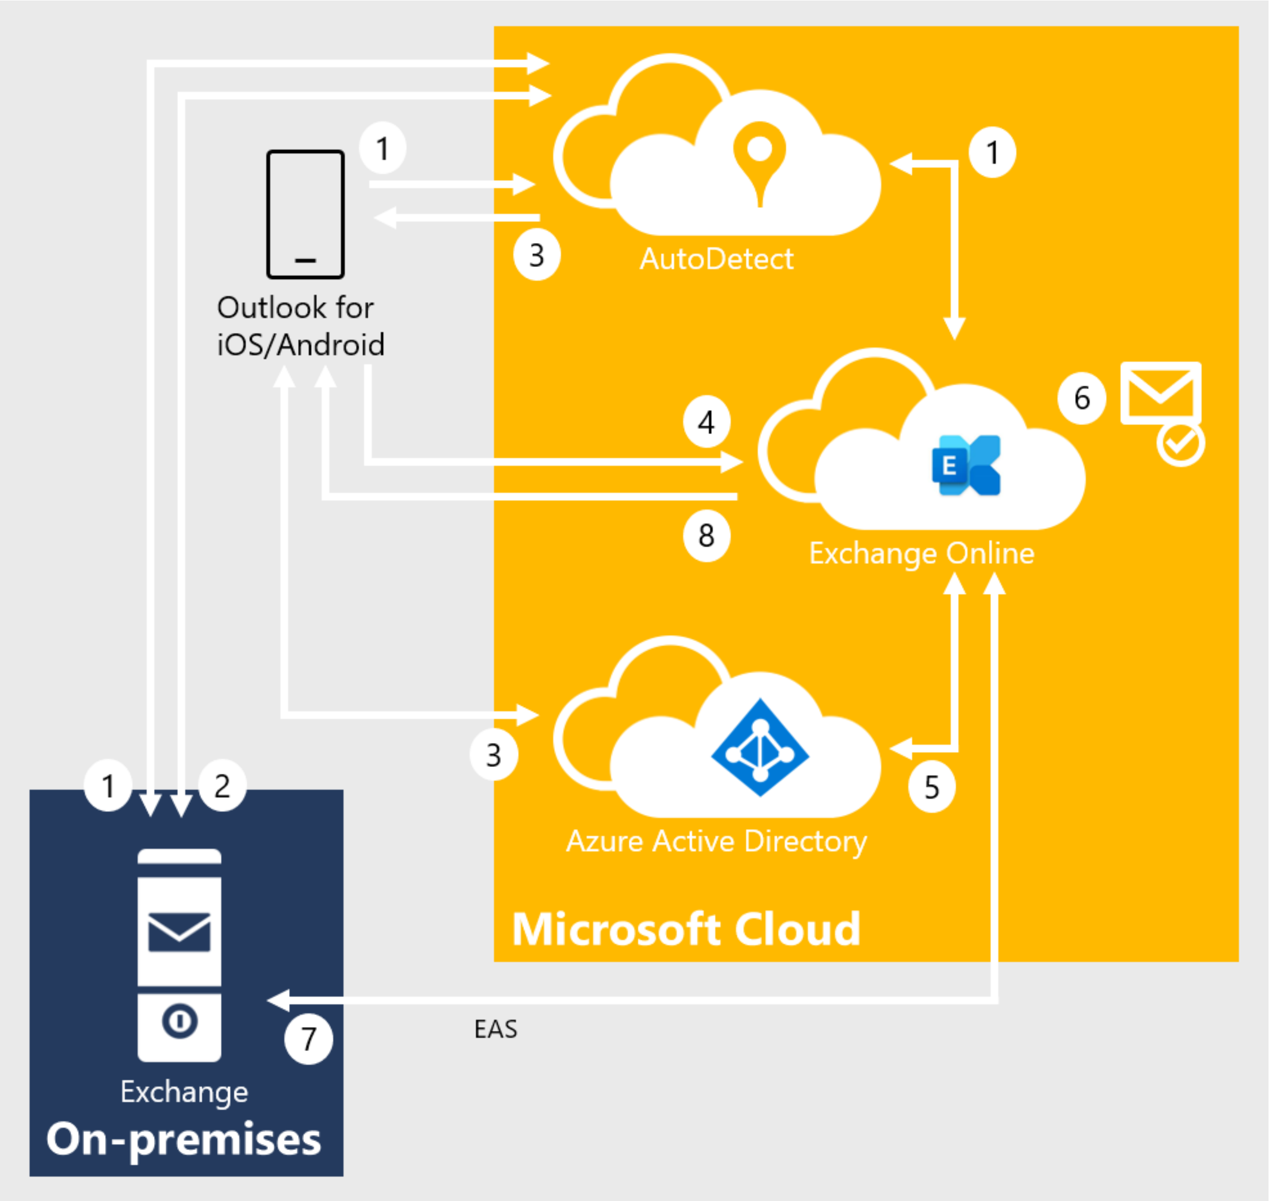

1. After the user enters their email address, Outlook for iOS and Android connects to the AutoDetect service. AutoDetect determines the mailbox type by initiating an AutoDiscover query to Exchange Online. Exchange Online determines that the user's mailbox is on-premises and returns a 302-redirect to AutoDetect with the on-premises Autodiscover URL. AutoDetect initiates a query against the on-premises AutoDiscover service to determine the ActiveSync endpoint for the email address. The URL attempted on-premises is similar to this example: https://autodiscover.contoso.com/autodiscover/autodiscover.json?Email=test%40contoso.com&Protocol=activesync&RedirectCount=3.

2. AutoDetect initiates a connection to the on-premises ActiveSync URL returned in Step 1 above with an empty bearer challenge. The empty bearer challenge tells the on-premises ActiveSync that the client supports Modern Authentication. On-premises ActiveSync responds with a 401-challenge response and includes the *WWW-Authenticate: Bearer* header. Within the WWW-Authenticate: Bearer header is the authorization_uri value that identifies the Azure Active Directory (AAD) endpoint that should be used to obtain an OAuth token.

3. AutoDetect returns the AAD endpoint to the client. The client begins the log-in flow and the user is presented with a Web form (or redirected to the Microsoft Authenticator app) and can enter credentials. Depending on the identity configuration, this may or may not involve a federated endpoint redirect to an on-premises identity provider. Ultimately, the client obtains an access-and-refresh token pair, which is named AT1/RT1. This access token is scoped to the Outlook for iOS and Android client with an audience of the Exchange Online endpoint.

4. Outlook for iOS and Android establishes a connection to Exchange Online and issues a provisioning request which includes the user's access token (AT1) and the on-premises ActiveSync endpoint.

5. The MRS provisioning API within Exchange Online utilizes AT1 as input and obtains a second access-and-refresh token pair (named AT2/RT2) to access the on-premises mailbox via an on-behalf-of call to Active Directory. This second access token is scoped with the client being Exchange Online and an audience of the on-premises ActiveSync namespace endpoint.

6. If the mailbox is not provisioned, then the provisioning API creates a mailbox.

7. The MRS provisioning API establishes a secure connection to the on-premises ActiveSync endpoint and synchronizes the user's messaging data using the AT2 access token as the authentication mechanism. RT2 is used periodically to generate a new AT2 so that data can be synchronized in the background without user intervention.

8. Data is returned to the client.

Technical and licensing requirements

The hybrid Modern Authentication architecture has the following technical requirements:

> **NOTE**
>
> On-premises accounts leveraging hybrid Modern Authentication with Outlook mobile are not supported with Office 365 US Government Community and Defense tenants, Office 365 Germany tenants, and Office 365 China operated by 21Vianet tenants.

1. **Exchange on-premises setup**:

   - Exchange Server 2019 Cumulative Update 1 (CU1) or later, Exchange Server 2016 Cumulative Update 8 (CU8) or later, or Exchange Server 2013 CU19 or later on all Exchange servers. In hybrid deployments (on-premises Exchange and Exchange Online) or in organizations that use Exchange Online Archiving (EOA) with their on-premises Exchange deployment, you need to deploy the most current CU or one CU prior to the most current.

   - All Exchange 2007 or Exchange 2010 servers must be removed from the environment. These versions of Exchange are out of mainstream support and will not work with Intune-managed Outlook for iOS and Android. In this architecture, Outlook for iOS and Android uses OAuth as the authentication mechanism. One of the on-premises configuration changes that occurs enables the OAuth endpoint to the Microsoft Cloud as the default authorization endpoint. When this change is made, clients can start negotiating the use of OAuth. Because this is an organization-wide change, Exchange 2010 mailboxes fronted by either Exchange 2013 or 2016 will incorrectly think they can perform OAuth and will end up in a disconnected state, since Exchange 2010 does not support OAuth as an authentication mechanism.

2. **Active Directory Synchronization**. Active Directory synchronization of the entire on-premises mail recipient directory with Azure Active Directory, via Azure AD Connect. If you have **Azure AD app and attribute filtering** enabled in Azure AD Connect configuration, ensure that the following applications are selected:

   - Office 365 ProPlus

   - Exchange Online

   - Azure RMS

   - Intune

If you do not have **Azure AD app and attribute filtering** enabled in Azure AD Connect configuration, all required applications are already selected by default.

> **IMPORTANT**
>
> Outlook for iOS and Android uses the tenant's Exchange Online Global Address List for on-premises mailboxes that leverage hybrid Modern Authentication. If all mail recipients are not synchronized into Azure Active Directory, users will experience mail flow issues.

3. **Exchange hybrid setup**: Requires full hybrid relationship between Exchange on-premises with Exchange Online.

    - A hybrid Microsoft 365 or Office 365 organization is configured in full hybrid configuration using Exchange Classic Hybrid Topology mode and is set up as specified in the Exchange Deployment Assistant.

        > **NOTE**
        >
        > Hybrid Modern Authentication is not supported with the Hybrid Agent.

    - Requires a Microsoft 365 or Office 365 Enterprise, Business, or Education organization.
    - The on-premises mailbox data is synchronized in the same datacenter region where that Microsoft 365 or Office 365 organization is set up. For more information about where Microsoft 365 and Office 365 data is located, visit the Microsoft Trust Center.
    - The external URL host names for Exchange ActiveSync and AutoDiscover must be published as service principals to Azure Active Directory through the Hybrid Configuration Wizard.
    - AutoDiscover and Exchange ActiveSync namespaces must be accessible from the Internet and cannot be fronted by a pre-authentication solution.
    - Ensure SSL or TLS offloading is not being used between the load balancer and your Exchange servers, as this will affect the use of the OAuth token. SSL and TLS bridging (termination and re-encryption) is supported.

4. **Intune setup**: Both Intune standalone and Co-Management deployments are supported (Basic Mobility and Security for Microsoft 365 is not supported).

5. **Microsoft 365 and Office 365 licensing**:

    - Outlook for iOS and Android is free for consumer usage from the iOS App store and from Google Play. However, commercial users require a Microsoft 365 or Office 365 subscription that includes the Office desktop applications: Microsoft 365 Apps for Business, Microsoft 365 Business Standard, Microsoft 365 Apps for enterprise, Office 365 Enterprise E3, Office 365 Enterprise E5, or the corresponding versions of those plans for Government or Education. Commercial users with the following subscriptions are allowed to use the Outlook mobile app on devices with integrated screens 10.1" diagonally or less: Office 365 Enterprise E1, Office 365 F1, Office 365 A1, Microsoft 365 Business Basic, and if you only have an Exchange Online license (without Office). If you only have an Exchange on-premises (Exchange Server) license, you are not licensed to use the app.
    - Use of advanced Exchange Online features (e.g., Service Encryption with Customer Key or Multi-Geo Capabilities) require the on-premises user to be assigned the applicable Office 365 or Microsoft 365 subscription license within the Microsoft 365 Admin Center.

    For more information on how to assign a license, see [Add users individually or in bulk] (https://docs.microsoft.com/microsoft-365/admin/add-users/add-users.

6. **EMS licensing**: Each on-premises user must have one of the following licenses:

   - Intune standalone + Azure Active Directory Premium 1 or Azure Active Directory Premium 2
   - Enterprise Mobility + Security E3, Enterprise Mobility + Security E5

Implementation steps

Enabling support for hybrid Modern Authentication in your organization requires each of the following steps, which are detailed in the following sections:

1. Create a conditional access policy
2. Create an Intune app protection policy
3. Enable hybrid Modern Authentication

Create a conditional access policy

When an organization decides to standardize how users access Exchange data, using Outlook for iOS and Android as the only email app for end users, they can configure a conditional access policy that blocks other mobile access methods. Outlook for iOS and Android authenticates via the Azure Active Directory identity object and then connects to Exchange Online. Therefore, you will need to create Azure Active Directory conditional access policies to restrict mobile device connectivity to Exchange Online. To do this, you will need two conditional access policies, with each policy targeting all potential users. Details on creating these policies can be found in Require app protection policy for cloud app access with Conditional Access.

1. Follow "Step 1: Configure an Azure AD Conditional Access policy for Microsoft 365 or Office 365" in Scenario 1: Microsoft 365 and Office 365 apps require approved apps with app protection policies, which allows Outlook for iOS and Android, but blocks OAuth capable Exchange ActiveSync clients from connecting to Exchange Online.

   > **NOTE**
   >
   > This policy ensures mobile users can access all Office endpoints using the applicable apps.

2. Follow "Step 2: Configure an Azure AD Conditional Access policy for Exchange Online with ActiveSync (EAS)" in Scenario 1: Microsoft 365 and Office 365 apps require approved apps with app protection policies, which prevents Exchange ActiveSync clients leveraging basic authentication from connecting to Exchange Online.

   The above policies leverage the grant control Require app protection policy, which ensures that an Intune App Protection Policy is applied to the associated account within Outlook for iOS and Android prior to granting access. If the user isn't assigned to an Intune App Protection Policy, isn't licensed for Intune, or the app isn't included in the Intune App Protection Policy, then the policy prevents the user from obtaining an access token and gaining access to messaging data.

3. Finally, follow How to: Block legacy authentication to Azure AD with Conditional Access to block legacy authentication for other Exchange protocols on iOS and Android devices; this policy should target only Microsoft 365 or Office 365 Exchange Online cloud app and iOS and Android device platforms. This ensures mobile apps using Exchange Web Services, IMAP4, or POP3 protocols with basic authentication cannot connect to Exchange Online.

IMPORTANT

To leverage app-based conditional access policies, the Microsoft Authenticator app must be installed on iOS devices. For Android devices, the Intune Company Portal app is required. For more information, see App-based conditional access with Intune.

In order to block other mobile device clients (such as the native mail client included in the mobile operating system) from connecting to your on-premises environment (which authenticate via basic authentication against on-premises Active Directory), you have two options:

1. You can leverage the built-in Exchange mobile device access rules and block all mobile devices from connecting by setting the following in the Exchange Management Shell:

   ```
   Set-ActiveSyncOrganizationSettings -DefaultAccessLevel Block
   ```

2. You can leverage an on-premises conditional access policy within Intune after installing the on-premises Exchange connector. For more information, see Create a conditional access policy for Exchange on-premises and legacy Exchange Online Dedicated.

   NOTE

   When implementing either of the above on-premises options, be aware that it may impact users connecting to Exchange on their mobile devices.

Create an Intune app protection policy

After hybrid Modern Authentication is enabled, all on-premises mobile users are able to leverage Outlook for iOS and Android using the Microsoft 365 or Office 365-based architecture. Therefore, it's important to protect corporate data with an Intune app protection policy.

Create Intune app protection policies for both iOS and Android using the steps documented in How to create and assign app protection policies. At a minimum, each policy must include the following:

1. They include all Microsoft mobile applications, such as Word, Excel, or PowerPoint, as this will ensure that users can access and manipulate corporate data within any Microsoft app in a secure fashion.

2. They mimic the security features that Exchange provides for mobile devices, including:

   - Requiring a PIN for access (which includes Select Type, PIN length, Allow Simple PIN, Allow fingerprint)

   - Encrypting app data

   - Blocking managed apps from running on "jailbroken" and rooted devices

3. They are assigned to all users. This ensures that all users are protected, regardless of whether they use Outlook for iOS and Android.

In addition to the above minimum policy requirements, you should consider deploying advanced protection policy settings like **Restrict cut, copy and paste with other apps** to further prevent corporate data leakage. For more information on the available settings, see Android app protection policy settings in Microsoft Intune and iOS app protection policy settings.

> **IMPORTANT**
>
> To apply Intune app protection policies against apps on Android devices that are not enrolled in Intune, the user must also install the Intune Company Portal. For more information, see What to expect when your Android app is managed by app protection policies.

Enable hybrid Modern Authentication

1. If you haven't enabled hybrid Modern Authentication, review the prerequisites as outlined in Hybrid Modern Authentication overview and prerequisites for using it with on-premises Skype for Business and Exchange servers. After you've completed the prerequisites, do the steps in How to configure Exchange Server on-premises to use hybrid Modern Authentication.

2. Create an Exchange on-premises device access allow rule to allow Exchange Online to connect to your on-premises environment using the ActiveSync protocol:

   ```
   If ((Get-ActiveSyncOrganizationSettings).DefaultAccessLevel -ne "Allow") {New-
   ActiveSyncDeviceAccessRule -Characteristic DeviceType -QueryString "OutlookService" -AccessLevel Allow}
   ```

   > **NOTE**
   >
   > Device management through the on-premises Exchange admin center is not possible. Intune is required to manage mobile devices.

3. Create an Exchange on-premises device access rule that prevents users from connecting to the on-premises environment with Outlook for iOS and Android with basic authentication over the Exchange ActiveSync protocol:

   ```
   New-ActiveSyncDeviceAccessRule -Characteristic DeviceModel -QueryString "Outlook for iOS and Android" -
   AccessLevel Block
   ```

   > **NOTE**
   >
   > Once this rule is created, Outlook for iOS and Android with Basic authentication users will be blocked.

4. Ensure your on-premises Exchange ActiveSync maxRequestLength is configured to match your transport configuration's MaxSendSize/MaxReceiveSize:

   - Path: `%ExchangeInstallPath%\FrontEnd\HttpProxy\Sync\web.config`
   - Property: `maxRequestLength`
   - Value: set in KB size (10MB is 10240, for example)

Client features that aren't supported

The following features are not supported for on-premises mailboxes using hybrid Modern Authentication with Outlook for iOS and Android.

- Draft folder and Draft messages synchronization
- Shared calendar access and delegate calendar access
- Shared and delegate mailbox data access

- Cortana Time to Leave / Travel Time
- Calendar attachments
- Rich meeting locations
- Task management with Microsoft To-Do
- Add-ins
- Interesting Calendars
- Play My Emails
- Sensitivity labeling
- S/MIME

Connection Flow FAQ

Q: My organization has a security policy that requires Internet inbound connections to be restricted to approved IP addresses or FQDNs. Is that possible with this architecture?

A: Microsoft recommends that the on-premises endpoints for AutoDiscover and ActiveSync protocols be opened and accessible from the Internet without any restrictions. In certain situations that may not be possible. For example, if you're in a co-existence period with another MDM solution, you may want to place restrictions on the ActiveSync protocol to prevent users from bypassing the third-party MDM solution while you migrate to Intune and Outlook for iOS and Android. If you must place restrictions on your on-premises firewall or gateway edge devices, Microsoft recommends filtering based on FQDN endpoints. If FQDN endpoints cannot be used, then filter on IP addresses. Make sure the following IP subnets and FQDNs are whitelisted:

- All Exchange Online FQDNs and IP subnet ranges as defined in Additional endpoints not included in the Microsoft 365 or Office 365 IP Address and URL Web service.
- The AutoDetect FQDNs and IP subnet ranges defined in Additional endpoints not included in the Microsoft 365 or Office 365 IP Address and URL Web service. This is required because the AutoDetect service establishes connections to the on-premises infrastructure.
- All Outlook iOS and Android and Office mobile app FQDNs as defined in Microsoft 365 and Office 365 URLs and IP address ranges.

Q: My organization currently uses a third-party MDM solution to control mobile device connectivity. If I expose the Exchange ActiveSync namespace on the Internet, that introduces a way for users to bypass the third-party MDM solution during the co-existence period. How can I prevent this?

A: There are three potential solutions to resolving this issue:

1. Implement Exchange mobile device access rules to control which devices are approved to connect.
2. Some third-party MDM solutions integrate with Exchange mobile device access rules, blocking unapproved access, while adding approved devices in the user's ActiveSyncAllowedDeviceIDs property.
3. Implement IP restrictions on the Exchange ActiveSync namespace.

Q: Can I leverage Azure ExpressRoute for managing traffic between the Microsoft Cloud and my on-premises environment?

A: Connectivity to the Microsoft Cloud requires Internet connectivity. Microsoft recommends exposing AutoDiscover and Exchange ActiveSync directly to the Internet; for more information, see Microsoft 365 and Office 365 Network Connectivity Principles. However, Azure ExpressRoute is supported for Exchange hybrid scenarios.

For more information, see Azure ExpressRoute for Microsoft 365 and Office 365.

With ExpressRoute, there is no private IP space for ExpressRoute connections, nor can there be "private" DNS resolution. That means that any endpoint your company wants to use over ExpressRoute must resolve in public DNS. If that endpoint resolves to an IP that is contained in the advertised prefixes associated with the ExpressRoute circuit (your company must configure those prefixes in the Azure portal when you enable Microsoft peering on the ExpressRoute connection), then the outbound connection from Exchange Online to your on-premises environment will route through the ExpressRoute circuit. Your company will have to ensure that the return traffic associated with these connections goes through the ExpressRoute circuit (avoiding asymmetric routing).

> **IMPORTANT**
>
> Because your company will be adding the Exchange AutoDiscover and ActiveSync namespaces to the advertised prefixes in the ExpressRoute circuit, the only way to reach the Exchange AutoDiscover and ActiveSync endpoints will be via the ExpressRoute. In other words, the only mobile device that will be able to connect to on-premises via the AutoDiscover and ActiveSync namespaces will be Outlook for iOS and Android. All other clients (such as mobile devices' native mail clients) will be unable to connect to the on-premises environment as the connection will not be established from the Microsoft Cloud. This is because there cannot be any overlaps of the public IP space advertised to Microsoft on the ExpressRoute circuit and the public IP space advertised on your Internet circuit(s).

Q: Given only four weeks of message data is synchronized to Exchange Online, does this mean that search queries executed in Outlook for iOS and Android cannot return information beyond the data available on the local device?

A: When a search query is performed in Outlook for iOS and Android, items that match the search query are returned if they are located on the device. In addition, the search query is passed to Exchange on-premises via Exchange Online. Exchange on-premises executes the search query against the on-premises mailbox and returns the results to Exchange Online, which relays the results to the client. The on-premises query results are stored in Exchange Online for one day before being deleted.

Q: How do I know that the email account is added correctly in Outlook for iOS and Android?

A: On-premises mailboxes that are added via hybrid Modern Authentication are labelled as **Exchange (Hybrid)** in the account settings in Outlook for iOS and Android, similar to the following example:

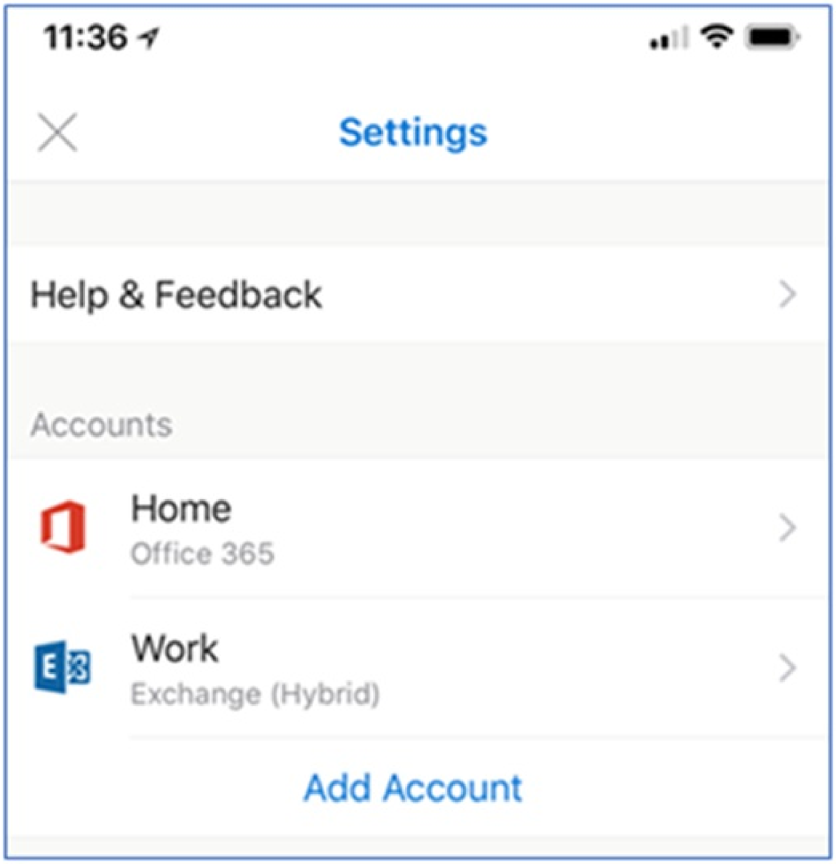

Authentication FAQ

Q: What identity configurations are supported with hybrid Modern Authentication and Outlook for iOS and Android?

A: The following identity configurations with Azure Active Directory are supported with hybrid Modern Authentication:

- Federated Identity with any on-premises identity provider that is supported by Azure Active Directory
- Password Hash Synchronization via Azure Active Directory Connect
- Pass-through Authentication via Azure Active Directory Connect

Q: What authentication mechanism is used for Outlook for iOS and Android? Are credentials stored in Microsoft 365 or Office 365?

A: See Account setup with modern authentication in Exchange Online.

Q: Do Outlook for iOS and Android and other Microsoft Office mobile apps support single sign-on?

A: See Account setup with modern authentication in Exchange Online.

Q: What is the lifetime of the tokens generated and used by the Active Directory Authentication Library (ADAL) in Outlook for iOS and Android?

A: See Account setup with modern authentication in Exchange Online.

Q: What happens to the access token when a user's password is changed?

A: See Account setup with modern authentication in Exchange Online.

Q: Is there a way for a user to bypass AutoDetect when adding their account to Outlook for iOS and Android?

A: Yes, a user can bypass AutoDetect at any time and manually configure the connection using Basic authentication over the Exchange ActiveSync protocol. To ensure that the user does not establish a connection to your on-premises environment via a mechanism that does not support Azure Active Directory Conditional Access or Intune app protection policies, the on-premises Exchange Administrator needs to configure an Exchange device access rule that blocks the ActiveSync connection. To do this, type the following command in the Exchange Management Shell:

```
New-ActiveSyncDeviceAccessRule -Characteristic DeviceModel -QueryString "Outlook for iOS and Android" -
AccessLevel Block
```

Troubleshooting

Below are the most common issues or errors with on-premises mailboxes using hybrid Modern Authentication with Outlook for iOS and Android.

AutoDiscover and ActiveSync

During profile creation, the user should be presented a Modern Authentication dialog similar to this:

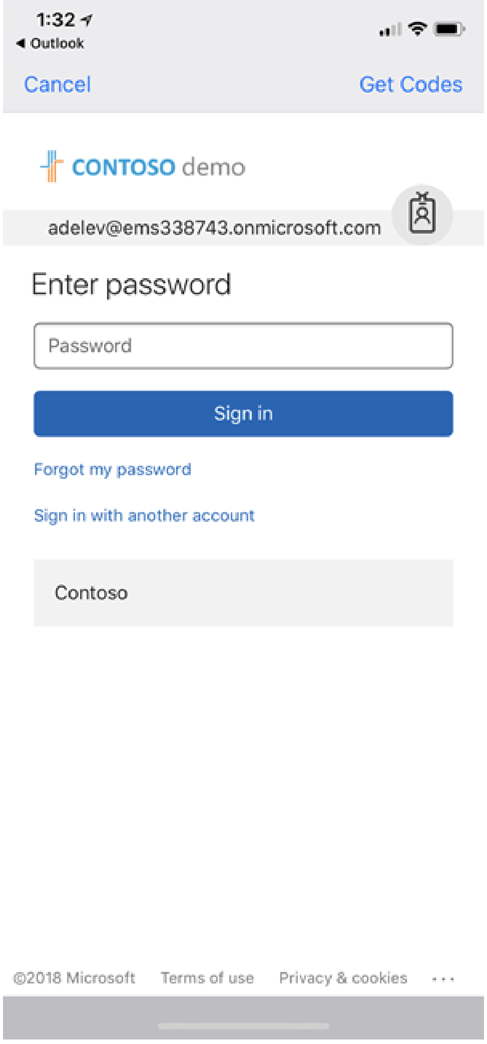

If, instead, the user is presented with one of the following dialogs, then there is an issue with either the Autodiscover or ActiveSync on-premises endpoints.

Here is an example of a user being presented with the legacy Basic authentication Exchange ActiveSync experience:

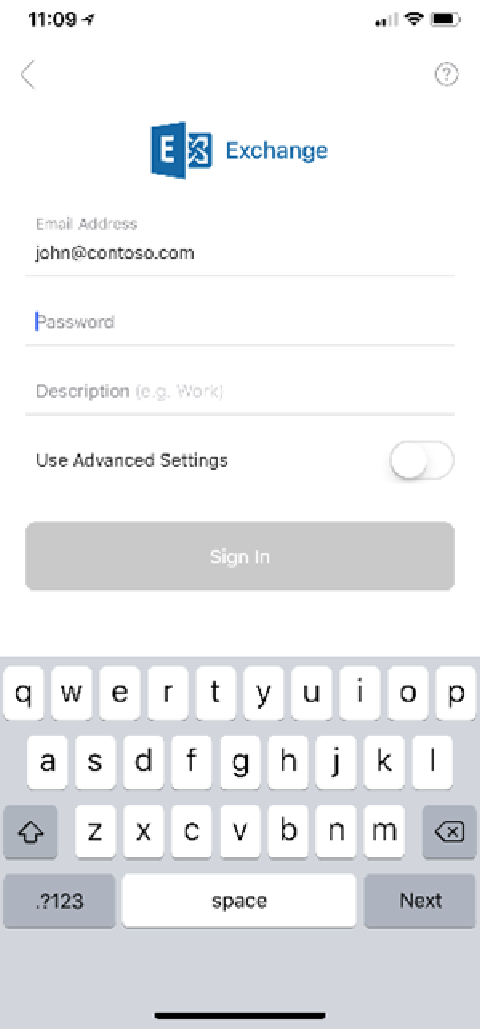

Email Address
john@contoso.com

And here's an example of what users see when AutoDetect isn't able to discover the configuration for users' on-premises mailboxes.

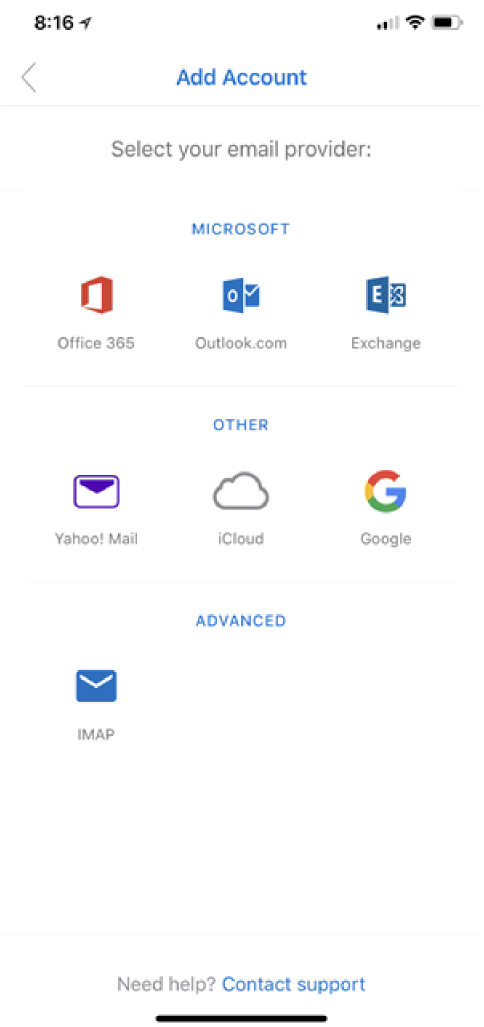

In either scenario, verify that your on-premises environment is correctly configured. To do this: from the TechNet Gallery, download and execute the script for Validating Hybrid Modern Authentication setup for Outlook for iOS and Android.

When you review the output from the script, you should be seeing the following from AutoDiscover:

```
{
    "Protocol": "activesync",
    "Url": "https://mail.contoso.com/Microsoft-Server-ActiveSync"
}
```

The on-premises ActiveSync endpoint should return the following response, where the WWW-Authenticate header includes an authorization_uri:

```
Content-Length →0
Date →Mon, 29 Jan 2018 19:51:46 GMT
Server →Microsoft-IIS/10.0 Microsoft-HTTPAPI/2.0
WWW-Authenticate →Bearer client_id="00000002-0000-0ff1-ce00-000000000000", trusted_issuers="00000001-0000-0000-c000-000000000000@5de110f8-2e0f-4d45-891d-bcf2218e253d,00000004-0000-0ff1-ce00-000000000000@contoso.com", token_types="app_asserted_user_v1 service_asserted_app_v1", authorization_uri="https://login.windows.net/common/oauth2/authorize"
Www-Authenticate →Basic realm="mail.contoso.com"
X-Powered-By →ASP.NET
request-id →5ca2c827-5147-474c-8457-63c4e5099c6e
```

If the AutoDiscover or ActiveSync responses are not similar to the above examples, you can investigate the following as possible causes:

1. If the AutoDiscover endpoint cannot be reached, then it's likely there's a firewall or load balancer configuration issue (for example, IP restrictions are configured and the required IP ranges are not present). Also, there may be a device in front of Exchange requiring pre-authentication to access the AutoDiscover endpoint.

2. If the AutoDiscover endpoint does not return the correct URL, then there is a configuration issue with the ActiveSync virtual directory's ExternalURL value.

3. If the ActiveSync endpoint cannot be reached, then there is a firewall or load balancer configuration issue. Again, one example is IP restrictions are configured and the required IP ranges are not present. Also, there may be a device in front of Exchange requiring pre-authentication to access the ActiveSync endpoint.

4. If the ActiveSync endpoint does not contain an authorization_uri value, verify that the EvoSTS authentication server is configured as the default endpoint using Exchange Management Shell:

   ```
   Get-AuthServer EvoSts | Format-List IsDefaultAuthorizationEndpoint
   ```

5. If the ActiveSync endpoint does not contain a WWW-Authenticate header, then a device in front of Exchange may be responding to the query.

Client synchronization issues

There are a few scenarios that can result in data being stale in Outlook for iOS and Android. Typically, this is due to an issue with the second access token (the token used by MRS in Exchange Online to synchronize the data with the on-premises environment). The two most common reasons for this issue are:

- SSL/TLS offloading on-premises.
- EvoSTS certificate metadata issues.

With SSL/TLS offloading, tokens are issued for a specific uri and that value includes the protocol value ("https://"). When the load balancer offloads SSL/TLS, the request Exchange receives comes in via HTTP, resulting in a claim mismatch due to the protocol value being http://. The following is an example of a response header from a Fiddler trace:

```
Content-Length →0
Date →Mon, 29 Jan 2018 19:51:46 GMT
Server →Microsoft-IIS/10.0 Microsoft-HTTPAPI/2.0
WWW-Authenticate →Bearer client_id="00000002-0000-0ff1-ce00-000000000000", trusted_issuers="00000001-0000-
0000-c000-000000000000@00c118a9-2de9-41d3-b39a-81648a7a5e4d",
authorization_uri="https://login.windows.net/common/oauth2/authorize", error="invalid_token"
WWW-Authenticate →Basic realm="mail.contoso.com"
X-Powered-By →ASP.NET
request-id →2323088f-8838-4f97-a88d-559bfcf92866
x-ms-diagnostics →2000003;reason="The hostname component of the audience claim value is invalid. Expected
'https://mail.contoso.com'. Actual 'http://mail.contoso.com'.";error_category="invalid_resource"
```

As specified above in the section *Technical and licensing requirements*, SSL/TLS offloading is not supported for OAuth flows.

For EvoSTS Certificate Metadata, the certificate metadata leveraged by EvoSTS is occasionally updated in Microsoft 365 or Office 365. The Exchange on-premises arbitration mailbox that has the organization capability of "OrganizationCapabilityManagement" is responsible for detecting the changes and for updating the corresponding metadata on-premises; this process executes every eight hours.

Exchange Administrators can find this mailbox by executing the following cmdlet using Exchange Management Shell:

```
$x=Get-mailbox -arbitration | ? {$_.PersistedCapabilities -like "OrganizationCapabilityManagement"};Get-
MailboxDatabaseCopyStatus $x.database.name
```

On the server hosting the database for the OrganizationCapabilityManagement arbitration mailbox, review the application event logs for events with a source of **MSExchange AuthAdmin**. The events should tell you if Exchange was able to refresh the metadata. If the metadata is out of date, you can manually refresh it with this ccmdlet:

```
Set-AuthServer EvoSts -RefreshAuthMetadata
```

You can also create a scheduled task that executes the above command every 24 hours.

Exchange Online statistics

You can use the following Exchange Online cmdlets to see statistical information for each synchronized on-premises mailbox.

1. First, obtain the location of the synchronized on-premises mailbox in the tenant, specifying the on-premises mailbox's identity (for example, jane@contoso.com).

   ```
   $m = Get-MailboxLocation <identity>
   ```

2. To see mailbox-related statistics, use

   ```
   Get-MailboxStatistics $m.id
   ```

3. To see mobile device statistics (like seeing when Outlook for iOS and Android last synchronized to Exchange Online), use

   ```
   Get-MobileDeviceStatistics -Mailbox $m.id
   ```

For more information, see Get-MailboxStatistics and Get-MobileDeviceStatistics.

Other issues

There are other issues that may prevent hybrid Modern Authentication from functioning correctly. For more information, see the troubleshooting section in Announcing Hybrid Modern Authentication for Exchange On-Premises.

Using Basic authentication with Outlook for iOS and Android

8/3/2020 • 4 minutes to read • Edit Online

The Outlook app for iOS and Android is designed to bring together email, calendar, contacts, and other files, enabling users in your organization to do more from their mobile devices. This article provides an overview of the architecture and the storage design of the app, so that Exchange administrators can deploy and maintain Outlook for iOS and Android in their Exchange organizations.

NOTE

This article is about using the app in an Exchange 2010, Exchange 2013, Exchange 2016 or Exchange 2019 environment where hybrid modern authentication is **not** enabled. For more information about using hybrid Modern Authentication for on-premises mailboxes with the app, see Using hybrid Modern Authentication with Outlook for iOS and Android. For information about using the app with Exchange Online, see Outlook for iOS and Android in Exchange Online.

Outlook for iOS and Android architecture

Outlook for iOS and Android is a cloud-backed application. This means your experience consists of a locally installed app powered by a secure and scalable service running in the Microsoft Cloud.

For Exchange Server mailboxes, Outlook for iOS and Android's architecture is built directly into the Microsoft Cloud, providing customers the additional benefits of security, privacy, built-in compliance, and transparent operations that Microsoft commits to in the Microsoft Trust Center.

The following environments will take advantage of this Microsoft 365 or Office 365-based architecture:

- In Exchange Server 2010 environments
- When a hybrid relationship between Exchange 2013, 2016, or 2019 on-premises and Microsoft 365 or Office 365 has not been configured
- When hybrid Modern Authentication has not been enabled between Exchange 2013, 2016, or 2019 on-premises and Microsoft 365 or Office 365

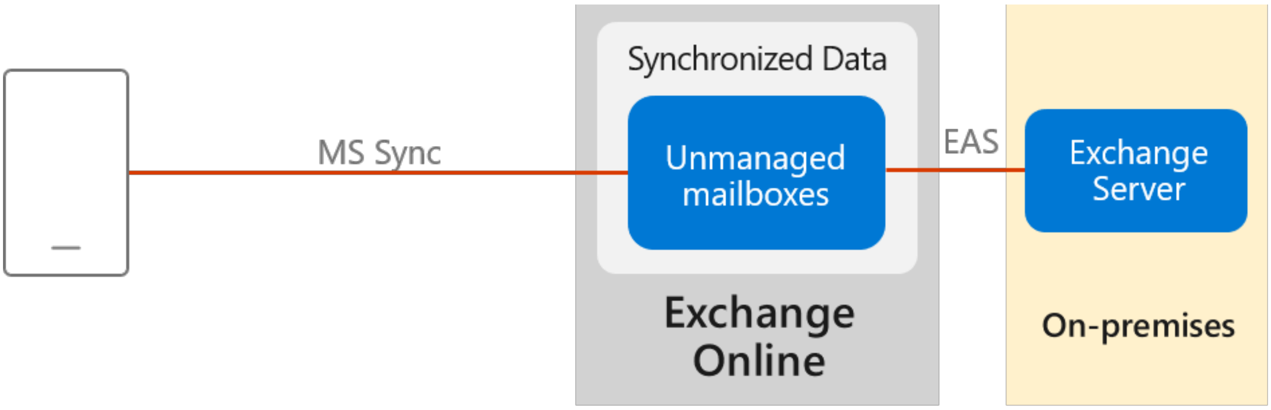

Within the Microsoft 365 or Office 365-based architecture, Outlook for iOS and Android utlilizes the native Microsoft sync technology for data synchronization which is protected by TLS-secured connections end-to-end, between Microsoft 365 or Office 365 and the app.

The Exchange ActiveSync (EAS) connection between Exchange Online and the on-premises environment enables synchronization of the users' on-premises data and includes four weeks of email, all calendar data, all contact data, and out-of-office status. The region in which this data is synchronized into depends on the IP address in use by the mobile device at the time synchronization is setup. If you have a hybrid setup with an Exchange Online tenant, the on-premises data is not synchronized into your tenant; instead, the data is synchronized into Outlook.com. If you want to control and manage your on-premises data from within your tenant, you need to enable hybrid Modern Authentication with Outlook for iOS and Android.

Data synchronization between the Exchange on-premises environment and Exchange Online happens independent of user behavior. This ensures that new messages are delivered to the devices very quickly. For more information on how the user authentication model enables data synchronization independently of user behavior, see Passwords and security in Outlook for iOS and Android for Exchange Server.

Processing information in the Microsoft Cloud enables advanced features and capabilities, such as the categorization of email for the Focused Inbox, customized experience for travel and calendar, and improved search speed. Relying on the cloud for intensive processing and minimizing the resources required from users' devices enhances the app's performance and stability. Lastly, it allows Outlook to build features that work across all email accounts, regardless of the technological capabilities of the underlying servers (such as different versions of Exchange Server, Microsoft 365, or Office 365).

IMPORTANT

On-premises mailboxes using basic authentication with Outlook for iOS and Android do not support Enterprise Mobility + Security features such as Azure Active Directory conditional access and Intune app protection policies. For support with these technologies, see Using hybrid Modern Authentication with Outlook for iOS and Android.

Data security, access, and auditing controls

With on-premises data being synchronized with Exchange Online, customers have questions about how the data is protected in Exchange Online. The white paper Encryption in the Microsoft Cloud discusses how BitLocker is used for volume-level encryption.

By default, Microsoft engineers have zero standing administrative privileges and zero standing access to customer content in Microsoft 365 and Office 365. The white paper Administrative Access Controls in Microsoft 365 discusses personnel screening, background checks, Lockbox and Customer Lockbox, and more.

ISO Audited Controls on Service Assurance documentation provides the status of audited controls from global information security standards and regulations that Microsoft 365 and Office 365 have implemented.

Connectivity Requirements

Microsoft recommends that the on-premises endpoints for AutoDiscover and ActiveSync protocols be opened and accessible from the Internet without any restrictions. In certain situations that may not be possible. If you must place restrictions on your on-premises firewall or gateway edge devices, Microsoft recommends filtering based on FQDN endpoints. If FQDN endpoints cannot be used, then filter on IP addresses. Make sure the following IP subnets and FQDNs are whitelisted:

- All Exchange Online FQDNs and IP subnet ranges as defined in Microsoft 365 and Office 365 URLs and IP address ranges.
- The AutoDetect FQDNs and IP subnet ranges defined in Additional Microsoft 365 or Office 365 IP Addresses and URLs not included in the web services. This is required because the AutoDetect service establishes connections to the on-premises infrastructure.
- All Outlook iOS and Android and Office mobile app FQDNs as defined in Microsoft 365 or Office 365 URLs

and IP address ranges.

Client features that aren't supported

The following features are not support for on-premises mailboxes using basic authentication with Outlook for iOS and Android:

- Draft folder and Draft messages synchronization
- Shared calendar access and Delegate calendar access
- Shared and delegate mailbox data access
- Cortana Time to Leave / Travel Time
- Calendar attachments
- Rich meeting locations
- Task management with Microsoft To-Do
- Favorite People with Notifications
- Add-ins
- Interesting Calendars
- Avatar support
- Play My Emails
- S/MIME
- Sensitivity labeling
- Discover Feed
- Privacy settings

Account setup in Outlook for iOS and Android using Basic authentication

8/3/2020 • 5 minutes to read • Edit Online

Outlook for iOS and Android offers Exchange administrators the ability to "push" account configurations to their on-premises users who use Basic authentication with the ActiveSync protocol. This capability works with any Mobile Device Management (MDM) provider who uses the Managed App Configuration channel for iOS or the Android in the Enterprise channel for Android.

For on-premises users enrolled in Microsoft Intune, you can deploy the account configuration settings using Intune in the Azure Portal.

Once an account configuration has been created and the user enrolls their device, Outlook for iOS and Android will detect that an account is "Found" and will then prompt the user to add the account. The only information the user needs to enter to complete the setup process is their password. Then, the user's mailbox content will load and the user can begin using the app.

The following images show an example of the end-user setup process after Outlook for iOS and Android has been configured in Intune in the Azure Portal.

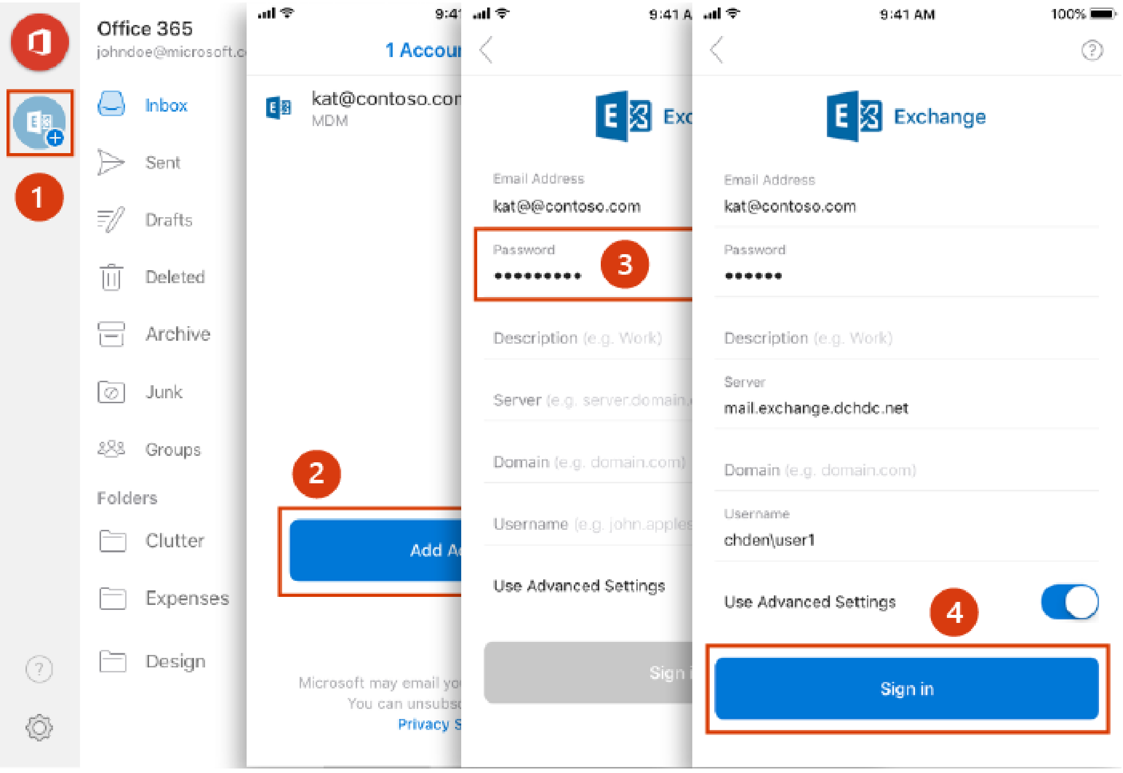

Create an app configuration policy for Outlook for iOS and Android using Microsoft Intune

If you're using Microsoft Intune as your mobile device management provider, the following steps will allow you to deploy account configuration settings for your on-premises mailboxes that leverage basic authentication with the ActiveSync protocol. Once the configuration is created, you can assign the settings to groups of users, as detailed in the next section, Assign configuration settings.

> **NOTE**
>
> If users in your organization use both iOS and Android for Work devices, you'll need to create a separate app configuration policy for each platform.

1. Sign into Microsoft Endpoint Manager.
2. Select **Apps** and then select **App configuration policies**.
3. On the **App Configuration policies** blade, choose **Add** and select **Managed devices**.
4. On the **Add app configuration** blade, enter a **Name**, and optional **Description** for the app configuration settings.
5. For **Platform**, choose either **iOS/iPadOS** or **Android**.
6. For **Associated app**, choose **Select the required app**, and then, on the **Targeted apps** blade, choose **Microsoft Outlook**.

   > **NOTE**
   >
   > If Outlook is not listed as an available app, then you must add it by following the instructions in Add Android store apps to Microsoft Intune and How to add iOS store apps to Microsoft Intune.

7. Click **OK** to return to the **Add app configuration** blade.
8. Choose **Configuration Settings**. On the **Configuration** blade, select **Use configuration designer** for the **Configuration settings format**. The key value pairs used in this section are defined in the section Key value pairs.
9. If you want to deploy account setup configuration, select **Yes** for **Configure email account settings** and configure appropriately:

   - For **Authentication type**, select **Basic authentication**. This is required for on-premises accounts that do not leverage hybrid modern authentication.
   - For **Username attribute from AAD**, select **User Principal Name** or **sAMAccountName**. If **sAMAccountName** is selected, enter the NetBIOS domain name in the **Account domain** field.
   - For **Email address attribute from AAD**, select **Primary SMTP Address**.
   - For **Email server**, enter the Exchange ActiveSync externally accessible domain name.
   - For **Email account name**, enter a descriptive value for the account.

10. If you want to deploy general app configuration settings, configure the desired settings accordingly:

    - For **Focused Inbox**, choose from the available options: **Not configured** (default), **On** (app default), **Off**.
    - For **Require Biometrics to access the app**, choose from the available options: **Not configured** (default), **On**, **Off** (app default). When selecting **On** or **Off**, administrators can choose to allow the user to change the app setting's value. Select **Yes** (app default) to allow the user to change the setting or choose **No** if you want to prevent the user from changing the setting's value. This setting is only available in Outlook for iOS.
    - For **Save Contacts**, choose from the available options: **Not configured** (default), **On**, **Off** (app default). When selecting **On** or **Off**, administrators can choose to allow the user to change the app

setting's value. Select **Yes** (app default) to allow the user to change the setting or choose **No** if you want to prevent the user from changing the setting's value.

- For **Default app signature**, choose from the available options: **Not configured** (default), **On** (app default), **Off**.
- For **Block external images**, choose from the available options: **Not configured** (default), **On**, **Off** (app default). When selecting **On** or **Off**, administrators can choose to allow the user to change the app setting's value. Select **Yes** (app default) to allow the user to change the setting or choose **No** if you want to prevent the user from changing the setting's value.
- For **Organize mail by thread**, choose from the available options: **Not configured** (default), **On** (app default), **Off**.

11. When you are done, choose **OK**.
12. On the **Add app configuration** blade, choose **Add**.

The newly created configuration policy is displayed on the **App configuration** blade.

Assign configuration settings

You assign the settings to groups of users in Azure Active Directory. When a user has the Microsoft Outlook app installed, the app is managed by the settings you have specified. To do this:

1. From the **Apps - App configuration policies** blade, select the app configuration policy you want to assign.
2. On the next blade, choose **Assignments**.
3. On the **Assignments** blade, select **Select groups to include** and choose the Azure AD group to which you want to assign the app configuration, and then choose **Select**.
4. Select **Save** to save and assign the app configuration policy.

Key value pairs

When you create an app configuration policy in the Azure Portal or through your MDM provider, you will need the following key value pairs:

KEY	VALUES
com.microsoft.outlook.EmailProfile.EmailAccountName	This value specifies the display name email account as it will appear to users on their devices. **Value type**: String **Accepted values**: Display Name **Default if not specified**: <blank> **Required**: Yes **Example**: user **Intune Token***: {{username}}
com.microsoft.outlook.EmailProfile.EmailAddress	This value specifies the email address to be used for sending and receiving mail. **Value type**: String **Accepted values**: Email address **Default if not specified**: <blank> **Required**: Yes **Example**: user@companyname.com **Intune Token***: {{mail}}

KEY	VALUES
com.microsoft.outlook.EmailProfile.EmailUPN	This value specifies the User Principal Name or username for the email profile that will be used to authenticate the account. **Value type**: String **Accepted values**: UPN Address or username **Default if not specified**: <blank> **Required**: Yes **Example**: userupn@companyname.com **Intune Token**[*]: {{userprincipalname}}
com.microsoft.outlook.EmailProfile.ServerAuthentication	This value specifies the authentication method for the user. **Value type**: String **Accepted values**: 'Username and Password' **Default if not specified**: 'Username and Password' **Required**: No **Example**: 'Username and Password'
com.microsoft.outlook.EmailProfile.ServerHostName	This value specifies the host name of your Exchange server. **Value type**: String **Accepted values**: ActiveSync FQDN **Default if not specified**: <blank> **Required**: Yes **Example**: mail.companyname.com
com.microsoft.outlook.EmailProfile.AccountDomain	This value specifies the user's account domain. **Value type**: String **Accepted values**: Domain **Default if not specified**: <blank> **Required**: No **Example**: companyname
com.microsoft.outlook.EmailProfile.AccountType	This value specifies the account type being configured based on the authentication model. **Value type**: String **Accepted values**: BasicAuth **Default if not specified**: BasicAuth **Required**: No **Example**: BasicAuth

[*] Microsoft Intune users can use tokens that will expand to the correct value according to the MDM enrolled user. See Add app configuration policies for managed iOS devices for more information.

Passwords and security in Outlook for iOS and Android for Exchange Server

8/3/2020 • 6 minutes to read • Edit Online

This article describes how passwords and security work in Outlook for iOS and Android with Exchange Server when using Basic authentication with the Exchange ActiveSync protocol.

> **IMPORTANT**
> Outlook for iOS and Android supports hybrid Modern Authentication for on-premises mailboxes which eliminates the need to leverage basic authentication. The information contained in this article only pertains to basic authentication. For more information, please see Using hybrid Modern Authentication with Outlook for iOS and Android.

Creating an account and protecting passwords

The first time the Outlook app for iOS and Android is run in an Exchange on-premises environment, Outlook generates a random AES-128 key. This key is known as the *device key* and is stored only on the user's device.

When a user logs onto Exchange with Basic authentication, the username, password, and a unique AES-128 device key are sent from the user's device to the Outlook cloud service over a TLS connection, where the device key is held in runtime compute memory. After verifying the password with the Exchange server, the Microsoft 365 or Office 365-based architecture uses the device key to encrypt the password, and the encrypted password is then stored in the service. The device key, meanwhile, is wiped from memory and never stored in the Microsoft 365 or Office 365-based architecture (the key is only stored on the user's device).

Next, when a user attempts to connect to Exchange to retrieve mailbox data, the device key is again passed from the device to the Microsoft 365 or Office 365-based architecture over a TLS-secured connection, where it is used to decrypt the password in runtime compute memory. Once decrypted, the password is never stored in the service or written to a local storage disk, and the device key is once again wiped from memory.

After the Microsoft 365 or Office 365-based architecture has decrypted the password at runtime, the service can then connect to the Exchange server to synchronize mail, calendar, and other mailbox data. As long as the user continues to open and use Outlook periodically, the Microsoft 365 or Office 365-based architecture will keep a copy of the user's decrypted password in memory to keep the connection to the Exchange server active.

Compliance considerations when sending passwords

Before you enable anything that allows for the transmission of passwords from your on-premises Exchange environment, be sure to consider the possible ramifications. For example, transmitting passwords to Microsoft 365 or Office 365-based architecture might result in your inability to meet the requirements of PCI-DSS or ISO/IEC 27001.

Furthermore, if you connect and synchronize email, calendars, and other email-related data, you might run into issues of compliance with GDPR, which restricts the private information that you can transmit without owner consent. This information might be contained in and found within emails, calendar items, and so on.

Account inactivity and flushing passwords from memory

After three days of inactivity, the Microsoft 365 or Office 365-based architecture will flush a decrypted password from memory. With the decrypted password flushed, the architecture is unable to access a user's mailbox on-premises. The encrypted password remains stored in the Microsoft 365 or Office 365-based architecture, but decrypting it again isn't possible without the device key, which is only available from the user's device.

There are three ways a user account can become inactive:

- Outlook for iOS and Android is uninstalled by the user.
- Background app refresh is disabled in the Settings options, and then a force-quit is applied to Outlook.
- No internet connection is available on the device, preventing Outlook from synchronizing with Exchange.

> **NOTE**
>
> Outlook will not become inactive simply because the user does not open the app for some time, such as over a weekend or while on vacation. As long as background app refresh is enabled (which is the default setting for Outlook for iOS and Android), functions like push notifications and background synchronization of email will count as activity.

Flushing encrypted password and synchronized mailbox data from Microsoft 365 or Office 365

The Microsoft 365 or Office 365-based architecture flushes, or deletes, inactive accounts on a weekly schedule. After a user account becomes inactive, the architecture will flush both the encrypted password and all of the user's synchronized mailbox content out of the service.

Device and service security combination

Each user's unique device key is never stored in the Microsoft 365 or Office 365-based architecture, and a user's Exchange password is never stored on the device. This architecture means that for a malicious party to gain access to a user's password, they would need both unauthorized access to the Microsoft 365 or Office 365-based architecture and physical access to that user's device.

By enforcing PIN policies and encryption on devices in your organization, the malicious party would also have to defeat a device's encryption to get access to the device key. This would all have to take place before the user noticed that the device was compromised and could request a remote wipe for the device.

Password security FAQ

The following are frequently asked questions regarding security design and settings for Outlook for iOS and Android when used with Basic authentication.

Are user credentials stored in the Microsoft 365 or Office 365-based architecture if I block Outlook from accessing my Exchange Server?

If you have chosen to block Outlook for iOS and Android from accessing your on-premises Exchange servers, the initial connection will be rejected by Exchange. User credentials will not be stored by the Outlook cloud service and the credentials presented in the failed connection attempt are immediately flushed from memory.

How is the unique device key and user password encrypted in transit to the Microsoft 365 or Office 365-based architecture?

All communication between the Outlook app and the Microsoft 365 or Office 365-based architecture is through an encrypted TLS connection. The Outlook app is capable of connecting with the Microsoft 365 or Office 365-based architecture and nothing else.

How do I remove a user's credentials and mailbox information from the Microsoft 365 or Office 365-based architecture?

Have the user uninstall Outlook for iOS and Android on all devices. All data will be removed from the Microsoft 365 or Office 365-based architecture in approximately 3-7 days.

The app is closed or uninstalled, but I still see it connecting to my Exchange server. How is this happening?

The Microsoft 365 or Office 365-based architecture decrypts user passwords in runtime compute memory and then uses the decrypted passwords to connect to Exchange. Since the architecture is connecting to Exchange on behalf of the device to fetch and cache mailbox data, it can continue for a short period until the service detects that Outlook is no longer requesting data.

If a user uninstalls the app from their device without first using the **Delete Account** option, the Microsoft 365 or Office 365-based architecture will stay connected to your Exchange server until the account becomes inactive, as described above in "Account inactivity and flushing passwords from memory." To stop this activity, follow Option 1 or Option 3 from the above FAQ, or block the app, as described in Blocking Outlook for iOS and Android.

Is a user password less secure in Outlook for iOS and Android than when using other Exchange ActiveSync clients?

No. EAS clients generally save user credentials locally on the user's device. This means a stolen or compromised device could result in a malicious party gaining access to the user's password. With the security design of Outlook for iOS and Android, a malicious party would need unauthorized access to the Microsoft 365 or Office 365-based architecture **and** have physical access to a user's device.

What happens if a user attempts to use Outlook for iOS and Android after their data has been deleted from the Outlook cloud service?

If a user account becomes inactive (such as by disabling background app refresh on the device or having their device disconnected from the Internet for some time), the Outlook app will reconnect to the Microsoft 365 or Office 365-based architecture the next time the app is launched, and the password encryption and email caching process will restart. This is all transparent to the user.

Is there a way to prevent the use of Basic authentication for on-premises mailboxes with Outlook for iOS and Android?

Yes, you can deploy hybrid Modern Authentication. For more information, see Using hybrid Modern Authentication with Outlook for iOS and Android.

Managing devices for Outlook for iOS and Android for Exchange Server

8/3/2020 • 8 minutes to read • Edit Online

IMPORTANT

Outlook for iOS and Android supports hybrid Modern Authentication for on-premises mailboxes which eliminates the need to leverage basic authentication. The information contained in this article only pertains to basic authentication. For more information, please see Using hybrid Modern Authentication with Outlook for iOS and Android.

Microsoft recommends Exchange ActiveSync for managing the mobile devices that are used to access Exchange mailboxes in your on-premises environment. Exchange ActiveSync is a Microsoft Exchange synchronization protocol that lets mobile phones access an organization's information on a server that's running Microsoft Exchange.

This article focuses on specific Exchange ActiveSync features and scenarios for mobile devices running Outlook for iOS and Android when authenticating with Basic authentication. Complete information about the Microsoft Exchange synchronization protocol is available in Exchange ActiveSync. In addition, there is information on the Office Blog detailing password enforcement and other benefits of using Exchange ActiveSync with devices running Outlook for iOS and Android.

Mobile device mailbox policy

Outlook for iOS and Android supports the following mobile device mailbox policy settings in Exchange on-premises:

- Device encryption enabled
- Min password length (only on Android)
- Password enabled
- Allow Bluetooth (used to manage the Outlook for Android wearable app)
  - When AllowBluetooth is enabled (default behavior) or configured for HandsfreeOnly, wearable synchronization between Outlook on the Android device and Outlook on the wearable is allowed for the work or school account.
  - When AllowBluetooth is disabled, Outlook for Android will disable synchronization between Outlook on the Android device and Outlook on the wearable for the specified work or school account (and delete any data previously synced for the account). Disabling the synchronization is controlled entirely within Outlook itself; Bluetooth is not disabled on the device or wearable nor is any other wearable app affected.

NOTE

Outlook for Android will roll out support for the AllowBluetooth setting beginning at the end of August.

For information on how to create or modify an existing mobile device mailbox policy, see Mobile device mailbox policies.

PIN lock and device encryption

If your organization's Exchange ActiveSync policy requires a password on mobile devices in order for users to synchronize email, Outlook will enforce this policy at the device level. This works differently between iOS devices and Android devices, based on the available controls provided by Apple and Google.

On iOS devices, Outlook checks to make sure a passcode or PIN is properly set. In the event a passcode is not set, Outlook prompts users to create a passcode in iOS settings. Until the passcode is setup, the user will be unable to access Outlook for iOS.

On Android devices, Outlook will enforce screen lock rules. In addition, Google provides controls that allow Outlook for Android to comply with Exchange policies regarding password length and complexity, and the number of allowable screen-unlock attempts before wiping the phone. Outlook for Android will also encourage storage encryption if it is not enabled, guiding users through this process with a step-by-step walkthrough.

iOS and Android devices that do not support these password security settings will not be able to connect to an Exchange mailbox.

Device encryption

iOS devices are shipped with built-in encryption, which Outlook uses once the passcode is enabled to encrypt all the data Outlook stores locally on the iOS device. Therefore, iOS devices with a PIN are encrypted whether or not this is required by an ActiveSync policy.

Outlook for Android supports device encryption via Exchange mobile device mailbox policies. However, prior to Android 7.0, the availability and implementation of this process varies by Android OS version and device manufacturer, which allow the user to cancel out during the encryption process. With changes that Google introduced to Android 7.0, Outlook for Android is now able to enforce encryption on devices running Android 7.0 or later. Users with devices running those operating systems will not be able to cancel out of the encryption process.

Even if the Android device is unencrypted and an attacker is in possession of the device, as long as a device PIN is enabled, the Outlook database remains inaccessible. This is true even with USB debugging enabled and the Android SDK installed. If an attacker attempts to root the device to bypass the PIN to gain access to this information, the rooting process wipes all device storage and removes all Outlook data. If the device is unencrypted and rooted by the user prior to being stolen, it is possible for an attacker to gain access to the Outlook database by enabling USB debugging on the device and plugging the device into a computer with the Android SDK installed.

Remote wipe with Exchange ActiveSync

Exchange ActiveSync enables administrators to remotely wipe devices, such as if they become compromised or lost/stolen. With Outlook for iOS and Android, a remote wipe only wipes data within the Outlook app itself and does not trigger a full device wipe.

See Perform a remote wipe on a mobile phone for more information.

Device access policy

Outlook for iOS and Android should be enabled by default, but in some existing Exchange on-premises environments the app may be blocked for a variety of reasons. Once an organization decides to standardize how users access Exchange data and use Outlook for iOS and Android as the only email app for end users, you can configure blocks for other email apps running on users' iOS and Android devices. You have two options for instituting these blocks within Exchange on-premises: the first option blocks all devices and only allows usage of Outlook for iOS and Android; the second option allows you to block individual devices from using the native Exchange ActiveSync apps.

NOTE

Because device IDs are not governed by any *physical device* ID, they can change without notice. When this happens, it can cause unintended consequences when device IDs are used for managing user devices, as existing 'allowed' devices may be unexpectedly blocked or quarantined by Exchange. Therefore, Microsoft recommends administrators only set mobile device access policies that allow/block devices based on device type or device model.

Option 1: Block all email apps except Outlook for iOS and Android

You can define a default block rule and then configure an allow rule for Outlook for iOS and Android, and for Windows devices, using the following Exchange on-premises PowerShell commands. This configuration will prevent any Exchange ActiveSync native app from connecting, and will only allow Outlook for iOS and Android.

1. Create the default block rule:

   ```
   Set-ActiveSyncOrganizationSettings -DefaultAccessLevel Block
   ```

2. Create an allow rule for Outlook for iOS and Android

   ```
   New-ActiveSyncDeviceAccessRule -Characteristic DeviceModel -QueryString "Outlook for iOS and Android" -
   AccessLevel Allow
   ```

3. **Optional**: Create rules that allow Outlook on Windows devices for Exchange ActiveSync connectivity (WindowsMail refers to the Mail app included in Windows 10):

   ```
   New-ActiveSyncDeviceAccessRule -Characteristic DeviceType -QueryString "WindowsMail" -AccessLevel Allow
   ```

Option 2: Block native Exchange ActiveSync apps on Android and iOS devices

Alternatively, you can block native Exchange ActiveSync apps on specific Android and iOS devices or other types of devices.

1. Confirm that there are no Exchange ActiveSync device access rules in place that block Outlook for iOS and Android:

   ```
   Get-ActiveSyncDeviceAccessRule | where {$_.AccessLevel -eq "Block" -and $_.QueryString -like "Outlook*"}
   | ft Name,AccessLevel,QueryString -auto
   ```

   If any device access rules that block Outlook for iOS and Android are found, type the following to remove them:

   ```
   Get-ActiveSyncDeviceAccessRule | where {$_.AccessLevel -eq "Block" -and $_.QueryString -like "Outlook*"}
   | Remove-ActiveSyncDeviceAccessRule
   ```

2. You can block most Android and iOS devices with the following commands:

   ```
   New-ActiveSyncDeviceAccessRule -Characteristic DeviceType -QueryString "Android" -AccessLevel Block
   New-ActiveSyncDeviceAccessRule -Characteristic DeviceType -QueryString "iPad" -AccessLevel Block
   New-ActiveSyncDeviceAccessRule -Characteristic DeviceType -QueryString "iPhone" -AccessLevel Block
   New-ActiveSyncDeviceAccessRule -Characteristic DeviceType -QueryString "iPod" -AccessLevel Block
   ```

3. Not all Android device manufacturers specify "Android" as the DeviceType. Manufacturers may specify a unique value with each release. In order to find other Android devices that are accessing your environment, execute the following command to generate a report of all devices that have an active Exchange ActiveSync

partnership:

```
Get-MobileDevice | Select-Object DeviceOS,DeviceModel,DeviceType | Export-CSV c:\temp\easdevices.csv
```

4. Create additional block rules, depending on your results from Step 3. For example, if you find your environment has a high usage of HTCOne Android devices, you can create an Exchange ActiveSync device access rule that blocks that particular device, forcing the users to use Outlook for iOS and Android. In this example, you would type:

```
New-ActiveSyncDeviceAccessRule -Characteristic DeviceType -QueryString "HTCOne" -AccessLevel Block
```

> **NOTE**
>
> The QueryString parameter does not accept wildcards or partial matches.

Additional resources:

- New-ActiveSyncDeviceAccessRule
- Get-MobileDevice
- Set-ActiveSyncOrganizationSettings

Blocking Outlook for iOS and Android

Every Exchange organization has different policies regarding security and device management. If an organization decides that Outlook for iOS and Android doesn't meet their needs or is not the best solution for them, administrators have the ability to block the app. Once the app is blocked, mobile Exchange users in your organization can continue accessing their mailboxes by using the built-in mail applications on iOS and Android.

The `New-ActiveSyncDeviceAccessRule` cmdlet has a `Characteristic` parameter, and there are three `Characteristic` options that administrators can use to block the Outlook for iOS and Android app. The options are UserAgent, DeviceModel, and DeviceType. In the two blocking options described in the following sections, you will use one or more of these characteristic values to restrict the access that Outlook for iOS and Android has to the mailboxes in your organization.

The values for each characteristic are displayed in the following table:

CHARACTERISTIC	STRING FOR IOS	STRING FOR ANDROID
DeviceModel	Outlook for iOS and Android	Outlook for iOS and Android
DeviceType	Outlook	Outlook
UserAgent	Outlook-iOS-Android/1.0	Outlook-iOS-Android/1.0

With the `New-ActiveSyncDeviceAccessRule` cmdlet, you can define a device access rule, using either the `DeviceModel` or `DeviceType` characteristic. In both cases, the access rule blocks Outlook for iOS and Android across all platforms, and will prevent any device, on both the iOS platform and Android platform, from accessing an Exchange mailbox via the app.

The following are two examples of a device access rule. The first example uses the `DeviceModel` characteristic; the second example uses the `DeviceType` characteristic.

```
New-ActiveSyncDeviceAccessRule -Characteristic DeviceType -QueryString "Outlook" -AccessLevel Block
```

```
New-ActiveSyncDeviceAccessRule -Characteristic DeviceModel -QueryString "Outlook for iOS and Android" -
AccessLevel Block
```

Default settings for Exchange virtual directories

8/3/2020 • 2 minutes to read • Edit Online

Exchange Server 2016 and Exchange Server 2019 automatically configure multiple Internet Information Services (IIS) virtual directories during the server installation. The tables in the following sections show the settings for the Client Access (frontend) services on Mailbox servers and the default IIS authentication and Secure Sockets Layer (SSL) settings.

Client Access services (frontend) on Mailbox servers

The following table lists the default settings in the Client Access services (the default web site) on Exchange Mailbox servers.

VIRTUAL DIRECTORY	AUTHENTICATION METHOD	SSL SETTINGS	MANAGEMENT METHOD
Default Web Site	Anonymous	Required	IIS management console
API[1]	Anonymous authentication Windows authentication	SSL required Requires 128-bit encryption	
aspnet_client	Anonymous authentication	SSL required Requires 128-bit encryption	IIS management console
Autodiscover	Anonymous authentication Basic authentication Windows authentication	SSL required Requires 128-bit encryption	EAC or Exchange Management Shell
ecp	Anonymous authentication Basic authentication	SSL required Requires 128-bit encryption	EAC or Exchange Management Shell
EWS	Anonymous authentication Windows authentication	SSL required Requires 128-bit encryption	EAC or Exchange Management Shell
MAPI	Windows authentication	SSL required Requires 128-bit encryption	EAC or Exchange Management Shell
Microsoft-Server-ActiveSync	Basic authentication	SSL required Requires 128-bit encryption	EAC or Exchange Management Shell
OAB	Windows authentication	SSL required Requires 128-bit encryption	EAC or Exchange Management Shell
owa	Basic authentication	SSL required Requires 128-bit encryption	EAC or Exchange Management Shell
PowerShell	By default, all authentication methods are disabled.	Not required	EAC or Exchange Management Shell
Rpc	Basic authentication Windows authentication	Not required	EAC or Exchange Management Shell

[1] The API virtual directory is available in Exchange 2016 CU3 or newer.

Back End Virtual Directories on Mailbox servers

The following table lists the default settings in the back end services on Exchange Mailbox servers.

VIRTUAL DIRECTORY	AUTHENTICATION METHOD	SSL SETTINGS	MANAGEMENT METHOD
Exchange Back End	Anonymous authentication	SSL required Requires 128-bit encryption	This virtual directory should not be configured by the user.
API[1]	Anonymous authentication Windows authentication	SSL required Requires 128-bit encryption	This virtual directory should not be configured by the user.
Autodiscover	Anonymous authentication Windows authentication	SSL required Requires 128-bit encryption	This virtual directory should not be configured by the user.
ecp	Anonymous authentication Windows authentication	SSL required Requires 128-bit encryption	This virtual directory should not be configured by the user.
EWS	Anonymous authentication Windows authentication	SSL required Requires 128-bit encryption	This virtual directory should not be configured by the user.
Microsoft-Server-ActiveSync	Basic authentication	SSL required Requires 128-bit encryption	This virtual directory should not be configured by the user.
OAB	Windows authentication	SSL required Requires 128-bit encryption	This virtual directory should not be configured by the user.
owa	Anonymous authentication Windows authentication	SSL required Requires 128-bit encryption	This virtual directory should not be configured by the user.
PowerShell	Windows authentication	SSL required Requires 128-bit encryption	This virtual directory should not be configured by the user.
Rpc	Windows authentication	Not required	This virtual directory should not be configured by the user.
RpcWithCert	Windows authentication	Not required	This virtual directory should not be configured by the user.

[1]The API virtual directory is available in Exchange 2016 CU3 or newer.

See also

Virtual directory management

Outlook on the web in Exchange Server

8/3/2020 • 2 minutes to read • Edit Online

The user interface in Outlook on the web (formerly known as Outlook Web App) for Exchange Server has been optimized and simplified for use with phones and tablets. Supported web browsers give users access to more Outlook features. Unsupported web browsers give users the light version of Outlook on the web that has less features. For more information about features and supported web browsers, see Outlook on the web (formerly Outlook Web App) and Outlook on the web (formerly Outlook Web App).

When you install Exchange Server, Outlook on the web is automatically available for internal users at `https://<ServerName>/owa` (for example, `https://mailbox01.contoso.com/owa`). But, you'll likely want to configure Outlook on the web for external access (for example, `https://mail.contoso.com/owa`). For more information, see Step 4: Configure external URLs in Configure mail flow and client access on Exchange servers.

In an Outlook 2010 or later installation that's connected to an Exchange mailbox, you can typically see the Outlook on the web URL at **File** > **Info** > **Account Information** in the **Account Settings** section.

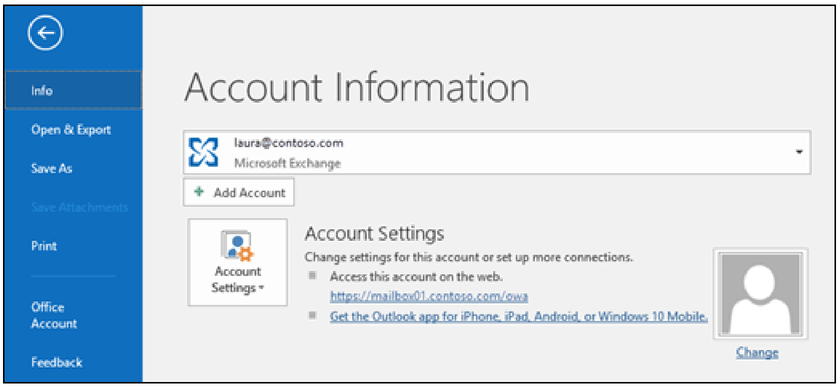

Account Information

laura@contoso.com
Microsoft Exchange

Account Settings

Change settings for this account or set up more connections.

- Access this account on the web.
  https://mailbox01.contoso.com/owa
- Get the Outlook app for iPhone, iPad, Android, or Windows 10 Mobile.

Outlook on the web is provided by the Client Access (frontend) services on Mailbox servers. In Exchange Server, Client Access services are part of the Mailbox server, so you can't configure a standalone Client Access server like you could in previous versions of Exchange. For more information, see Client access protocol architecture.

If you're looking for information about Outlook on the web in Microsoft 365 or Office 365, see Using email in Outlook on the web.

Administrative tasks for managing Outlook on the web

The configuration and management tasks that are documented for Outlook on the web in Outlook 2016 are listed in the following table.

TOPIC	DESCRIPTION
View or configure Outlook on the web virtual directories in Exchange Server	View and configure the properties of Outlook on the web for all users that connect to the server.
Configure http to https redirection for Outlook on the web in Exchange Server	Redirect Outlook on the web unencrypted http requests to https.
Create a theme for Outlook on the web in Exchange Server	Outlook on the web comes with built-in themes that define the colors and icons that are used in Outlook on the web, but you can also create your own themes.

TOPIC	DESCRIPTION
Customize the Outlook on the web sign-in, language selection, and error pages in Exchange Server	Customize key pages in Outlook on the web.
Use AD FS claims-based authentication with Outlook on the web	Centralize Outlook on the web authentication by using Active Directory Federation Services.

Enable or disable Outlook on the web access to mailboxes in Exchange Server

8/3/2020 • 5 minutes to read • Edit Online

Administrators can use the Exchange admin center (EAC) or the Exchange Management Shell to enable or disable Outlook on the web access to a mailbox. By default, users can access their mailboxes by using Outlook on the web. When you disable Outlook on the web access to mailboxes, users can still access their mailboxes by using Outlook or other email clients.

For additional management tasks related to user access to mailboxes, see these topics:

- Enable or disable Exchange ActiveSync access to mailboxes in Exchange Server
- Enable or disable POP3 or IMAP4 access to mailboxes in Exchange Server
- Enable or disable MAPI access to mailboxes in Exchange Server

What do you need to know before you begin?

- Estimated time to complete: 5 minutes.
- For more information about accessing and using the EAC, see Exchange admin center in Exchange Server.
- To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Client Access user settings" entry in the Clients and mobile devices permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Enable or disable Outlook on the web access to a single mailbox

Use the EAC to Enable or disable Outlook on the web access to a single mailbox

1. In the EAC, go to **Recipients** > **Mailboxes**.
2. In the list of mailboxes, find the mailbox that you want to modify. You can:
   - Scroll through the list of mailboxes.
   - Click **Search** 🔍 and enter part of the user's name, email address, or alias.
   - Click **More options** ••• > **Advanced search** to find the mailbox.

   Once you've found the mailbox that you want to modify, select it, and then click **Edit** ✎.

3. On the mailbox properties page that opens, click **Mailbox features**.

4. In the **Email Connectivity** section, configure one of these settings:

   - If you see **Outlook on the web: Enabled**, click **Disable** to disable it, and then click **Yes** in the warning message that appears.

   - If you see **Outlook on the web: Disabled**, click **Enable** to enable it.

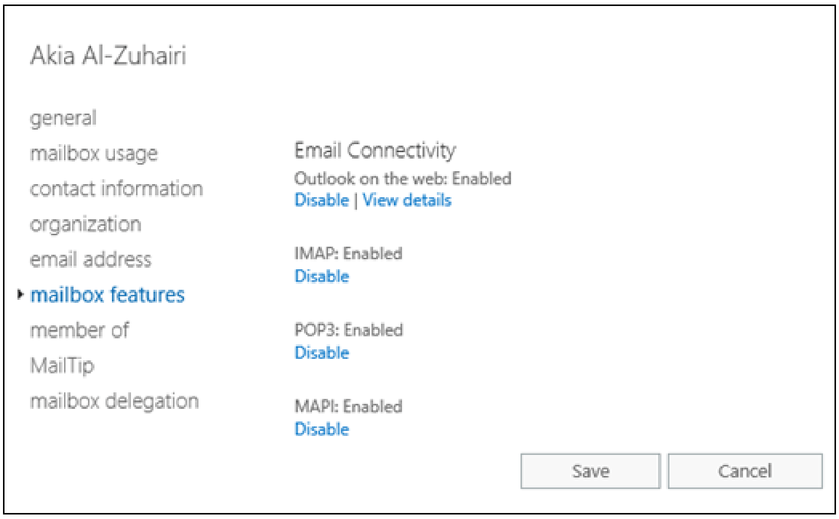

Akia Al-Zuhairi

Email Connectivity

Outlook on the web: Enabled

IMAP: Enabled

POP3: Enabled

MAPI: Enabled

When you're finished, click **Save**.

Use the Exchange Management Shell to enable or disable Outlook on the web access to a mailbox

To enable or disable Outlook on the web access to a single mailbox, use this syntax:

```
Set-CasMailbox -Identity <MailboxIdentity> -OWAEnabled <$true | $false>
```

This example disables Outlook on the web access to the mailbox named Yan Li.

```
Set-CasMailbox -Identity "Yan Li" -OWAEnabled $false
```

This example enables Outlook on the web access to the mailbox named Elly Nkya.

```
Set-CasMailbox -Identity "Elly Nkya" -OWAEnabled $true
```

For detailed syntax and parameter information, see Set-CASMailbox.

Enable or disable Outlook on the web access to multiple mailboxes

Use the EAC to enable or disable Outlook on the web access to multiple mailboxes

1. In the EAC, go to **Recipients** > **Mailboxes**.

2. In the list of mailboxes, find the mailboxes that you want to modify. You can:

   - Scroll through the list of mailboxes.

   - Click **Search** 🔍 and enter part of the user's name, email address, or alias.

   - Click **More options** ••• > **Advanced search** to find the mailbox.

3. In the list of mailboxes, select multiple mailboxes of the same type (for example, **User**) from the list. For example:

- Select a mailbox, hold down the Shift key, and select another mailbox that's farther down in the list.
- Hold down the CTRL key as you select each mailbox.

After you select multiple mailboxes of the same type, the title of the details pane changes to **Bulk Edit**.

4. In the details pane, scroll down to **Outlook on the web**, click **Enable** or **Disable**, and then click **OK** in the warning message that appears.

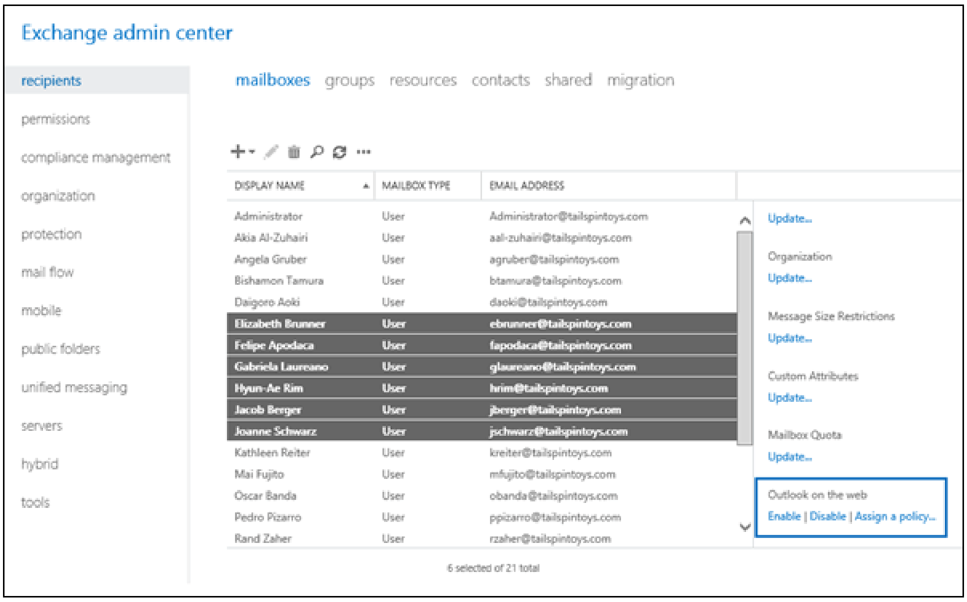

DISPLAY NAME	MAILBOX TYPE	EMAIL ADDRESS
Administrator	User	Administrator@tailspintoys.com
Akia Al-Zuhairi	User	aal-zuhairi@tailspintoys.com
Angela Gruber	User	agruber@tailspintoys.com
Bishamon Tamura	User	btamura@tailspintoys.com
Daigoro Aoki	User	daoki@tailspintoys.com
Elizabeth Brunner	User	ebrunner@tailspintoys.com
Felipe Apodaca	User	fapodaca@tailspintoys.com
Gabriela Laureano	User	glaureano@tailspintoys.com
Hyun-Ae Rim	User	hrim@tailspintoys.com
Jacob Berger	User	jberger@tailspintoys.com
Joanne Schwarz	User	jschwarz@tailspintoys.com
Kathleen Reiter	User	kreiter@tailspintoys.com
Mai Fujito	User	mfujito@tailspintoys.com
Oscar Banda	User	obanda@tailspintoys.com
Pedro Pizarro	User	ppizarro@tailspintoys.com
Rand Zaher	User	rzaher@tailspintoys.com

6 selected of 21 total

Use the Exchange Management Shell to enable or disable Outlook on the web access to multiple mailboxes

You can use the **Get-Mailbox**, **Get-User** or **Get-Content** cmdlets to identify the mailboxes that you want to modify. For example:

- Use the *OrganizationalUnit* parameter to filter the mailboxes by organizational unit (OU).
- Use the *Filter* parameter to create OPATH filters that identify the mailboxes. For more information, see Filterable Properties for the -Filter Parameter.
- Use a text file to specify the mailboxes. The text file contains one mailbox (email address, name, or other unique identifier) on each line like this:

```
ebrunner@tailspintoys.com
fapodaca@tailspintoys.com
glaureano@tailspintoys.com
hrim@tailspintoys.com
```

This example disables Outlook on the web access to all user mailboxes in the North America\Finance OU.

```
$NAFinance = Get-Mailbox -OrganizationalUnit "OU=Marketing,OU=North America,DC=contoso,DC=com" -Filter
"RecipientTypeDetails -eq 'UserMailbox'" -ResultSize Unlimited; $NAFinance | foreach  {Set-CasMailbox
$_.Identity -OWAEnabled $false}
```

This example disables Outlook on the web access to all user mailboxes in the Engineering department in Washington state.

```
Get-User -Filter "RecipientType -eq 'UserMailbox' -and Department -like 'Engineering*' -and StateOrProvince -
eq 'WA'" | Set-CasMailbox -OWAEnabled $false
```

This example uses the text file C:\My Documents\Accounts.txt to disable Outlook on the web access to the specified mailboxes.

```
Get-Content "C:\My Documents\Accounts.txt" | foreach {Set-CasMailbox $_ -OWAEnabled $false}
```

For detailed syntax and parameter information, see Get-Mailbox and Get-User.

How do you know this worked?

To verify that you've successfully enabled or disabled Outlook on the web access to a mailbox, do any of these steps:

- In the EAC, go to **Recipients** > **Mailboxes** > select the mailbox > click **Edit** ✎ > **Mailbox features** and verify the **Outlook on the web** value in the **Email Connectivity** section.

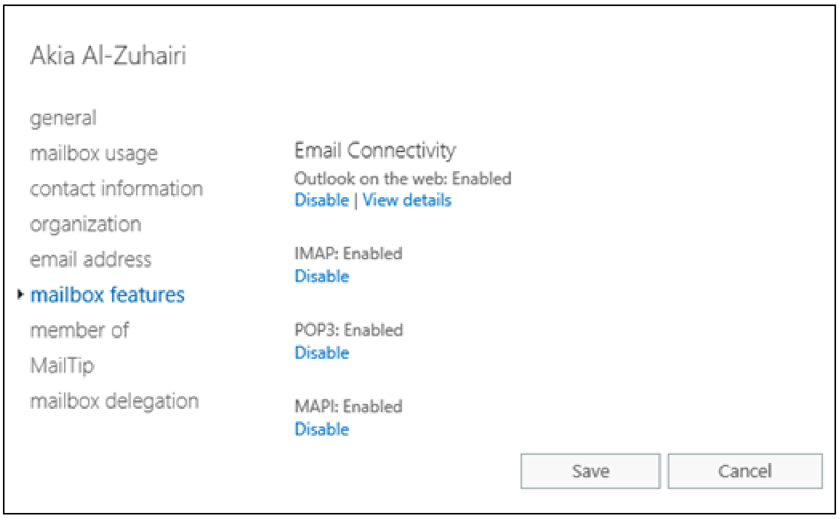

Akia Al-Zuhairi

Email Connectivity

Outlook on the web: Enabled

IMAP: Enabled

POP3: Enabled

MAPI: Enabled

- In the Exchange Management Shell, replace *<MailboxIdentity>* with the identity of the mailbox (for example, name, alias, or email address), and run this command:

```
Get-CasMailbox -Identity "<MailboxIdentity>"
```

- Use the same filter that you used to identify the mailboxes, but use the **Get-CasMailbox** cmdlet instead of **Set-CasMailbox**. For example:

```
Get-User -Filter "RecipientType -eq 'UserMailbox' -and Department -like 'Engineering*' -and
StateOrProvince -eq 'WA'" | Get-CasMailbox
```

- In the Exchange Management Shell, run this command to show all mailboxes where Outlook on the web access is disabled:

```
Get-CasMailbox -ResultSize unlimited -Filter "OWAEnabled -eq `$false"
```

View or configure Outlook on the web virtual directories in Exchange Server

8/3/2020 • 10 minutes to read • Edit Online

You can use the Exchange admin center (EAC) or the Exchange Management Shell to view or modify the properties of an Outlook on the web (formerly known as Outlook Web App) virtual directory. Although the name has changed to Outlook on the web, the name of the virtual directory is still "owa".

What do you need to know before you begin?

- Estimated time to complete each procedure: 10 minutes.
- For more information about the EAC, see .Exchange admin center in Exchange Server. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- Secure Sockets Layer (SSL) is being replaced by Transport Layer Security (TLS) as the protocol that's used to encrypt data sent between computer systems. They're so closely related that the terms "SSL" and "TLS" (without versions) are often used interchangeably. Because of this similarity, references to "SSL" in Exchange topics, the Exchange admin center, and the Exchange Management Shell have often been used to encompass both the SSL and TLS protocols. Typically, "SSL" refers to the actual SSL protocol only when a version is also provided (for example, SSL 3.0). To find out why you should disable the SSL protocol and switch to TLS, check out Protecting you against the SSL 3.0 vulnerability.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Outlook on the web virtual directories" entry in the Clients and mobile devices permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to view or configure Outlook on the web virtual directory properties

1. In the EAC, go to **Servers** > **Virtual directories**.

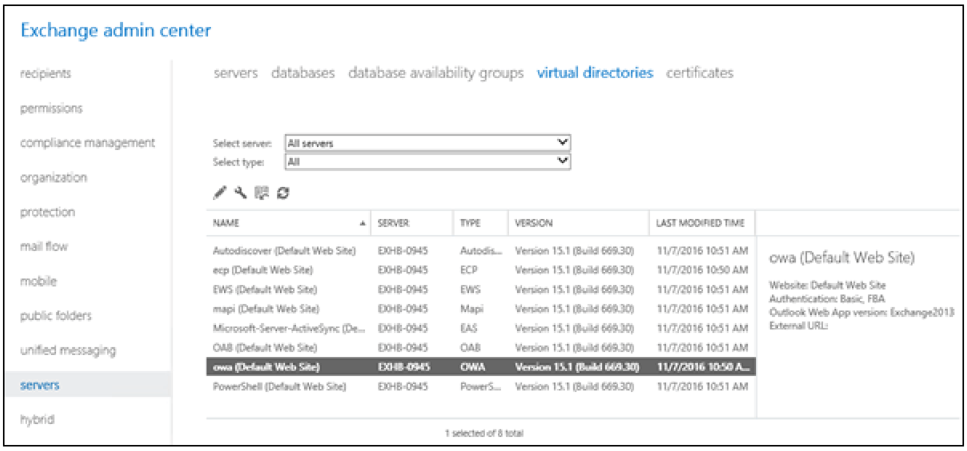

Select server: All servers

Select type: All

NAME	SERVER	TYPE	VERSION	LAST MODIFIED TIME
Autodiscover (Default Web Site)	EXHB-0945	Autodis...	Version 15.1 (Build 669.30)	11/7/2016 10:51 AM
ecp (Default Web Site)	EXHB-0945	ECP	Version 15.1 (Build 669.30)	11/7/2016 10:50 AM
EWS (Default Web Site)	EXHB-0945	EWS	Version 15.1 (Build 669.30)	11/7/2016 10:51 AM
mapi (Default Web Site)	EXHB-0945	Mapi	Version 15.1 (Build 669.30)	11/7/2016 10:51 AM
Microsoft-Server-ActiveSync (De...	EXHB-0945	EAS	Version 15.1 (Build 669.30)	11/7/2016 10:51 AM
OAB (Default Web Site)	EXHB-0945	OAB	Version 15.1 (Build 669.30)	11/7/2016 10:51 AM
owa (Default Web Site)	EXHB-0945	OWA	Version 15.1 (Build 669.30)	11/7/2016 10:50 A...
PowerShell (Default Web Site)	EXHB-0945	PowerS...	Version 15.1 (Build 669.30)	11/7/2016 10:51 AM

owa (Default Web Site)

Website: Default Web Site
Authentication: Basic, FBA
Outlook Web App version: Exchange2013
External URL:

1 selected of 8 total

2. Select the Outlook on the web virtual directory you want to view or configure.

   - You can use the **Select server** drop down list to filter the Exchange servers by name.

   - To only display Outlook on the web virtual directories, select **OWA** in the **Select type** drop down list.

   After you select the virtual directory, you can see the following properties and values in the feature pane:

   - **Website** (read-only): The default web site is named **Default Web Site**.

   - **Authentication**: The default authentication methods are **Basic** and **FBA** (forms-based authentication).

   - **Outlook on the web version**: The default version is `Exchange2013`.

   - **External URL**: The default value is blank (not configured).

3. To see more properties, or to modify the settings that aren't read only, click **Edit** (✎). The following tabs and settings are available:

   - **General** tab:

     - **Internal URL**: The URL that's used to access Outlook on the web from the internal network. This value is configured automatically during Exchange Server setup, and the default value is https:// *<Server FQDN>*/owa (for example, https://mailbox01.contoso.com/owa).

     - **External URL**: The URL that's used to access Outlook on the web from the Internet. The default value is blank.

       For Internet-facing Exchange servers, this is the value that clients use to access Outlook on the web. To configure this setting, see the Use the EAC to configure the external URL for Outlook on the web section in this topic.

       For Exchange servers that don't have an Internet presence, the leave the **External URL** value blank.

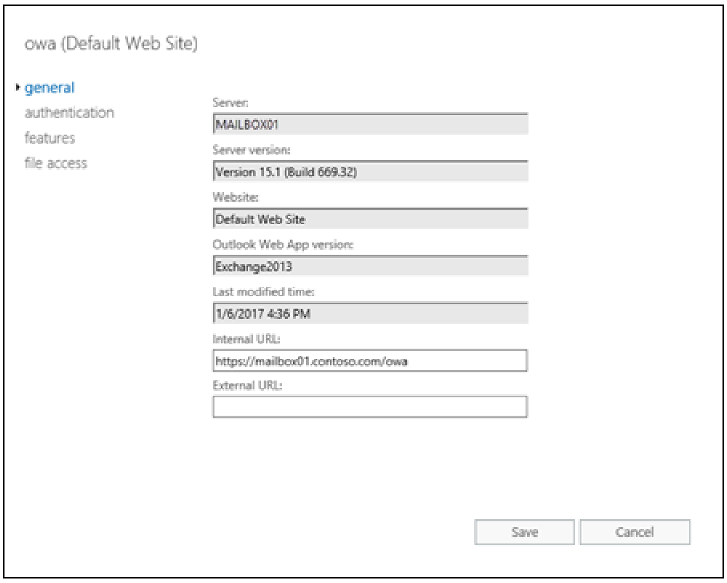

owa (Default Web Site)

Server:
MAILBOX01

Server version:
Version 15.1 (Build 669.32)

Website:
Default Web Site

Outlook Web App version:
Exchange2013

Last modified time:
1/6/2017 4:36 PM

Internal URL:
https://mailbox01.contoso.com/owa

External URL:

- **Authentication** tab:
  - **Use one or more standard authentication methods**: Select this option to use one or more of the following authentication methods:
  - **Integrated Windows authentication**: This method requires that users have a valid Active Directory user account, and the client computer is a member of the same domain as the Exchange server (or a domain that's trusted by the Exchange server's domain). Users aren't prompted for their account names and passwords. Instead, the server negotiates with the Windows security packages that are installed on the client computer. No unencrypted information is transmitted over the network.
  - **Digest authentication for Windows domain servers**: This method requires that users have a valid Active Directory user account. Passwords are transmitted over the network as a hash value for additional security.
  - **Basic authentication (password is sent in clear text)**: This is the default value. When you use basic authentication, you should require TLS encrypted connections between client computers and the Exchange server.
  - **Use forms-based authentication**: Forms-based authentication provides enhanced security and allows you to configure the type of prompt that's used to sign-in. However, forms-based authentication won't provide a secure channel unless TLS is enabled.

    Select one of the following logon formats to use with forms-based authentication. The examples use the account for the user named Valeria Barrios in the contoso.com domain.

    - **Domain\user name** For example, CONTOSO\VBarrios. This is the default value.
    - **User principal name (UPN)** For example, vbarrios@contoso.com. Note that if the UPN doesn't match the email address, users can't access Outlook on the web by using this method.
    - **Username only** For example, VBarrios. This setting requires you to configure the default domain that's used with all user names. Click **Browse** in the **Logon Domain** property to select the default Active Directory domain. If the user isn't a member of the specified domain, they're required to enter the domain and username when they sign in.

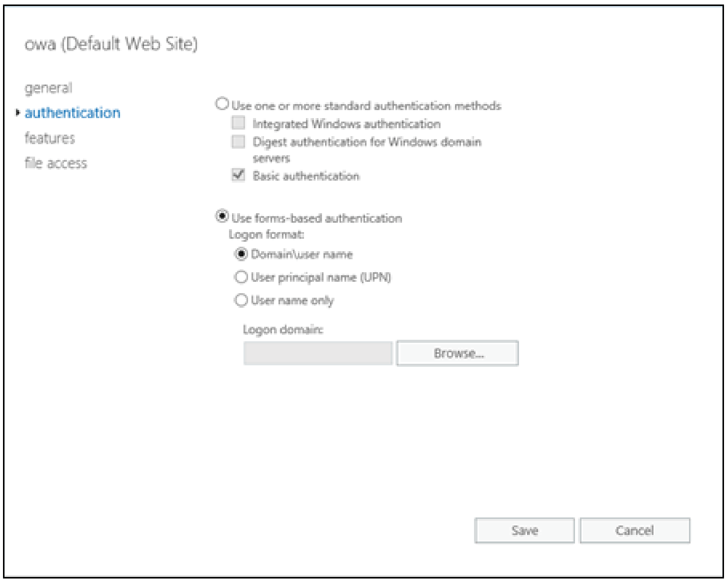

owa (Default Web Site)

○ 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:

- **Features** tab:

  These settings affect all users who connect to the Outlook on the web virtual directory. You can configure custom Outlook on the web settings for specific users or groups of users by using Outlook on the web mailbox policies. For more information, see View or configure Outlook on the web mailbox policy properties.

  - **Communication management**
  - **Instant messaging**
  - **Text messaging**
  - **Unified Messaging**: (In Exchange 2016 only; not available in Exchange 2019)
  - **Exchange ActiveSync**
  - **Contacts**
  - **All address lists***
  - **Information management**
  - **Journaling**
  - **Inbox rules***
  - **Recover deleted items***: Disabling this setting doesn't affect the deleted item retention for mailboxes; it prevents users from viewing or recovering deleted items in Outlook on the web.
  - **Security**
  - **Change password**
  - **Junk email-filtering**: This setting doesn't enable or disable the junk email rule in mailboxes; it controls the *availability* of the junk email settings for users in Outlook on the web. For more information about the junk email rule and junk email filtering in mailboxes, see Configure Exchange antispam settings on mailboxes.
  - **User experience**
  - **Themes**
  - **Premium client**: If you uncheck this setting, The standard version of Outlook on the web (formerly known as the premium version of Outlook Web App) is disabled, and all clients are

  forced to use the light version of Outlook on the web.

  - **Email signature***
  - **Time management***
  - **Calendar***
  - **Tasks***
  - **Reminders and notifications***

  * These settings are available after you click **More options**.

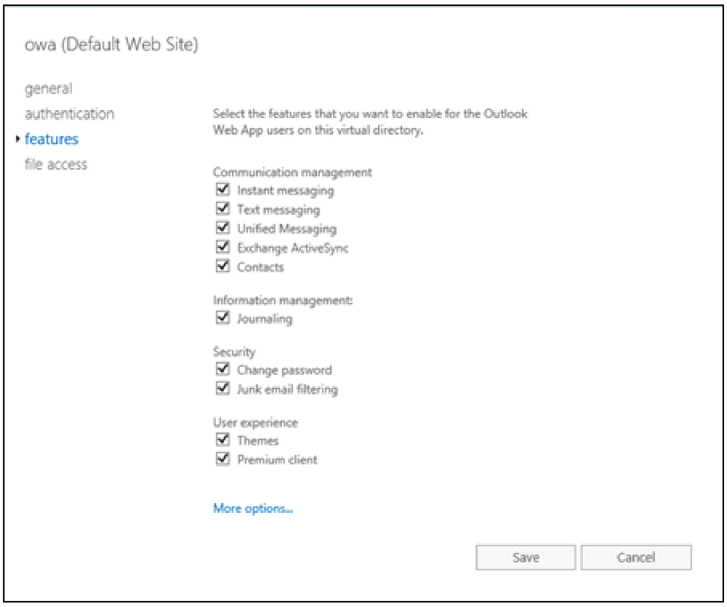

owa (Default Web Site)

general
authentication
features
file access

Select the features that you want to enable for the Outlook Web App users on this virtual directory.

Communication management
☑ Instant messaging
☑ Text messaging
☑ Unified Messaging
☑ Exchange ActiveSync
☑ Contacts

Information management:
☑ Journaling

Security
☑ Change password
☑ Junk email filtering

User experience
☑ Themes
☑ Premium client

- **File access** tab:

  The direct file access settings on this page affect traditional file attachments that you click on to open or save, or MIME files (typically, image files) that are embedded directly in the message. Disabling direct file access doesn't affect file access in other email clients (for example, in Outlook), or by using other access methods in Outlook on the web (for example, web document access that's provided by Office Online Server, or links to files in the cloud).

  Note that users can select public or private computer access in Outlook on the web only when the virtual directory is configured for forms-based authentication. All other authentication methods automatically use private computer access.

  - **Direct file access** for public or shared computers.
  - **Direct file access** for private computers.

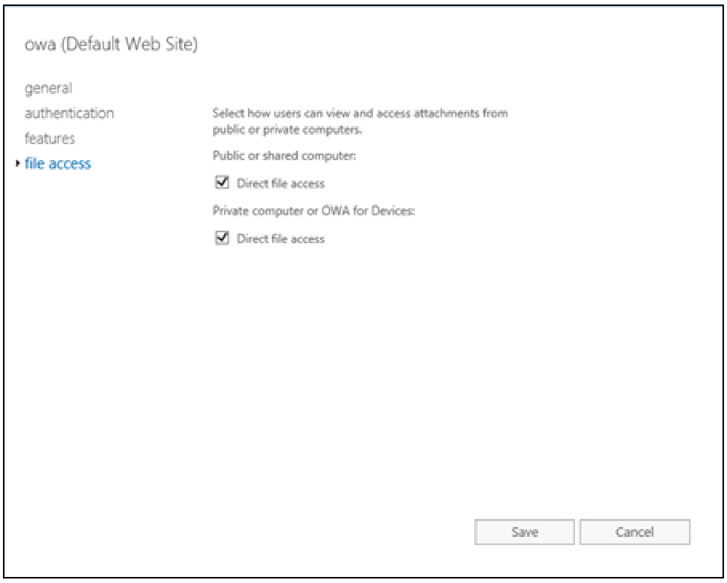

Select how users can view and access attachments from public or private computers.

Public or shared computer:

☑ Direct file access

Private computer or OWA for Devices:

☑ Direct file access

4. If you changed any of the virtual directory settings, click **Save**. If you're just browsing, click **Cancel**.

Use the EAC to configure the external URL for Outlook on the web

1. In the EAC, go to **Servers** > **Virtual directories**, select the Outlook on the web virtual directory you want to view or configure, and then click **Configure** (🔧).

   - You can use the **Select server** drop down list to filter the Exchange servers by name.

   - To only display Outlook on the web virtual directories, select **OWA** in the **Select type** drop down list.

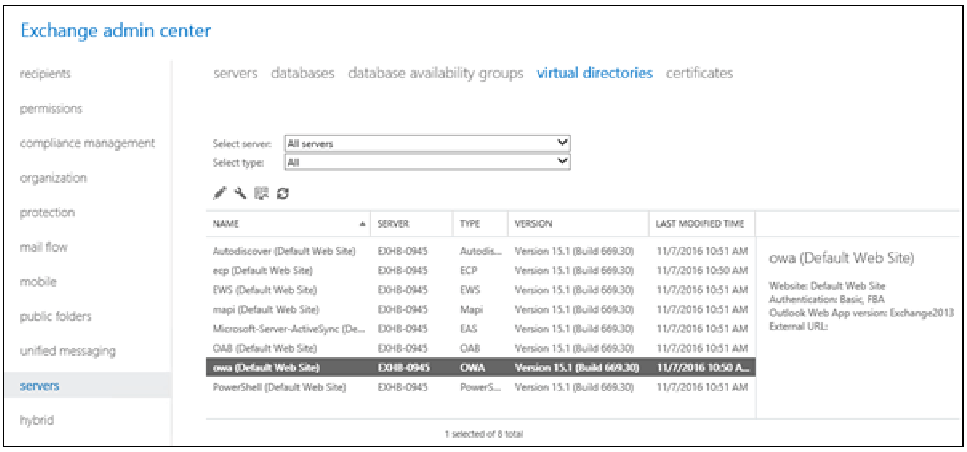

Select server: All servers

Select type: All

NAME	SERVER	TYPE	VERSION	LAST MODIFIED TIME
Autodiscover (Default Web Site)	EXHB-0945	Autodis...	Version 15.1 (Build 669.30)	11/7/2016 10:51 AM
ecp (Default Web Site)	EXHB-0945	ECP	Version 15.1 (Build 669.30)	11/7/2016 10:50 AM
EWS (Default Web Site)	EXHB-0945	EWS	Version 15.1 (Build 669.30)	11/7/2016 10:51 AM
mapi (Default Web Site)	EXHB-0945	Mapi	Version 15.1 (Build 669.30)	11/7/2016 10:51 AM
Microsoft-Server-ActiveSync (De...	EXHB-0945	EAS	Version 15.1 (Build 669.30)	11/7/2016 10:51 AM
OAB (Default Web Site)	EXHB-0945	OAB	Version 15.1 (Build 669.30)	11/7/2016 10:51 AM
owa (Default Web Site)	EXHB-0945	OWA	Version 15.1 (Build 669.30)	11/7/2016 10:50 A...
PowerShell (Default Web Site)	EXHB-0945	PowerS...	Version 15.1 (Build 669.30)	11/7/2016 10:51 AM

owa (Default Web Site)

Website: Default Web Site
Authentication: Basic, FBA
Outlook Web App version: Exchange2013
External URL:

1 selected of 8 total

2. In the **Configure external access domain** page that opens, configure the following settings:

   - **Select the servers to use with the external URL**: Click **Add** (+) and select one or more Exchange servers that external clients will use to connect to Outlook on the web (don't select internal only servers).

   - **Enter the domain name you will use with your external servers**: Enter the FQDN that external clients will use to connect to Outlook on the web (for example, mail.contoso.com). Note that this value needs to be configured and resolvable in your organization's public DNS.

   When you're finished, click **Save**.

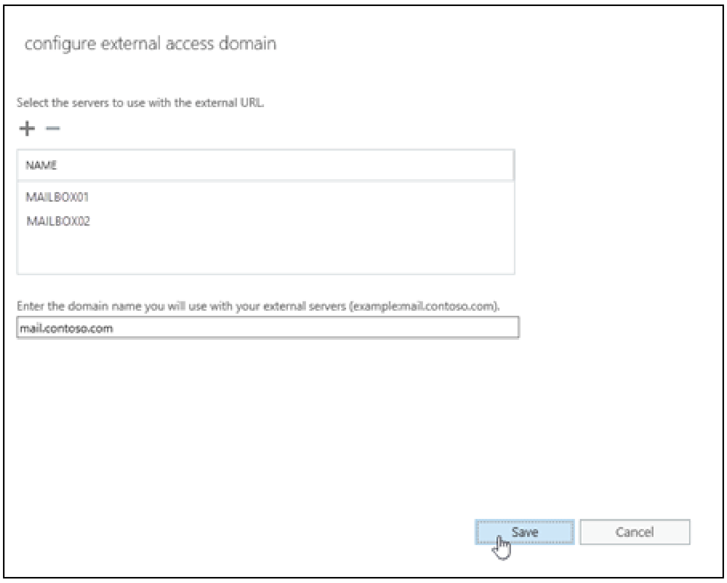

configure external access domain

Select the servers to use with the external URL.

NAME
MAILBOX01
MAILBOX02

Enter the domain name you will use with your external servers (example:mail.contoso.com).

mail.contoso.com

Reset an Outlook on the web virtual directory

If an Outlook on the web virtual directory isn't working the way you expect, you can reset it. The virtual directory is deleted and recreated with the default settings. Although any customized settings are lost, you're forced to select a location for a text document to backup the current settings.

1. In the EAC, go to **Servers** > **Virtual directories**, select the Outlook on the web virtual directory you want to view or configure, and then click **Reset** (▦).

   - You can use the **Select server** drop down list to filter the Exchange servers by name.
   - To only display Outlook on the web virtual directories, select **OWA** in the **Select type** drop down list.

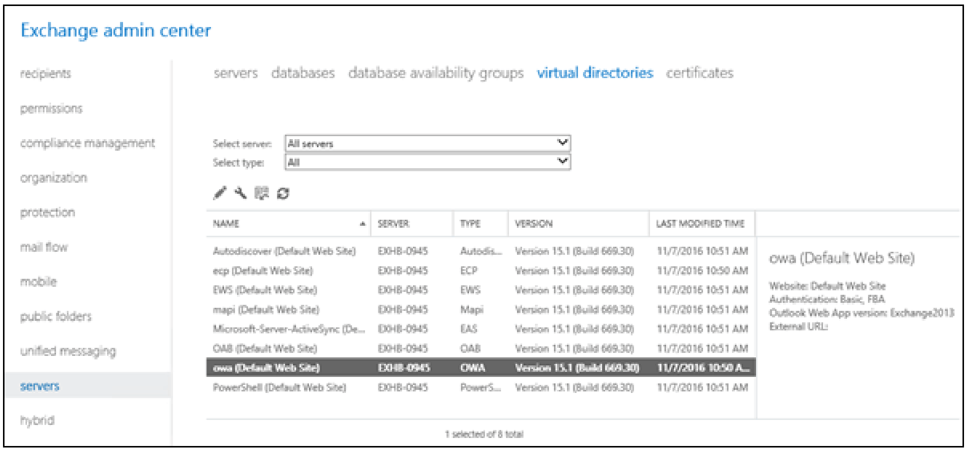

Select server: All servers

Select type: All

NAME	SERVER	TYPE	VERSION	LAST MODIFIED TIME
Autodiscover (Default Web Site)	EXHB-0945	Autodis...	Version 15.1 (Build 669.30)	11/7/2016 10:51 AM
ecp (Default Web Site)	EXHB-0945	ECP	Version 15.1 (Build 669.30)	11/7/2016 10:50 AM
EWS (Default Web Site)	EXHB-0945	EWS	Version 15.1 (Build 669.30)	11/7/2016 10:51 AM
mapi (Default Web Site)	EXHB-0945	Mapi	Version 15.1 (Build 669.30)	11/7/2016 10:51 AM
Microsoft-Server-ActiveSync (De...	EXHB-0945	EAS	Version 15.1 (Build 669.30)	11/7/2016 10:51 AM
OAB (Default Web Site)	EXHB-0945	OAB	Version 15.1 (Build 669.30)	11/7/2016 10:51 AM
owa (Default Web Site)	EXHB-0945	OWA	Version 15.1 (Build 669.30)	11/7/2016 10:50 A...
PowerShell (Default Web Site)	EXHB-0945	PowerS...	Version 15.1 (Build 669.30)	11/7/2016 10:51 AM

owa (Default Web Site)

Website: Default Web Site
Authentication: Basic, FBA
Outlook Web App version: Exchange2013
External URL:

1 selected of 8 total

2. In the **Warning** page that opens, specify the UNC path of the file to save the current virtual directory settings (for example, \ *<Server>*\ *<Share>*\owavdir.txt or \ *<LocalServerName>*\c$\owavdir.txt).

   When you're finished, click **Reset**.

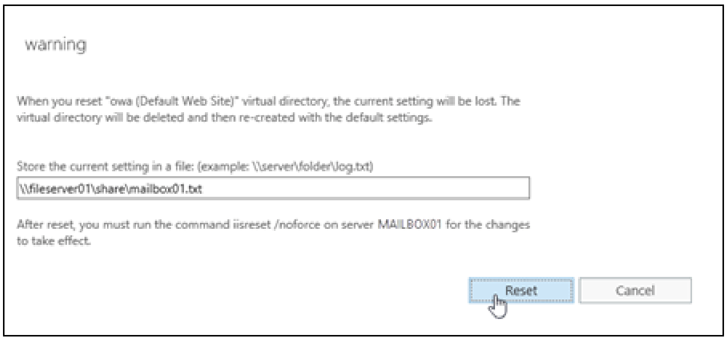

warning

When you reset "owa (Default Web Site)" virtual directory, the current setting will be lost. The virtual directory will be deleted and then re-created with the default settings.

Store the current setting in a file: (example: \\server\folder\log.txt)

\\fileserver01\share\mailbox01.txt

After reset, you must run the command iisreset /noforce on server MAILBOX01 for the changes to take effect.

3. Restart IIS by using either of the following methods:

- IIS Manager:

   a. Open IIS Manager on the Exchange server. An easy way to do this in Windows Server 2012 or later is to press Windows key + Q, type inetmgr, and select **Internet Information Services (IIS) Manager** in the results.

   b. In IIS Manager, select the server.

   c. In the **Actions** pane, click **Restart**.

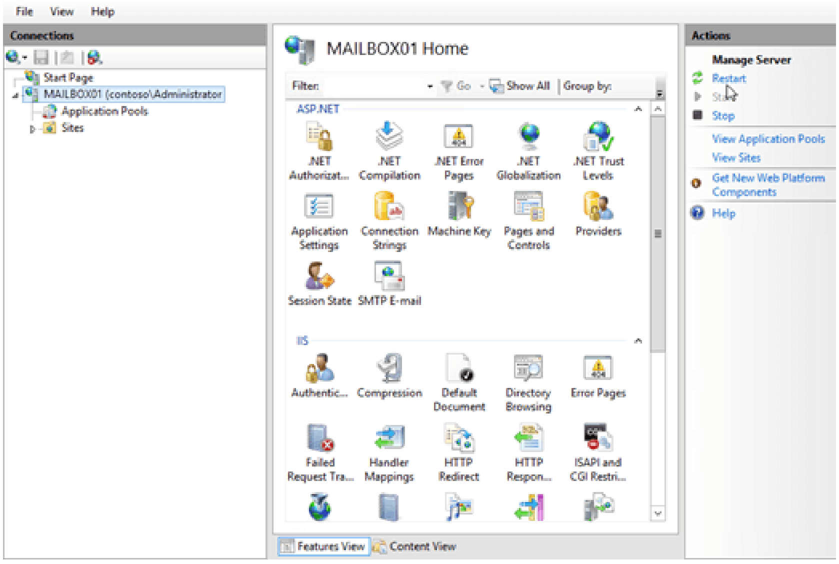

- Command prompt:

Open an elevated command prompt on the Exchange server (a Command Prompt window you open by selecting **Run as administrator**) and run the following commands:

```
net stop w3svc /y
```

```
net start w3svc
```

Use the Exchange Management Shell to view Outlook on the web virtual directory properties

To use the Exchange Management Shell to view the properties of Outlook on the web virtual directories, use the following syntax:

```
Get-OWAVirtualDirectory [-Identity "<ExchangeServer>\owa <Website>"]
```

This example returns a summary list of all Outlook on the web virtual directories on all Exchange servers in the organization.

```
Get-OWAVirtualDirectory
```

This example returns detailed information for the Outlook on the web virtual directory in the default website on the Exchange server named Mailbox01.

```
Get-OWAVirtualDirectory -Identity "Mailbox01\owa (Default Web Site)" | Format-List
```

This example returns the authentication methods and settings for the same virtual directory:

```
Get-OWAVirtualDirectory -Identity "Mailbox01\owa (Default Web Site)" | Format-List *Authentication*
```

Note: Not every setting is applicable to Exchange 2016 or Exchange 2019 (for example, **SpellCheckerEnabled**).

For detailed syntax and parameter information, see Get-OWAVirtualDirectory.

Use the Exchange Management Shell to configure Outlook on the web virtual directory settings

There are many more configuration settings available for Outlook on the web virtual directories in the Exchange Management Shell (the **Set-OwaVirtualDirectory** cmdlet) than in the EAC. Hare are some of the Outlook on the web virtual directory settings that are only available in the Exchange Management Shell:

PARAMETER	FUNCTION
AllowedFileTypes *BlockedFileTypes* *ForceSaveFileTypes* *AllowedMimeTypes* *BlockedMimeTypes* *ForceSaveMimeTypes* *ActionForUnknownFileAndMIMETypes*	Defines the file types for direct file access (traditional file attachments an embedded MIME files) in Outlook on the web (not in other email clients).
DefaultTheme	Specifies the default theme that's used in Outlook on the web.
LogonAndErrorLanguage *OutboundCharset* *UseGB18030* *UseISO885915*	Configures the various language settings for Outlook on the web.
DisplayPhotosEnabled *SetPhotoEnabled* *SetPhotoURL*	Configures the user photo settings in Outlook on the web.

Note: Not all of the available parameters apply to Exchange 2016 or Exchange 2019 (for example, *SpellCheckerEnabled*).

To use the Exchange Management Shell to configure the properties of Outlook on the web virtual directories, use the following syntax:

```
Set-OWAVirtualDirectory -Identity "<ExchangeServer>\owa <Website>" <Settings>
```

This example enables configures direct file access in Outlook on the web to block file types that aren't specifically defined in the Allow list (the default action is allow).

```
Set-OwaVirtualDirectory -Identity "Contoso\owa (Default Web Site)" -ActionForUnknownFileAndMIMETypes Block
```

For detailed syntax and parameter information, see Set-OwaVirtualDirectory.

Configure http to https redirection for Outlook on the web in Exchange Server

8/3/2020 • 7 minutes to read • Edit Online

By default in Exchange Server, the URL https://*<ServerName>* redirects users to https://*<ServerName>*/owa. But, if anyone tries to access Outlook on the web (formerly known as Outlook Web App) by using http://*<ServerName>* or http://*<ServerName>*/owa, they'll get an error.

You can configure http redirection for Outlook on the web so that requests for http://*<ServerName>* or http://*<ServerName>*/owa are automatically redirected to https://*<ServerName>*/owa. This requires the following configuration steps in Internet Information Services (IIS):

1. Remove the **Require SSL** setting from the default website.

2. Restore the **Require SSL** setting on other virtual directories in the default website that had it enabled by default (except for /owa).

3. Configure the default website to redirect http requests to the /owa virtual directory.

4. Remove http redirection from all virtual directories in the default website (including /owa).

5. Reset IIS for the changes to take effect.

For the default SSL and http redirect settings on all virtual directories in the default website, see the Default Require SSL and HTTP Redirect settings in the default website on an Exchange server section at the end of this topic.

What do you need to know before you begin?

- Estimated time to complete this procedure: 15 minutes.

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "IIS Manager" entry in the Outlook on the web permissions section of the Clients and mobile devices permissions topic.

- The procedures in this topic might cause a web.config file to be created in the folder `%ExchangeInstallPath%ClientAccess\OAB`. If you later remove http redirection for Outlook on the web, Outlook might freeze when users click **Send and Receive**. To prevent Outlook from freezing after you remove http redirection, delete the web.config file in `%ExchangeInstallPath%ClientAccess\OAB`.

- Secure Sockets Layer (SSL) is being replaced by Transport Layer Security (TLS) as the protocol that's used to encrypt data sent between computer systems. They're so closely related that the terms "SSL" and "TLS" (without versions) are often used interchangeably. Because of this similarity, references to "SSL" in Exchange topics, the Exchange admin center, and the Exchange Management Shell have often been used to encompass both the SSL and TLS protocols. Typically, "SSL" refers to the actual SSL protocol only when a version is also provided (for example, SSL 3.0). To find out why you should disable the SSL protocol and switch to TLS, check out Protecting you against the SSL 3.0 vulnerability.

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Step 1: Use IIS Manager to remove the Require SSL setting from the default website

1. Open IIS Manager on the Exchange server. An easy way to do this in Windows Server 2012 or later is to press Windows key + Q, type inetmgr, and select **Internet Information Services (IIS) Manager** in the results.

2. Expand the server, and expand **Sites**.

3. Select **Default Web Site**. and verify **Features View** is selected at the bottom of the page.

4. In the **IIS** section, double-click **SSL Settings**.

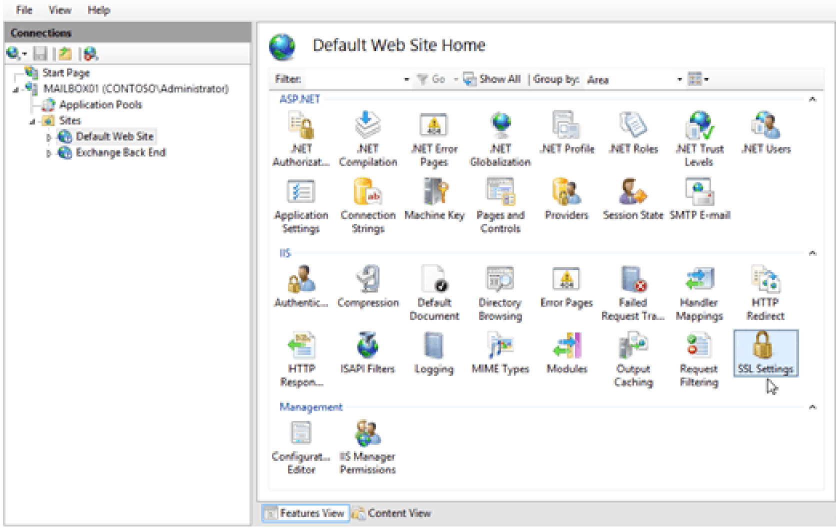

5. On the **SSL Settings** page, clear the **Require SSL** check box, and in the **Actions** pane, click **Apply**.

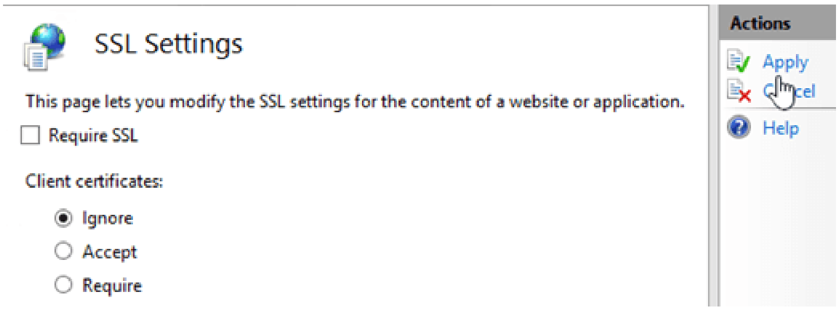

SSL Settings

This page lets you modify the SSL settings for the content of a website or application.

□ Require SSL

Client certificates:

● Ignore

○ Accept

○ Require

Note: To perform this procedure on the command line, open an elevated command prompt on the Exchange server (a Command Prompt window you open by selecting **Run as administrator**) and run the following command:

```
%windir%\system32\inetsrv\appcmd.exe set config "Default Web Site" -section:access -sslFlags:None -commit:APPHOST
```

Step 2: Use IIS Manager to restore the Require SSL setting on other virtual directories in the default website

When you change the **Require SSL** setting on a website in IIS, the setting is automatically inherited by all virtual directories in the website. Because we're only interested in configuring Outlook on the web, you need to restore the **Require SSL** setting for other virtual directories that had it enabled by default.

Based on the information in the Default Require SSL and HTTP Redirect settings in the default website on an Exchange server section, use the following procedure to restore the setting on the other virtual directories where **Require SSL** was enabled by default:

1. In IIS Manager, expand the server, expand **Sites**, and expand **Default Web Site**.

2. Select the virtual directory, and verify **Features View** is selected at the bottom of the page.

3. In the **IIS** section, double-click **SSL Settings**.

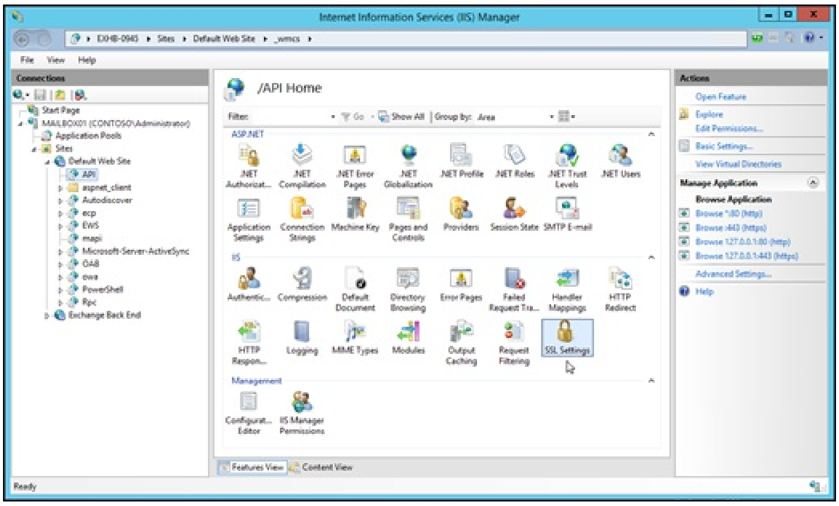

4. On the **SSL Settings** page, select the **Require SSL** check box, and in the **Actions** pane, click **Apply**.

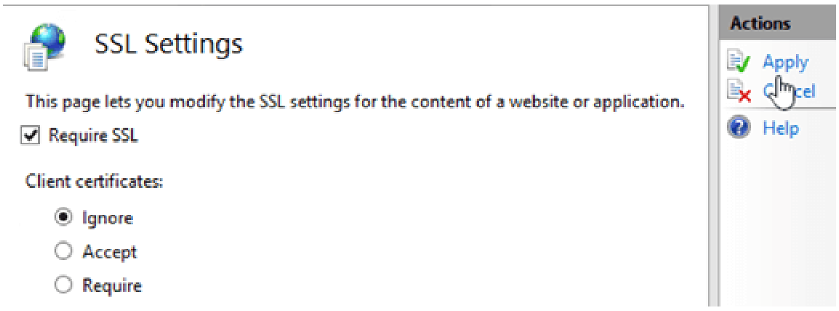

This page lets you modify the SSL settings for the content of a website or application.

☑ Require SSL

Client certificates:

● Ignore
○ Accept
○ Require

5. Repeat the previous steps on each virtual directory in the default website that had **Require SSL** enabled by default (except for /owa). The only virtual directories that don't have **Require SSL** enabled by default are /PowerShell and /Rpc.

Note: To perform these procedures on the command line, replace *<VirtualDirectory>* with the name of the virtual directory, and run the following command in an elevated command prompt:

```
%windir%\system32\inetsrv\appcmd.exe set config "Default Web Site/<VirtualDirectory>" -section:Access -
sslFlags:Ssl,Ssl128 -commit:APPHOST
```

Step 3: Use IIS Manager to configure the default website to redirect to the /owa virtual directory.

1. In IIS Manager, expand the server, and expand **Sites**.

2. Select **Default Web Site**. and verify **Features View** is selected at the bottom of the page.

3. In the **IIS** section, double-click **HTTP Redirect**.

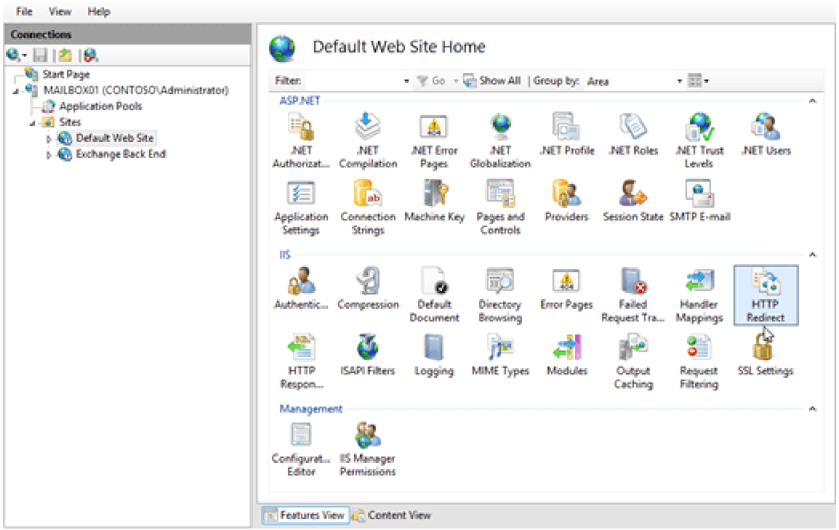

4. On the **HTTP Redirect** page, configure the following settings:

5. Select the **Redirect requests to this destination** check box, and enter the value https:// *<OWAUrl>*/owa (For example, https://webmail.contoso.com/owa).

6. In the **Redirect Behavior** section, select the **Only redirect requests to content in this directory (not subdirectories)** check box.

7. In the **Status code** list, verify **Found (302)** is selected.

   When you're finished, click **Apply** in the **Actions** pane.

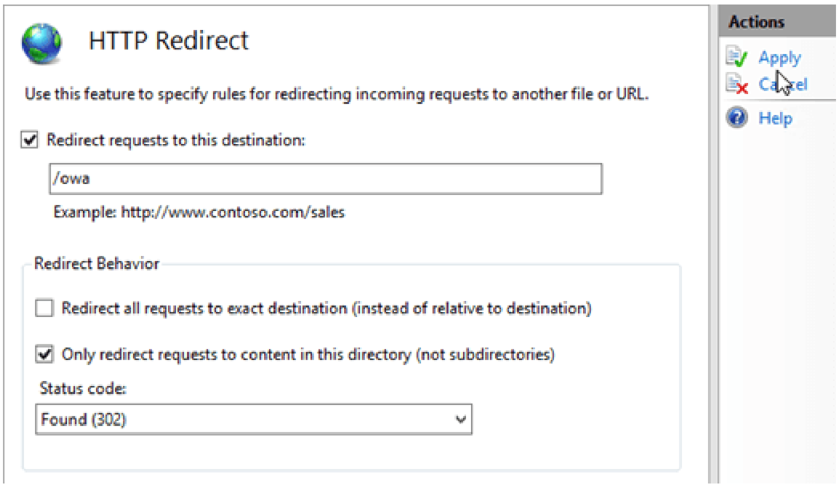
HTTP Redirect

Use this feature to specify rules for redirecting incoming requests to another file or URL.

☑ Redirect requests to this destination:

/owa

Example: http://www.contoso.com/sales

Redirect Behavior

□ Redirect all requests to exact destination (instead of relative to destination)

☑ Only redirect requests to content in this directory (not subdirectories)

Status code:

Found (302)

Note: To perform this procedure on the command line, replace *<OWAUrl>* with the URL of the OWA virtual directory, open an elevated command prompt and run the following command:

```
%windir%\system32\inetsrv\appcmd.exe set config "Default Web Site" -section:httpredirect -enabled:true -
destination:"https://<OWAUrl>/owa" -childOnly:true
```

Step 4: Use IIS Manager to remove http redirection from all virtual directories in the default website

When you enable redirection on a website in IIS, the setting is automatically inherited by all virtual directories in the website. Because we're only interested in configuring redirection for the default website, you need to remove the redirect setting from all virtual directories. By default, no directories or virtual directories in the default website are enabled for redirection. For more information, see the Default Require SSL and HTTP Redirect settings in the default website on an Exchange server section.

Use the following procedure to remove the redirect setting from all virtual directories in the default website (including /owa):

1. In IIS Manager, expand the server, expand **Sites**, and expand **Default Web Site**.

2. Select the virtual directory, and verify **Features View** is selected at the bottom of the page.

3. In the **IIS** section, double-click **HTTP Redirect**.

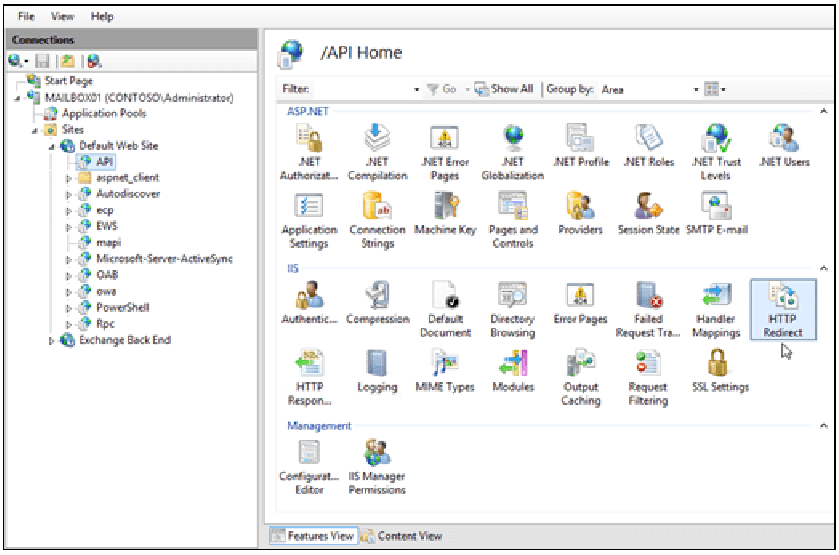

4. On the **HTTP Redirect** page, change the following settings:

5. Clear the **Only redirect requests to content in this directory (not subdirectories)** check box.

6. Clear the **Redirect requests to this destination** check box.

7. In the **Actions** pane, click **Apply**.

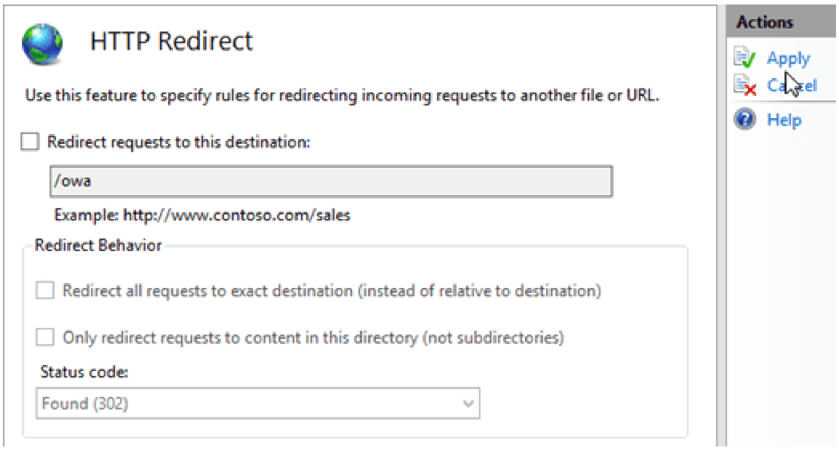

HTTP Redirect

Use this feature to specify rules for redirecting incoming requests to another file or URL.

□ Redirect requests to this destination:

/owa

Example: http://www.contoso.com/sales

Redirect Behavior

□ Redirect all requests to exact destination (instead of relative to destination)

□ Only redirect requests to content in this directory (not subdirectories)

Status code:

Found (302)

8. Repeat the previous steps on each virtual directory in the default website.

Note: To perform these procedures on the command line, replace *<VirtualDirectory>* with the name of the virtual directory, and run the following command in an elevated command prompt:

```
%windir%\system32\inetsrv\appcmd.exe set config "Default Web Site/<VirtualDirectory>" -section:httpredirect -
enabled:false -destination:"" -childOnly:false
```

Step 5: Use IIS Manager to restart IIS

1. In IIS Manager, select the server.

2. In the **Actions** pane, click **Restart**.

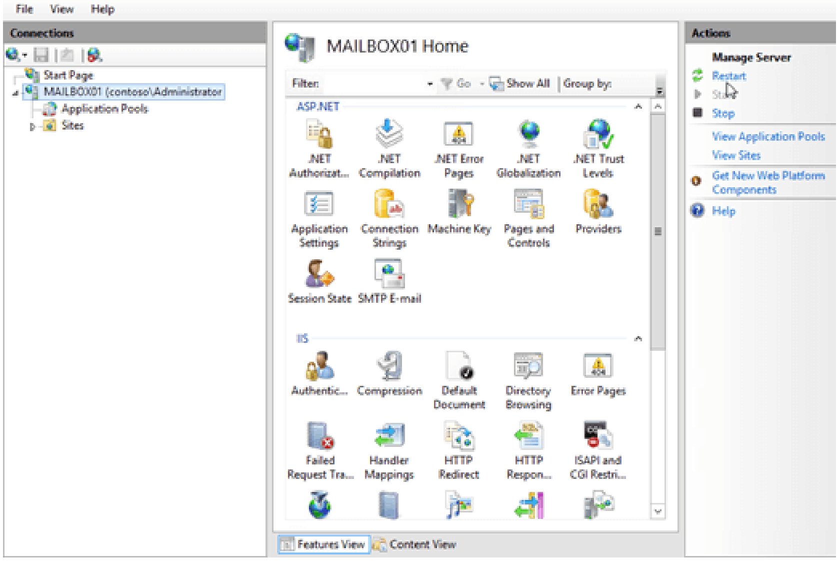

Note: To perform this procedure on the command line, open an elevated command prompt on the Exchange server and run the following commands:

```
net stop w3svc /y
```

```
net start w3svc
```

How do you know this worked?

To verify that you have successfully configured http to https redirection for Outlook on the web, perform the following steps:

1. On a client computer, open a web browser and enter the URL http://*<ServerName>*. On the local server, you can use the value http://127.0.0.1 or http://localhost.

2. Verify that you're redirected to Outlook on the web in https, and verify that you can log in successfully.

3. Open the URL http://*<ServerName>*/owa (or http://127.0.0.1/owa or http://localhost/owa).

4. Verify that you're redirected to Outlook on the web in https, and verify that you can log in successfully.

Default Require SSL and HTTP Redirect settings in the default website on an Exchange server

The default **Require SSL** and **HTTP Redirect** settings for the default website and all virtual directories in the default website on an Exchange server are described in the following table.

WEBSITE	VIRTUAL DIRECTORY	REQUIRE SSL	HTTP REDIRECT
Default Web Site	n/a	yes	none
Default Web Site	API	yes	none
Default Web Site	aspnet_client (directory)	yes	none

WEBSITE	VIRTUAL DIRECTORY	REQUIRE SSL	HTTP REDIRECT
Default Web Site	Autodiscover	yes	none
Default Web Site	ecp	yes	none
Default Web Site	EWS	yes	none
Default Web Site	mapi	yes	none
Default Web Site	Microsoft-Server-ActiveSync	yes	none
Default Web Site	OAB	yes	none
Default Web Site	owa	yes Subdirectories: • auth: yes • Calendar: no • Integrated: yes • oma: yes	none
Default Web Site	PowerShell	no	none
Default Web Site	Rpc	no	none

View or configure Outlook on the web mailbox policy properties

8/3/2020 • 3 minutes to read • Edit Online

You can configure mailbox policies in Exchange Server for Outlook on the web through the Exchange admin center (EAC) or Exchange Management Shell. After you create an Outlook on the web mailbox policy, you can then configure a variety of options to control the features available to users in Outlook on the web. For example, you can enable or disable Inbox rules or create a list of allowed file types for attachments.

What do you need to know before you begin?

- Estimated time to complete each procedure: 3 minutes.
- To open the EAC, see Exchange admin center in Exchange Server. To open the Exchange Management Shell, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Outlook on the web mailbox policies" entry in the Clients and mobile devices permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to view or configure Outlook on the web mailbox policies

1. In the EAC, click **Permissions** > **Outlook Web App policies**.
2. In the result pane, click to select the mailbox policy you want to view or configure.
3. Click **Edit**.
4. On the **General** tab, you can view and edit the name of the policy.
5. On the **Features** tab, use the check boxes to enable or disable features. By default, the most common features are displayed. To see all features that can be enabled or disabled, click **More options**.

   Notes:

   - Features settings for Outlook on the web mailbox policies override Outlook on the web virtual directory settings. You can change segmentation settings for individual users by using the **Set-CASMailbox** cmdlet in the Exchange Management Shell.
   - The option to enable or disable the standard version of Outlook on the web by using the **Premium client** check box has been deprecated and will be removed from the settings. The standard version of Outlook on the web is always enabled.
6. On the **File Access** tab, use the check boxes to configure the file access and viewing options for users. File

access lets a user open or view the contents of files attached to an email message.

File access can be controlled based on whether a user has signed in on a public or private computer. The option for users to select private computer access or public computer access is available only when you're using forms-based authentication. All other forms of authentication default to private computer access.

- **Direct file access**: Select this check box if you want to enable direct file access. Direct file access lets users open files attached to email messages.

- **WebReady Document Viewing**: Select this check box if you want to enable supported documents to be converted to HTML and displayed in a web browser.

- **Force WebReady Document Viewing when a converter is available**: Select this check box if you want to force documents to be converted to HTML and displayed in a web browser before users can open them in the viewing application. Documents can be opened in the viewing application only if direct file access has been enabled.

7. On the **Offline access** tab, use the option buttons to configure offline access availability.

8. Click **Save** to update the policy.

Use the Exchange Management Shell to view Outlook on the web mailbox policies

This example retrieves the properties of the Outlook on the web mailbox policy `Executives` in the organization `Fabrikam`.

```
Get-OwaMailboxPolicy -Identity Fabrikam\Executives
```

For more information about syntax and parameters, see Get-OwaMailboxPolicy.

Use the Exchange Management Shell to configure Outlook on the web mailbox policies

This example enables calendar access in the default mailbox policy.

```
Set-OwaMailboxPolicy -Identity Default -CalendarEnabled $true
```

For more information about syntax and parameters, see Set-OwaMailboxPolicy.

How do you know this worked?

To verify that you've successfully edited an Outlook on the web mailbox policy:

1. In the EAC, click **Permissions** > **Outlook Web App Policies**, and then choose a specific Outlook on the web mailbox policy.

2. Click **Edit** to view the properties of the mailbox policy.

3. Click **Save** or **Cancel** to close the properties page.

See also

Outlook Web App mailbox policy procedures in Exchange 2013

Create a theme for Outlook on the web in Exchange Server

8/3/2020 • 11 minutes to read • Edit Online

A *theme* defines the colors, fonts, and images that are displayed to users in Outlook on the web (formerly known as Outlook Web App) in Exchange Server. Each theme is a collection of files that are stored on the Exchange server. The built-in themes are described in the Default Outlook on the web themes in Exchange Server section at the end of this topic.

The basic steps to create a new theme for Outlook on the web are:

1. Copy the folders and files of an existing theme, and rename the copied folders and files.
2. Configure the display name and sort order of the new theme.
3. Customize the new theme.
4. (Optional) Set the new theme as the default, and prevent users from selecting themes.
5. (Optional) Allow users to see and select the new theme
6. Restart IIS for the changes to take affect.

If you use multiple Exchange servers for Outlook on the web client connections, you need to copy the new theme to each server. You should also create a backup copy of the new theme so you can copy the files back after you reinstall or upgrade the Exchange server.

After you create a theme, you may also want to customize elements that are common to all themes. For more information, see Customize the Outlook on the web sign-in, language selection, and error pages in Exchange Server.

What do you need to know before you begin?

- Estimated time to complete this task: 45 minutes.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Outlook on the web virtual directories" entry in the Clients and mobile devices permissions topic. The account you use also needs to be a member of the local Administrators group on the Exchange server.
- The light version of Outlook on the web doesn't support themes.
- To replace an existing color with a new color, you need the HTML RGB value of the new color. You can find HTML RGB values at Color Table. If you can't find the color there, you can use an image editing tool or an HTML color codes web site to determine its HTML RGB value.
- Don't delete the folder `%ExchangeInstallPath%ClientAccess\OWA\prem\<ExchangeVersion>\resources\themes\base`, or any files in it.
- If you decide to directly edit an existing theme (not a copy of the theme), make a backup copy of the original files before you modify them.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server.

Step 1: Use File Explorer to copy the folders and files of an existing theme, and rename the copied folders and files

You can inspect the built-in themes by opening a mailbox in Outlook on the web, selecting **Settings**, and then selecting **Change theme**.

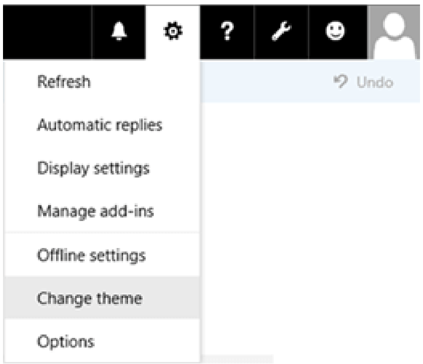

You can use the information in the Default Outlook on the web themes in Exchange Server section at the end of this topic to match the display name of the theme in Outlook on the web to the name of the theme folder on the Exchange server.

The theme files and folders are stored in the following locations:

- `%ExchangeInstallPath%ClientAccess\OWA\prem\<ExchangeVersion>\resources\themes\` contains the theme folder that holds the header image, theme preview image, and theme description text.

- `%ExchangeInstallPath%ClientAccess\OWA\prem\<ExchangeVersion>\resources\styles\` contains the `_fabric.color.variables.theme.<ThemeFolderName>.less` and `fabric.color.theme.<ThemeFolderName>.css` files that define the colors that are used in the theme.

  Note: The *<ExchangeVersion>* subfolder uses the syntax 15.1. *nnn*. *nn*, and indicates the Exchange Cumulative Update (CU) that's installed.

After you've identified the theme that's closest to what you want (for example, with or without a header image), you need to copy the theme folder and the corresponding files, and then rename the copied folders and files

1. In File Explorer, browse to `%ExchangeInstallPath%ClientAccess\OWA\prem\<ExchangeVersion>\resources\themes`.

2. Select an existing theme folder in the `\themes` folder, copy it, and then paste it back into the `\themes` folder. This results in a new folder named `<ThemeFolderName> - Copy`.

   Note: An easy way to copy and paste the theme folder is to select the folder, press the Control key + C, and then press the Control key + V.

3. Rename the new theme folder that you created in the previous step. For example, `fourthcoffee`.

   Note: An easy way to rename the folder is to select it, and then press the F2 key.

4. In File Explorer, browse to `%ExchangeInstallPath%ClientAccess\OWA\prem\<ExchangeVersion>\resources\styles\`.

5. Locate the files named `_fabric.color.variables.theme.<ThemeFolderName>.less` and `fabric.color.theme.<ThemeFolderName>.css` that correspond to the theme folder you copied in step 2. Select each file, copy it, and paste it back into the `\styles` folder. This results in new files named

`_fabric.color.variables.theme.<ThemeFolderName> - Copy.less` and `fabric.color.theme.<ThemeFolderName> - Copy.css`.

6. Rename the new files that you created in the previous step. The *<ThemeFolderName>* value must match the folder name from step 3. For example, `_fabric.color.variables.theme.fourthcoffee.less` and `fabric.color.theme.fourthcoffee.css`.

Step 2: Use Notepad to configure the display name and sort order of the new theme

You need to configure a unique display name and sort order for the new theme, because the new theme has the same display name and sort order as the theme you copied. The theme's display name appears in the **Change theme** panel in Outlook on the web. The sort order determines where the theme appears in the list of themes.

1. Use Notepad to open the file named `themeinfo.xml` in the new theme folder `%ExchangeInstallPath%ClientAccess\OWA\prem\<ExchangeVersion>\resources\themes\<NewThemeFolder>` that you created in Step 1. The contents of the file look like this:

   ```
   <theme displayname="__<CopiedThemeName>__" sortorder="<CopiedThemeSortOrder>"/>
   ```

2. Change the `displayname="__<CopiedThemeName>__"` value to the value you want. For example `displayname = "Fourth Coffee Corporate Theme"`.

   Note: The theme display name value `"__<ThemeName>__"` is a code string that's localized into different languages. The text value that you specify for the new theme isn't localized into different languages.

3. Change the `sortorder="<CopiedThemeSortOrder>"` integer value to the unique value you want. A lower value appears earlier in the list of themes. You can use the information in the Default Outlook on the web themes in Exchange Server section at the end of this topic to find the sort order values for the built-in themes. The Default theme has `sortorder="0"`, and appears first in the list.

   - If you want to insert your new theme among the list of built-in themes, change the number to a unique value that isn't already in use. For example, if you want your new theme to appear second in the list, you can use the value `sortorder="5"`.

   - If you want to *replace* the position of a built-in theme in the list, set the number to the same value as built-in theme, and then change the sort order for the built-in theme. For example, if you want your new theme to appear first in the list, you need to set your new theme to `sortorder="0"`. But, you also need to open the `themeinfo.xml` file in the `\base` folder (the Default theme) to change the value `sortorder="0"` to something else (for example, `sortorder="5")`.

4. When you're finished, save and close the `themeinfo.xml` file.

Step 3: Customize the new theme

Image files

Theme image files are stored in the following folders in `%ExchangeInstallPath%ClientAccess\OWA\prem\<ExchangeVersion>\resources\themes\<ThemeFolderName>`:

- `\images\0`: These files are used in left-to-right languages.

- `\images\rtl`: These files are used in right-to-left languages. Depending on the image, the file might be exactly the same as the left-to-right version, or it might be reversed (right-to-left instead of left-to-right).

The image files that exist in these folders are described in the following table:

FILE NAME	DIMENSIONS (WIDTH X HEIGHT IN PIXELS)	BIT DEPTH	DESCRIPTION
headerbgmaing2.png	2000 x 50	32	The header image for themes that use a static header image. The size of the file varies. If the theme doesn't use a static header image, the file is 1 x 1, and the size is 2815 bytes.
headerbgmaing2.gif	2000 x 50	24	The header image for themes that use an animated header image. The size of the file varies. If the theme doesn't use an animated header image, the file is 1 x 1, and the size is 43 bytes.
themepreview.png	64 x 64	24 or 8	The small square image that represents the theme in the **Change theme** panel in Outlook on the web. For the Default theme and the Black theme, this file 1 x 1, and the preview image is a black square.

You can edit the existing image file, or replace the file with a new file that has the same name and dimensions.

Colors

Theme colors are defined in the following files in the `%ExchangeInstallPath%ClientAccess\OWA\prem\<ExchangeVersion>\resources\styles` folder:

- `fabric.color.theme.<ThemeFolderName>.css`
- `_fabric.color.variables.<ThemeFolderName>.less`

If you change a color value, you need to change all references to the color in both files.

Step 4: (Optional) Set the default theme and prevent users from selecting a theme

Setting a new default theme only affects users who haven't manually selected their theme. To force all users to use the default theme, you also need to disable theme selection in Outlook on the web. These settings affect all users who connect to Outlook on the web through the Exchange server.

To set the default theme and prevent users from changing their theme in Outlook on the web, use the following syntax:

```
Set-OwaVirtualDirectory -Identity <VirtualDirectoryIdentity> -DefaultTheme <ThemeFolderName> -
ThemeSelectionEnabled $false
```

This example configures the theme folder named `fourthcoffee` as the default theme in Outlook on the web for the default website on the server named Mailbox01.

```
Set-OwaVirtualDirectory -Identity "Mailbox01\owa (Default Web Site)" -DefaultTheme fourthcoffee -
ThemeSelectionEnabled $false
```

Notes:

- By default, the value of the *DefaultTheme* parameter is blank (`$null`). This value indicates that no default theme is specified, and the theme named Default is used if the user hasn't manually selected a theme.
- Exchange doesn't validate the value that you specify for the *DefaultTheme* parameter. Make sure that the theme exists.
- To specify a default theme for specific users that overrides the default theme setting on the Outlook on the web virtual directory, use the *DefaultTheme* parameter on the **Set-OwaMailboxPolicy** cmdlet.

Step 5: (Optional) Allow users to select the new theme

If you don't want to force all users to use the new theme, you need to add the new theme to the `stylemanifest.xml` file so users can find and select it in the list of themes. The `stylemanifest.xml` file is located in `%ExchangeInstallPath%ClientAccess\OWA\prem\<ExchangeVersion>\manifests` .

This example adds a new line in the `stylemanifest.xml` file for the new `fourthcoffee` theme.

```
<themeVariables themeName="fourthcoffee" fileName="_fabric.color.variables.theme.fourthcoffee.less" />
```

Step 6: Restart IIS

You need to restart Internet Information Services (IIS) for the changes to take effect.

1. Open IIS Manager on the Exchange server. An easy way to do this in Windows Server 2012 or later is to press Windows key + Q, type inetmgr, and select **Internet Information Services (IIS) Manager** in the results.
2. In IIS Manager, select the server.
3. In the **Actions** pane, click **Restart**.

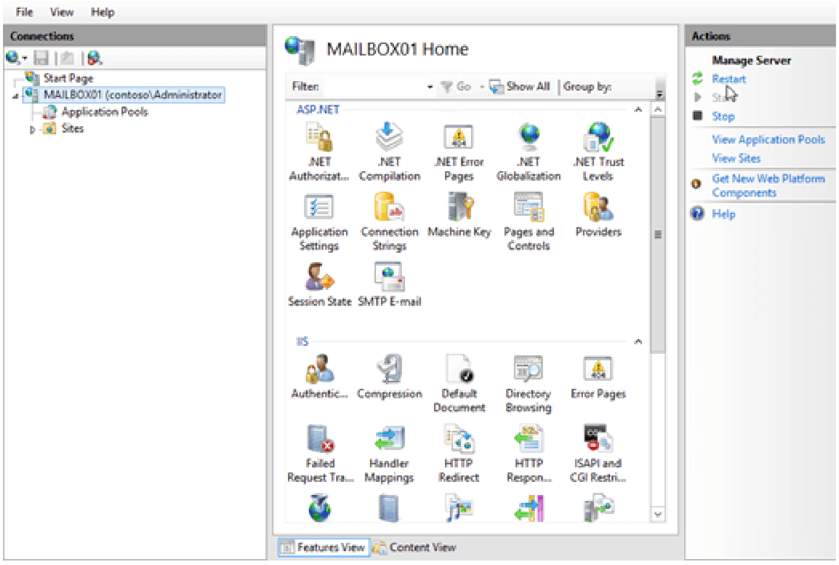

Note: To perform this procedure on the command line, open an elevated command prompt on the Exchange server (a Command Prompt window you open by selecting **Run as administrator**) and run the following command:

```
net stop w3svc /y
```

```
net start w3svc
```

How do you know this worked?

To verify that you've successfully created an Outlook on the web theme, perform the following steps:

1. Open a mailbox in Outlook on the web. On the Exchange server, you can test your theme by opening the URL https://localhost/owa or https://127.0.0.1/owa.

2. Depending on the settings you configured, verify the new theme is used by default, or verify that you can see and select the new theme at **Settings** > **Change theme**.

3. If you don't see your changes after you restart IIS, clear your browsing history (delete temporary Internet files), and refresh the browser window.

Default Outlook on the web themes in Exchange Server

The built-in Outlook on the web themes are located in the folder `%ExchangeInstallPath%ClientAccess\OWA\prem\<ExchangeVersion>\resources\themes`, and are described in the following table.

FOLDER NAME	DISPLAY NAME IN OUTLOOK ON THE WEB	SORT ORDER IN OUTLOOK ON THE WEB (LOWER LISTED FIRST)	HEADER IMAGE TYPE
angular	Angular 80's	110	Static
balloons	Balloons	240	Static
base	Default	0	None
beach	Beach Sunset	40	Animated
black	Black	670	None
blueberry	Blueberry	600	None
blueprint	Blueprint	120	Static
bricks	Bricks	20	Static
cats	Cats	300	Static
chevron	Chevron	80	Static
circuit	Circuit	130	Static
comic	Comic Book	170	Static

FOLDER NAME	DISPLAY NAME IN OUTLOOK ON THE WEB	SORT ORDER IN OUTLOOK ON THE WEB (LOWER LISTED FIRST)	HEADER IMAGE TYPE
contrast	Contrast	500	None
cordovan	Cordovan	650	None
crayon	Crayon	140	Static
cubes	3D Cubes	190	Static
cubism	Cubism	310	Static
darkcordovan	Dark Cordovan	660	None
darkorange	Dark Orange	620	None
diamonds	Floating Diamonds	160	Static
far	Far, Far Away	150	Animated
grape	Grape	610	None
jelly	Jelly Fish	70	Animated
lightblue	Light Blue	530	None
lightgreen	Light Green	540	None
lite	Lite	510	None
mediumdarkblue	Dark Blue	640	None
minimal	Minimal	520	None
modern	20th Century Modern	280	Static
mountain	Mountain Peak	50	Static
orange	Orange	580	None
paint	Finger paints	290	Static
pink	Pink	550	None
pixel	Pixel Pop	60	Static
polka	Polka Dot	200	Static
pomegranate	Pomegranate	590	None
primary	Primary	180	Static

FOLDER NAME	DISPLAY NAME IN OUTLOOK ON THE WEB	SORT ORDER IN OUTLOOK ON THE WEB (LOWER LISTED FIRST)	HEADER IMAGE TYPE
raspberry	Raspberry	570	None
robot	Robot	100	Animated
simple	Simple Facets	230	Static
spectrum	Spectrum Facets	90	Static
strawberry	Strawberry	250	Static
super	Super sparkle happy	10	Static
teagarden	Tea Garden	210	Static
teal	Teal	550	None
watermelon	Watermelon	630	None
whale	Whale of a Time	30	Animated
whimsical	Whimsical	220	Static
wntrlnd	Winterland	260	Static
wrld	One World	270	Static

Customize the Outlook on the web sign-in, language selection, and error pages in Exchange Server

8/3/2020 • 3 minutes to read • Edit Online

The Outlook on the web (formerly known as Outlook Web App) sign-in, language selection, and error pages are based on image and content style sheet (CSS) files in the themes resources folder in the Client Access (front end) services on an Exchange Server 2016 or Exchange 2019 server. Outlook on the web uses only one set of sign-in, language selection, and error pages for all themes. Any modifications to those pages will be seen by all users who connect to the Exchange server for Outlook on the web.

Notes:

- Backup the default Outlook on the web files before you make any changes.
- Create a back-up copy of your customized files so you can reapply them after a reinstallation or upgrade of the Exchange server.
- If you use multiple Exchange servers for Outlook on the web connections, you need to copy the modified files to each server.

For more information about Outlook on the web, see Outlook on the web in Exchange Server. For information about creating a custom theme, see Create a theme for Outlook on the web in Exchange Server.

What do you need to know before you begin?

- Estimated time to complete this task: 30 minutes.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Graphics editor" entry under "Outlook on the web Permissions" in the Clients and mobile devices permissions topic.
- To replace the existing color with a new color, you need the HTML RGB value of the new color. You can find HTML RGB values in the topic Color Table. If you can't find the color there, you can use an image editing tool to sample a color and determine its HTML RGB value.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Customize the color of the Outlook on the web sign-in page

1. Use Notepad to open the file `%ExchangeInstallPath%FrontEnd\HttpProxy\owa\auth\<ExchangeVersion>\themes\resources\logon.css`.

   Note: The *<ExchangeVersion>* subfolder uses the syntax 15.1. *nnn*. *nn*, and changes every time you install an Exchange Cumulative Update (CU).

2. In the `logon.css` file, replace the default blue color value #0072c6 with the HTML RGB value that you want

to use.

3. When you're finished, save and close the file.

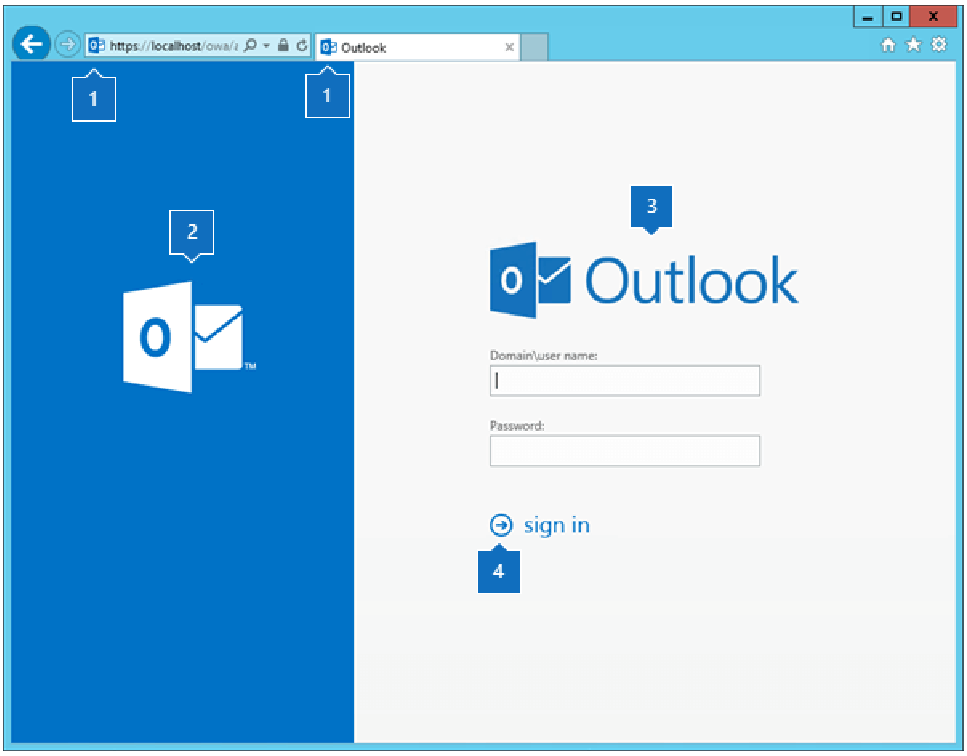

Customize the color of the Outlook on the web error page

1. Use Notepad to open the file `%ExchangeInstallPath%FrontEnd\HttpProxy\owa\auth\<ExchangeVersion>\themes\resources\errorFE.css` .

2. In the `errorFE.css` file, replace the default blue color value #0072c6 with the HTML RGB value that you want to use.

3. When you're finished, save and close the file.

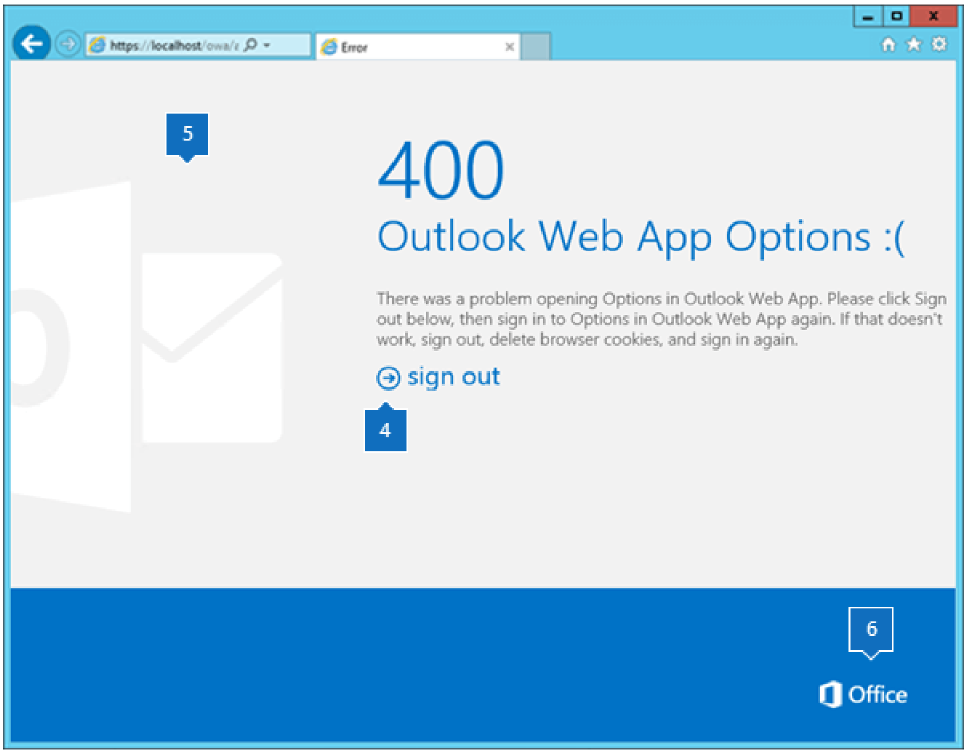

400

Outlook Web App Options :(

There was a problem opening Options in Outlook Web App. Please click Sign out below, then sign in to Options in Outlook Web App again. If that doesn't work, sign out, delete browser cookies, and sign in again.

sign out

Customize the color of the Outlook on the web language selection page

1. Use Notepad to open the file `%ExchangeInstallPath%ClientAccess\Owa\prem\<ExchangeVersion>\resources\styles\languageselection.css`.
2. In the `languageselection.css` file, replace the default blue color value #0072c6 with the HTML RGB value that you want to use.
3. When you're finished, save and close the file.

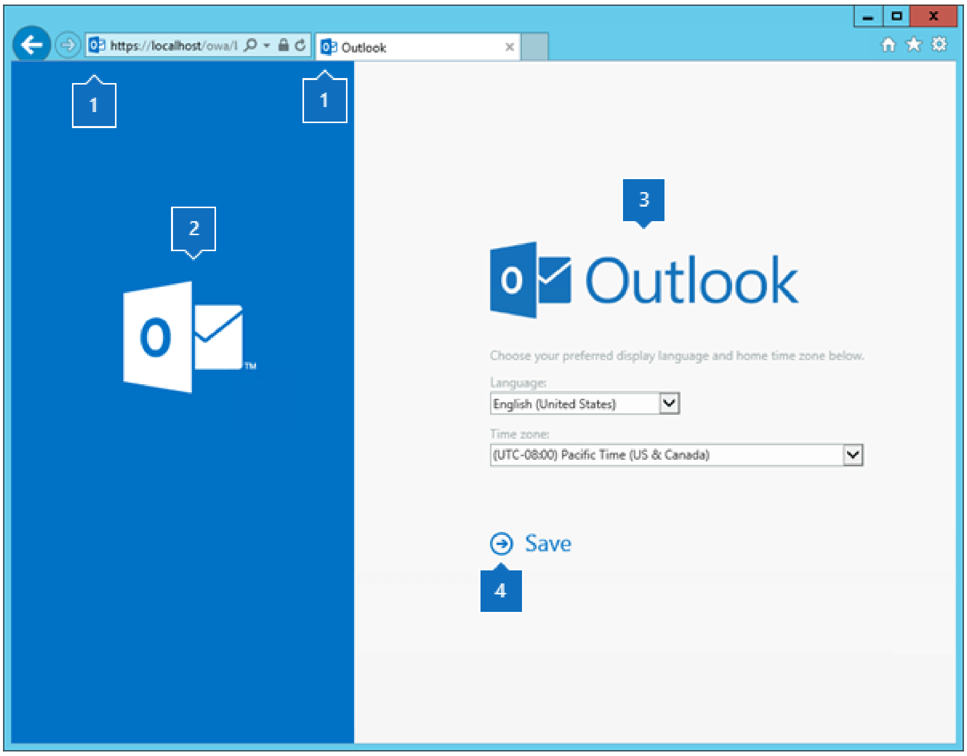

Choose your preferred display language and home time zone below.

Language:
English (United States)

Time zone:
(UTC-08:00) Pacific Time (US & Canada)

Save

Customize the images on the Outlook on the web sign-in, language selection, and error pages

You can edit the existing image files, or replace the files with new files that have the same names and dimensions. The images are described in the following table:

IMAGE	FILE NAME	LOCATION	DIMENSIONS (WIDTH X HEIGHT IN PIXELS)	BIT DEPTH
1	favicon.ico	`%ExchangeInstallPath%FrontEnd\HttpProxy\owa\auth\<ExchangeVersion>\themes\resources`	16 x 16	32
2	olk_logo_white.png	`%ExchangeInstallPath%ClientAccess\Owa\prem\<ExchangeVersion>\resources\images\0`	128 x 108	32
3	owa_text_blue.png	`%ExchangeInstallPath%ClientAccess\Owa\prem\<ExchangeVersion>\resources\images\0`	300 x 76	32

IMAGE	FILE NAME	LOCATION	DIMENSIONS (WIDTH X HEIGHT IN PIXELS)	BIT DEPTH
4	Sign_in_arrow.png (for left-to-right languages) Sign_in_arrow_rtl.png (for right-to-left languages)	`%ExchangeInstallPath%FrontEnd\HttpProxy\owa\auth\<ExchangeVersion>\themes\resources`	22 x 22	32
5	olk_logo_white_cropped.png	`%ExchangeInstallPath%FrontEnd\HttpProxy\owa\auth\<ExchangeVersion>\themes\resources`	265 x 310	32
6	office_logo_white_small.png	`%ExchangeInstallPath%ClientAccess\Owa\prem\<ExchangeVersion>\resources\images\0` (for left-to-right languages) `%ExchangeInstallPath%ClientAccess\Owa\prem\<ExchangeVersion>\resources\images\rtl` (for right-to-left languages)	81 x 26	8

How do you know this worked?

To verify that you've successfully customized the Outlook on the web sign-in, language selection, and error pages, perform the following steps:

1. Open the Outlook on the web sign-in page in a web browser. On the Exchange server that hosts the Outlook on the web virtual directory, you can test your changes by opening the URL https://localhost/owa or https://127.0.0.1/owa.

2. If you don't see your changes, clear your browsing history (delete temporary Internet files), and refresh the browser window.

   Note: To see the effects of your changes, you can keep the .css file open and refresh the browser window after you save each change.

Use AD FS claims-based authentication with Outlook on the web

8/3/2020 • 31 minutes to read • Edit Online

Installing and configuring Active Directory Federation Services (AD FS) in Exchange Server organizations allows clients to use AD FS claims-based authentication to connect to Outlook on the web (formerly known as Outlook Web App) and the Exchange admin center (EAC). Claims-based identity is another approach to authentication that removes authentication management from the application, and makes it easier for you to manage accounts by centralizing authentication. When claims-based authentication is enabled, Outlook on the web and the EAC aren't responsible for authenticating users, storing user accounts and passwords, looking up user identity details, or integrating with other identity systems. Centralizing authentication helps make it easier to upgrade authentication methods in the future.

AD FS claims-based authentication replaces the traditional authentication methods that are available for Outlook on the web and the EAC. For example:

- Active Directory client certificate authentication
- Basic authentication
- Digest authentication
- Forms authentication
- Windows authentication

Setting up AD FS claims-based authentication for Outlook on the web and the EAC in Exchange Server involves the following additional servers:

- A Windows Server 2012 or later domain controller (Active Directory Domain Services server role).
- A Windows Server 2012 or later AD FS server (Active Directory Federation Services server role). Windows Server 2012 uses AD FS 2.1, and Windows Server 2012 R2 uses AD FS 3.0. You need to be a member of the Domain Admins, Enterprise Admins, or local Administrators security group to install AD FS, and to create the required relying party trusts and claim rules on the AD FS server.
- Optionally, a Windows Server 2012 R2 or later Web Application Proxy server (Remote Access server role, Web Application Proxy role service).
  - Web Application Proxy is a reverse proxy server for web applications that are inside the corporate network. Web Application Proxy allows users on many devices to access published web applications from outside the corporate network. For more information, see Installing and Configuring Web Application Proxy for Publishing Internal Applications.
  - Although Web Application Proxy is typically recommended when AD FS is accessible to external clients, offline access in Outlook on the web isn't supported when using AD FS authentication through Web Application Proxy.
  - Installing Web Application Proxy on a Windows Server 2012 R2 server requires local administrator permissions.
  - You need to deploy and configure the AD FS server before you configure the Web Application Proxy server, and you can't install Web Application Proxy on the same server where AD FS is installed.

What do you need to know before you begin?

- Estimated time to complete this procedure: 45 minutes.
- The procedures in this topic are based on Windows Server 2012 R2.
- Outlook on the web for devices doesn't support AD FS claims-based authentication.
- For the procedures in the Exchange organization, you need to have Organization Management permissions.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Step 1: Review the certificate requirements for AD FS

AD FS requires two basic types of certificates:

- A service communication Secure Sockets Layer (SSL) certificate for encrypted web services traffic between the AD FS server, clients, Exchange servers, and the optional Web Application Proxy server. We recommend that you use a certificate that's issued by an internal or commercial certification authority (CA), because all clients need to trust this certificate.
- A token-signing certificate for encrypted communication and authentication between the AD FS server, Active Directory domain controllers, and Exchange servers. We recommend that you use the default self-signed AD FS token signing certificate.

For more information about creating and importing SSL certificates in Windows, see Server Certificates.

Here's a summary of the certificates that we'll be using in this scenario:

COMMON NAME (CN) IN THE CERTIFICATE (IN THE SUBJECT, SUBJECT ALTERNATIVE NAME, OR A WILDCARD CERTIFICATE MATCH)	TYPE	REQUIRED ON SERVERS	COMMENTS
`adfs.contoso.com`	Issued by a CA	AD FS server Web Application Proxy server	This is the host name that's visible to clients, so clients need to trust the issuer of this certificate.

COMMON NAME (CN) IN THE CERTIFICATE (IN THE SUBJECT, SUBJECT ALTERNATIVE NAME, OR A WILDCARD CERTIFICATE MATCH)	TYPE	REQUIRED ON SERVERS	COMMENTS
`ADFS Signing - adfs.contoso.com`	Self-signed	AD FS server Exchange servers Web Application Proxy server	The default self-signed certificate is automatically copied over during the configuration of the optional Web Application Proxy server, but you'll need to manually import it into the Trusted Root Certificate store on all Exchange servers in your organization. By default, the self-signed token-signing certificates are valid for one year. The AD FS server is configured to automatically renew (replace) its self-signed certificates before they expire, but you'll need to re-import the certificate on the Exchange servers. You can increase the default certificate expiration period by running this command in Windows PowerShell on the AD FS server: `Set-AdfsProperties -CertificateDuration <Days>` (the default value is 365). For more information, see Set-AdfsProperties. To export the certificate from the AD FS Management console, select **Service** > **Certificates** > right-click on the token-signing certificate > select **View Certificate** > click the **Details** tab > click **Copy to File**.
`mail.contoso.com`	Issued by a CA	Exchange servers Web Application Proxy server	This is the typical certificate that's used to encrypt external client connections to Outlook on the web (and likely other Exchange IIS services). For more information, see Certificate requirements for Exchange services.

For more information, see the "Certificate requirements" section in AD FS Requirements.

NOTE

Secure Sockets Layer (SSL) is being replaced by Transport Layer Security (TLS) as the protocol that's used to encrypt data sent between computer systems. They're so closely related that the terms "SSL" and "TLS" (without versions) are often used interchangeably. Because of this similarity, references to "SSL" in Exchange topics, the Exchange admin center, and the Exchange Management Shell have often been used to encompass both the SSL and TLS protocols. Typically, "SSL" refers to the actual SSL protocol only when a version is also provided (for example, SSL 3.0). To find out why you should disable the SSL protocol and switch to TLS, check out Protecting you against the SSL 3.0 vulnerability.

Step 2: Deploy an AD FS server

You can use Server Manager or Windows PowerShell to install the Active Directory Federation Services role service on the target server.

To use Server Manager to install AD FS, follow these steps:

1. On the target server, open **Server Manager**, click **Manage**, and then select **Add Roles and Features**.

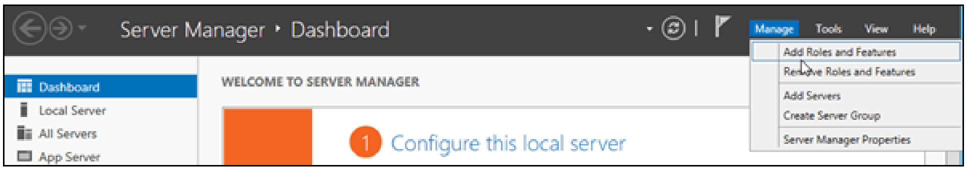
Server Manager • Dashboard
Manage
Tools
View
Help
Add Roles and Features
Remove Roles and Features
Add Servers
Create Server Group
Server Manager Properties
Dashboard
Local Server
All Servers
App Server
WELCOME TO SERVER MANAGER
1
Configure this local server

2. The **Add Roles and Features Wizard** opens. You'll start on the **Before you begin** page unless you previously selected **Skip this page by default**. Click **Next**.

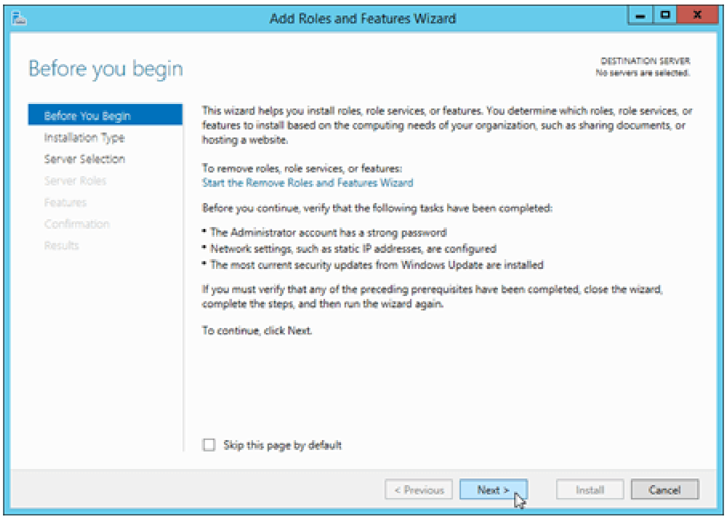

This wizard helps you install roles, role services, or features. You determine which roles, role services, or features to install based on the computing needs of your organization, such as sharing documents, or hosting a website.

To remove roles, role services, or features:
Start the Remove Roles and Features Wizard

Before you continue, verify that the following tasks have been completed:

- The Administrator account has a strong password
- Network settings, such as static IP addresses, are configured
- The most current security updates from Windows Update are installed

If you must verify that any of the preceding prerequisites have been completed, close the wizard, complete the steps, and then run the wizard again.

To continue, click Next.

□ Skip this page by default

3. On the **Select installation type** page, verify that **Role-based or feature-based installation** is selected, and then click **Next**.

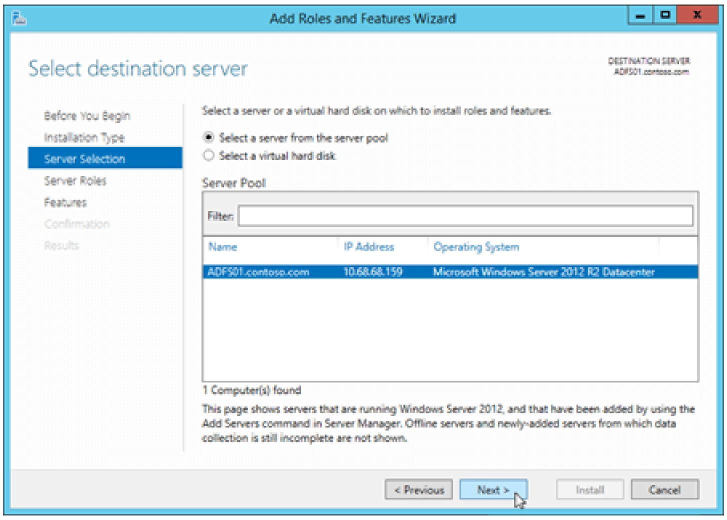

Select destination server

DESTINATION SERVER
ADFS01.contoso.com

Select a server or a virtual hard disk on which to install roles and features.

● Select a server from the server pool
○ Select a virtual hard disk

Server Pool

Filter:

Name	IP Address	Operating System
ADFS01.contoso.com	10.68.68.159	Microsoft Windows Server 2012 R2 Datacenter

1 Computer(s) found

This page shows servers that are running Windows Server 2012, and that have been added by using the Add Servers command in Server Manager. Offline servers and newly-added servers from which data collection is still incomplete are not shown.

4. On the **Select destination server** page, verify the server selection, and then click **Next**.

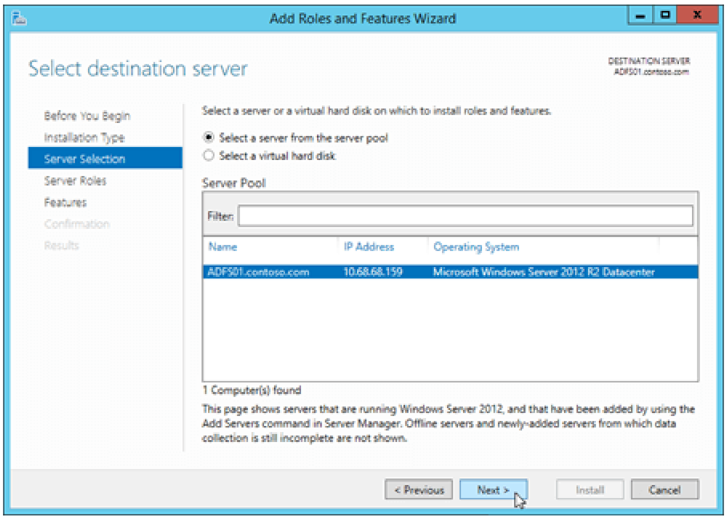

Select destination server

DESTINATION SERVER
ADFS01.contoso.com

Select a server or a virtual hard disk on which to install roles and features.

● Select a server from the server pool
○ Select a virtual hard disk

Server Pool

Filter:

Name	IP Address	Operating System
ADFS01.contoso.com	10.68.68.159	Microsoft Windows Server 2012 R2 Datacenter

1 Computer(s) found

This page shows servers that are running Windows Server 2012, and that have been added by using the Add Servers command in Server Manager. Offline servers and newly-added servers from which data collection is still incomplete are not shown.

5. On the **Select server roles** page, select **Active Directory Federation Services** from the list, and then click **Next**.

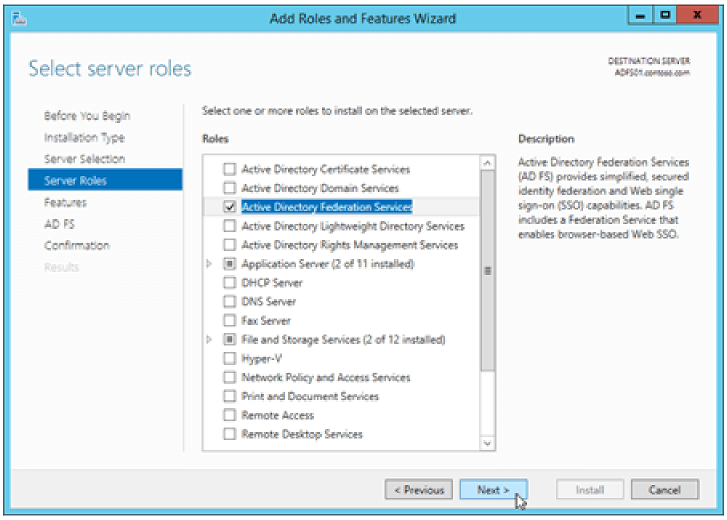

Select server roles

DESTINATION SERVER
ADFS01.contoso.com

Select one or more roles to install on the selected server.

Roles

- □ Active Directory Certificate Services
- □ Active Directory Domain Services
- ☑ Active Directory Federation Services
- □ Active Directory Lightweight Directory Services
- □ Active Directory Rights Management Services
- ▨ Application Server (2 of 11 installed)
- □ DHCP Server
- □ DNS Server
- □ Fax Server
- ▨ File and Storage Services (2 of 12 installed)
- □ Hyper-V
- □ Network Policy and Access Services
- □ Print and Document Services
- □ Remote Access
- □ Remote Desktop Services

Description

Active Directory Federation Services (AD FS) provides simplified, secured identity federation and Web single sign-on (SSO) capabilities. AD FS includes a Federation Service that enables browser-based Web SSO.

6. On the **Select features** page, click **Next** (accept the default feature selections).

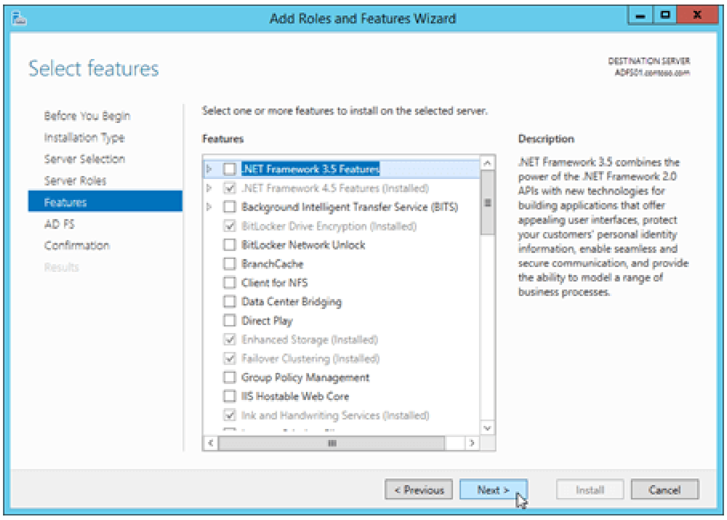

Select features

Select one or more features to install on the selected server.

Features

- ☐ .NET Framework 3.5 Features
- ☑ .NET Framework 4.5 Features (Installed)
- ☐ Background Intelligent Transfer Service (BITS)
- ☑ BitLocker Drive Encryption (Installed)
- ☐ BitLocker Network Unlock
- ☐ BranchCache
- ☐ Client for NFS
- ☐ Data Center Bridging
- ☐ Direct Play
- ☑ Enhanced Storage (Installed)
- ☑ Failover Clustering (Installed)
- ☐ Group Policy Management
- ☐ IIS Hostable Web Core
- ☑ Ink and Handwriting Services (Installed)

Description

.NET Framework 3.5 combines the power of the .NET Framework 2.0 APIs with new technologies for building applications that offer appealing user interfaces, protect your customers' personal identity information, enable seamless and secure communication, and provide the ability to model a range of business processes.

7. On the **Active Directory Federation Services (AD FS)** page, click **Next**.

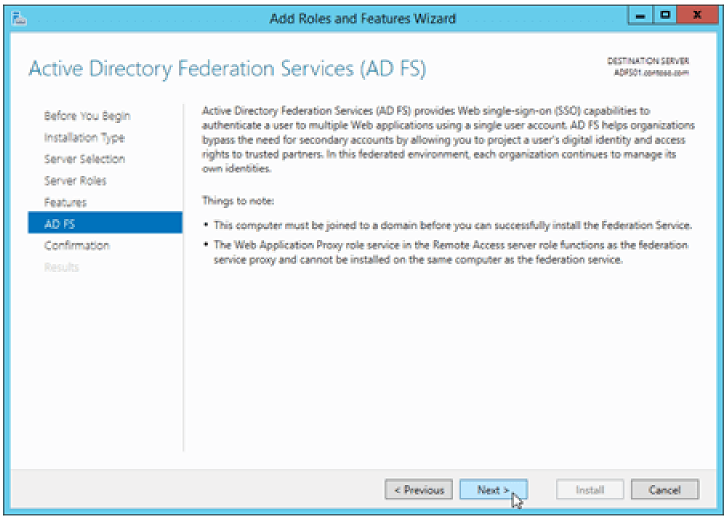

Active Directory Federation Services (AD FS)

Active Directory Federation Services (AD FS) provides Web single-sign-on (SSO) capabilities to authenticate a user to multiple Web applications using a single user account. AD FS helps organizations bypass the need for secondary accounts by allowing you to project a user's digital identity and access rights to trusted partners. In this federated environment, each organization continues to manage its own identities.

Things to note:

- This computer must be joined to a domain before you can successfully install the Federation Service.
- The Web Application Proxy role service in the Remote Access server role functions as the federation service proxy and cannot be installed on the same computer as the federation service.

8. **Windows Server 2012 only**: On the **Select role services** page, click **Next** (accept the default role service selections).

9. On the **Confirm installation selections** page, click **Install**.

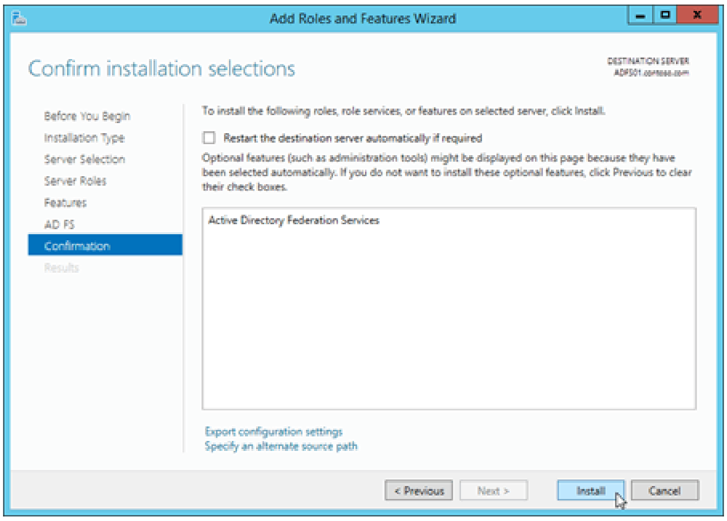

Confirm installation selections

To install the following roles, role services, or features on selected server, click Install.

☐ Restart the destination server automatically if required

Optional features (such as administration tools) might be displayed on this page because they have been selected automatically. If you do not want to install these optional features, click Previous to clear their check boxes.

Active Directory Federation Services

10. On the **Installation progress** page, you can watch the progress bar to verify that the installation was successful. When the installation is finished, leave the wizard open so you can click **Configure the federation service on this server** in Step 3b: Configure the AD FS server.

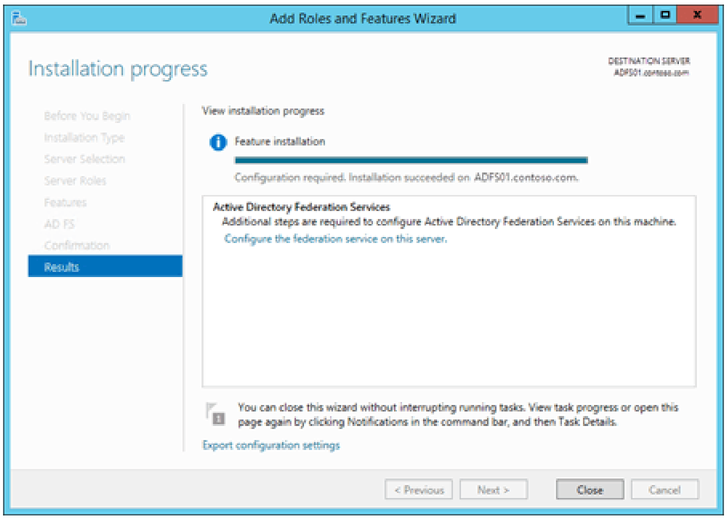

Configuration required. Installation succeeded on ADFS01.contoso.com.

Active Directory Federation Services

Additional steps are required to configure Active Directory Federation Services on this machine.

Configure the federation service on this server.

You can close this wizard without interrupting running tasks. View task progress or open this page again by clicking Notifications in the command bar, and then Task Details.

Export configuration settings

To use Windows PowerShell to install AD FS, run the following command:

```
Install-WindowsFeature ADFS-Federation -IncludeManagementTools
```

Step 3: Configure and test the AD FS server

You can also refer to this checklist to help you configure AD FS: Checklist: Setting Up a Federation Server.

Step 3a: Create a gMSA on a domain controller

Before you configure the AD FS server, you need to create a group Managed Service Account (gMSA) on a Windows Server 2012 or later domain controller. You do this in an elevated Windows PowerShell window on the domain controller (a Windows PowerShell window you open by selecting **Run as administrator**).

1. Run the following command:

   ```
   Add-KdsRootKey -EffectiveTime (Get-Date).AddHours(-10)
   ```

   If the command is successful, a GUID value is returned. For example:

   > Guid
   > ----
   > 2570034b-ab50-461d-eb80-04e73ecf142b

2. To create a new gMSA account for the AD FS server, use the following syntax:

   ```
   New-ADServiceAccount -Name <AccountName> -DnsHostName <FederationServiceName> -ServicePrincipalNames
   http/<FederationServiceName>
   ```

   This example creates a new gMSA account named FSgMSA for the Federation Service named adfs.contoso.com. The Federation Service name is the value that's visible to clients.

   ```
   New-ADServiceAccount -Name FSgMSA -DnsHostName adfs.contoso.com -ServicePrincipalNames
   http/adfs.contoso.com
   ```

Step 3b: Configure the AD FS server

To configure the AD FS server, you can use Server Manager or Windows PowerShell.

To use Server Manager, following these steps:

1. If you left the **Add Roles and Features Wizard** open on the AD FS server from Step 2: Deploy an AD FS server, you can click the **Configure the federation service on this server** link on the **Installation progress** page.

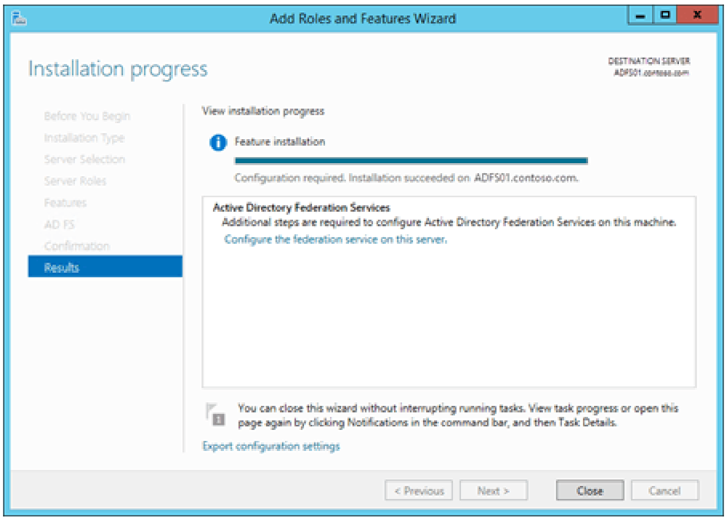

Add Roles and Features Wizard

Installation progress

DESTINATION SERVER
ADFS01.contoso.com

View installation progress

Feature installation

Configuration required. Installation succeeded on ADFS01.contoso.com.

Active Directory Federation Services
Additional steps are required to configure Active Directory Federation Services on this machine.
Configure the federation service on this server.

You can close this wizard without interrupting running tasks. View task progress or open this page again by clicking Notifications in the command bar, and then Task Details.

Export configuration settings

If you closed the **Add Roles and Features Wizard** or you used Windows PowerShell to install AD FS, you can get to the same place in Server Manager by clicking **Notifications**, and then clicking **Configure the federation service on this server** in the **Post-deployment Configuration** warning.

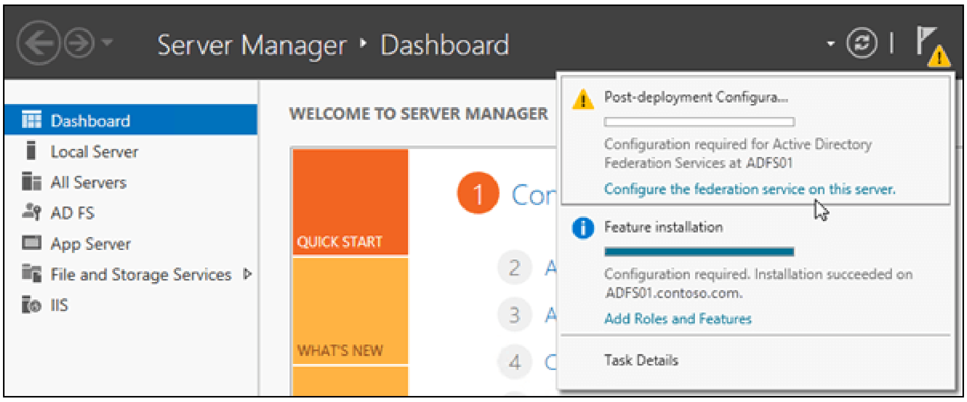

Server Manager ▸ Dashboard

Post-deployment Configura...

Configuration required for Active Directory Federation Services at ADFS01

Configure the federation service on this server.

Feature installation

Configuration required. Installation succeeded on ADFS01.contoso.com.

Add Roles and Features

Task Details

2. The **Active Directory Federation Services Wizard** opens. On the **Welcome** page, verify **Create the first federation server in a federation server farm** is selected, and then click **Next**.

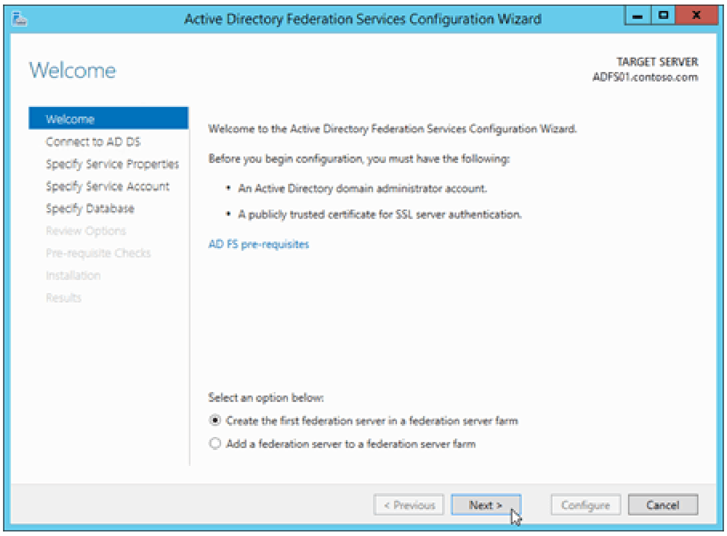

Active Directory Federation Services Configuration Wizard

Welcome

TARGET SERVER
ADFS01.contoso.com

Welcome to the Active Directory Federation Services Configuration Wizard.

Before you begin configuration, you must have the following:

- An Active Directory domain administrator account.
- A publicly trusted certificate for SSL server authentication.

AD FS pre-requisites

Select an option below:

● Create the first federation server in a federation server farm

○ Add a federation server to a federation server farm

3. On the **Connect to Active Directory Federation Services** page, select a domain administrator account in the domain where the AD FS server resides (your current credentials are selected by default). If you need to select a different user, click **Change**. When you're finished, click **Next**.

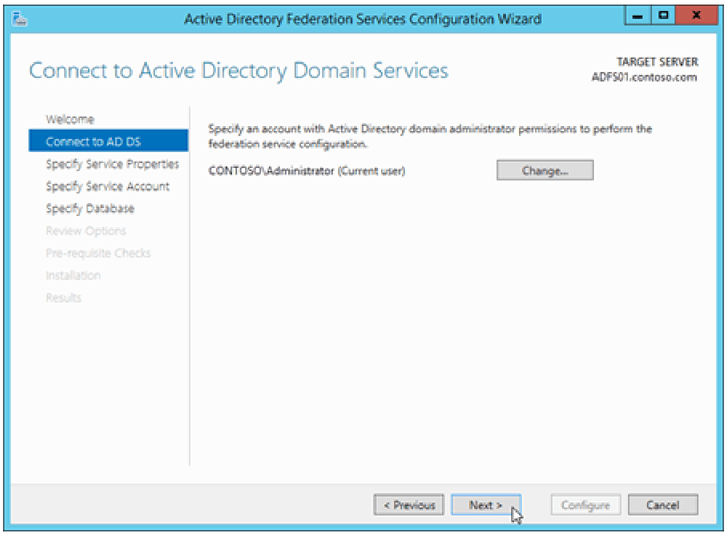

Active Directory Federation Services Configuration Wizard

Connect to Active Directory Domain Services

TARGET SERVER
ADFS01.contoso.com

Specify an account with Active Directory domain administrator permissions to perform the federation service configuration.

CONTOSO\Administrator (Current user)

4. On the **Specify Service Properties** page, configure the following settings:

   - **SSL Certificate**: Import or select the SSL certificate that contains the federation service name that you configured in Step 3a: Create a gMSA on a domain controller (for example `adfs.contoso.com`). When you import a certificate that isn't already installed on the server, you need to import a .pfx file (likely, a password-protected file that contains the certificate's private key). The common name (CN) value in the certificate's Subject field is displayed here.

   - **Federation Service Name**: This field is automatically populated based on the type of SSL certificate that you select or import:

     - **Single subject certificate**: The CN value of the certificate's Subject field is displayed, and you can't change it (for example, `adfs.contoso.com`).

     - **SAN certificate**: If the certificate contains the required federation service name, that value is displayed (for example, `adfs.contoso.com`). You can use the drop down list to see other CN values in the certificate.

     - **Wildcard certificate**: The CN value of the certificate's Subject field is displayed (for example, `*.contoso.com`), but you need to change it to the required federation service name (for example, `adfs.contoso.com`).

     Note: If the certificate you select doesn't contain the required federation service name (the **Federation Service Name** field doesn't contain the required value), you'll receive the following error:

     ```
     The federation service name does not match any of the subject names found in the certificate.
     ```

   - **Federation Service Display Name**: Enter the name of your organization. For example, Contoso, Ltd..

   When you're finished, click **Next**.

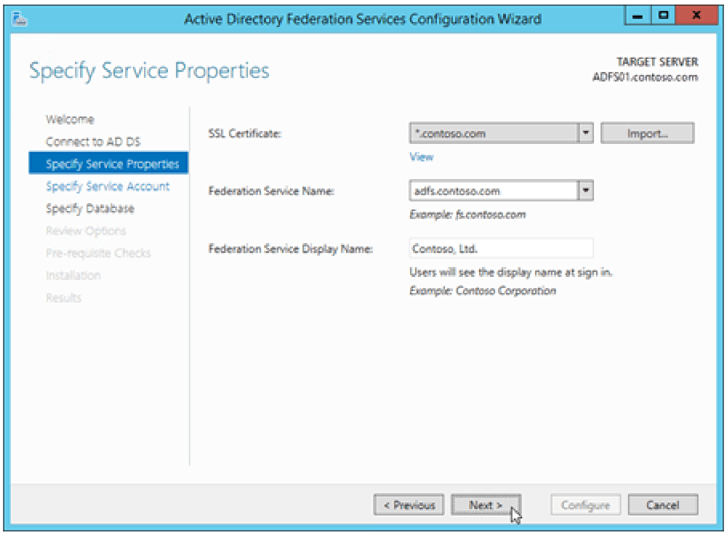

Specify Service Properties

TARGET SERVER
ADFS01.contoso.com

SSL Certificate: *.contoso.com

Federation Service Name: adfs.contoso.com

Example: fs.contoso.com

Federation Service Display Name: Contoso, Ltd.

Users will see the display name at sign in.

Example: Contoso Corporation

5. On the **Specify Service Account** page, configure the following settings:

   - Select **Use an existing domain user account or group Managed Service Account**.
   - **Account Name**: Click **Select** and enter the gMSA account that you created in Step 3a: Create a gMSA on a domain controller (for example, `FSgMSA`). Note that after you select it, the value that's displayed is `<Domain>\<gMSAAccountName>$` (for example, `CONTOSO\FSgMSA$`).

   When you're finished, click **Next**.

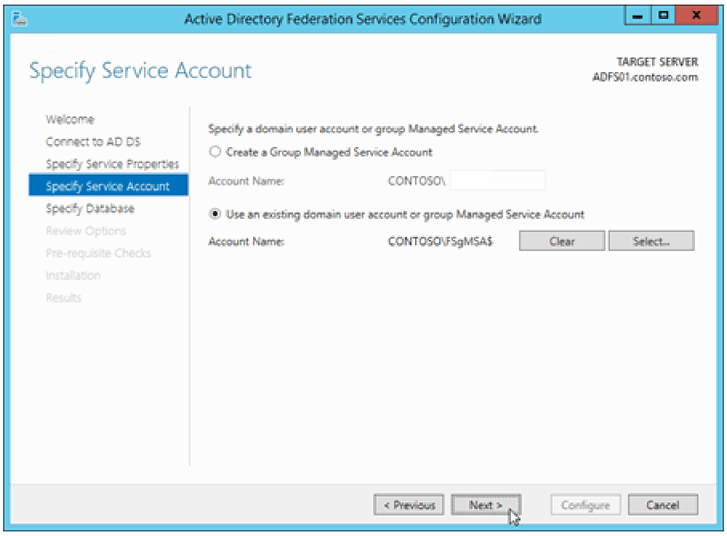

Specify Service Account

TARGET SERVER
ADFS01.contoso.com

Specify a domain user account or group Managed Service Account.

○ Create a Group Managed Service Account

Account Name: CONTOSO\

● Use an existing domain user account or group Managed Service Account

Account Name: CONTOSO\FSgMSA$

6. On the **Specify Configuration Database** page, verify that **Create a database on this server using Windows Internal Database** is selected, and then click **Next**.

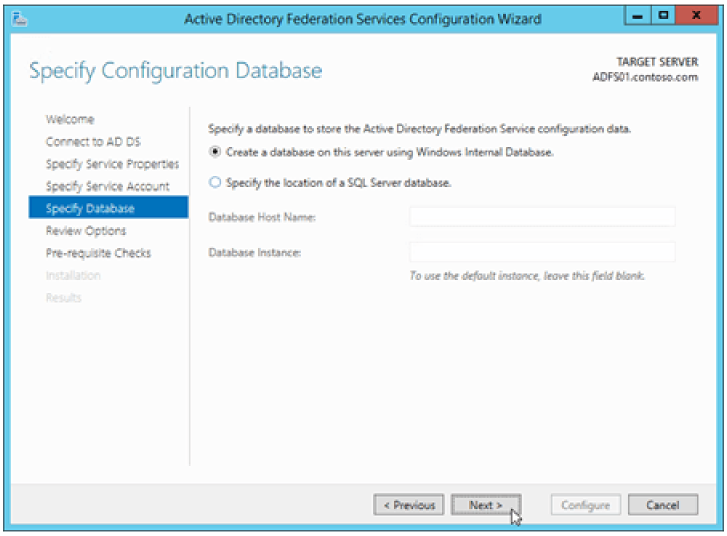

Specify Configuration Database

TARGET SERVER
ADFS01.contoso.com

Specify a database to store the Active Directory Federation Service configuration data.

● Create a database on this server using Windows Internal Database.

○ Specify the location of a SQL Server database.

Database Host Name:

Database Instance:

To use the default instance, leave this field blank.

7. On the **Review Options** page, verify your selections. You can click **View Script** button to copy the Windows PowerShell equivalent of the selections that you made for future use. When you're finished, click **Next**.

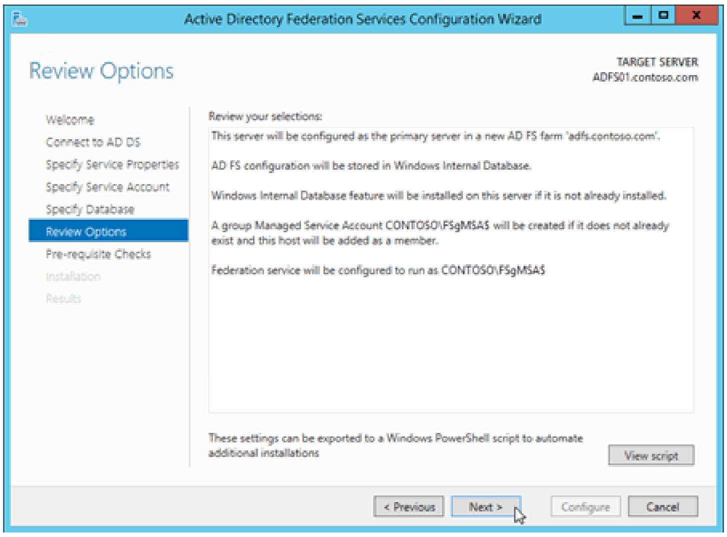

Active Directory Federation Services Configuration Wizard

Review Options

TARGET SERVER
ADFS01.contoso.com

Review your selections:

This server will be configured as the primary server in a new AD FS farm 'adfs.contoso.com'.

AD FS configuration will be stored in Windows Internal Database.

Windows Internal Database feature will be installed on this server if it is not already installed.

A group Managed Service Account CONTOSO\FSgMSA$ will be created if it does not already exist and this host will be added as a member.

Federation service will be configured to run as CONTOSO\FSgMSA$

These settings can be exported to a Windows PowerShell script to automate additional installations

8. On the **Pre-requisite Checks** page, verify that all the prerequisite checks were successfully completed, and then click **Configure**.

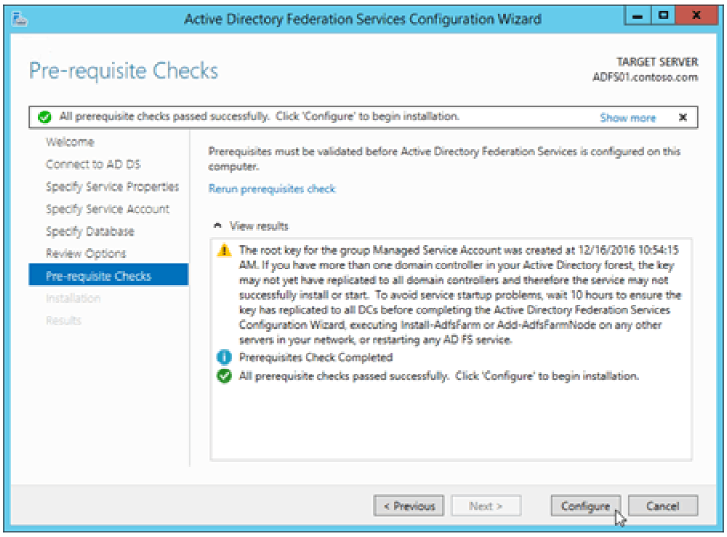

Active Directory Federation Services Configuration Wizard

Pre-requisite Checks

TARGET SERVER
ADFS01.contoso.com

✓ All prerequisite checks passed successfully. Click 'Configure' to begin installation.

Prerequisites must be validated before Active Directory Federation Services is configured on this computer.

Rerun prerequisites check

View results

⚠ The root key for the group Managed Service Account was created at 12/16/2016 10:54:15 AM. If you have more than one domain controller in your Active Directory forest, the key may not yet have replicated to all domain controllers and therefore the service may not successfully install or start. To avoid service startup problems, wait 10 hours to ensure the key has replicated to all DCs before completing the Active Directory Federation Services Configuration Wizard, executing Install-AdfsFarm or Add-AdfsFarmNode on any other servers in your network, or restarting any AD FS service.

ℹ Prerequisites Check Completed

✓ All prerequisite checks passed successfully. Click 'Configure' to begin installation.

9. On the **Results** page, review the results, verify that the configuration completed successfully. You can click **Next steps required for completing your federation service deployment** if you want to read about the next steps (for example, configuring DNS). When you're finished, click **Close**.

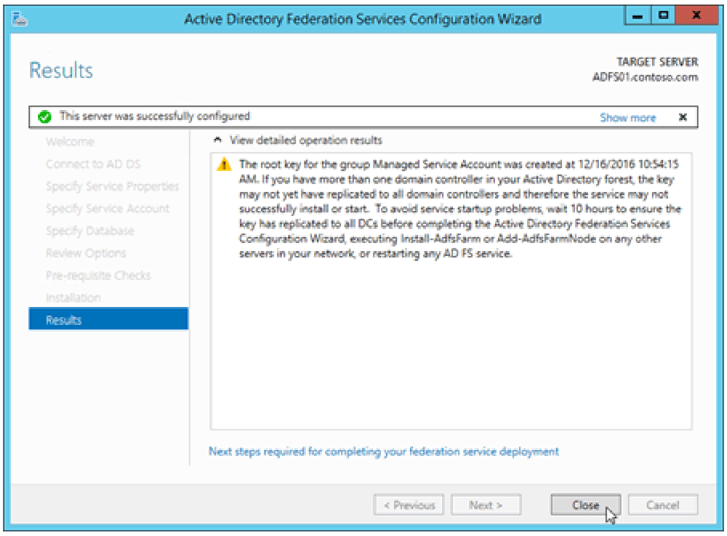

Active Directory Federation Services Configuration Wizard

Results

TARGET SERVER
ADFS01.contoso.com

✓ This server was successfully configured

View detailed operation results

⚠ The root key for the group Managed Service Account was created at 12/16/2016 10:54:15 AM. If you have more than one domain controller in your Active Directory forest, the key may not yet have replicated to all domain controllers and therefore the service may not successfully install or start. To avoid service startup problems, wait 10 hours to ensure the key has replicated to all DCs before completing the Active Directory Federation Services Configuration Wizard, executing Install-AdfsFarm or Add-AdfsFarmNode on any other servers in your network, or restarting any AD FS service.

Next steps required for completing your federation service deployment

To use Windows PowerShell to configure AD FS, follow these steps:

1. Run the following command on the AD FS server to find the thumbprint value of the installed certificate that contains `adfs.contoso.com`:

   ```
   Set-Location Cert:\LocalMachine\My; Get-ChildItem | Format-List FriendlyName,Subject,Thumbprint
   ```

2. Run the following command:

   ```
   Import-Module ADFS
   ```

3. Use the following syntax:

   ```
   Install-AdfsFarm -CertificateThumbprint <ThumbprintValue> -FederationServiceName <FederationServiceName>
   -FederationServiceDisplayName <FederationServiceDisplayName> -GroupServiceAccountIdentifier <gMSA>
   ```

This example configures AD FS with the following settings:

- **adfs.contoso.com certificate thumbprint**: The `*.contoso.com` certificate that has the thumbprint value `5AE82C737900B29C2BAC3AB6D8C44D249EE05609`.

- **Federation service name**: `adfs.contoso.com`

- **Federation service display name**: `Contoso, Ltd.`

- **Federation gMSA SAM account name and domain**: For example, for the gMSA account named `FSgMSA` in the `contoso.com` domain, the required value is `contoso\FSgMSA$`.

```
Install-AdfsFarm -CertificateThumbprint 5AE82C737900B29C2BAC3AB6D8C44D249EE05609 -FederationServiceName
adfs.contoso.com -FederationServiceDisplayName "Contoso, Ltd." -GroupServiceAccountIdentifier
"contoso\FSgMSA`$"
```

Notes:

- When you create the gMSA, the `$` is automatically appended to the **Name** value to create the **SamAccountName** value, which is required here.

- The escape character (``) is required for the `$` in the **SamAccountName**.

For details and syntax, see Install-AdfsFarm.

Step 3c: Test the AD FS server

After you configure AD FS, you can verify the installation on the AD FS server by successfully opening the URL of the federation metadata in a web browser. The URL uses the syntax `https://<FederationServiceName>/federationmetadata/2007-06/federationmetadata.xml`. For example, `https://adfs.contoso.com/federationmetadata/2007-06/federationmetadata.xml`.

Step 4: Create a relying party trust and custom claim rules in AD FS for Outlook on the web and the EAC

- On the Exchange server, Outlook on the web uses the virtual directory named `owa` and the EAC uses the virtual directory named `ecp`.

- The trailing slash (`/`) that's used in the Outlook on the web and EAC URL values is intentional. It's important that the AD FS relying party trusts and Exchange Audience URI's **are identical**. They **both must have** or **both must omit** the trailing slashes in their URLs. The examples in this section contain the trailing slashes after the owa and ecp URLs (`owa/` and `ecp/`).

- In organizations with multiple Active Directory sites that use separate namespaces (for example, `eu.contoso.com` and `na.contoso.com`), you need to configure relying party trusts for each namespace for both Outlook on the web and the EAC.

Step 4a: Create relying party trusts in AD FS for Outlook on the web and the EAC

To create the relying party trusts on the AD FS server, you can use the AD FS Management console or Windows PowerShell.

To use the AD FS Management console to create the relying party trusts, follow these steps:

Note: You need to go through these steps twice: once for Outlook on the web, and once for the EAC. The only difference is the values that you enter in steps 5 and 8 (the **Specify Display Name** and **Configure URL** pages in the wizard).

1. In **Server Manager**, click **Tools**, and then select **AD FS Management**.

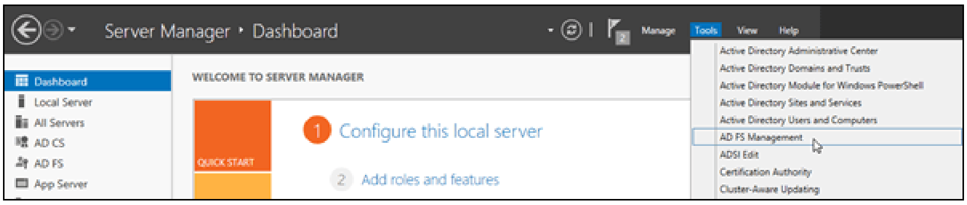
Server Manager • Dashboard
Manage
Tools
View
Help
Dashboard
Local Server
All Servers
AD CS
AD FS
App Server
WELCOME TO SERVER MANAGER
QUICK START
1 Configure this local server
2 Add roles and features
Active Directory Administrative Center
Active Directory Domains and Trusts
Active Directory Module for Windows PowerShell
Active Directory Sites and Services
Active Directory Users and Computers
AD FS Management
ADSI Edit
Certification Authority
Cluster-Aware Updating

2. In the AD FS Management console, expand **Trust Relationships** and then select **Relying Party Trusts**. In the **Actions** pane, select **Add Relying Party Trust**.

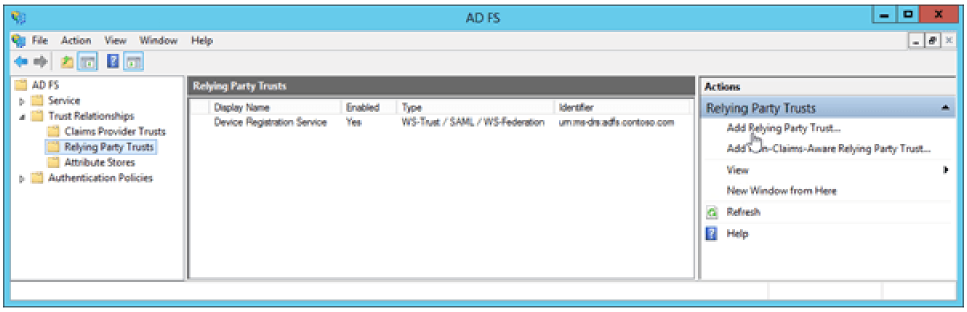

Display Name	Enabled	Type	Identifier
Device Registration Service	Yes	WS-Trust / SAML / WS-Federation	urn:ms-drs:adfs.contoso.com

3. The **Add Relying Party Trust Wizard** opens. On the **Welcome** page, click **Start**.

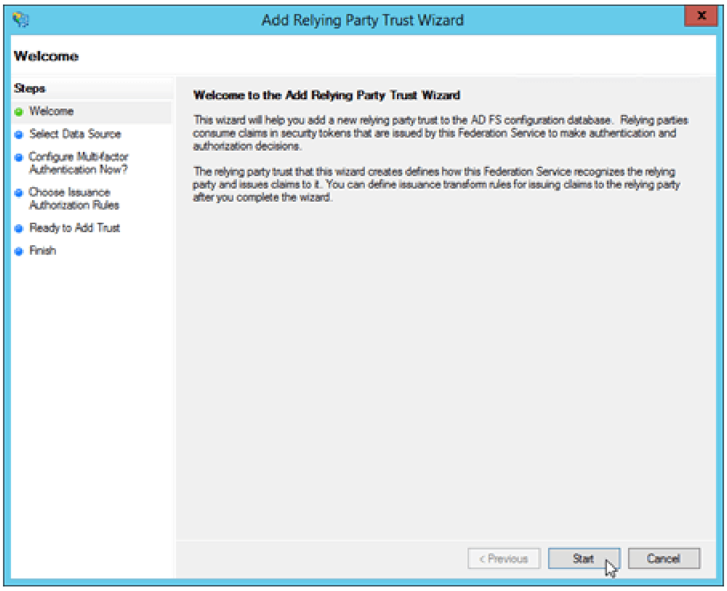
Add Relying Party Trust Wizard
Welcome
Steps
Welcome
Select Data Source
Configure Multi-factor Authentication Now?
Choose Issuance Authorization Rules
Ready to Add Trust
Finish
< Previous
Start
Cancel

Welcome to the Add Relying Party Trust Wizard

This wizard will help you add a new relying party trust to the AD FS configuration database. Relying parties consume claims in security tokens that are issued by this Federation Service to make authentication and authorization decisions.

The relying party trust that this wizard creates defines how this Federation Service recognizes the relying party and issues claims to it. You can define issuance transform rules for issuing claims to the relying party after you complete the wizard.

4. On the **Select Data Source** page, select **Enter data about the relying party manually**, and then click

Next.

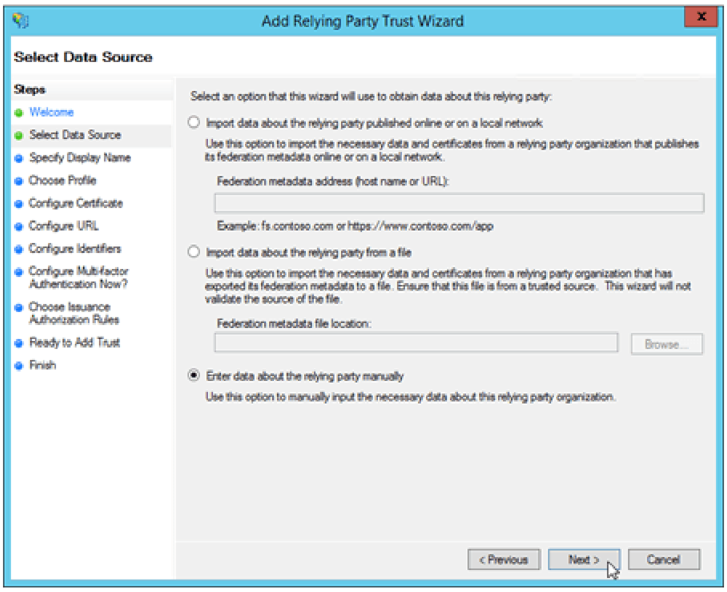

Select Data Source

Select an option that this wizard will use to obtain data about this relying party:

○ Import data about the relying party published online or on a local network

Use this option to import the necessary data and certificates from a relying party organization that publishes its federation metadata online or on a local network.

Federation metadata address (host name or URL):

Example: fs.contoso.com or https://www.contoso.com/app

○ Import data about the relying party from a file

Use this option to import the necessary data and certificates from a relying party organization that has exported its federation metadata to a file. Ensure that this file is from a trusted source. This wizard will not validate the source of the file.

Federation metadata file location:

● Enter data about the relying party manually

Use this option to manually input the necessary data about this relying party organization.

5. On the **Specify Display Name** page, configure the following settings:

- **For Outlook on the web**:

- **Display Name**: Type Outlook on the web.

- **Notes**: Enter a description. For example, This is a trust for https://mail.contoso.com/owa/.

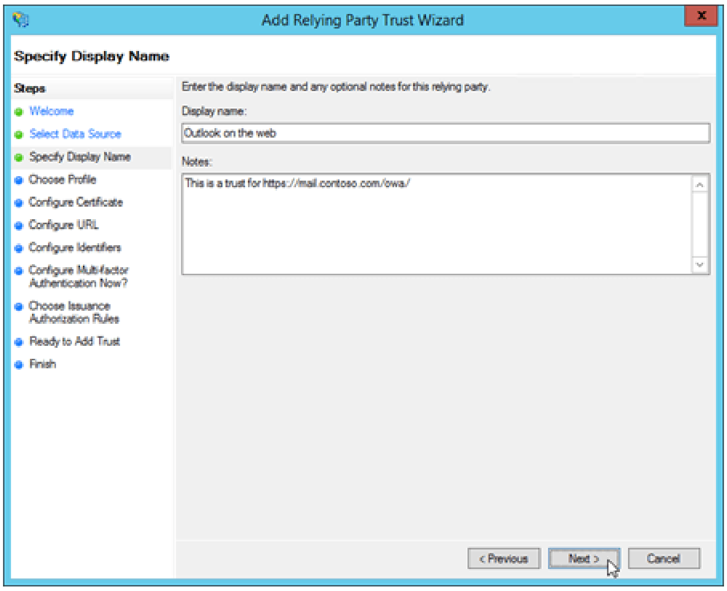

Specify Display Name

Enter the display name and any optional notes for this relying party.

Display name:

Outlook on the web

Notes:

This is a trust for https://mail.contoso.com/owa/

- **For the EAC**:

- **Display Name**: Type EAC.

- **Notes**: Enter a description. For example, This is a trust for https://mail.contoso.com/ecp/.

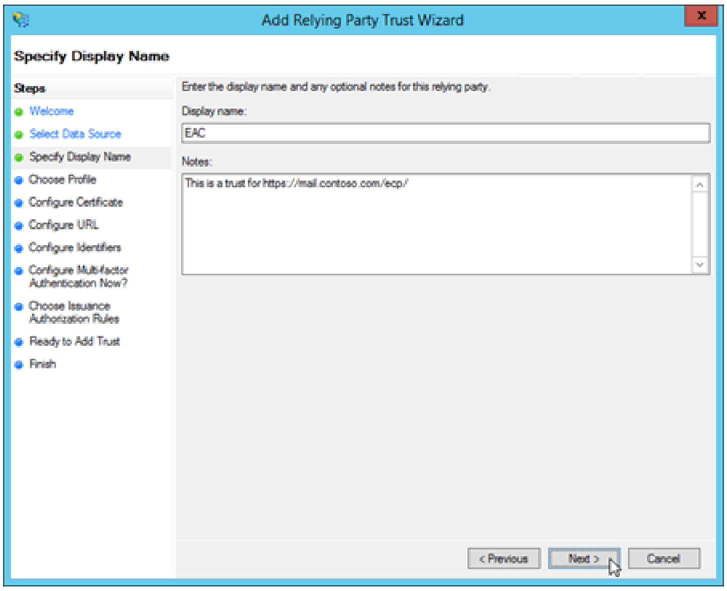

Add Relying Party Trust Wizard

Specify Display Name

Steps
- Welcome
- Select Data Source
- Specify Display Name
- Choose Profile
- Configure Certificate
- Configure URL
- Configure Identifiers
- Configure Multi-factor Authentication Now?
- Choose Issuance Authorization Rules
- Ready to Add Trust
- Finish

Enter the display name and any optional notes for this relying party.

Display name:
EAC

Notes:
This is a trust for https://mail.contoso.com/ecp/

< Previous | Next > | Cancel

When you're finished, click **Next**.

6. On the **Choose Profile** page, verify **AD FS profile** is selected, and then click **Next**.

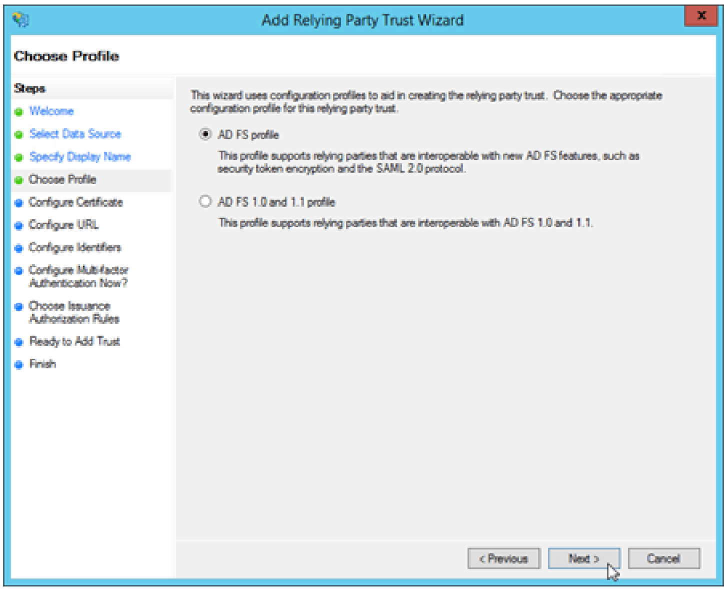

Add Relying Party Trust Wizard

Choose Profile

Steps
- Welcome
- Select Data Source
- Specify Display Name
- Choose Profile
- Configure Certificate
- Configure URL
- Configure Identifiers
- Configure Multi-factor Authentication Now?
- Choose Issuance Authorization Rules
- Ready to Add Trust
- Finish

This wizard uses configuration profiles to aid in creating the relying party trust. Choose the appropriate configuration profile for this relying party trust.

● AD FS profile
This profile supports relying parties that are interoperable with new AD FS features, such as security token encryption and the SAML 2.0 protocol.

○ AD FS 1.0 and 1.1 profile
This profile supports relying parties that are interoperable with AD FS 1.0 and 1.1.

< Previous | Next > | Cancel

7. On the **Configure Certificate** page, click **Next** (don't specify an optional token encryption certificate).

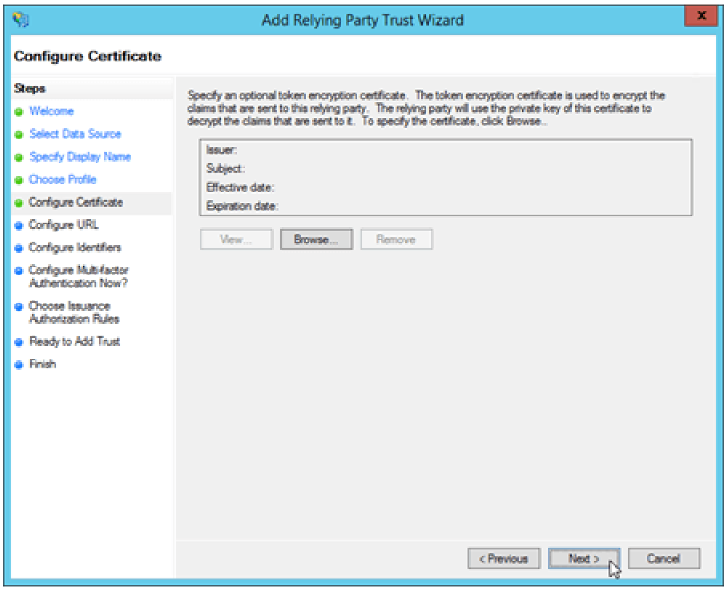

Add Relying Party Trust Wizard

Configure Certificate

Steps
- Welcome
- Select Data Source
- Specify Display Name
- Choose Profile
- Configure Certificate
- Configure URL
- Configure Identifiers
- Configure Multi-factor Authentication Now?
- Choose Issuance Authorization Rules
- Ready to Add Trust
- Finish

Specify an optional token encryption certificate. The token encryption certificate is used to encrypt the claims that are sent to this relying party. The relying party will use the private key of this certificate to decrypt the claims that are sent to it. To specify the certificate, click Browse...

Issuer:
Subject:
Effective date:
Expiration date:

View... | Browse... | Remove

< Previous | Next > | Cancel

8. On the **Configure URL** page, select **Enable support for the WS-Federation Passive protocol**, and in

Relying party WS-Federation Passive protocol URL, enter the following information:

- **Outlook on the web**: Type your external Outlook on the web URL (for example, https://mail.contoso.com/owa/).

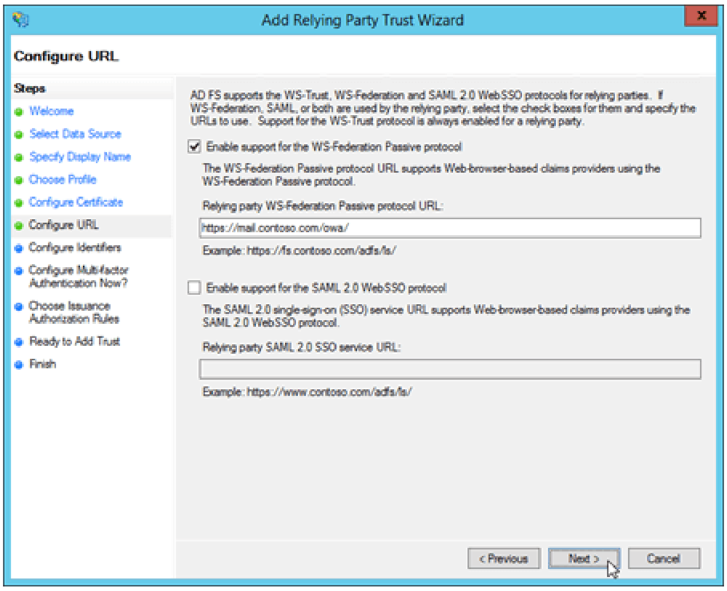

Configure URL

AD FS supports the WS-Trust, WS-Federation and SAML 2.0 WebSSO protocols for relying parties. If WS-Federation, SAML, or both are used by the relying party, select the check boxes for them and specify the URLs to use. Support for the WS-Trust protocol is always enabled for a relying party.

☑ Enable support for the WS-Federation Passive protocol

The WS-Federation Passive protocol URL supports Web-browser-based claims providers using the WS-Federation Passive protocol.

Relying party WS-Federation Passive protocol URL:

https://mail.contoso.com/owa/

Example: https://fs.contoso.com/adfs/ls/

☐ Enable support for the SAML 2.0 WebSSO protocol

The SAML 2.0 single-sign-on (SSO) service URL supports Web-browser-based claims providers using the SAML 2.0 WebSSO protocol.

Relying party SAML 2.0 SSO service URL:

Example: https://www.contoso.com/adfs/ls/

- **EAC**: Type your external EAC URL (for example, https://mail.contoso.com/ecp/).

When you're finished, click **Next**.

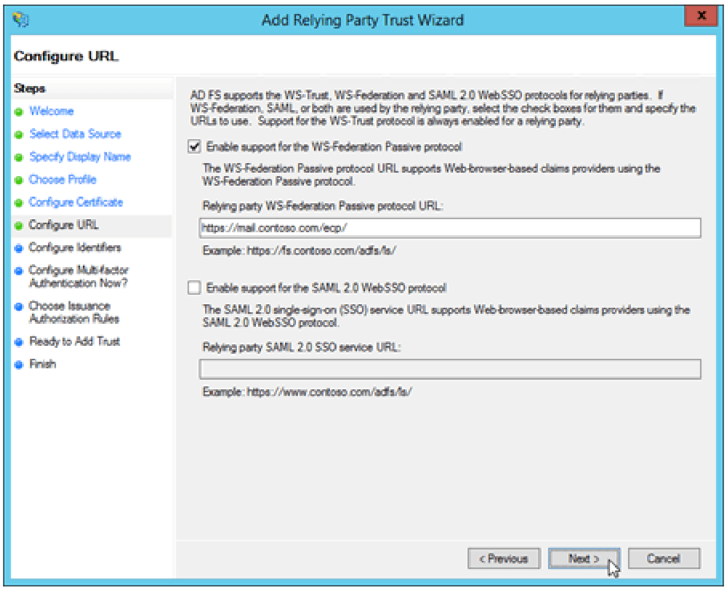

Configure URL

AD FS supports the WS-Trust, WS-Federation and SAML 2.0 WebSSO protocols for relying parties. If WS-Federation, SAML, or both are used by the relying party, select the check boxes for them and specify the URLs to use. Support for the WS-Trust protocol is always enabled for a relying party.

☑ Enable support for the WS-Federation Passive protocol

The WS-Federation Passive protocol URL supports Web-browser-based claims providers using the WS-Federation Passive protocol.

Relying party WS-Federation Passive protocol URL:

https://mail.contoso.com/ecp/

Example: https://fs.contoso.com/adfs/ls/

☐ Enable support for the SAML 2.0 WebSSO protocol

The SAML 2.0 single-sign-on (SSO) service URL supports Web-browser-based claims providers using the SAML 2.0 WebSSO protocol.

Relying party SAML 2.0 SSO service URL:

Example: https://www.contoso.com/adfs/ls/

9. On the **Configure Identifiers** page, click **Next** (the URL from the previous step is listed in **Relying party trust identifiers**).

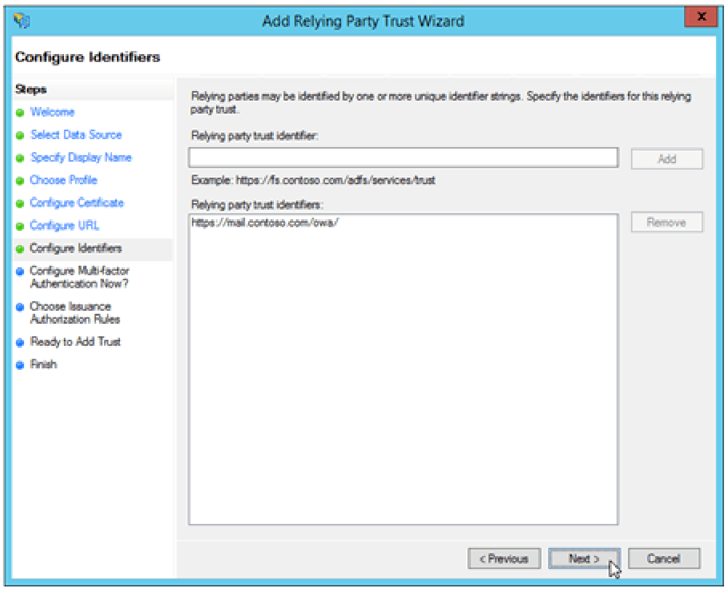

Configure Identifiers

Relying parties may be identified by one or more unique identifier strings. Specify the identifiers for this relying party trust.

Relying party trust identifier:

Example: https://fs.contoso.com/adfs/services/trust

Relying party trust identifiers:

https://mail.contoso.com/owa/

10. On the **Configure Multi-factor Authentication Now?** page, verify that **I do not want to configure multi-factor authentication settings for this relying party trust at this time** is selected, and then click **Next**.

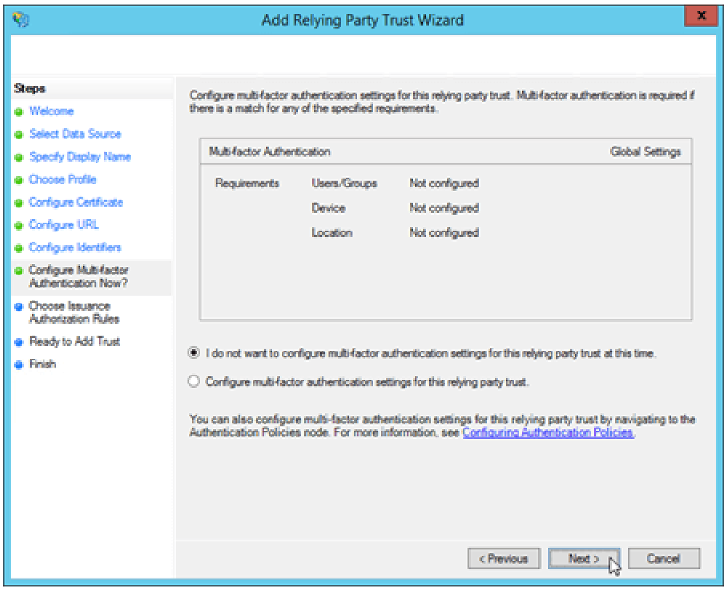

Configure multi-factor authentication settings for this relying party trust. Multi-factor authentication is required if there is a match for any of the specified requirements.

Multi-factor Authentication		Global Settings
Requirements	Users/Groups	Not configured
	Device	Not configured
	Location	Not configured

● I do not want to configure multi-factor authentication settings for this relying party trust at this time.

○ Configure multi-factor authentication settings for this relying party trust.

You can also configure multi-factor authentication settings for this relying party trust by navigating to the Authentication Policies node. For more information, see Configuring Authentication Policies.

11. On the **Choose Issuance Authorization Rules** page, verify **Permit all users to access this relying party** is selected, and then click **Next**.

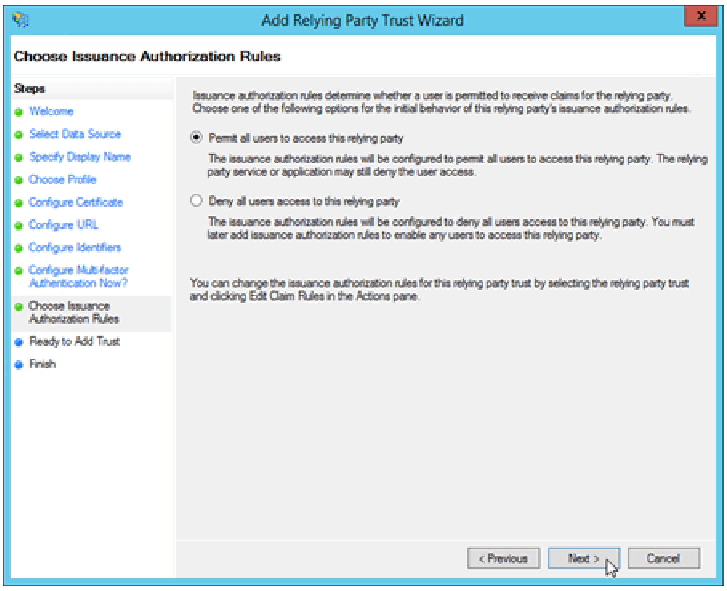

Choose Issuance Authorization Rules

Issuance authorization rules determine whether a user is permitted to receive claims for the relying party. Choose one of the following options for the initial behavior of this relying party's issuance authorization rules.

● Permit all users to access this relying party

The issuance authorization rules will be configured to permit all users to access this relying party. The relying party service or application may still deny the user access.

○ Deny all users access to this relying party

The issuance authorization rules will be configured to deny all users access to this relying party. You must later add issuance authorization rules to enable any users to access this relying party.

You can change the issuance authorization rules for this relying party trust by selecting the relying party trust and clicking Edit Claim Rules in the Actions pane.

12. On the **Ready to Add Trust** page, review the settings, and then click **Next** to save your relying party trust information.

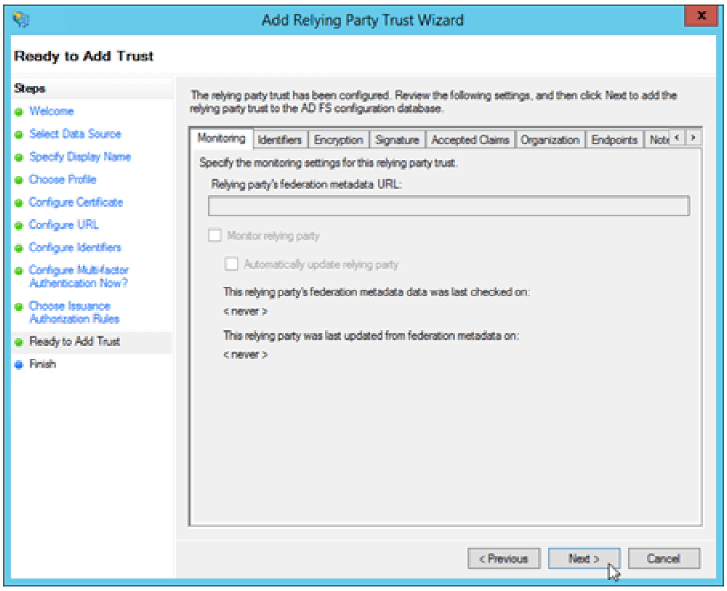

The relying party trust has been configured. Review the following settings, and then click Next to add the relying party trust to the AD FS configuration database.

Specify the monitoring settings for this relying party trust.

Relying party's federation metadata URL:

□ Monitor relying party

□ Automatically update relying party

This relying party's federation metadata data was last checked on:
< never >

This relying party was last updated from federation metadata on:
< never >

13. On the **Finish** page, uncheck **Open the Edit Claim Rules dialog for this relying party trust when the wizard closes**, and then click **Close**.

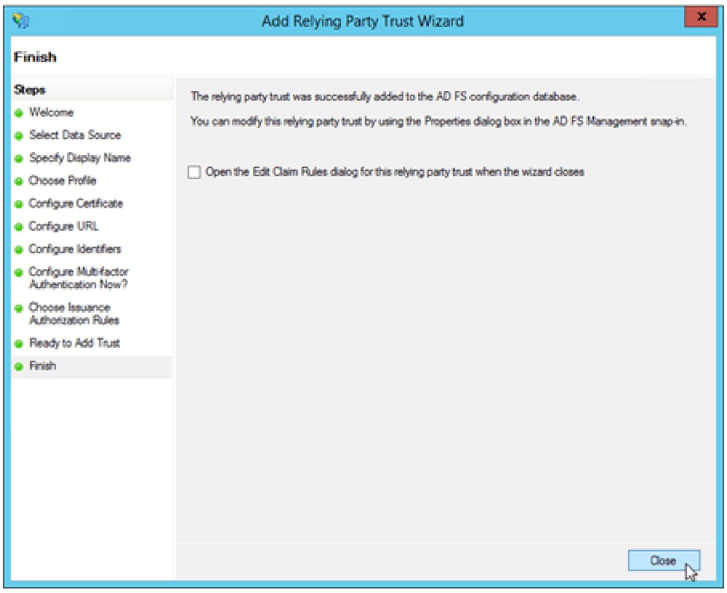

The relying party trust was successfully added to the AD FS configuration database.

You can modify this relying party trust by using the Properties dialog box in the AD FS Management snap-in.

□ Open the Edit Claim Rules dialog for this relying party trust when the wizard closes

To use Windows PowerShell prompt to create the relying party trusts, follow these steps:

1. In an elevated Windows PowerShell window, run the following command:

```
Import-Module ADFS
```

2. Use the following syntax:

```
Add-AdfsRelyingPartyTrust -Name <"Outlook on the web" | EAC> -Notes "This is a trust for <OotwURL |
EACURL>" -Identifier <OotwURL | EACURL> -WSFedEndpoint <OotwURL | EACURL> -IssuanceAuthorizationRules
'@RuleTemplate = "AllowAllAuthzRule" => issue(Type =
"http://schemas.microsoft.com/authorization/claims/permit", Value = "true");' -IssueOAuthRefreshTokensTo
NoDevice
```

This example creates a relying party trust for Outlook on the web using the following values:

- **Name**: Outlook on the web

- **Notes**: This is a trust for https://mail.contoso.com/owa/
- **Identifier**: https://mail.contoso.com/owa/
- **WSFedEndpoint**: https://mail.contoso.com/owa/

```
Add-AdfsRelyingPartyTrust -Name "Outlook on the web" -Notes "This is a trust for
https://mail.contoso.com/owa/" -Identifier https://mail.contoso.com/owa/ -WSFedEndpoint
https://mail.contoso.com/owa/ -IssuanceAuthorizationRules '@RuleTemplate = "AllowAllAuthzRule" => issue(Type =
"http://schemas.microsoft.com/authorization/claims/permit", Value = "true");' -IssueOAuthRefreshTokensTo
NoDevice
```

This example creates a relying party trust for the EAC using the following values:

- **Name**: EAC
- **Notes**: This is a trust for https://mail.contoso.com/ecp/
- **Identifier**: https://mail.contoso.com/ecp/
- **WSFedEndpoint**: https://mail.contoso.com/ecp/

```
Add-AdfsRelyingPartyTrust -Name EAC -Notes "This is a trust for https://mail.contoso.com/ecp/" -Identifier
https://mail.contoso.com/ecp/ -WSFedEndpoint https://mail.contoso.com/ecp/ -IssuanceAuthorizationRules
'@RuleTemplate = "AllowAllAuthzRule" => issue(Type =
"http://schemas.microsoft.com/authorization/claims/permit", Value = "true");' -IssueOAuthRefreshTokensTo
NoDevice
```

Step 4b: Create custom claim rules in AD FS for Outlook on the web and the EAC

For both Outlook on the web and the EAC, you need to create two claim rules:

- Active Directory user SID
- Active Directory UPN

To create the claim rules on the AD FS server, you can use the AD FS Management console or Windows PowerShell.

To use the AD FS Management console to create the claim rules, follow these steps:

Note: You need to go through these steps twice: once for Outlook on the web, and once for EAC. The only difference is the relying party trust that you select in the first step. All other values in the procedure are identical.

To add the required claims rules:

1. In the AD FS Management console, expand **Trust Relationships** select **Relying Party Trusts**, and then select the Outlook on the web or EAC relying party trust. In the **Actions** pane, select **Edit Claim Rules**.

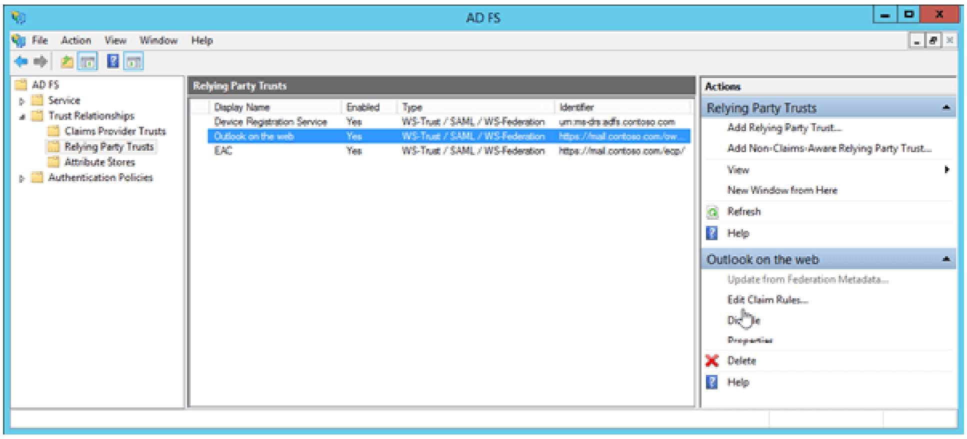

Display Name	Enabled	Type	Identifier
Device Registration Service	Yes	WS-Trust / SAML / WS-Federation	urn:ms-drs:adfs.contoso.com
Outlook on the web	Yes	WS-Trust / SAML / WS-Federation	https://mail.contoso.com/ow...
EAC	Yes	WS-Trust / SAML / WS-Federation	https://mail.contoso.com/ecp/

2. In the **Edit Claim Rules for <RuleName>** window that opens, verify that the **Issuance Transform Rules** tab is selected, and then click **Add Rule**.

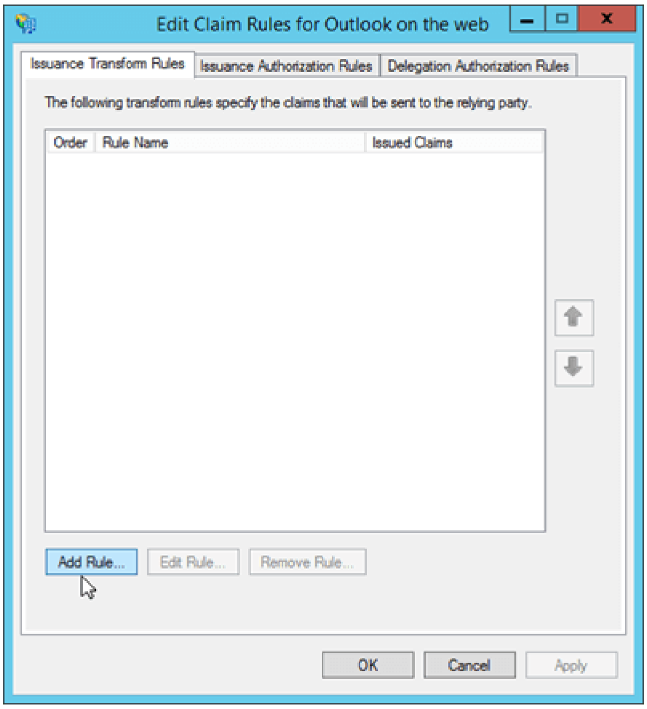

The following transform rules specify the claims that will be sent to the relying party.

3. The **Add Transform Claim Rule Wizard** opens. On the **Select Rule Template** page, click the **Claim rule template** drop down, and then select **Send Claims Using a Custom Rule**. When you're finished, click **Next**.

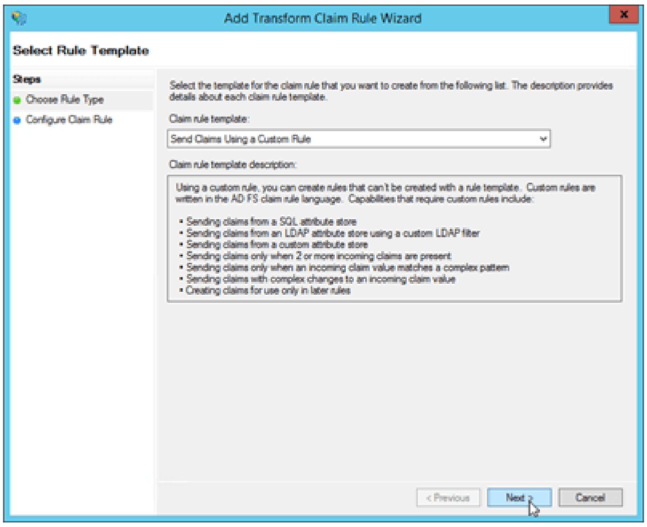

Select the template for the claim rule that you want to create from the following list. The description provides details about each claim rule template.

Claim rule template: Send Claims Using a Custom Rule

Claim rule template description:

Using a custom rule, you can create rules that can't be created with a rule template. Custom rules are written in the AD FS claim rule language. Capabilities that require custom rules include:

- Sending claims from a SQL attribute store
- Sending claims from an LDAP attribute store using a custom LDAP filter
- Sending claims from a custom attribute store
- Sending claims only when 2 or more incoming claims are present
- Sending claims only when an incoming claim value matches a complex pattern
- Sending claims with complex changes to an incoming claim value
- Creating claims for use only in later rules

4. On the **Configure Rule** page, enter the following information:

   - **Claim rule name**: Enter a descriptive name for the claim rule. For example, ActiveDirectoryUserSID.

   - **Custom rule**: Copy and paste the following text:

```
c:[Type == "http://schemas.microsoft.com/ws/2008/06/identity/claims/windowsaccountname", Issuer
== "AD AUTHORITY"] => issue(store = "Active Directory", types =
("http://schemas.microsoft.com/ws/2008/06/identity/claims/primarysid"), query = ";objectSID;{0}",
param = c.Value);
```

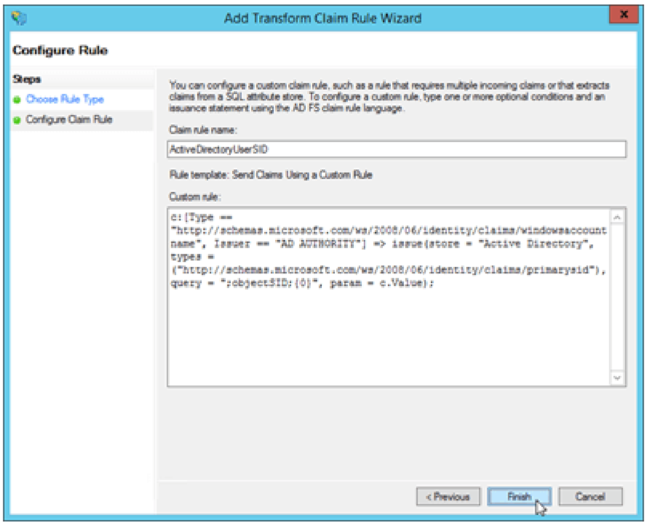

You can configure a custom claim rule, such as a rule that requires multiple incoming claims or that extracts claims from a SQL attribute store. To configure a custom rule, type one or more optional conditions and an issuance statement using the AD FS claim rule language.

Claim rule name:

ActiveDirectoryUserSID

Rule template: Send Claims Using a Custom Rule

Custom rule:

```
c:[Type ==
"http://schemas.microsoft.com/ws/2008/06/identity/claims/windowsaccount
name", Issuer == "AD AUTHORITY"] => issue(store = "Active Directory",
types =
("http://schemas.microsoft.com/ws/2008/06/identity/claims/primarysid"),
query = ";objectSID;{0}", param = c.Value);
```

When you're finished, click **Finish**.

5. Back on the **Edit Claim Rules for <RuleName>** window, verify that the **Issuance Transform Rules** tab is selected, and then click **Add Rule**.

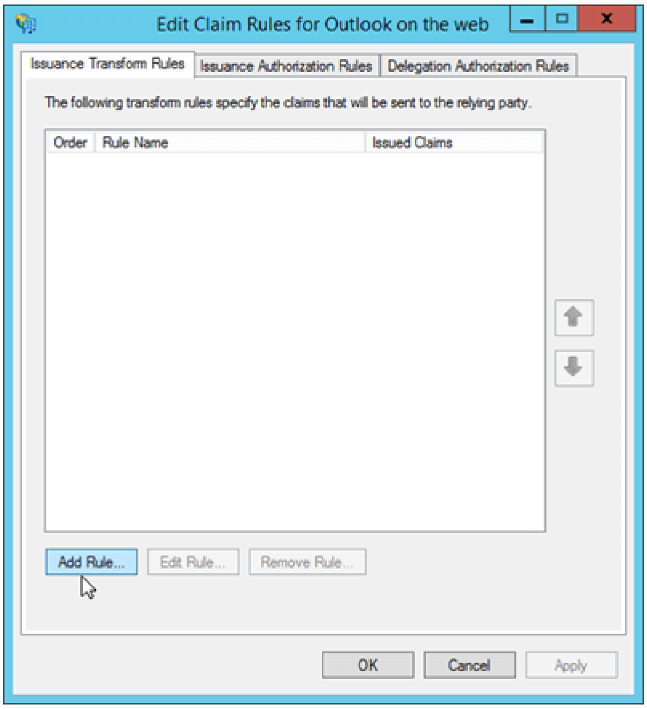
Edit Claim Rules for Outlook on the web
Issuance Transform Rules
Issuance Authorization Rules
Delegation Authorization Rules

The following transform rules specify the claims that will be sent to the relying party.

Order	Rule Name	Issued Claims

Add Rule...
Edit Rule...
Remove Rule...
OK
Cancel
Apply

6. The **Add Transform Claim Rule Wizard** opens. On the **Select Rule Template** page, click the **Claim rule template** drop down, and then select **Send Claims Using a Custom Rule**. When you're finished, click **Next**.

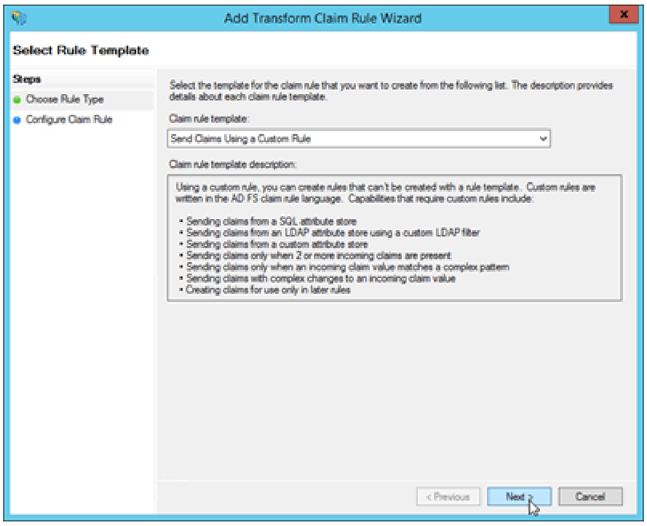

Select the template for the claim rule that you want to create from the following list. The description provides details about each claim rule template.

Claim rule template:

Send Claims Using a Custom Rule

Claim rule template description:

Using a custom rule, you can create rules that can't be created with a rule template. Custom rules are written in the AD FS claim rule language. Capabilities that require custom rules include:

- Sending claims from a SQL attribute store
- Sending claims from an LDAP attribute store using a custom LDAP filter
- Sending claims from a custom attribute store
- Sending claims only when 2 or more incoming claims are present
- Sending claims only when an incoming claim value matches a complex pattern
- Sending claims with complex changes to an incoming claim value
- Creating claims for use only in later rules

7. On the **Configure Rule** page, enter the following information:

   - **Claim rule name**: Enter a descriptive name for the claim rule. For example, ActiveDirectoryUPN.

   - **Custom rule**: Copy and paste the following text:

   ```
   c:[Type == "http://schemas.microsoft.com/ws/2008/06/identity/claims/windowsaccountname", Issuer
   == "AD AUTHORITY"] => issue(store = "Active Directory", types =
   ("http://schemas.xmlsoap.org/ws/2005/05/identity/claims/upn"), query = ";userPrincipalName;{0}",
   param = c.Value);
   ```

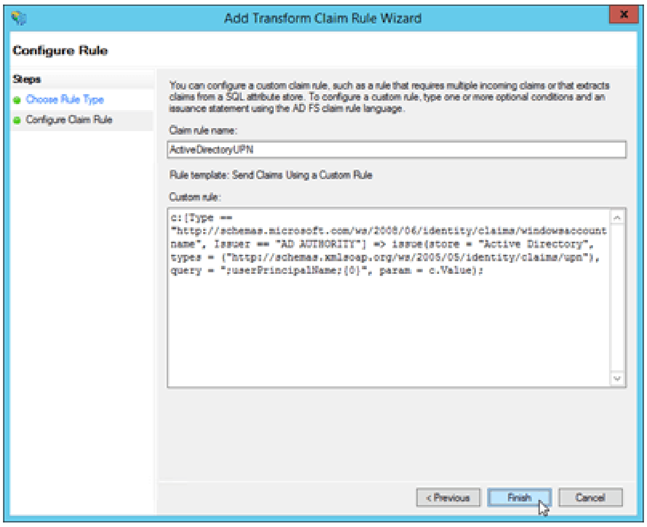

You can configure a custom claim rule, such as a rule that requires multiple incoming claims or that extracts claims from a SQL attribute store. To configure a custom rule, type one or more optional conditions and an issuance statement using the AD FS claim rule language.

Claim rule name:

ActiveDirectoryUPN

Rule template: Send Claims Using a Custom Rule

Custom rule:

```
c:[Type ==
"http://schemas.microsoft.com/ws/2008/06/identity/claims/windowsaccount
name", Issuer == "AD AUTHORITY"] => issue(store = "Active Directory",
types = ("http://schemas.xmlsoap.org/ws/2005/05/identity/claims/upn"),
query = ";userPrincipalName;{0}", param = c.Value);
```

   When you're finished, click **Finish**.

8. Back on the **Edit Claim Rules for <RuleName>** window, click **OK**.

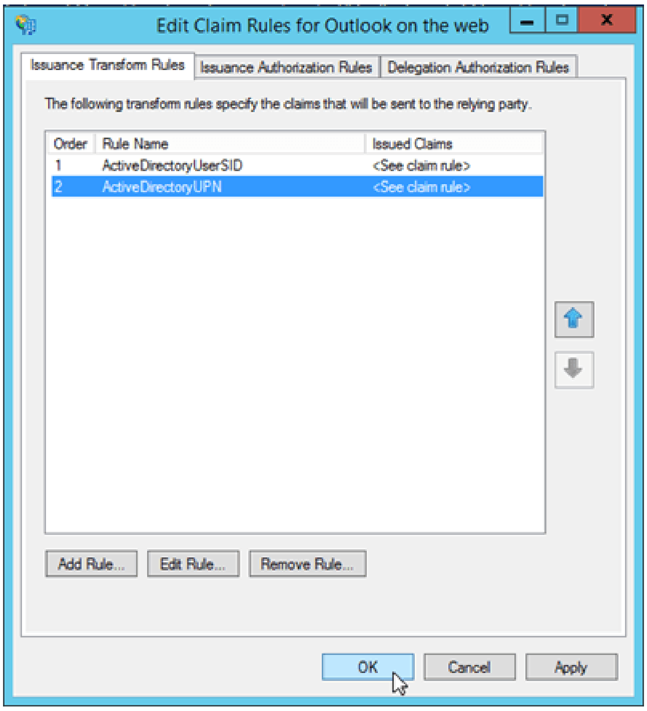
Edit Claim Rules for Outlook on the web
Issuance Transform Rules
Issuance Authorization Rules
Delegation Authorization Rules
Add Rule...
Edit Rule...
Remove Rule...
OK
Cancel
Apply

The following transform rules specify the claims that will be sent to the relying party.

Order	Rule Name	Issued Claims
1	ActiveDirectoryUserSID	<See claim rule>
2	ActiveDirectoryUPN	<See claim rule>

To use Windows PowerShell to create the custom claim rules, follow these steps:

1. Open an elevated Windows PowerShell window, and run the following command:

```
Import-Module ADFS
```

2. Use the following syntax:

```
Set-AdfsRelyingPartyTrust -TargetName <OotwRelyingPartyTrust | EACRelyingPartyTrust> -
IssuanceTransformRules '@RuleName = "ActiveDirectoryUserSID" c:[Type ==
"http://schemas.microsoft.com/ws/2008/06/identity/claims/windowsaccountname", Issuer == "AD AUTHORITY"]
=> issue(store = "Active Directory", types =
("http://schemas.microsoft.com/ws/2008/06/identity/claims/primarysid"), query = ";objectSID;{0}", param
= c.Value); @RuleName = "ActiveDirectoryUPN" c:[Type ==
"http://schemas.microsoft.com/ws/2008/06/identity/claims/windowsaccountname", Issuer == "AD AUTHORITY"]
=> issue(store = "Active Directory", types =
("http://schemas.xmlsoap.org/ws/2005/05/identity/claims/upn"), query = ";userPrincipalName;{0}", param =
c.Value);'
```

To create the custom claim rules in the existing relying party trust named Outlook on the web, run the following command:

```
Set-AdfsRelyingPartyTrust -TargetName "Outlook on the web" -IssuanceTransformRules '@RuleName =
"ActiveDirectoryUserSID" c:[Type ==
"http://schemas.microsoft.com/ws/2008/06/identity/claims/windowsaccountname", Issuer == "AD AUTHORITY"] =>
issue(store = "Active Directory", types =
("http://schemas.microsoft.com/ws/2008/06/identity/claims/primarysid"), query = ";objectSID;{0}", param =
c.Value); @RuleName = "ActiveDirectoryUPN" c:[Type ==
"http://schemas.microsoft.com/ws/2008/06/identity/claims/windowsaccountname", Issuer == "AD AUTHORITY"] =>
issue(store = "Active Directory", types = ("http://schemas.xmlsoap.org/ws/2005/05/identity/claims/upn"), query
= ";userPrincipalName;{0}", param = c.Value);'
```

To create the custom claim rules in the existing relying party trust named EAC, run the following command:

```
Set-AdfsRelyingPartyTrust -TargetName EAC -IssuanceTransformRules '@RuleName = "ActiveDirectoryUserSID" c:
[Type == "http://schemas.microsoft.com/ws/2008/06/identity/claims/windowsaccountname", Issuer == "AD
AUTHORITY"] => issue(store = "Active Directory", types =
("http://schemas.microsoft.com/ws/2008/06/identity/claims/primarysid"), query = ";objectSID;{0}", param =
c.Value); @RuleName = "ActiveDirectoryUPN" c:[Type ==
"http://schemas.microsoft.com/ws/2008/06/identity/claims/windowsaccountname", Issuer == "AD AUTHORITY"] =>
issue(store = "Active Directory", types = ("http://schemas.xmlsoap.org/ws/2005/05/identity/claims/upn"), query
= ";userPrincipalName;{0}", param = c.Value);'
```

Step 5: (Optional) Deploy and configure a Windows Server 2012 R2 Web Application Proxy server

The steps in this section are required only if you want to publish Outlook on the web and the EAC using Web Application Proxy, and you want Web Application Proxy perform the AD FS authentication. Remember:

- You can't use offline access in Outlook on the web if you use AD FS authentication through Web Application Proxy.
- You can't install Web Application Proxy on the same server where AD FS is installed.

If you aren't going to use Web Application Proxy, skip to Step 6.

Step 5a: Install Web Application Proxy

To use Server Manager to install Web Application Proxy, follow these steps:

1. On the target server, open **Server Manager**, click **Manage**, and then select **Add Roles and Features**.

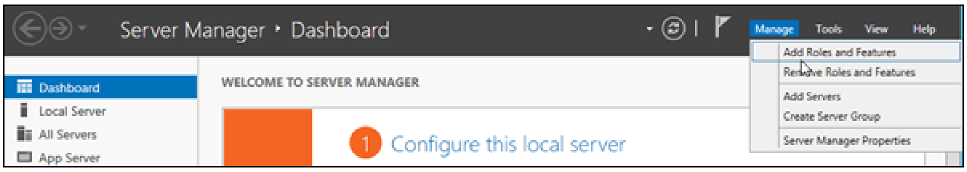
Server Manager • Dashboard
Manage
Tools
View
Help
Add Roles and Features
Remove Roles and Features
Add Servers
Create Server Group
Server Manager Properties
Dashboard
Local Server
All Servers
App Server
WELCOME TO SERVER MANAGER
1
Configure this local server

2. The **Add Roles and Features Wizard** opens. You'll start on the **Before you begin** page unless you previously selected **Skip this page by default**. Click **Next**.

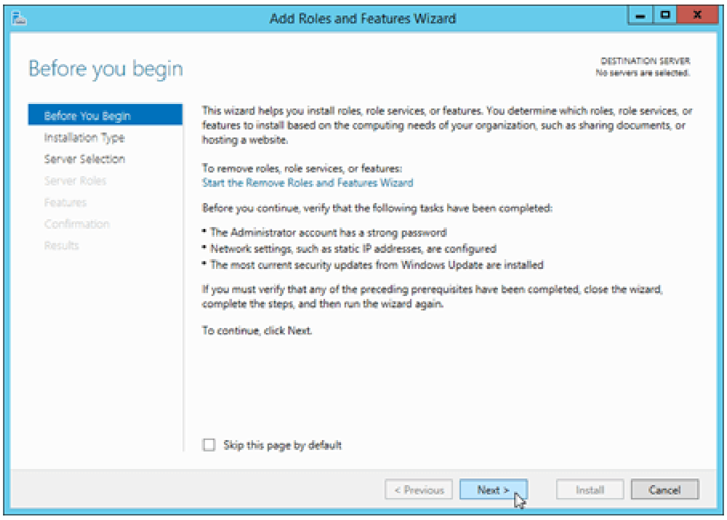

Before you begin

DESTINATION SERVER
No servers are selected.

This wizard helps you install roles, role services, or features. You determine which roles, role services, or features to install based on the computing needs of your organization, such as sharing documents, or hosting a website.

To remove roles, role services, or features:
Start the Remove Roles and Features Wizard

Before you continue, verify that the following tasks have been completed:

- The Administrator account has a strong password
- Network settings, such as static IP addresses, are configured
- The most current security updates from Windows Update are installed

If you must verify that any of the preceding prerequisites have been completed, close the wizard, complete the steps, and then run the wizard again.

To continue, click Next.

☐ Skip this page by default

3. On the **Select installation type** page, verify that **Role-based or feature-based installation** is selected, and then click **Next**.

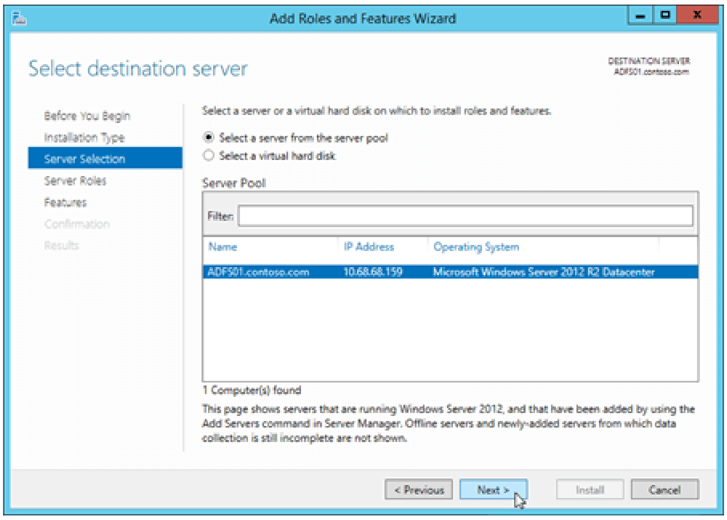

Select destination server

DESTINATION SERVER
ADFS01.contoso.com

Select a server or a virtual hard disk on which to install roles and features.

● Select a server from the server pool

○ Select a virtual hard disk

Server Pool

Filter:

Name	IP Address	Operating System
ADFS01.contoso.com	10.68.68.159	Microsoft Windows Server 2012 R2 Datacenter

1 Computer(s) found

This page shows servers that are running Windows Server 2012, and that have been added by using the Add Servers command in Server Manager. Offline servers and newly-added servers from which data collection is still incomplete are not shown.

4. On the **Select destination server** page, verify the server selection, and then click **Next**.

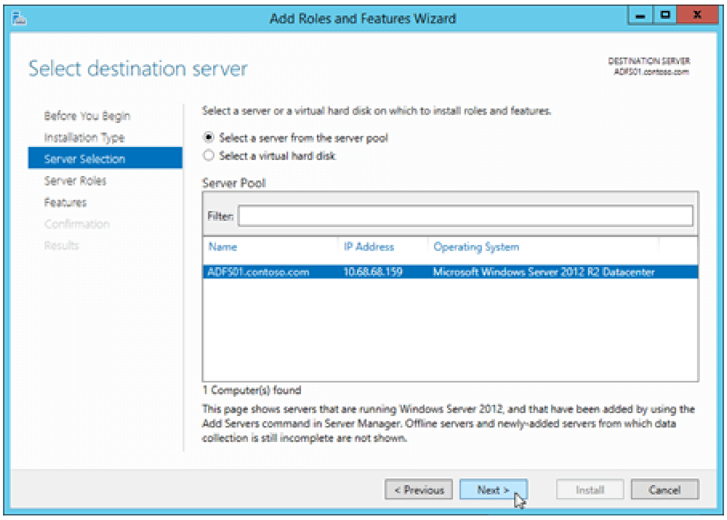

Select destination server

DESTINATION SERVER
ADFS01.contoso.com

Select a server or a virtual hard disk on which to install roles and features.

● Select a server from the server pool

○ Select a virtual hard disk

Server Pool

Filter:

Name	IP Address	Operating System
ADFS01.contoso.com	10.68.68.159	Microsoft Windows Server 2012 R2 Datacenter

1 Computer(s) found

This page shows servers that are running Windows Server 2012, and that have been added by using the Add Servers command in Server Manager. Offline servers and newly-added servers from which data collection is still incomplete are not shown.

5. On the **Select server roles** page, select **Remote Access** in the list of roles, and then click **Next**.

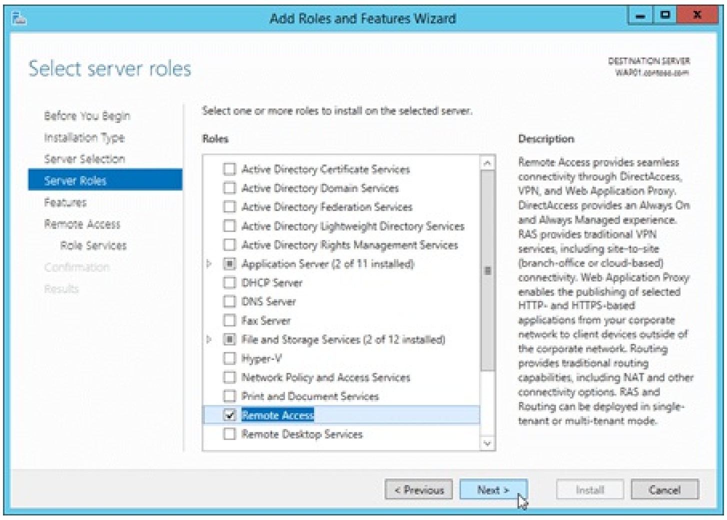

Select server roles

DESTINATION SERVER
WAP01.contoso.com

Select one or more roles to install on the selected server.

Roles

- □ Active Directory Certificate Services
- □ Active Directory Domain Services
- □ Active Directory Federation Services
- □ Active Directory Lightweight Directory Services
- □ Active Directory Rights Management Services
- ▣ Application Server (2 of 11 installed)
- □ DHCP Server
- □ DNS Server
- □ Fax Server
- ▣ File and Storage Services (2 of 12 installed)
- □ Hyper-V
- □ Network Policy and Access Services
- □ Print and Document Services
- ☑ Remote Access
- □ Remote Desktop Services

Description

Remote Access provides seamless connectivity through DirectAccess, VPN, and Web Application Proxy. DirectAccess provides an Always On and Always Managed experience. RAS provides traditional VPN services, including site-to-site (branch-office or cloud-based) connectivity. Web Application Proxy enables the publishing of selected HTTP- and HTTPS-based applications from your corporate network to client devices outside of the corporate network. Routing provides traditional routing capabilities, including NAT and other connectivity options. RAS and Routing can be deployed in single-tenant or multi-tenant mode.

6. On the **Features** page, click **Next** (accept the default feature selections).

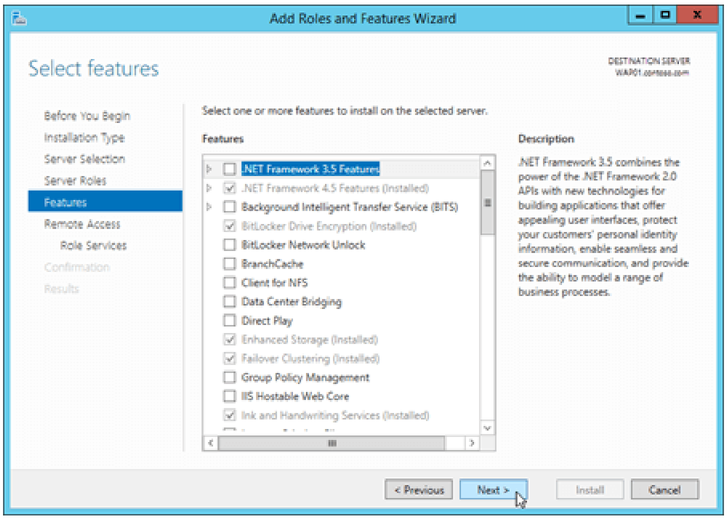

Select one or more features to install on the selected server.

Features

- ☐ .NET Framework 3.5 Features
- ☑ .NET Framework 4.5 Features (Installed)
- ☐ Background Intelligent Transfer Service (BITS)
- ☑ BitLocker Drive Encryption (Installed)
- ☐ BitLocker Network Unlock
- ☐ BranchCache
- ☐ Client for NFS
- ☐ Data Center Bridging
- ☐ Direct Play
- ☑ Enhanced Storage (Installed)
- ☑ Failover Clustering (Installed)
- ☐ Group Policy Management
- ☐ IIS Hostable Web Core
- ☑ Ink and Handwriting Services (Installed)

Description

.NET Framework 3.5 combines the power of the .NET Framework 2.0 APIs with new technologies for building applications that offer appealing user interfaces, protect your customers' personal identity information, enable seamless and secure communication, and provide the ability to model a range of business processes.

7. On the **Remote Access** page, read the information, and then click **Next**.

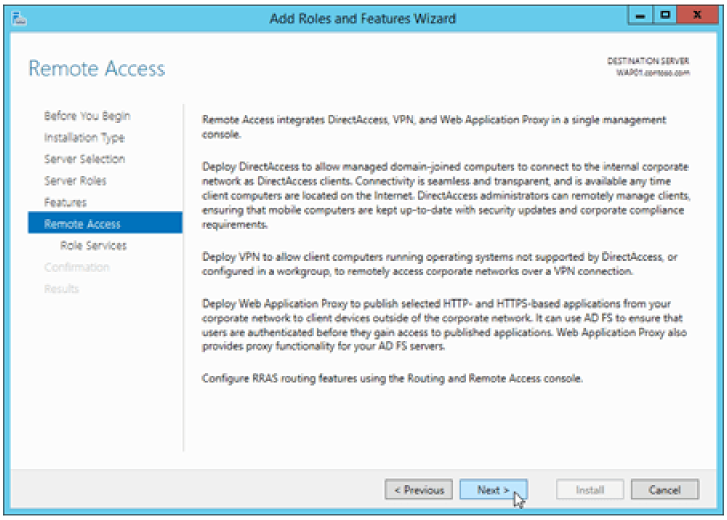

Remote Access integrates DirectAccess, VPN, and Web Application Proxy in a single management console.

Deploy DirectAccess to allow managed domain-joined computers to connect to the internal corporate network as DirectAccess clients. Connectivity is seamless and transparent, and is available any time client computers are located on the Internet. DirectAccess administrators can remotely manage clients, ensuring that mobile computers are kept up-to-date with security updates and corporate compliance requirements.

Deploy VPN to allow client computers running operating systems not supported by DirectAccess, or configured in a workgroup, to remotely access corporate networks over a VPN connection.

Deploy Web Application Proxy to publish selected HTTP- and HTTPS-based applications from your corporate network to client devices outside of the corporate network. It can use AD FS to ensure that users are authenticated before they gain access to published applications. Web Application Proxy also provides proxy functionality for your AD FS servers.

Configure RRAS routing features using the Routing and Remote Access console.

8. On the **Select role services** page, select **Web Application Proxy**. In the add features dialog box that opens, click **Add Features** to accept the default values and close the dialog box. Back on the **Select role services** page, click **Next**.

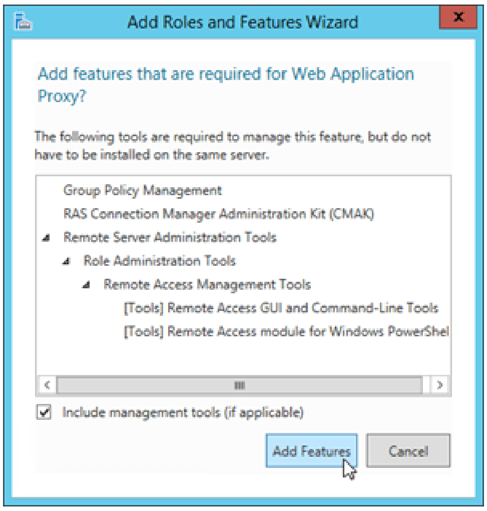

Add features that are required for Web Application Proxy?

The following tools are required to manage this feature, but do not have to be installed on the same server.

- Group Policy Management
- RAS Connection Manager Administration Kit (CMAK)
- Remote Server Administration Tools
  - Role Administration Tools
    - Remote Access Management Tools
      - [Tools] Remote Access GUI and Command-Line Tools
      - [Tools] Remote Access module for Windows PowerShel

☑ Include management tools (if applicable)

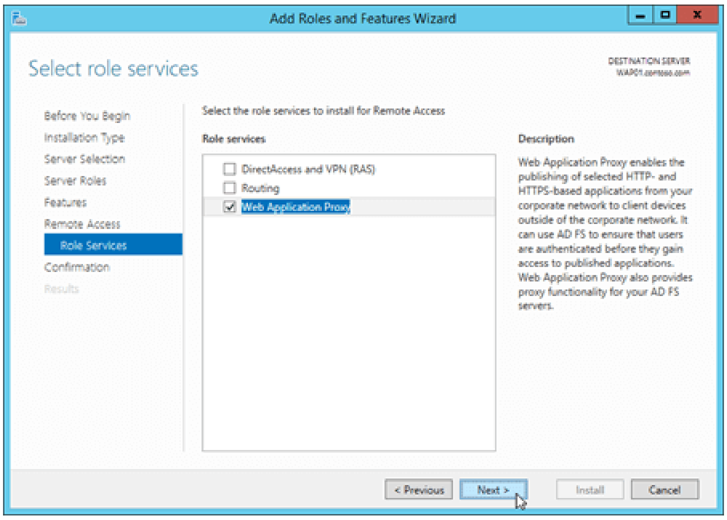

Select role services

DESTINATION SERVER
WAP01.contoso.com

Select the role services to install for Remote Access

Role services

☐ DirectAccess and VPN (RAS)
☐ Routing
☑ Web Application Proxy

Description

Web Application Proxy enables the publishing of selected HTTP- and HTTPS-based applications from your corporate network to client devices outside of the corporate network. It can use AD FS to ensure that users are authenticated before they gain access to published applications. Web Application Proxy also provides proxy functionality for your AD FS servers.

9. On the **Confirm installation selections** page, click **Install**.

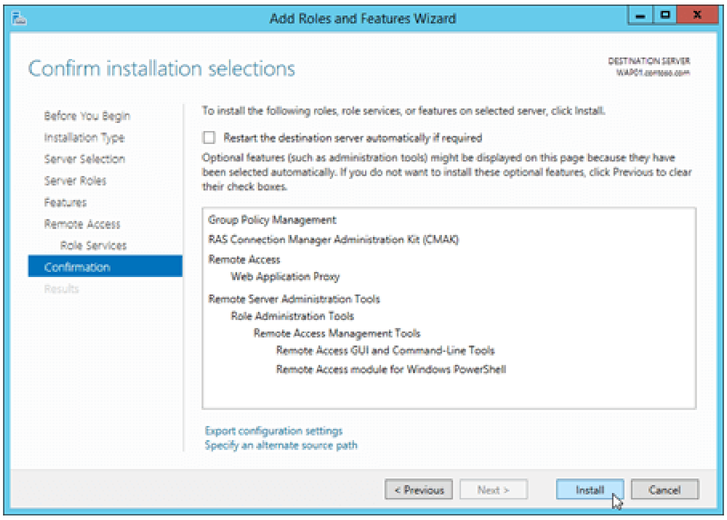

Confirm installation selections

DESTINATION SERVER
WAP01.contoso.com

To install the following roles, role services, or features on selected server, click Install.

☐ Restart the destination server automatically if required

Optional features (such as administration tools) might be displayed on this page because they have been selected automatically. If you do not want to install these optional features, click Previous to clear their check boxes.

```
Group Policy Management
RAS Connection Manager Administration Kit (CMAK)
Remote Access
    Web Application Proxy
Remote Server Administration Tools
    Role Administration Tools
        Remote Access Management Tools
            Remote Access GUI and Command-Line Tools
            Remote Access module for Windows PowerShell
```

10. On the **Installation progress** page, watch the progress bar to verify that the installation was successful. When the installation is finished, leave the wizard open so you can click **Open the Web Application Proxy Wizard** in the next step (5b).

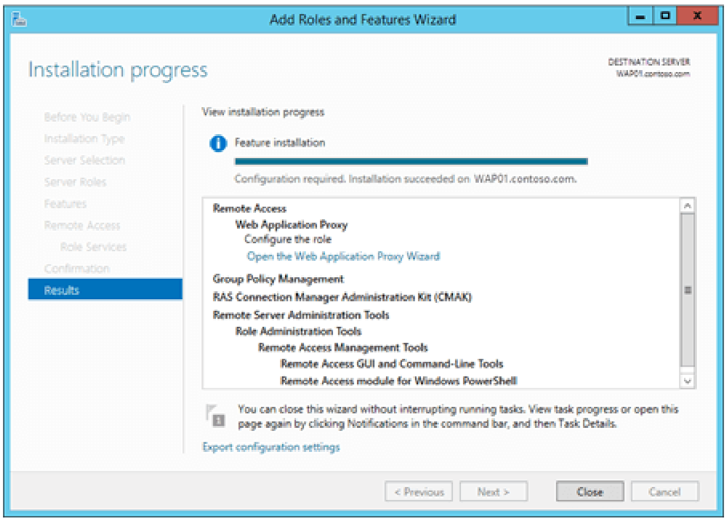

Installation progress

DESTINATION SERVER
WAP01.contoso.com

View installation progress

Feature installation

Configuration required. Installation succeeded on WAP01.contoso.com.

```
Remote Access
    Web Application Proxy
        Configure the role
        Open the Web Application Proxy Wizard
Group Policy Management
RAS Connection Manager Administration Kit (CMAK)
Remote Server Administration Tools
    Role Administration Tools
        Remote Access Management Tools
            Remote Access GUI and Command-Line Tools
            Remote Access module for Windows PowerShell
```

You can close this wizard without interrupting running tasks. View task progress or open this page again by clicking Notifications in the command bar, and then Task Details.

To use Windows PowerShell to install Web Application Proxy, run the following command:

```
Install-WindowsFeature Web-Application-Proxy -IncludeManagementTools
```

Step 5b: Configure the Web Application Proxy server

After you deploy the Web Application Proxy server, you need to configure the following Web Application Proxy settings:

- **Federation service name**: For example, `adfs.contoso.com`.
- **Federation service trust credential**: The username and password of a local administrator account on the AD FS server.
- **AD FS Proxy Certificate**: A certificate that's installed on the Web Application Proxy server that identifies the server to clients as a proxy for the Federation Service, and therefore contains the federation service name (for example, `adfs.contoso.com`). Also, the federation service name must be accessible to the Web Application Proxy server (resolvable in DNS).

You can use Server Manager or Windows PowerShell to configure the Web Application Proxy server.

To use Server Manager to configure Web Application Proxy, follow these steps:

1. If you left the **Add Roles and Features Wizard** open on the Web Application Proxy server from the previous step, you can click the **Open the Web Application Proxy Wizard** link on the **Installation progress** page.

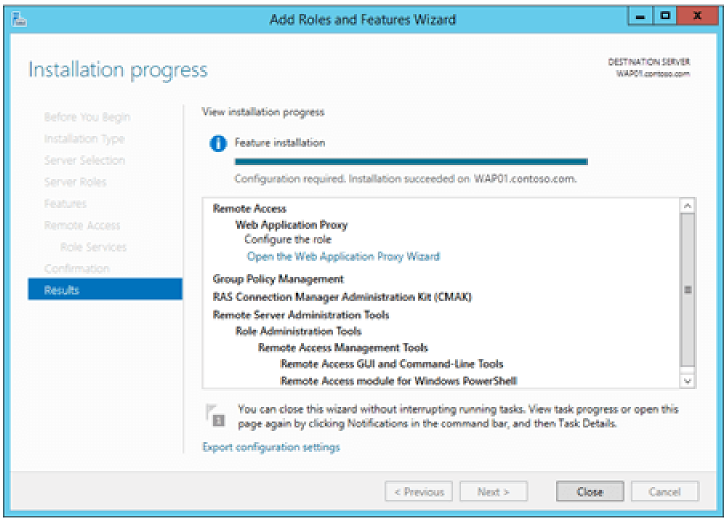

Add Roles and Features Wizard

Installation progress

DESTINATION SERVER
WAP01.contoso.com

View installation progress

Feature installation

Configuration required. Installation succeeded on WAP01.contoso.com.

Remote Access
  Web Application Proxy
    Configure the role
    Open the Web Application Proxy Wizard
Group Policy Management
RAS Connection Manager Administration Kit (CMAK)
Remote Server Administration Tools
  Role Administration Tools
    Remote Access Management Tools
      Remote Access GUI and Command-Line Tools
      Remote Access module for Windows PowerShell

You can close this wizard without interrupting running tasks. View task progress or open this page again by clicking Notifications in the command bar, and then Task Details.

Export configuration settings

If you closed the **Add Roles and Features Wizard** or you used Windows PowerShell to install Web Application Proxy, you can get to the same place by clicking **Notifications**, and then clicking **Open the Web Application Proxy Wizard** in the **Post-deployment Configuration** warning.

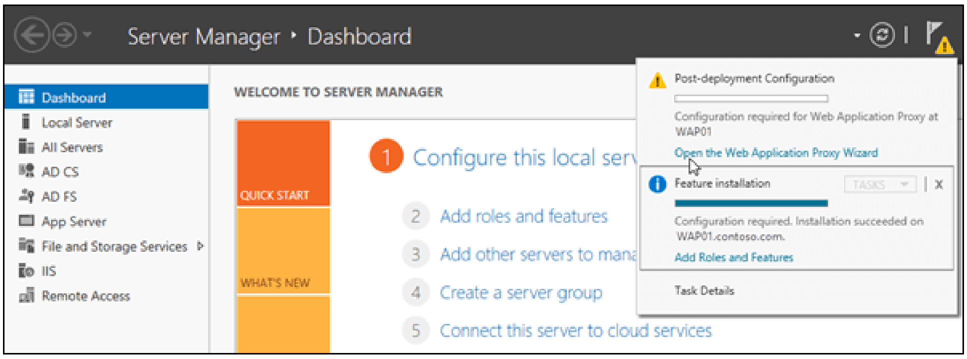

Server Manager • Dashboard

WELCOME TO SERVER MANAGER

Post-deployment Configuration

Configuration required for Web Application Proxy at WAP01

Open the Web Application Proxy Wizard

Feature installation

Configuration required. Installation succeeded on WAP01.contoso.com.

Add Roles and Features

Task Details

2. The **Web Application Proxy Configuration Wizard** opens. On the **Welcome** page, click **Next**.

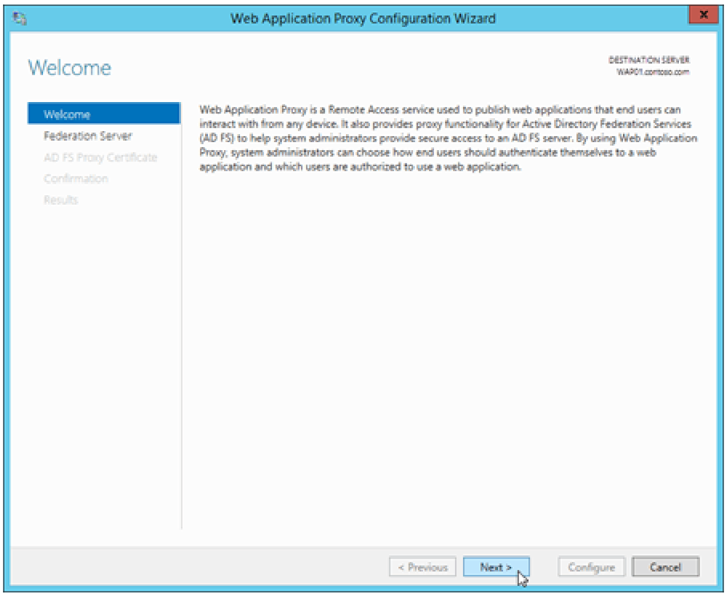

Web Application Proxy is a Remote Access service used to publish web applications that end users can interact with from any device. It also provides proxy functionality for Active Directory Federation Services (AD FS) to help system administrators provide secure access to an AD FS server. By using Web Application Proxy, system administrators can choose how end users should authenticate themselves to a web application and which users are authorized to use a web application.

3. On the **Federation Server** page, enter the following information:

   - **Federation service name**: For example, `adfs.contoso.com`.

   - **User name** and **Password**: Type the credentials of a local administrator account on the AD FS server.

   When you're finished, click **Next**.

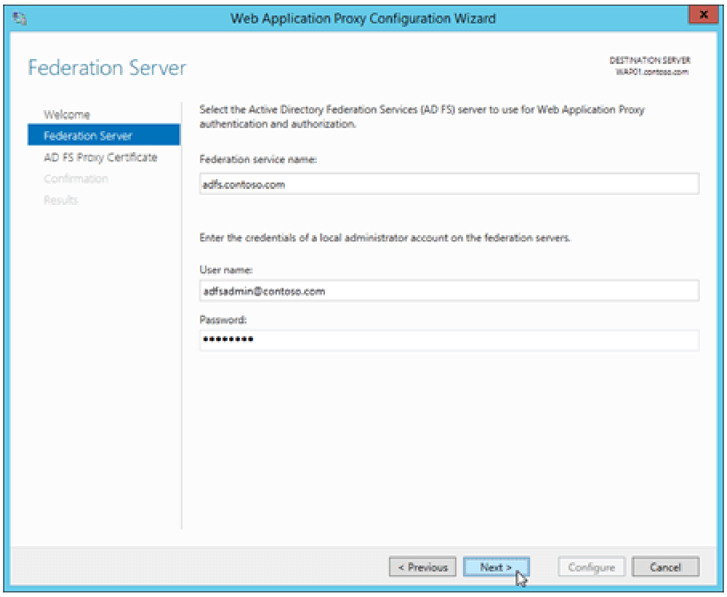

Select the Active Directory Federation Services (AD FS) server to use for Web Application Proxy authentication and authorization.

Federation service name:
adfs.contoso.com

Enter the credentials of a local administrator account on the federation servers.

User name:
adfsadmin@contoso.com

Password:
•••••••••

4. On the **AD FS Proxy Certificate** page, select an installed certificate that contains the federation service name (for example `adfs.contoso.com`). You can select a certificate in the drop down list, and then click **View > Details** to see more information about the certificate. When you're finished, click **Next**.

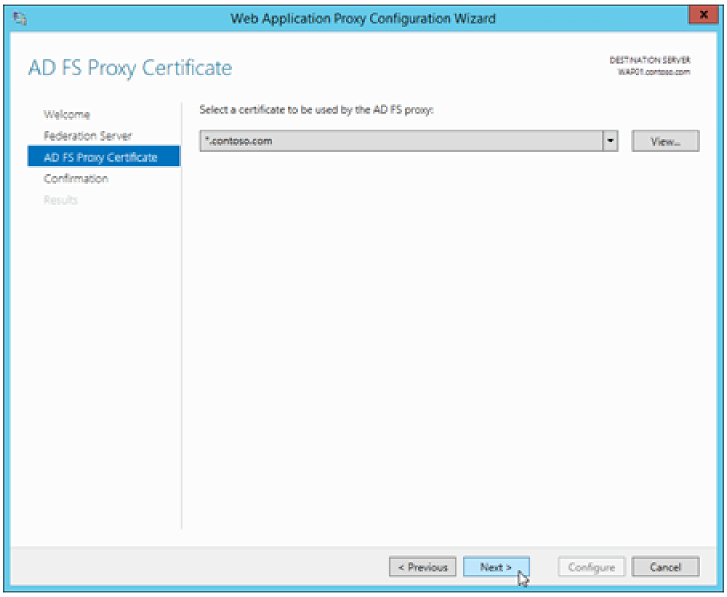

Select a certificate to be used by the AD FS proxy:

*.contoso.com

5. On the **Confirmation** page, review the settings. You can copy the Windows PowerShell command to automate additional installations (in particular, the certificate thumbprint value). When you're finished, click **Configure**.

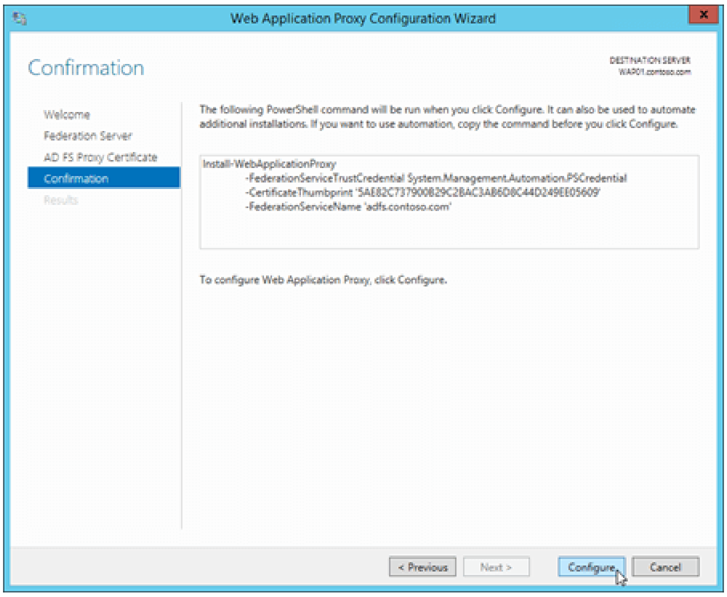

The following PowerShell command will be run when you click Configure. It can also be used to automate additional installations. If you want to use automation, copy the command before you click Configure.

```
Install-WebApplicationProxy
        -FederationServiceTrustCredential System.Management.Automation.PSCredential
        -CertificateThumbprint '5AE82C737900B29C2BAC3AB6D8C44D249EE05609'
        -FederationServiceName 'adfs.contoso.com'
```

To configure Web Application Proxy, click Configure.

6. On the **Results** page, verify that the configuration was successful, and then click **Close**.

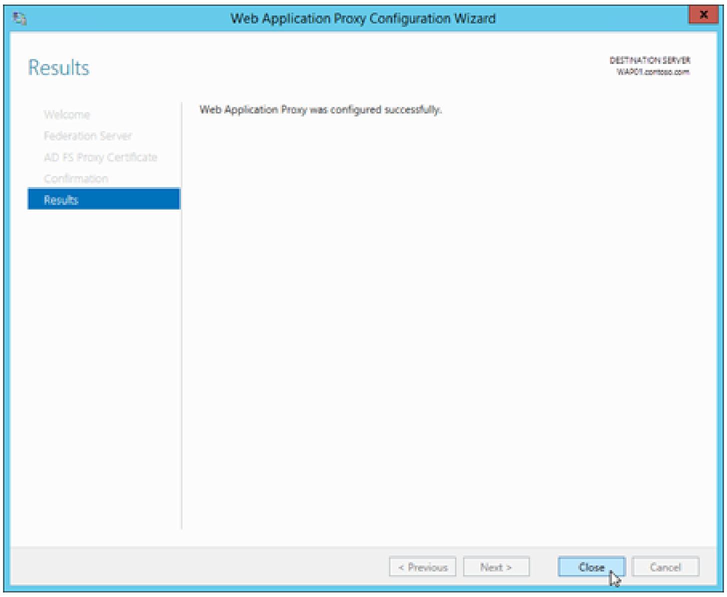

Web Application Proxy was configured successfully.

To use Windows PowerShell to configure Web Application Proxy, follow these steps:

1. Run the following command on the Web Application Proxy server to find the thumbprint value of the installed certificate that contains `adfs.contoso.com`:

   ```
   Set-Location Cert:\LocalMachine\My; Get-ChildItem | Format-List FriendlyName,Subject,Thumbprint
   ```

2. Run the following command, and enter the username and password of a local administrator account on the AD FS server.

   ```
   $ADFSServerCred = Get-Credential
   ```

3. Use the following syntax:

   ```
   Install-WebApplicationProxy -FederationServiceName <FederationServiceName> -
   FederationServiceTrustCredential $ADFSServerCred -CertificateThumprint <ADFSCertThumbprint>
   ```

   This example configure the Web Application Proxy server with the following settings:

   - **Federation service name**: `adfs.contoso.com`

   - **AD FS SSL certificate thumbprint**: The `*.contoso.com` certificate that has the thumbprint value `5AE82C737900B29C2BAC3AB6D8C44D249EE05609`.

   ```
   Install-WebApplicationProxy -FederationServiceName adfs.contoso.com -FederationServiceTrustCredential
   $ADFSServerCred -CertificateThumprint 5AE82C737900B29C2BAC3AB6D8C44D249EE05609
   ```

Step 5c: Publish the claims relying party trusts for Outlook on the web and the EAC in Web Application Proxy

To publish the relying party trusts in Web Application Proxy, you can use the Remote Access Management console or Windows PowerShell.

To use the Remote Access Management console, follow these steps:

Note: You need to go through these steps twice: once for Outlook on the web, and once for EAC. The required settings are described in the procedure.

1. Open the Remote Access Management console on the Web Application Proxy server: in Server Manager, click **Tools** > **Remote Access Management**.

2. In the Remote Access Management console, under **Configuration**, click **Web Application Proxy**, and then in the **Tasks** pane, click **Publish**.

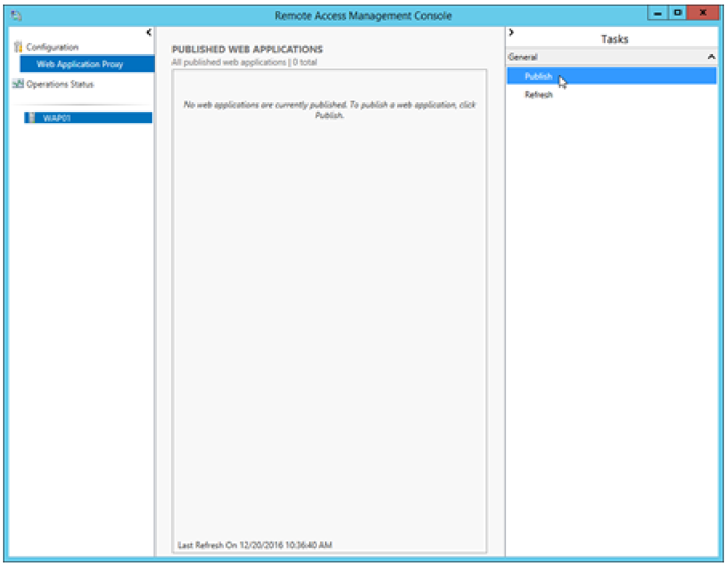
Remote Access Management Console
Configuration
Web Application Proxy
Operations Status
WAP01
PUBLISHED WEB APPLICATIONS
All published web applications | 0 total
Tasks
General
Publish
Refresh

No web applications are currently published. To publish a web application, click Publish.

Last Refresh On 12/20/2016 10:36:40 AM

3. The **Publish New Application Wizard** opens. On the **Welcome** page, click **Next**.

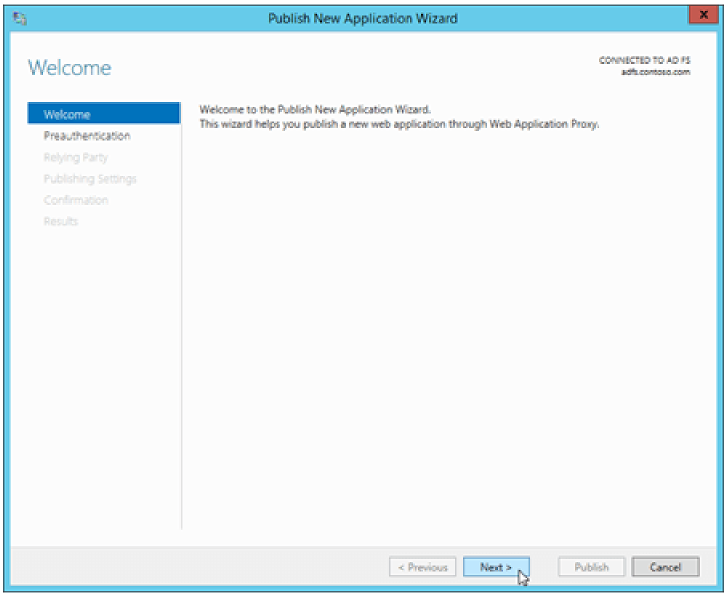

Welcome to the Publish New Application Wizard.
This wizard helps you publish a new web application through Web Application Proxy.

4. On the **Preauthentication** page, verify **Active Directory Federation Services (AD FS)** is selected, and then click **Next**.

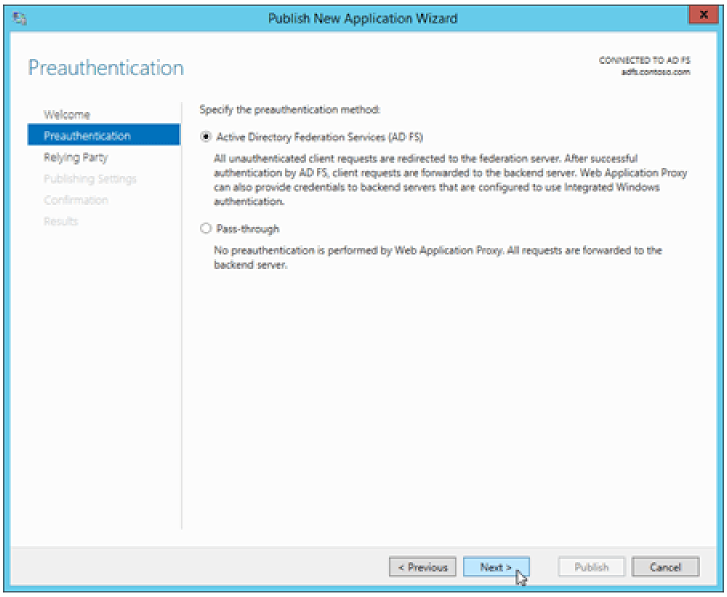

Specify the preauthentication method:

● Active Directory Federation Services (AD FS)

All unauthenticated client requests are redirected to the federation server. After successful authentication by AD FS, client requests are forwarded to the backend server. Web Application Proxy can also provide credentials to backend servers that are configured to use Integrated Windows authentication.

○ Pass-through

No preauthentication is performed by Web Application Proxy. All requests are forwarded to the backend server.

5. On the **Relying Party** page, select the relying party that you created on the AD FS server in Step 4: Create a relying party trust and custom claim rules in AD FS for Outlook on the web and the EAC:

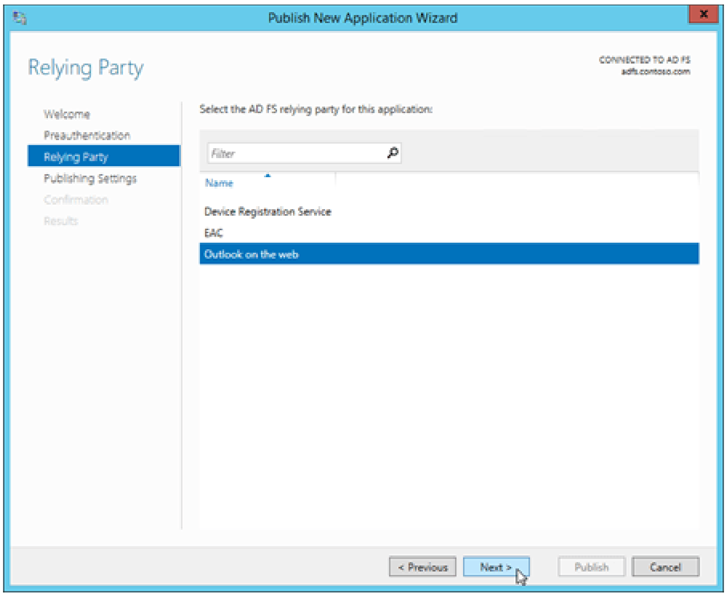

Relying Party

CONNECTED TO AD FS
adfs.contoso.com

Select the AD FS relying party for this application:

Name
Device Registration Service
EAC
Outlook on the web

- **For Outlook on the web**: Select Outlook on the web.
- **For the EAC**: Select EAC.

When you're finished, click **Next**.

6. On the **Publishing Settings** page, enter the following information:

- **For Outlook on the web**
  - **Name**: For example, `Outlook on the web`. This name is only visible in the Remote Access Management console.
  - **External URL**: For example, `https://mail.contoso.com/owa/`.
  - **External certificate**: Select an installed certificate that contains the host name of the external URL for Outlook on the web (for example, `mail.contoso.com`). You can select a certificate in the drop down list, and then click **View** > **Details** to see more information about the certificate.
  - **Backend server URL**: This value is automatically populated by the **External URL**. You only need to change it if the backend server URL is different from the external URL. For example, `https://server01.contoso.com/owa/`. Note that the paths in the external URL and backend server URL must match (`/owa/`), but the host name values can be different (for example, `mail.contoso.com` and `server01.contoso.com`).

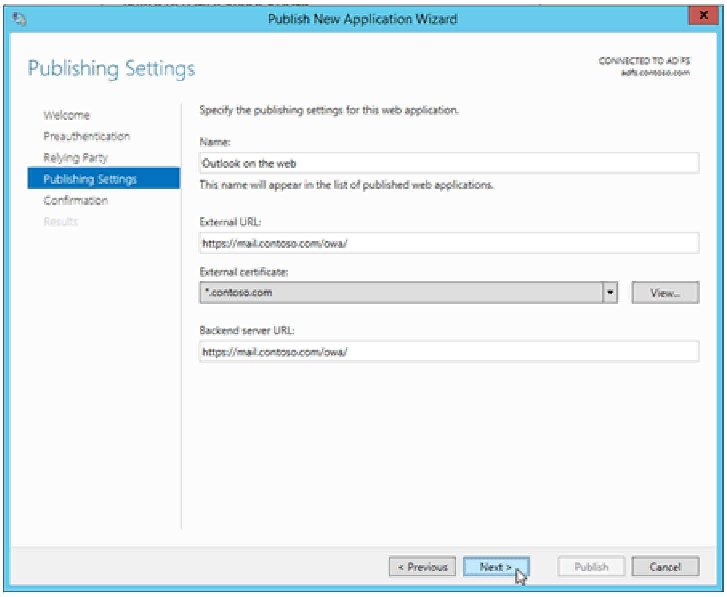

Publishing Settings

CONNECTED TO AD FS
adfs.contoso.com

Specify the publishing settings for this web application.

Name:
Outlook on the web
This name will appear in the list of published web applications.

External URL:
https://mail.contoso.com/owa/

External certificate:
*.contoso.com

Backend server URL:
https://mail.contoso.com/owa/

- For the EAC

    - **Name**: For example, `EAC`. This name is only visible in the Remote Access Management console.

    - **External URL**: The external URL for the EAC. For example, https://mail.contoso.com/ecp/.

    - **External certificate**: Select an installed certificate that contains the host name of the external URL for the EAC (for example, `mail.contoso.com`). The certificate is likely a wildcard certificate or SAN certificate. You can select a certificate in the drop down list, and then click **View** > **Details** to see more information about the certificate.

    - **Backend server URL**: This value is automatically populated by the **External URL**. You only need to change it if the backend server URL is different from the external URL. For example, `https://server01.contoso.com/ecp/`. Note that the paths in the external URL and backend server URL must match (`/ecp/`), but the host name values can be different (for example, `mail.contoso.com` and `server01.contoso.com`).

When you're finished, click **Next**.

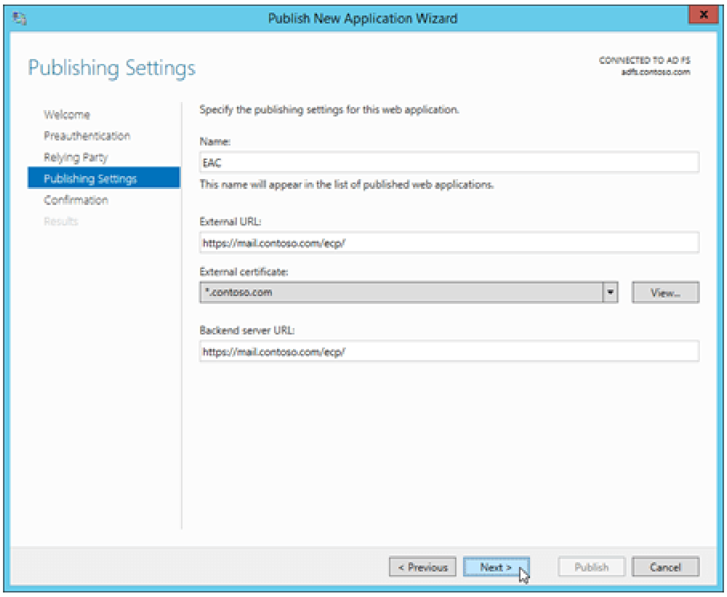

Specify the publishing settings for this web application.

Name:
EAC

This name will appear in the list of published web applications.

External URL:
https://mail.contoso.com/ecp/

External certificate:
*.contoso.com

Backend server URL:
https://mail.contoso.com/ecp/

7. On the **Confirmation** page, review the settings. You can copy the Windows PowerShell command to automate additional installations (in particular, the certificate thumbprint value). When you're finished, click **Publish**.

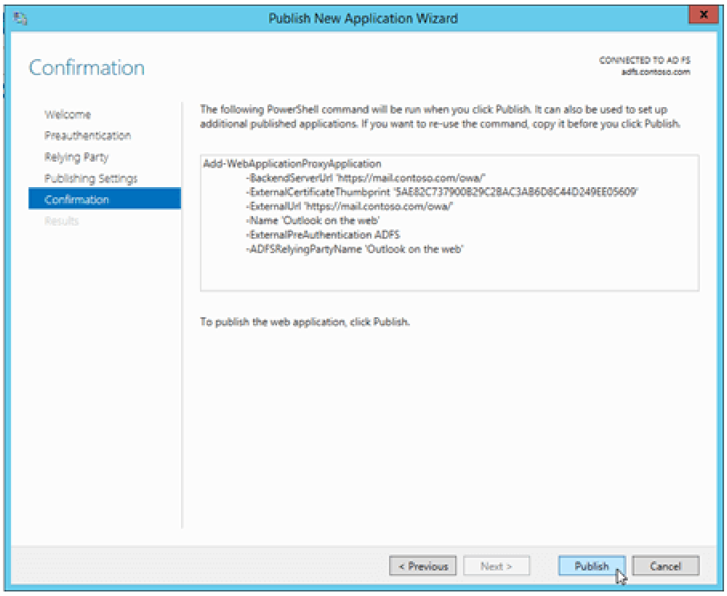

The following PowerShell command will be run when you click Publish. It can also be used to set up additional published applications. If you want to re-use the command, copy it before you click Publish.

```
Add-WebApplicationProxyApplication
        -BackendServerUrl 'https://mail.contoso.com/owa/'
        -ExternalCertificateThumbprint '5AE82C737900B29C28AC3AB6D8C44D249EE05609'
        -ExternalUrl 'https://mail.contoso.com/owa/'
        -Name 'Outlook on the web'
        -ExternalPreAuthentication ADFS
        -ADFSRelyingPartyName 'Outlook on the web'
```

To publish the web application, click Publish.

8. On the **Results** page, verify that the application published successfully, and then click **Close**.

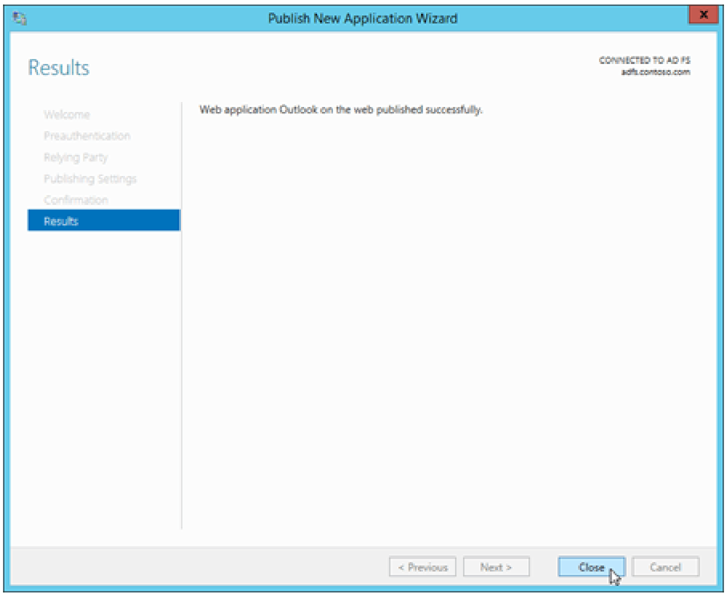

Publish New Application Wizard

Results

CONNECTED TO AD FS
adfs.contoso.com

Web application Outlook on the web published successfully.

To use Windows PowerShell to publish the relying party trusts, follow these steps:

1. Run the following command on the Web Application Proxy server to find the thumbprint of the installed certificate that contains the host name of the Outlook on the web and EAC URLs (for example, `mail.contoso.com`):

```
Set-Location Cert:\LocalMachine\My; Get-ChildItem | Format-List FriendlyName,Subject,Thumbprint
```

2. Use the following syntax:

```
Add-WebApplicationProxyApplication -ExternalPreAuthentication ADFS -ADFSRelyingPartyName
<OotwRelyingParty | EACRelyingParty> -Name "<Outlook on the web  | EAC>" -ExternalUrl <OotwURL | EACURL>
-ExternalCertificateThumbprint <Thumbprint> -BackendServerUrl <OotwURL | EACURL>
```

This example publishes Outlook on the web in Web Application Proxy with the following settings:

- **AD FS relying party**: Outlook on the web
- **Name**: Outlook on the web
- **External URL**: https://mail.contoso.com/owa/
- **External certificate thumbprint**: The `*.contoso.com` certificate that has the thumbprint value `5AE82C737900B29C2BAC3AB6D8C44D249EE05609`.
- **Backend server URL**: https://mail.contoso.com/owa/

```
Add-WebApplicationProxyApplication -ExternalPreAuthentication ADFS -ADFSRelyingPartyName "Outlook on the
web" -Name "Outlook on the web" -ExternalUrl https://mail.contoso.com/owa/ -
ExternalCertificateThumbprint 5AE82C737900B29C2BAC3AB6D8C44D249EE056093 -BackendServerUrl
https://mail.contoso.com/owa/
```

This example publishes the EAC in Web Application Proxy with the following settings:

- **Name**: EAC
- **External URL**: https://external.contoso.com/ecp/
- **External certificate thumbprint**: The `*.contoso.com` certificate that has the thumbprint value

`5AE82C737900B29C2BAC3AB6D8C44D249EE05609`.

- **Backend server URL**: https://mail.contoso.com/ecp/

```
Add-WebApplicationProxyApplication -ExternalPreAuthentication ADFS -ADFSRelyingPartyName EAC -Name EAC -
ExternalUrl https://external.contoso.com/ecp/ -ExternalCertificateThumbprint
5AE82C737900B29C2BAC3AB6D8C44D249EE05609 -BackendServerUrl https://mail.contoso.com/ecp/
```

Note: All AD FS endpoints that you want to publish through Web Application Proxy need to be proxy enabled. You do this in the AD FS Management console at **Service** > **Endpoints** (verify that **Proxy Enabled** is **Yes** for the specified endpoint).

Step 6: Configure the Exchange organization to use AD FS authentication

To configure the Exchange organization to use AD FS authentication, you need to use the Exchange Management Shell. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.

1. Run the following command to find the thumbprint value of the imported AD FS token signing certificate:

```
Set-Location Cert:\LocalMachine\Root; Get-ChildItem | Sort-Object Subject
```

Look for the Subject value `CN=ADFS Signing - <FederationServiceName>` (for example, `CN=ADFS Signing - adfs.contoso.com`).

You can confirm this thumbprint value on the AD FS server in an elevated Windows PowerShell window by running the command `Import-Module ADFS`, and then running the command `Get-AdfsCertificate -CertificateType Token-Signing`.

2. Use the following syntax:

```
Set-OrganizationConfig -AdfsIssuer https://<FederationServiceName>/adfs/ls/ -AdfsAudienceUris "
<OotwURL>","<EACURL>" -AdfsSignCertificateThumbprint "<Thumbprint>"
```

This example uses the following values:

- **AD FS URL**: `https://adfs.contoso.com/adfs/ls/`
- **Outlook on the web URL**: `https://mail.contoso.com/owa/`
- **EAC URL**: `https://mail.contoso.com/ecp/`
- **AD FS token-signing certificate thumbprint**: The `ADFS Signing - adfs.contoso.com` certificate that has the thumbprint value `88970C64278A15D642934DC2961D9CCA5E28DA6B`.

```
Set-OrganizationConfig -AdfsIssuer https://adfs.contoso.com/adfs/ls/ -AdfsAudienceUris
"https://mail.contoso.com/owa/","https://mail.contoso.com/ecp/" -AdfsSignCertificateThumbprint
"88970C64278A15D642934DC2961D9CCA5E28DA6B"
```

Note: The *AdfsEncryptCertificateThumbprint* parameter isn't supported in these scenarios.

Step 7: Configure AD FS authentication on the Outlook on the web and EAC virtual directories

For the Outlook on the web and EAC virtual directories, you need to configure AD FS authentication as the only available authentication method by disabling all other authentication methods.

- You need to configure the EAC virtual directory before you configure the Outlook on the web virtual directory.
- You'll likely want to configure AD FS authentication only on Internet-facing Exchange servers that clients use to connect to Outlook on the web and the EAC.
- By default, only Basic and Forms authentication are enabled for the Outlook on the web and EAC virtual directories.

To use the Exchange Management Shell to configure an EAC or Outlook on the web virtual directory to only accept AD FS authentication, use the following syntax:

```
Set-EcpVirtualDirectory -Identity <VirtualDirectoryIdentity> -AdfsAuthentication $true -BasicAuthentication
$false -DigestAuthentication $false -FormsAuthentication $false -OAuthAuthentication $false -
WindowsAuthentication $false
```

This example configures the EAC virtual directory in the default web site on the server named Mailbox01:

```
Set-EcpVirtualDirectory -Identity "Mailbox01\ecp (Default Web Site)" -AdfsAuthentication $true -
BasicAuthentication $false -DigestAuthentication $false -FormsAuthentication $false -OAuthAuthentication
$false -WindowsAuthentication $false
```

This example configures the Outlook on the web virtual directory in the default we site on the server named Mailbox01:

```
Set-OwaVirtualDirectory -Identity "Mailbox01\owa (Default Web Site)" -AdfsAuthentication $true -
BasicAuthentication $false -DigestAuthentication $false -FormsAuthentication $false -OAuthAuthentication
$false -WindowsAuthentication $false
```

Note: To configure all EAC and Outlook on the web virtual directories on every Exchange server in your organization, run the following commands:

```
Get-EcpVirtualDirectory | Set-EcpVirtualDirectory -AdfsAuthentication $true -BasicAuthentication $false -
DigestAuthentication $false -FormsAuthentication $false -OAuthAuthentication $false -WindowsAuthentication
$false
```

```
Get-OwaVirtualDirectory | Set-OwaVirtualDirectory -AdfsAuthentication $true -BasicAuthentication $false -
DigestAuthentication $false -FormsAuthentication $false -OAuthAuthentication $false -WindowsAuthentication
$false
```

Step 8: Restart IIS on the Exchange server

1. Open IIS Manager on the Exchange server. An easy way to do this in Windows Server 2012 or later is to press Windows key + Q, type inetmgr, and select **Internet Information Services (IIS) Manager** in the results.
2. In IIS Manager, select the server.
3. In the **Actions** pane, click **Restart**.

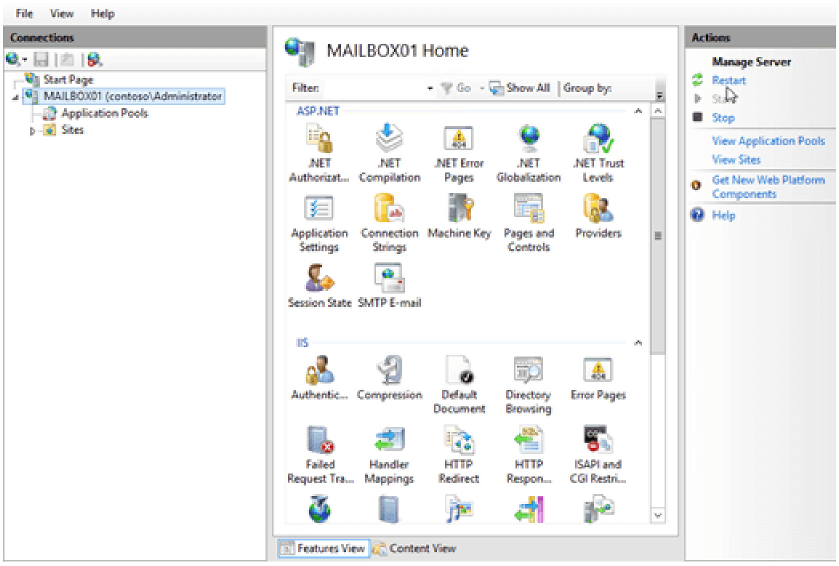

Note: To perform this procedure on the command line, open an elevated command prompt on the Exchange server (a Command Prompt window you open by selecting **Run as administrator**) and run the following commands:

```
net stop w3svc /y
```

```
net start w3svc
```

How do you know this worked?

To test the AD FS claims for Outlook on the web:

1. In a web browser, open Outlook on the web (for example, https://mail.contoso.com/owa).
2. If you get a certificate error in the web browser, just continue on to the Outlook on the web site. You should be redirected to the AD FS sign-in page or the AD FS prompt for credentials.
3. Type your username (domain\user) and password, and then click **Sign in**.
4. Outlook on the web will load in the window.

To test the AD FS claims for EAC:

1. In a web browser, open EAC (for example, https://mail.contoso.com/ecp).
2. If you get a certificate error in the web browser, just continue on to the EAC web site. You should be redirected to the AD FS sign-in page or the AD FS prompt for credentials.
3. Type your username (domain\user) and password, and then click **Sign in**.
4. EAC will load in the window.

Additional considerations

Multifactor authentication

Deploying and configuring AD FS for claims-based authentication allows Outlook on the web and the EAC to support multifactor authentication, such as certificate-based authentication, authentication or security tokens, and fingerprint authentication. Multifactor authentication requires two of these three authentication factors:

- Something only the user knows (for example, the password, PIN, or pattern).

- Something only the user has (for example, an ATM card, security token, smart card, or mobile phone).
- Something only the user is (for example, a biometric characteristic, such as a fingerprint).

For example, a password and a security code that's sent to a mobile phone, or a PIN and a fingerprint.

For details on multifactor authentication in Windows Server 2012 R2, see Overview: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications and Walkthrough Guide: Manage Risk with Additional Multi-Factor Authentication for Sensitive Applications.

On the AD FS server, the federation service functions as a security token service, and provides the security tokens that are used with claims. The federation service issues tokens based on the credentials that are presented. After the account store verifies a user's credentials, the claims for the user are generated according to the rules of the trust policy and then added to a security token that is issued to the client. For more information about claims, see Understanding Claims.

Co-existence with other versions of Exchange

You can use AD FS authentication for Outlook on the web and the EAC when you have more than one version of Exchange deployed in your organization. This scenario is supported only if all clients are connecting through Exchange servers, **and** all of those servers have been configured for AD FS authentication.

In Exchange 2016 organizations, users with mailboxes on Exchange 2010 servers can access their mailboxes through an Exchange 2016 server that's configured for AD FS authentication. The initial client connection to the or Exchange 2016 server uses AD FS authentication. However, the proxied connection to Exchange 2010 uses Kerberos. There's no supported way to configure Exchange 2010 for direct AD FS authentication.

Client Access Rules in Exchange 2019

8/3/2020 • 7 minutes to read • Edit Online

Client Access Rules help you control access to your Exchange 2019 organization in the Exchange admin center (EAC) and remote PowerShell based on client properties or client access requests. Client Access Rules are like mail flow rules (also known as transport rules) for EAC and remote PowerShell connections to your Exchange organization. You can prevent EAC and remote PowerShell clients from connecting to Exchange based on their IP address, authentication type, and user property values. For example:

- Prevent client access using remote PowerShell (which also includes the Exchange Management Shell).
- Block access to the EAC for users in a specific country or region.

For Client Access Rule procedures, see Procedures for Client Access Rules in Exchange Server.

Client Access Rule components

A rule is made of conditions, exceptions, an action, and a priority value.

- **Conditions**: Identify the client connections to apply the action to. For a complete list of conditions, see the Client Access Rule conditions and exceptions section later in this topic. When a client connection matches the conditions of a rule, the action is applied to the client connection, and rule evaluation stops (no more rules are applied to the connection).
- **Exceptions**: Optionally identify the client connections that the action shouldn't apply to. Exceptions override conditions and prevent the rule action from being applied to a connection, even if the connection matches all of the configured conditions. Rule evaluation continues for client connections that are allowed by the exception, but a subsequent rule could still affect the connection.
- **Action**: Specifies what to do to client connections that match the conditions in the rule, and don't match any of the exceptions. Valid actions are:
  - Allow the connection (the `AllowAccess` value for the *Action* parameter).
  - Block the connection (the `DenyAccess` value for the *Action* parameter).

  Note: When you block connections for a specific protocol, other applications that rely on the same protocol might also be affected.
- **Priority**: Indicates the order that the rules are applied to client connections (a lower number indicates a higher priority). The default priority is based on when the rule is created (older rules have a higher priority than newer rules), and higher priority rules are processed before lower priority rules. Remember, rule processing stops once the client connection matches the conditions in the rule.

  For more information about setting the priority value on rules, see Use the Exchange Management Shell to set the priority of Client Access Rules.

How Client Access Rules are evaluated

How multiple rules with the same condition are evaluated, and how a rule with multiple conditions, condition values, and exceptions are evaluated are described in the following table.

COMPONENT	LOGIC	COMMENTS
Multiple rules that contain the same condition	The first rule is applied, and subsequent rules are ignored	For example, if your highest priority rule blocks remote PowerShell connections, and you create another rule that allows remote PowerShell connections for a specific IP address range, all remote PowerShell connections are still blocked by the first rule. Instead of creating another rule for remote PowerShell, you need to add an exception to the existing remote PowerShell rule to allow connections from the specified IP address range.
Multiple conditions in one rule	AND	A client connection must match all conditions in the rule. For example, EAC connections from users in the Accounting department.
One condition with multiple values in a rule	OR	For conditions that allow more than one value, the connection must match any one (not all) of the specified conditions. For example, EAC or remote PowerShell connections.
Multiple exceptions in one rule	OR	If a client connection matches any one of the exceptions, the actions are not applied to the client connection. The connection doesn't have to match all the exceptions. For example, IP address 19.2.168.1.1 or Basic authentication.

You can test how a specific client connection would be affected by Client Access Rules (which rules would match and therefore affect the connection). For more information, see Use the Exchange Management Shell to test Client Access Rules.

Important notes

Client connections from your internal network

Connections from your local network aren't automatically allowed to bypass Client Access Rules. Therefore, when you create Client Access Rules that block client connections to Exchange, you need to consider how connections from your internal network might be affected. The preferred method to allow internal client connections to bypass Client Access Rules is to create a highest priority rule that allows client connections from your internal network (all or specific IP addresses). That way, the client connections are always allowed, regardless of any other blocking rules that you create in the future.

Client Access Rules and middle-tier applications

Many applications that access Exchange use a middle-tier architecture (clients talk to the middle-tier application and the middle-tier application talks to Exchange). A Client Access Rule that only allows access from your local network might block middle-tier applications. So, your rules need to allow the IP addresses of middle-tier applications.

Middle-tier applications owned by Microsoft (for example, Outlook for iOS and Android) will bypass blocking by Client Access Rules, and will always be allowed. To provide additional control over these applications, you need to use the control capabilities that are available in the applications.

Timing for rule changes

To improve overall performance, Client Access Rules use a cache, which means changes to rules don't immediately take effect. The first rule that you create in your organization can take up to 24 hours to take effect. After that, modifying, adding, or removing rules can take up to one hour to take effect.

Administration

You can only use the Exchange Management Shell (remote PowerShell) to manage Client Access Rules, so you need to be careful about rules that block your access to remote PowerShell.

As a best practice, create a Client Access Rule with the highest priority to preserve your access to remote PowerShell. For example:

```
New-ClientAccessRule -Name "Always Allow Remote PowerShell" -Action Allow -AnyOfProtocols RemotePowerShell -
Priority 1
```

Authentication types and protocols

Not all authentication types are supported for all protocols. The supported authentication types per protocol in Exchange Server are described in this table:

	ADFSAUTHENTICATION	BASICAUTHENTICATION	CERTIFICATEBASEDAUTHENTICATION	NONBASICAUTHENTICATION	OAUTHAUTHENTICATION
`ExchangeAdminCenter`	supported	supported	n/a	n/a	n/a
`RemotePowerShell`	n/a	supported	n/a	supported	n/a

Client Access Rule conditions and exceptions

Conditions and exceptions in Client Access Rules identify the client connections that the rule is applied to or not applied to. For example, if the rule blocks access by remote PowerShell clients, you can configure the rule to allow remote PowerShell connections from a specific range of IP addresses. The syntax is the same for a condition and the corresponding exception. The only difference is conditions specify client connections to include, while exceptions specify client connections to exclude.

This table describes the conditions and exceptions that are available in Client Access Rules:

CONDITION PARAMETER IN THE EXCHANGE MANAGEMENT SHELL	EXCEPTION PARAMETER IN THE EXCHANGE MANAGEMENT SHELL	DESCRIPTION
AnyOfAuthenticationTypes	*ExceptAnyOfAuthenticationTypes*	Valid values in Exchange Server are: • For the EAC: `AdfsAuthentication` and `BasicAuthentication` • For remote PowerShell: `BasicAuthentication` and `NonBasicAuthentication` You can specify multiple values separated by commas. You can use quotation marks around each individual value ("*value1*","*value2*"), but not around all values (don't use "*value1*,*value2*").

CONDITION PARAMETER IN THE EXCHANGE MANAGEMENT SHELL	EXCEPTION PARAMETER IN THE EXCHANGE MANAGEMENT SHELL	DESCRIPTION
AnyOfClientIPAddressesOrRanges	*ExceptAnyOfClientIPAddressesOrRanges*	Valid values are: • **A single IP address**: For example, `192.168.1.1`. • **An IP address range**: For example, `192.168.0.1-192.168.0.254`. • **Classless Inter-Domain Routing (CIDR) IP**: For example, `192.168.3.1/24`. You can specify multiple values separated by commas.
AnyOfProtocols	*ExceptAnyOfProtocols*	Valid values in Exchange Server are: • `ExchangeAdminCenter` • `RemotePowerShell` You can specify multiple values separated by commas. You can use quotation marks around each individual value (" *value1*","*value2*"), but not around all values (don't use "*value1*,*value2*"). **Note**: If you don't use this condition in a rule, the rule is applied to both protocols.
Scope	n/a	Specifies the type of connections that the rule applies to. Valid values are: • `Users`: The rule only applies to end-user connections. • `All`: The rule applies to all types of connections (end-users and middle-tier apps).
UsernameMatchesAnyOfPatterns	*ExceptUsernameMatchesAnyOfPatterns*	Accepts text and the wildcard character (*) to identify the user's account name in the format `<Domain>\<UserName>` (for example, `contoso.com\jeff` or `*jeff*`, but not `jeff*`). Non-alphanumeric characters don't require an escape character. You can specify multiple values separated by commas.

CONDITION PARAMETER IN THE EXCHANGE MANAGEMENT SHELL	EXCEPTION PARAMETER IN THE EXCHANGE MANAGEMENT SHELL	DESCRIPTION
UserRecipientFilter	n/a	Uses OPath filter syntax to identify the user that the rule applies to. For example, `"City -eq 'Redmond'"`. The filterable attributes are: • `City` • `Company` • `CountryOrRegion` • `CustomAttribute1` to `CustomAttribute15` • `Department` • `Office` • `PostalCode` • `StateOrProvince` • `StreetAddress` The search criteria uses the syntax `"<Property> -<Comparison operator> '<Value>'"` . • `<Property>` is a filterable property. • `-<Comparison Operator>` is an OPATH comparison operator. For example `-eq` for exact matches (wildcards are not supported) and `-like` for string comparison (which requires at least one wildcard in the property value). For more information about comparison operators, see about_Comparison_Operators. • `<Value>` is the property value. Text values with or without spaces or values with wildcards (*) need to be enclosed in quotation marks (for example, `'<Value>'` or `'*<Value>'`). Don't use quotation marks with the system value `$null` (for blank values). You can chain multiple search criteria together using the logical operators `-and` and `-or`. For example, `"<Criteria1> -and <Criteria2>"` or `"(<Criteria1> -and <Criteria2>) -or <Criteria3>"` . For more information about OPATH filter syntax, see Additional OPATH syntax information.

Procedures for Client Access Rules in Exchange 2019

8/3/2020 • 6 minutes to read • Edit Online

Client Access Rules allow or block Exchange admin center (EAC) or remote PowerShell connections to your Exchange 2019 organization based on the properties of the connection. For more information about Client Access Rules, see Client Access Rules in Exchange Server.

> **TIP**
>
> Verify that your rules work the way you expect. Be sure to thoroughly test each rule and the interactions between rules. For more information, see the Use the Exchange Management Shell to test Client Access Rules section later in this topic.

What do you need to know before you begin?

- Estimated time to complete each procedure: less than 5 minutes.
- The procedures in this topic are only available in the Exchange Management Shell. For more information, see Open the Exchange Management Shell or Connect to Exchange servers using remote PowerShell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Mail flow" entry in Mail flow permissions.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at Exchange Server.

Use the Exchange Management Shell to view Client Access Rules

To return a summary list of all Client Access Rules, run this command:

```
Get-ClientAccessRule
```

To return detailed information about a specific rule, use this syntax:

```
Get-ClientAccessRule -Identity "<RuleName>" | Format-List [<Specific properties to view>]
```

This example returns all the property values for the rule named "Block Client Connections from 192.168.1.0/24".

```
Get-ClientAccessRule -Identity "Block Client Connections from 192.168.1.0/24" | Format-List
```

This example returns only the specified properties for the same rule.

```
Get-ClientAccessRule -Identity "Block Client Connections from 192.168.1.0/24" | Format-List
Name,Priority,Enabled,Scope,Action
```

For detailed syntax and parameter information, see Get-ClientAccessRule.

Use the Exchange Management Shell to create Client Access Rules

To create Client Access Rules in the Exchange Management Shell, use this syntax:

```
New-ClientAccessRule -Name "<RuleName>" [-Priority <PriorityValue>] [-Enabled <$true | $false>] -Action
<AllowAccess | DenyAccess> [<Conditions>] [<Exceptions>]
```

This example creates a new Client Access Rule named Block PowerShell that blocks remote PowerShell access, except for clients in the IP address range 192.168.10.1/24.

```
New-ClientAccessRule -Name "Block PowerShell" -Action DenyAccess -AnyOfProtocols RemotePowerShell -
ExceptAnyOfClientIPAddressesOrRanges 192.168.10.1/24
```

Notes:

- As a best practice, create a Client Access Rule with the highest priority to preserve your administrator access to remote PowerShell. For example:
  ```
  New-ClientAccessRule -Name "Always Allow Remote PowerShell" -Action Allow -AnyOfProtocols
  RemotePowerShell -Priority 1
  ```
  .

- The rule has the default priority value, because we didn't use the *Priority* parameter. For more information, see the Use the Exchange Management Shell to set the priority of Client Access Rules section later in this topic.

- The rule is enabled, because we didn't use the *Enabled* parameter, and the default value is `$true`.

This example creates a new Client Access Rule named Restrict EAC Access that blocks access for the Exchange admin center, except if the client is coming from an IP address in the 192.168.10.1/24 range or if the user account name contains "tanyas".

```
New-ClientAccessRule -Name "Restrict EAC Access" -Action DenyAccess -AnyOfProtocols ExchangeAdminCenter -
ExceptAnyOfClientIPAddressesOrRanges 192.168.10.1/24 -ExceptUsernameMatchesAnyOfPatterns *tanyas*
```

For detailed syntax and parameter information, see New-ClientAccessRule.

How do you know this worked?

To verify that you've successfully created a Client Access Rule, use any of these procedures:

- Run this command in the Exchange Management Shell to see the new rule in the list of rules:

  ```
  Get-ClientAccessRule
  ```

- Replace *<RuleName>* with the name of the rule, and run this command to see the details of the rule:

  ```
  Get-ClientAccessRule -Identity "<RuleName>" | Format-List
  ```

- See which Client Access Rules would affect a specific client connection to Exchange by using the **Test-ClientAccessRule** cmdlet. For more information, see the Use the Exchange Management Shell to test Client Access Rules section later in this topic.

Use the Exchange Management Shell to modify Client Access Rules

No additional settings are available when you modify a Client Access Rule. They're the same settings that were available when you created the rule.

To modify a Client Access Rule in the Exchange Management Shell, use this syntax:

```
Set-ClientAccessRule -Identity "<RuleName>" [-Name "<NewName>"] [-Priority <PriorityValue>] [-Enabled <$true |
$false>] -Action <AllowAccess | DenyAccess> [<Conditions>] [<Exceptions>]
```

This example disables the existing Client Access Rule named Allow EAC.

```
Set-ClientAccessRule -Identity "Allow EAC" -Enabled $false
```

An important consideration when you modify Client Access Rules is modifying conditions or exceptions that accept multiple values:

- The values that you specify will *replace* any existing values.
- To add or remove values without affecting other existing values, use this syntax: `@{Add="<Value1>","<Value2>"...; Remove="<Value1>","<Value2>"...}`

This example adds the IP address range 172.17.17.27/16 to the existing Client Access Rule named Allow EAC without affecting the existing IP address values.

```
Set-ClientAccessRule -Identity "Allow EAC" -AnyOfClientIPAddressesOrRanges @{Add="172.17.17.27/16"}
```

For detailed syntax and parameter information, see Set-ClientAccessRule.

How do you know this worked?

To verify that you've successfully modified a Client Access Rule, use any of these procedures:

- Replace *<RuleName>* with the name of the rule, and run this command to see the details of the rule:

  ```
  Get-ClientAccessRule -Identity "<RuleName>" | Format-List
  ```

- See which Client Access Rules would affect a specific client connection to Exchange by using the **Test-ClientAccessRule** cmdlet. For more information, see the Use the Exchange Management Shell to test Client Access Rules section later in this topic.

Use the Exchange Management Shell to set the priority of Client Access Rules

By default, Client Access Rules are given a priority that's based on the order they were created in (newer rules are lower priority than older rules). A lower priority number indicates a higher priority for the rule, and rules are processed in priority order (higher priority rules are processed before lower priority rules). No two rules can have the same priority.

The highest priority you can set on a rule is 1. The lowest value you can set depends on the number of rules. For example, if you have five rules, you can use the priority values 1 through 5. Changing the priority of an existing rule can have a cascading effect on other rules. For example, if you have five rules (priorities 1 through 5), and you change the priority of a rule from 5 to 2, the existing rule with priority 2 is changed to priority 3, the rule with priority 3 is changed to priority 4, and the rule with priority 4 is changed to priority 5.

To set the priority of a Client Access Rule in the Exchange Management Shell, use this syntax:

```
Set-ClientAccessRule -Identity "<RuleName>" -Priority <Number>
```

This example sets the priority of the rule named Disable PowerShell to 3. All existing rules that have a priority less than or equal to 3 are decreased by 1 (their priority numbers are increased by 1).

```
Set-ClientAccessRule -Identity "Disable PowerShell" -Priority 4
```

Note: To set the priority of a new rule when you create it, use the *Priority* parameter on the **New-ClientAccessRule** cmdlet.

How do you know this worked?

To verify that you've successfully set the priority of a Client Access Rule, use either of these procedures:

- Run the this command in the Exchange Management Shell to see the list of rules and their **Priority** values:

  ```
  Get-ClientAccessRule
  ```

- Replace *<RuleName>* with the name of the rule, and run this command:

  ```
  Get-ClientAccessRule -Identity "<RuleName>" | Format-List Name,Priority
  ```

Use the Exchange Management Shell to remove Client Access Rules

To remove Client Access Rules in the Exchange Management Shell, use this syntax:

```
Remove-ClientAccessRule -Identity "<RuleName>"
```

This example removes the Client Access Rule named Block EAC.

```
Remove-ClientAccessRule -Identity "Block EAC"
```

Note: To disable a Client Access Rule without deleting it, use the *Enabled* parameter with the value `$false` on the **Set-ClientAccessRule** cmdlet.

For detailed syntax and parameter information, see Remove-ClientAccessRule.

How do you know this worked?

To verify that you've successfully removed a Client Access Rule, run this command in the Exchange Management Shell to verify that the rule is no longer listed:

```
Get-ClientAccessRule
```

Use the Exchange Management Shell to test Client Access Rules

To see which Client Access Rules would affect a specific client connection to Exchange, use this syntax:

```
Test-ClientAccessRule -User <MailboxIdentity> -AuthenticationType <AuthenticationType> -Protocol <Protocol> -
RemoteAddress <ClientIPAddress> -RemotePort <TCPPortNumber>
```

This example returns the Client Access Rules that would match a client connection to Exchange that has these properties:

- **Authentication type**: Basic
- **Protocol**: `ExchangeAdminCenter`
- **Remote address**: 172.17.17.26
- **Remote port**: 443
- **User**: julia@contoso.com

```
Test-ClientAccessRule -User julia@contoso.com -AuthenticationType BasicAuthentication -Protocol
ExchangeAdminCenter -RemoteAddress 172.17.17.26 -RemotePort 443
```

For detailed syntax and parameter information, see Test-ClientAccessRule.

Mail flow and the transport pipeline

8/3/2020 • 10 minutes to read • Edit Online

In Exchange Server, mail flow occurs through the transport pipeline. The *transport pipeline* is a collection of services, connections, components, and queues that work together to route all messages to the categorizer in the Transport service on an Exchange Mailbox server inside the organization.

For information about how to configure mail flow in a new Exchange 2016 or Exchange 2019 organization, see Configure mail flow and client access.

Understanding the transport pipeline

The transport pipeline consists of the following services:

- **Front End Transport service on Mailbox servers**: This service acts as a stateless proxy for all inbound and (optionally) outbound external SMTP traffic for the Exchange Server organization. The Front End Transport service doesn't inspect message content, doesn't communicate with the Mailbox Transport service, and doesn't queue any messages locally.

- **Transport service on Mailbox servers**: This service is virtually identical to the Hub Transport server role in Exchange Server 2010. The Transport service handles all SMTP mail flow for the organization, performs message categorization, and performs message content inspection. Unlike Exchange 2010, the Transport service never communicates directly with mailbox databases. That task is now handled by the Mailbox Transport service. The Transport service routes messages among the Mailbox Transport service, the Transport service, the Front End Transport service, and (depending on your configuration) the Transport service on Edge Transport servers. The Transport service on Mailbox servers is described in more detail later in this topic.

- **Mailbox Transport service on Mailbox servers**: This service consists of two separate services:

  - **Mailbox Transport Submission service**: This service connects to the local mailbox database using an Exchange remote procedure call (RPC) to retrieve messages. The service submits the messages over SMTP to the Transport service on the local Mailbox server or on other Mailbox servers. The Mailbox Transport Submission service has access to the same routing topology information as the Transport service.

  - **Mailbox Transport Delivery service**: This service receives SMTP messages from the Transport service on the local Mailbox server or on other Mailbox servers and connects to the local mailbox database using RPC to deliver the messages.

  The Mailbox Transport service doesn't communicate with the Front End Transport service, the Mailbox Transport service, or mailbox databases on other Mailbox servers. It also doesn't queue any messages locally.

- **Transport service on Edge Transport servers**: This service is very similar to the Transport service on Mailbox servers. If you have an Edge Transport server installed in the perimeter network, all mail coming from the Internet or going to the Internet flows through the Transport service Edge Transport server. This service is described in more detail later in this topic.

The following diagram shows the relationships among the components in the Exchange transport pipeline.

NOTE

Although the diagrams in this topic show the components on a single Exchange server, communication also occurs between those components on different Exchange servers. The only communication that always occurs on the local Exchange server is between the Mailbox Transport service and the local mailbox database.

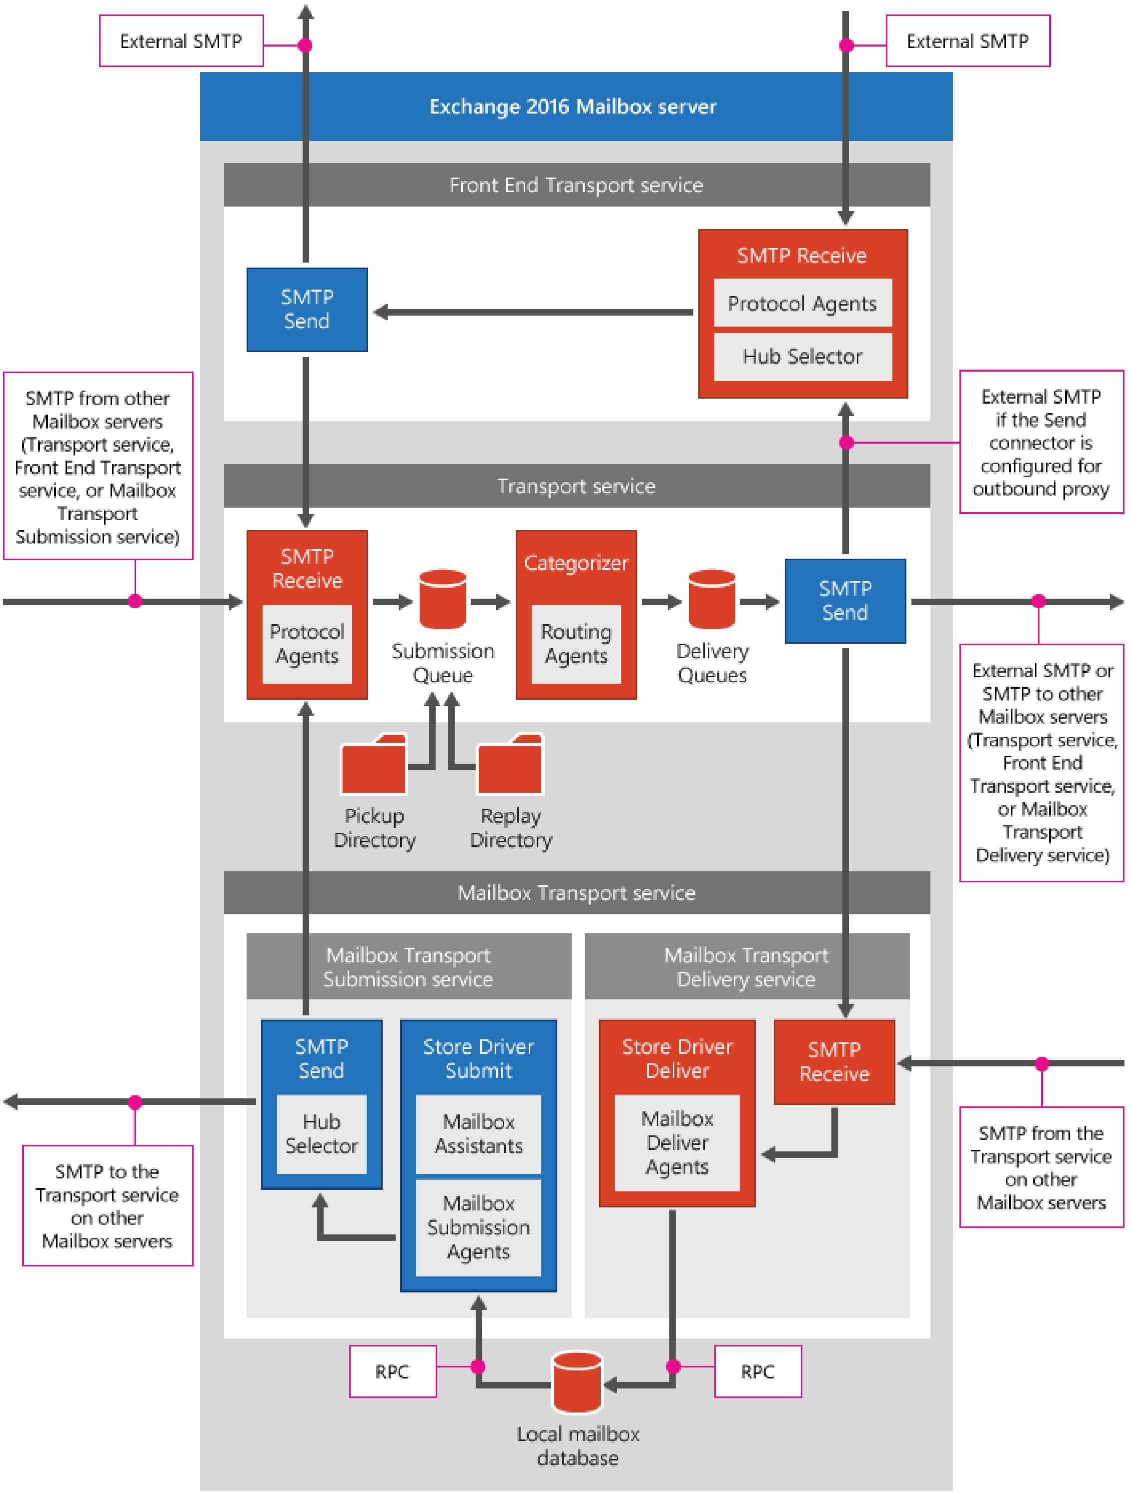

External SMTP
External SMTP
Exchange 2016 Mailbox server
Front End Transport service
SMTP Send
SMTP Receive
Protocol Agents
Hub Selector
SMTP from other Mailbox servers (Transport service, Front End Transport service, or Mailbox Transport Submission service)
External SMTP if the Send connector is configured for outbound proxy
Transport service
SMTP Receive
Protocol Agents
Submission Queue
Categorizer
Routing Agents
Delivery Queues
SMTP Send
External SMTP or SMTP to other Mailbox servers (Transport service, Front End Transport service, or Mailbox Transport Delivery service)
Pickup Directory
Replay Directory
Mailbox Transport service
Mailbox Transport Submission service
Mailbox Transport Delivery service
SMTP Send
Hub Selector
Store Driver Submit
Mailbox Assistants
Mailbox Submission Agents
Store Driver Deliver
Mailbox Deliver Agents
SMTP Receive
SMTP to the Transport service on other Mailbox servers
SMTP from the Transport service on other Mailbox servers
RPC
RPC
Local mailbox database

How messages from external senders enter the transport pipeline

The way messages from outside the Exchange organization enter the transport pipeline depends on whether you have a subscribed Edge Transport server deployed in your perimeter network.

Inbound mail flow (no Edge Transport servers)

The following diagram and list describe inbound mail flow with only Exchange Mailbox servers.

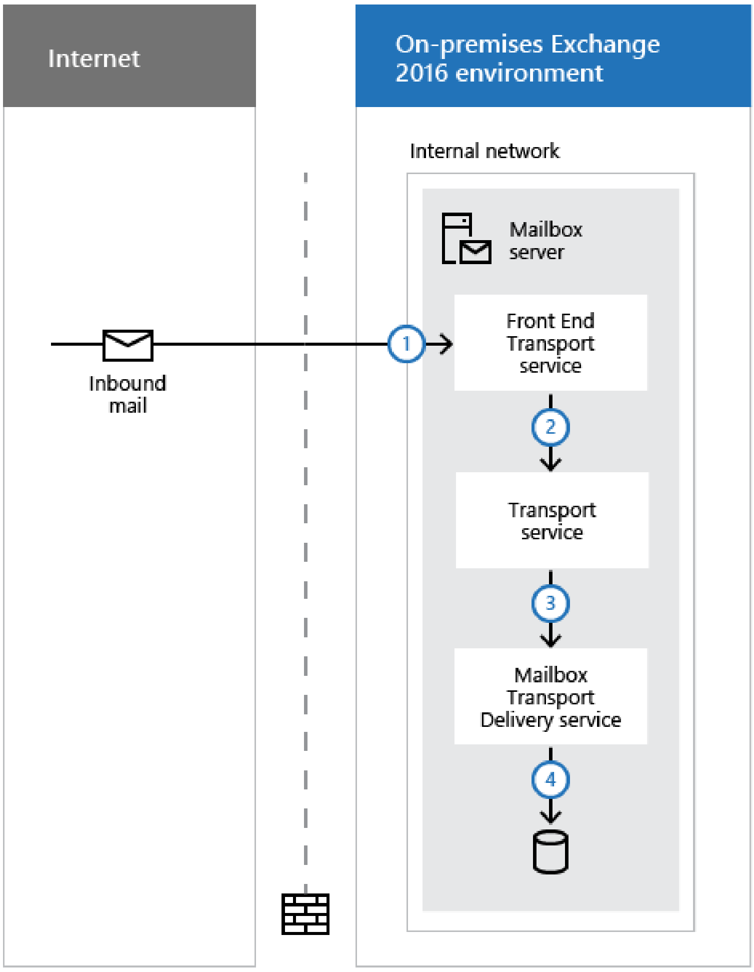
Internet
On-premises Exchange 2016 environment
Internal network
Mailbox server
Inbound mail
1
Front End Transport service
2
Transport service
3
Mailbox Transport Delivery service
4

1. A message from outside the organization enters the transport pipeline through the default Receive connector named "Default Frontend *<Mailbox server name>*" in the Front End Transport service.

2. The message is sent to the Transport service on the local Mailbox server or on a different Mailbox server. The Transport service listens for messages on the default Receive connector named "Default *<Mailbox server name>*".

3. The message is sent from the Transport service to the Mailbox Transport Delivery service on the local Mailbox server or on a different Mailbox server.

4. The Mailbox Transport Delivery service uses RPC to deliver the message to the local mailbox database.

Inbound mail flow with Edge Transport servers

The following diagram and list describe inbound mail flow with an Edge Transport server installed in the perimeter network

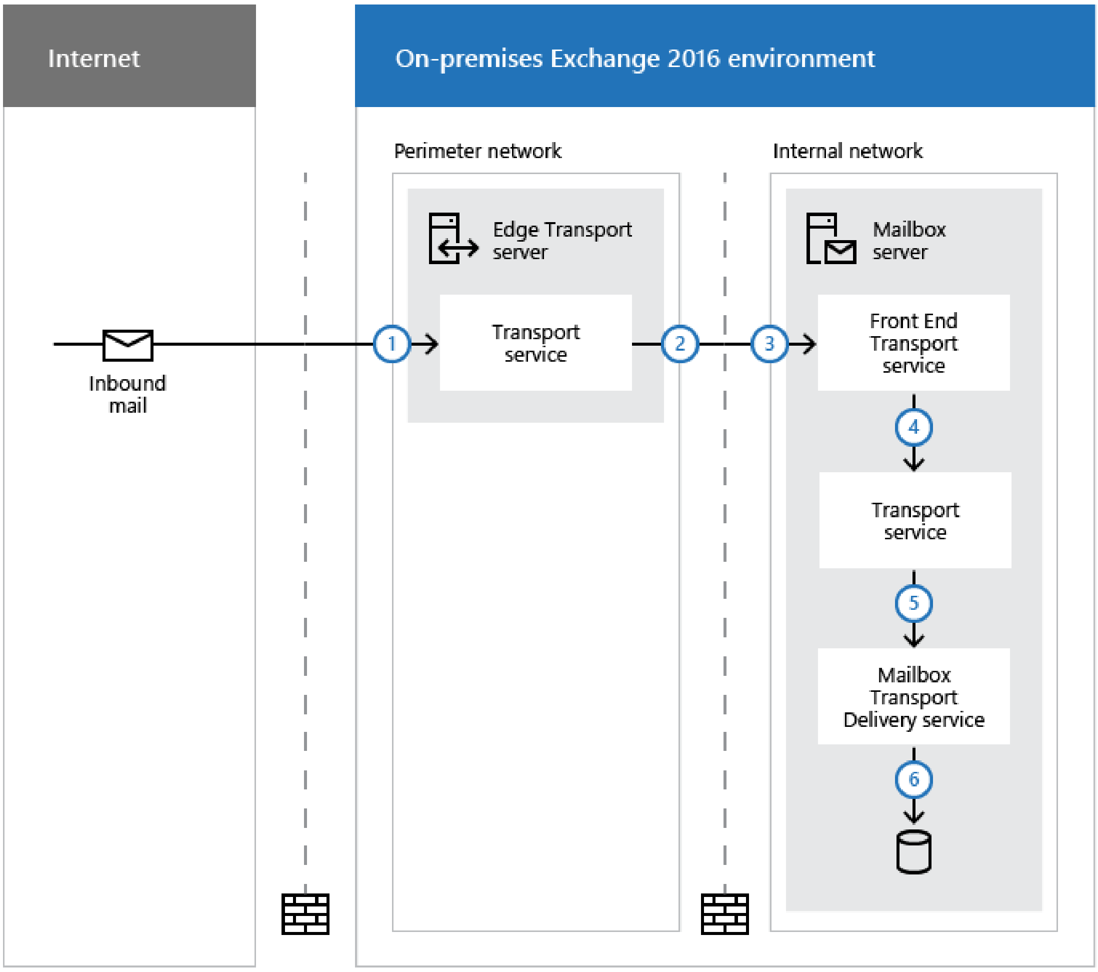

1. A message from outside the Exchange organization enters the transport pipeline through the default Receive connector named "Default internal Receive connector *<Edge Transport server name>*" in the Transport service on the Edge Transport server.

2. In the Transport service on the Edge Transport server, the default Send connector named "EdgeSync - Inbound to *<Active Directory site name>*" sends the message to a Mailbox server in the subscribed Active Directory site.

3. In the Front End Transport service on the Mailbox server, the default Receive connector named "Default Frontend *<Mailbox server name>*" accepts the message.

4. The message is sent from the Front End Transport service to the Transport service on the local Mailbox server or on a different Mailbox server. The Transport service listens for messages on the default Receive connector named "Default *<Mailbox server name>*".

5. The message is sent from the Transport service to the Mailbox Transport Delivery service on the local Mailbox server, or on a different Mailbox server.

6. The Mailbox Transport Delivery service uses RPC to deliver the message to the local mailbox database.

How messages from internal senders enter the transport pipeline

SMTP messages from inside the organization enter the transport pipeline through the Transport service on a Mailbox server in one of the following ways:

- Through a Receive connector.

- From the Pickup directory or the Replay directory.

- From the Mailbox Transport Submission service.

- Through agent submission.

The message is routed based on the routing destination or delivery group.

Outbound mail flow (no Edge Transport servers)

By default, in a new Exchange Server organization, there's no Send connector that's configured to send messages to the Internet. You need to create the Send connector yourself. After you do that, Outbound mail flow occurs as described in the following diagram and list.

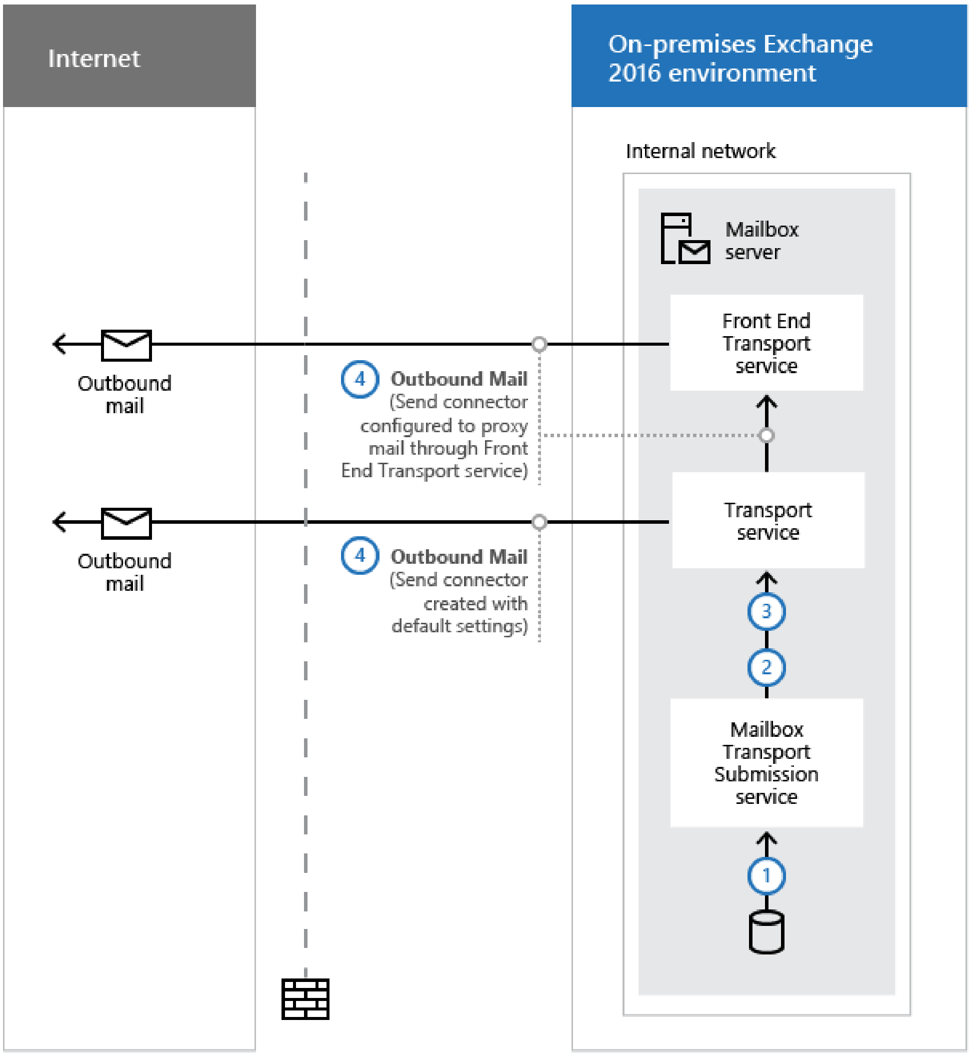

Internet
On-premises Exchange 2016 environment
Internal network
Mailbox server
Front End Transport service
Outbound mail
4 Outbound Mail (Send connector configured to proxy mail through Front End Transport service)
Transport service
Outbound mail
4 Outbound Mail (Send connector created with default settings)
3
2
Mailbox Transport Submission service
1

1. The Mailbox Transport Submission service uses RPC to retrieve the outbound message from the local mailbox database.

2. The Mailbox Transport Submission service uses SMTP to send the message to the Transport service on the local Mailbox server or on a different Mailbox server.

3. In the Transport service, the default Receive connector named "Default *<Mailbox server name>*" accepts the message.

4. What happens next depends on the configuration of the Send connector:

   - **Default**: The Transport service uses the Send connector you created to send the message to the Internet.

   - **Outbound proxy**: The Transport service uses the Send connector you created to send the message to the Front End Transport service on the local Mailbox server or on a remote Mailbox server. In the Front End Transport service, the default Receive connector named "Outbound Proxy Frontend *<Mailbox server name>*" accepts the message. The Front End Transport services sends the message to the Internet.

Outbound mail flow with Edge Transport servers

If you have an Edge Transport server installed in the perimeter network, outbound mail never flows through the Front End Transport service. Outbound mail flow with an Edge Transport server is described in the following diagram and list.

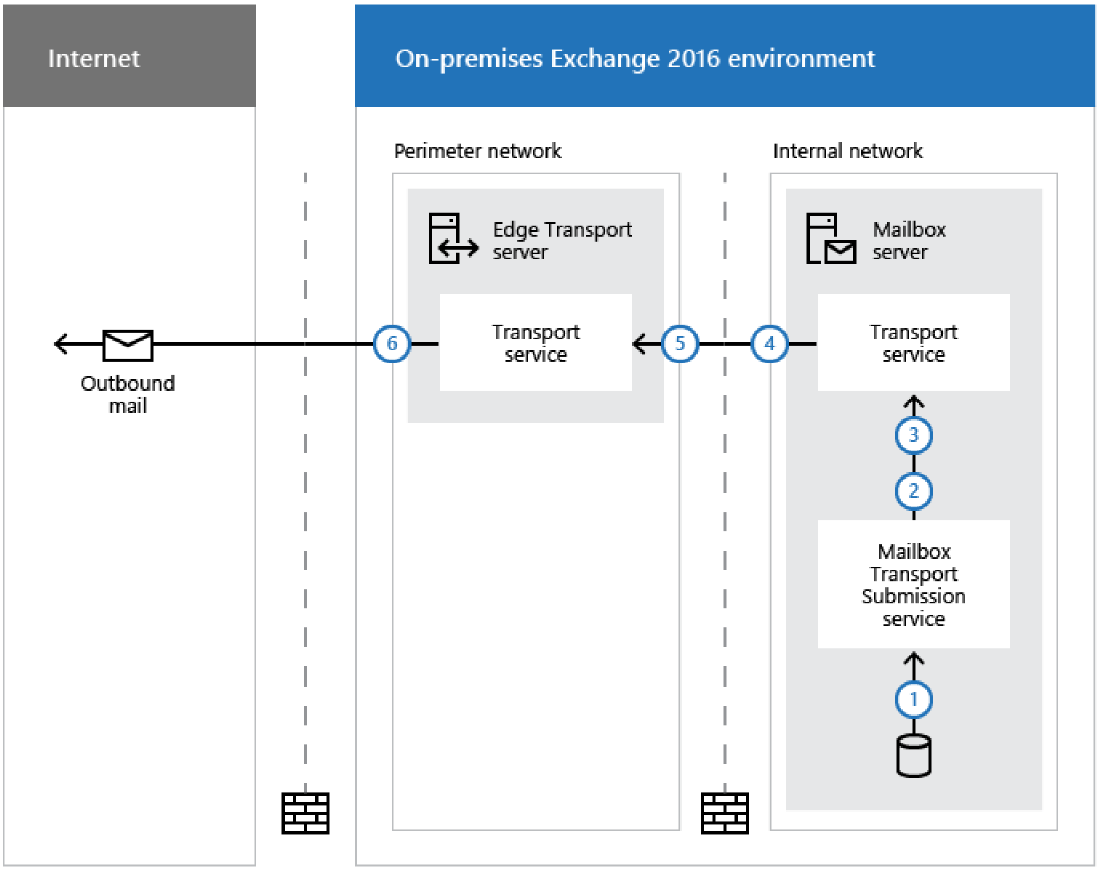

1. The Mailbox Transport Submission service uses RPC to retrieve the outbound message from the local mailbox database.

2. The Mailbox Transport Submission service uses SMTP to send the message to the Transport service on the local Mailbox server or on a different Mailbox server.

3. In the Transport service on a Mailbox server in the subscribed Active Directory site, the default Receive connector named "Default *<Mailbox server name>*" accepts the message.

4. The message is sent to the Edge Transport server using the implicit and invisible intra-organization Send connector that automatically sends mail between Exchange servers in the same organization.

5. In the Transport service on the Edge Transport server, the default Receive connector named "Default internal Receive connector *<Edge Transport server name>*" accepts the message.

6. In the Transport service on the Edge Transport server, the default Send connector named "EdgeSync - *<Active Directory site name>* to Internet" sends the message to the Internet.

Understanding the Transport service on Mailbox servers

Every message that's sent or received in an Exchange Server organization must be categorized in the Transport service on a Mailbox server before it can be routed and delivered. After a message has been categorized, it's put in a delivery queue for delivery to the destination mailbox database, the destination database availability group (DAG), Active Directory site or Active Directory forest, or to the destination domain outside the organization.

The Transport service on a Mailbox server consists of the following components and processes:

- **SMTP Receive**: When messages are received by the Transport service, message content inspection is performed and antispam inspection is performed if is enabled. The SMTP session has a series of events that work together in a specific order to validate the contents of a message before it's accepted. After a message has passed completely through SMTP Receive and isn't rejected by receive events, or by an antispam agent, it's put in the Submission queue.

- **Submission**: Submission is the process of putting messages into the Submission queue. The categorizer picks up one message at a time for categorization. Submission happens in three ways:

  - From SMTP Receive through a Receive connector.

  - Through the Pickup directory or the Replay directory. These directories exist on Mailbox servers and Edge Transport servers. Correctly formatted message files that are copied into the Pickup directory or the Replay directory are put directly into the Submission queue.

  - Through a transport agent.

- **Categorizer**: The categorizer picks up one message at a time from the Submission queue. The categorizer completes the following steps:

  - Recipient resolution, which includes top-level addressing, distribution group expansion, and message bifurcation.

  - Routing resolution.

  - Content conversion.

    Additionally, mail flow rules that the organization defined are applied. After messages have been categorized, they're put into a delivery queue that's based on the destination of the message. Messages are queued by the destination mailbox database, DAG, Active Directory site, Active Directory forest, or external domain.

- **SMTP Send**: How messages are routed from the Transport service depends on the location of the message recipients relative to the Mailbox server where categorization occurred. The message could be routed to one of the following locations:

  - To the Mailbox Transport Delivery service on the same Mailbox server.

  - To the Mailbox Transport Delivery service on a different Mailbox server that's part of the same DAG.

  - To the Transport service on a Mailbox server in a different DAG, Active Directory site, or Active Directory forest.

  - For delivery to the Internet through:

  - A Send connector on the same Mailbox server.

  - The Transport service on a different Mailbox server.

  - The Front End Transport service on the same Mailbox server or a different Mailbox server (if outbound proxy is configured).

  - The Transport service on an Edge Transport server in the perimeter network.

Understanding the Transport service on Edge Transport servers

The components of the Transport service on Edge Transport servers are identical to the components of the

Transport service on Mailbox servers. However, what actually happens during each stage of processing on Edge Transport servers is different. The differences are described in the following list.

- **SMTP Receive**: When an Edge Transport server is subscribed to an internal Active Directory site, the default Receive connector named "Default <Edge Transport server name>" is automatically configured to accept mail from internal Mailbox servers and from the Internet. When Internet messages arrive at the Edge Transport server, antispam agents filter connections and message contents and help identify the sender and the recipient while the message is being accepted into the organization. The antispam agents are installed and enabled by default. Additional attachment filtering and connection filtering features are available, but built-in malware filtering is not. Also, mail flow rules (also known as transport rules) are controlled by the Edge Rule agent. Compared to the Transport Rule agent on Mailbox servers, only a small subset of mail flow rule conditions are available on Edge Transport servers. But, there are unique mail flow rule actions related to SMTP connections that are available only on Edge Transport servers.

- **Submission**: On an Edge Transport server, messages typically enter the Submission queue through a Receive connector. However, the Pickup directory and the Replay directory are also available.

- **Categorizer**: On an Edge Transport server, categorization is a short process in which the message is put directly into a delivery queue for delivery to internal or external recipients.

- **SMTP Send**: When an Edge Transport server is subscribed to an internal Active Directory site, two Send connectors are automatically created and configured. One named "EdgeSync - <Active Directory site name> to Internet" is responsible for sending outbound mail to Internet recipients; the other named "EdgeSync - Inbound to <Active Directory site name>" is responsible for sending inbound mail from the Internet to internal recipients. Inbound mail is sent to the Front End Transport service on an available Mailbox server in the subscribed Active Directory site.

Accepted domains in Exchange Server

8/3/2020 • 6 minutes to read • Edit Online

Accepted domains are the SMTP name spaces (also known as address spaces) that you configure in an Exchange organization to receive email messages. For example, if your company registered the domain contoso.com, and you configured a mail exchanger (MX) record in your Internet DNS for contoso.com, you need to configure contoso.com as an accepted domain in your Exchange organization to accept messages that are addressed to @contoso.com recipients.

Accepted domains in Exchange 2016 and Exchange 2019 are basically unchanged from Exchange Server 2010, and consist of the following types:

- **Authoritative domains**: Recipients (in particular, mailboxes) are configured with email addresses in these domains. The Exchange organization accepts messages that are addressed to recipients in these domains, and is responsible for generating non-delivery reports (also known as NDRs or bounce messages) for non-existent recipients.

- **Relay domains**: The Exchange organization accepts messages that are addressed to recipients in relay domains, but isn't responsible for generating NDRs for non-existent recipients. Instead, Exchange (with additional configuration) relays the messages to messaging servers that are external to the Exchange organization. Relay domains can be **internal** (for domains that you control) or **external** (for domains that you don't control).

An accepted domain can be a single domain (contoso.com) or a domain with subdomains (*.contoso.com). Accepted domains are a global setting for the Exchange organization, and you can have multiple accepted domains of the same or different types.

To configure accepted domains, see Procedures for accepted domains in Exchange Server.

> **NOTE**
> If you have a subscribed Edge Transport server in your perimeter network, you configure accepted domains on a Mailbox server in your Exchange organization. The accepted domains configuration is replicated to the Edge Transport server during EdgeSync synchronization. For more information, see Edge Subscriptions.

Authoritative domains

You configure an accepted domain as an authoritative domain when all recipients in that domain exist in your Exchange organization.

By default, when you install the first Exchange Mailbox server, the fully qualified domain name (FQDN) of your forest root domain in Active Directory is configured as an authoritative domain. If you don't want to use this domain for email, you need to add another authoritative domain. For instructions, see Create accepted domains.

An organization can be configured with multiple authoritative domains. The set of email domains for an organization are the authoritative domains. You can use authoritative domains in email address policies, and Exchange is responsible for generating NDRs for non-existent recipients in authoritative domains.

Relay domains

You configure an accepted domain as a relay domain (also known as non-authoritative domain) when some or none of the recipients in that domain exist in your Exchange organization (for example, partners or subsidiaries). Exchange isn't responsible for generating NDRs for non-existent recipients in a relay domain. Instead, you configure a Send connector with the address space of the relay domain, and you configure this Send connector to use smart host routing to relay messages to their destination (directly or to the next hop). For more information about creating Send connectors that use smart host routing, see Create a Send connector to route outbound mail through a smart host.

You configure a relay domain as an internal relay domain or as an external relay domain. The differences are described in the following list:

- **Internal relay domains**
  - Some of the recipients in the internal relay domain don't exist in the Exchange organization. For example:
    - You share the domain between the Exchange organization and a third-party messaging system.
    - You share the domain between Exchange organizations in different Active Directory forests.
  - Recipients in the internal relay domain can be represented as mail contacts or mail users in the Exchange organization (manually created or automatically created by using directory synchronization).
  - The Send connector that you configure for the internal relay domain is sourced on an internal Mailbox server.

    Note: By default, you can't configure a Send connector for an internal relay domain on a subscribed Edge Transport server. Messages sent to recipients in the internal relay domain are automatically forwarded to internal Mailbox servers in the subscribed Active Directory site by using the default "EdgeSync - Inbound to *<Active Directory site name>*" Send connector. This Send connector is automatically configured to route mail for all authoritative domains and internal relay domains (the address space value is `--`). For more information, see Send connectors created automatically by the Edge Subscription.
  - You can use internal relay domains in email address policies.
- **External relay domains**
  - None of the recipients in the external relay domain exist in the Exchange organization (including mail contacts or mail users). For example, your Exchange organization is the central location for accepting Internet email for a group of separate organizations.
  - The Send connector that you configure for the external relay domain is sourced on an Edge Transport server or Internet-facing Mailbox server.
  - You can't use external relay domains in email address policies.

Accepted domains and email address policies

Email address policies assign email addresses to recipients. You need to add an authoritative domain or an internal relay domain before you can use that domain in an email address policy. For more information about email address policies, see Email address policies in Exchange Server.

Recipient Lookup in accepted domains

Recipient filtering on a subscribed Edge Transport server can block messages that are addressed to non-existent recipients in your Exchange organization. This feature is known as *Recipient Lookup*. For more information about recipient filtering, see Recipient filtering on Edge Transport servers.

You can enable or disable Recipient Lookup for an accepted domain by using the *AddressBookEnabled* parameter on the **Set-AcceptedDomain** cmdlet. The default value for each accepted domain type is described in the following table:

ACCEPTED DOMAIN TYPE	DEFAULT RECIPIENT LOOKUP (*ADDRESSBOOKENABLED* PARAMETER) VALUE	COMMENTS
Authoritative domain	Enabled (`$true`)	All recipients in an authoritative domain exist in the Exchange organization, so Recipient Lookup for the domain is enabled by default.
Internal relay domain	Disabled (`$false`)	If all recipients in the internal relay domain exist in the Exchange organization (including mail contacts and mail users), you can enable Recipient Lookup for the domain. If some or none of the recipients in the internal relay domain exist in the Exchange organization, you shouldn't enable Recipient Lookup for the domain.
External relay domain	Disabled (`$false`)	No recipients in the authoritative domain exist in the Exchange organization, so you shouldn't enable Recipient Lookup for the domain.

For configuration instructions, see Modify accepted domains.

Default domain

Because the forest root FQDN is automatically configured as the first accepted domain in your organization, that accepted domain is also configured as the *default domain*. However, after you add additional accepted domains, you can configure one of them as the default domain. Here's some information about the default domain:

- You can't delete the default domain. You need to configure another accepted domain as the default domain (one accepted domain is always configured as the default domain).
- The default domain is used in the external postmaster address: `postmaster@<default domain>` .
- The default domain is used in encapsulated non-SMTP email addresses (Internet Mail Connector Encapsulated Address or IMCEA encapsulation).
- The first default domain is used as the primary address for all recipients in the default email address policy. If you configure another accepted domain as the default domain, the default email address policy isn't automatically updated.
- Although you can configure any accepted domain as the default domain, you typically specify an authoritative domain.

Procedures for accepted domains in Exchange Server

8/3/2020 • 8 minutes to read • Edit Online

Accepted domains are the SMTP name spaces (also known as address spaces) that you configure in an Exchange organization to receive email messages. You use the Exchange admin center (EAC) or the Exchange Management Shell to configure accepted domains in Exchange Server.

For more information about accepted domains, see Accepted domains in Exchange Server. The types of accepted domains are summarized in the following list:

- **Authoritative domains**
  - All recipients in the authoritative domain exist in the Exchange organization.
  - Exchange is responsible for generating non-delivery reports (also known as NDRs or bounce messages) for non-existent recipients in an authoritative domain.
- **Internal relay domains**
  - Some recipients in the internal relay domain might exist in the Exchange organization.
  - Exchange isn't responsible for generating NDRs for non-existent recipients in an internal relay domain. Instead, you create a Send connector with the address space of the internal relay domain. You source this Send connector on an internal Mailbox server to relay messages for the non-existent recipients in the domain.
- **External relay domains**
  - None of the recipients in the external relay domain exist in the Exchange organization.
  - Exchange isn't responsible for generating NDRs for non-existent recipients in an external relay domain. Instead, you create a Send connector with the address space of the external relay domain. You source this Send connector on an Edge Transport server or Internet-facing Mailbox server to relay messages for all the recipients in the domain.

What do you need to know before you begin?

- Estimated time to complete each task: 5 minutes.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Accepted domains" entry in the Mail flow permissions topic and the "Email address policies" entry in the Recipients Permissions topic.
- If you have a subscribed Edge Transport server in your perimeter network, you configure accepted domains on a Mailbox server in your Exchange organization. The accepted domains configuration is replicated to the Edge Transport server during EdgeSync synchronization. For more information, see Edge Subscriptions.
- If Exchange accepts mail for recipients in an accepted domain from the Internet, you need to configure an MX record for the domain in your Internet-facing (public) DNS servers. Each MX record should resolve to the Internet-facing server that receives email for your organization.
- You need to create a Send connector to route mail for non-existent recipients in internal or external relay domains. For more information, see Create a Send connector to route outbound mail through a smart host.
- To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see

Open the Exchange Management Shell.

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Create accepted domains

After you create an accepted domain, you can't change the domain value (for example, from contoso.com to *.contoso.com, or vice-versa).

Use the EAC to create accepted domains

1. In the EAC, go to **Mail flow** > **Accepted domains**, and then click **Add** (+).

2. In the **New accepted domain** window that opens, configure the following settings:

   - **Name**: Enter a unique, descriptive name.

   - **Accepted domain**: Enter a single domain (for example, contoso.com) or a domain with subdomains (for example, *.contoso.com).

   - **This domain is**: Select **Authoritative**, **Internal Relay**, or **External Relay**.

   When you're finished, click **Save**.

Use the Exchange Management Shell to create accepted domains

To create an accepted domain, use the following syntax:

```
New-AcceptedDomain -Name <Name> -DomainName <DomainOrDomainWithSubdomains> -DomainType <Authoritative |
InternalRelay | ExternalRelay>
```

This example creates a new authoritative domain named Contoso Corp for contoso.com.

```
New-AcceptedDomain -Name "Contoso Corp" -DomainName contoso.com
```

Note: We didn't need to use the *DomainType* parameter, because the default value is `Authoritative`.

For detailed syntax and parameter information, see New-AcceptedDomain.

How do you know this worked?

To verify that you've successfully created an accepted domain, use either of the following procedures:

- In the EAC, go to **Mail flow** > **Accepted domains**, verify that the accepted domain is listed, and the details are correct.

- In the Exchange Management Shell, run the following command to verify the property values:

  ```
  Get-AcceptedDomain | Format-Table -Auto Name,DomainName,DomainType,Default,AddressBookEnabled
  ```

Modify accepted domains

- You can only *replace* the default domain with a new default domain (one accepted domain is always configured as the default domain). For more information about the default domain, see Default domain.

- You can enable and disable Recipient Lookup for an accepted domain only in the Exchange Management Shell. For more information, see Recipient Lookup in accepted domains.

Use the EAC to modify accepted domains

1. In the EAC, go to **Mail flow** > **Accepted domains**, select the accepted domain from the list, and then click **Edit** (✎).

2. In the properties window that opens, you can configure the following settings:

   - **Name**

   - **This domain is**: Authoritative, Internal Relay, or External Relay.

   - **Make this the default domain**: If the check box is cleared, select it to configure the accepted domain as the default domain.

   When you're finished, click **Save**.

Use the Exchange Management Shell to modify accepted domains

To modify an accepted domain, use the following syntax:

```
Set-AcceptedDomain -Identity <AcceptedDomainIdentity> [-Name <Name>]  [-DomainType <Authoritative |
InternalRelay | ExternalRelay>] [-AddressBookEnabled <$true | $false>] [-MakeDefault $true]
```

This example configures the authoritative domain named Contoso Corp as the default domain.

```
Set-AcceptedDomain -Identity "Contoso Corp" -MakeDefault $true
```

This example enables Recipient Lookup on a Edge Transport server for the internal relay domain named Fabrikam Corp. All external recipients in the fabrikam.com domain are represented in Exchange as mail users.

```
Set-AcceptedDomain -Identity "Fabrikam Corp" -AddressBookEnabled $true
```

For detailed syntax and parameter information, see Set-AcceptedDomain.

How do you know this worked?

To verify that you've successfully modified an accepted domain, use either of the following procedures:

- In the EAC, go to **Mail flow** > **Accepted domains**, and verify the property values.

  Notes:

  - To verify that the accepted domain is the default domain, you need to select the accepted domain from the list, and then click **Edit** (✎). If **Make this the default domain** is selected, it's the default domain.

  - You can't use the EAC to verify that Recipient Lookup is enabled or disabled for the accepted domain. You need to use the Exchange Management Shell.

- In the Exchange Management Shell, run the following command to verify the property values:

  ```
  Get-AcceptedDomain | Format-Table -Auto Name,DomainName,DomainType,Default,AddressBookEnabled
  ```

Remove accepted domains

- You can't remove the default domain. First, you need to configure another accepted domain as the default domain.
- You can't remove an accepted domain that's defined anywhere in an email address policy (including in the disabled email address templates). To see all the domains that are used in email address policies, run the following command in the Exchange Management Shell:

  ```
  Get-EmailAddressPolicy | Format-List Name,*EmailAddressTemplate*
  ```

  For more information about modifying email address policies, see Modify email address policies[Modify email address policies]).

Use the EAC to remove accepted domains

1. In the EAC, go to **Mail flow** > **Accepted domains**, select the accepted domain from the list, and then click **Remove** (−).
2. In the **Warning** dialog that appears, click **Yes** to confirm.

Use the Exchange Management Shell to remove accepted domains

To remove an accepted domain, use the following syntax:

```
Remove-AcceptedDomain -Identity <AcceptedDomainIdentity>
```

This example removes the accepted domain named Fabrikam Corp.

```
Remove-AcceptedDomain -Identity "Fabrikam Corp"
```

For detailed syntax and parameter information, see remove-AcceptedDomain.

How do you know this worked?

To verify that you've successfully removed an accepted domain, use either of the following procedures:

- In the EAC, go to **Mail flow** > **Accepted domains**, and verify that the accepted domain is no longer listed.
- In the Exchange Management Shell, run the following command to verify that the accepted domain isn't listed:

  ```
  Get-AcceptedDomain
  ```

Configure Exchange to accept mail for multiple authoritative domains

These are some scenarios that require multiple authoritative domains:

- Your organization is changing its SMTP domain name.
- You want to provision different email addresses for different business units within your organization.
- You provide email hosting services, and you need to accept email for more than one SMTP domain.

After you configure one or more authoritative domains, you need to decide how to use those domains in your organization. For example:

- Do you want to replace the existing primary (**Reply-To:**) address for the recipients, or add the new email address as a proxy address?

- Do you want to keep old email addresses as a proxy addresses so the recipients continue to receive mail that's sent to their old email addresses?
- Do you want the email addresses in the new authoritative domain to apply to all recipients and all types of recipients, or only to specific types of recipients or specific recipients based on their user properties (for example, only users in the Finance department)?

These are the steps that are required to configure Exchange to accept mail for multiple authoritative domains:

1. Create one or more authoritative domains as described in the Create accepted domains section.

   For example, if you already have contoso.com configured as an authoritative domain, add fourthcoffee.com as an authoritative domain.

2. Create or modify an email address policy that uses the authoritative domains to meet your requirements. For example:

   - Modify the default email address policy so all recipients get the required primary and proxy email addresses.

   For example, modify the default policy so *<alias>*@fourthcoffee.com is the primary SMTP email address, and *<alias>*@contoso.com is kept as a proxy address. For instructions, see Modify email address policies.

   - Create a new email address policy that applies the required primary and proxy email addresses to a filtered set of recipients.

   For example, create a new policy named Fourth Coffee Recipients with the following settings:

   - **Precanned recipient filter**: All users with mailboxes where the **Company** value is Fourth Coffee.
   - **Primary SMTP email address**: *<alias>*@fourthcoffee.com.
   - **Additional proxy email addresses**: None. The affected recipients can no longer receive messages sent to their old @contoso.com primary email address.
   - **Priority**: 1. The first email address policy that identifies a recipient configures the recipient's email addresses. All other policies are ignored, even if the first policy is unapplied and can't configure the recipient's email addresses.

     For instructions, see Create email address policies.

3. Apply the new or updated email address policy to the affected recipients. For instructions, see Apply email address policies to recipients.

To verify that you've configured Exchange to accept mail for multiple authoritative domains, perform the following procedures:

1. Send test messages to an affected recipient from a mailbox that's outside of your Exchange organization, and verify the email addresses that accept or reject mail.

2. Send test messages from an affected mailbox to an external recipient, and verify the From address of the message.

Connectors

8/3/2020 • 2 minutes to read • Edit Online

Exchange uses connectors to enable incoming and outgoing mail flow on Exchange servers, and also between services in the transport pipeline on the local Exchange server.

These are the types of connectors that are available in Exchange.

CONNECTOR	DESCRIPTION
Receive connectors	Receive connectors control incoming SMTP mail flow. They listen for incoming connections that match the configuration of the connector. Multiple default Receive connectors are created when you install Exchange. For more information, see Receive connectors.
Send connectors	Send connectors control outgoing SMTP mail flow. A Send connector is chosen based on the message recipients and the configuration of the connector. No default Send connectors for external mail flow are created when you install Exchange, but implicit and invisible Send connectors exist, and are used to route mail between internal Exchange servers. For more information, see Send connectors.
Delivery agents and Delivery Agent Connectors	Delivery agents and Delivery Agent connectors control outgoing mail flow to non-SMTP systems. Outgoing messages are put into message queues for delivery to the non-SMTP system. Delivery agents and Delivery agent connectors are preferred over Foreign connectors due to their improved performance and management. For more information, see Delivery Agents and Delivery Agent Connectors.
Foreign connectors	Foreign connectors control outgoing mail flow to non-SMTP systems. Outgoing messages are written to files in a location called the Drop directory to be picked up by the non-SMTP system. For information, see Foreign Connectors.

Receive connectors

8/3/2020 • 17 minutes to read • Edit Online

Exchange servers use Receive connectors to control inbound SMTP connections from:

- Messaging servers that are external to the Exchange organization.
- Services in the transport pipeline on the local Exchange server or on remote Exchange servers.
- Email clients that need to use authenticated SMTP to send messages.

You can create Receive connectors in the Transport service on Mailbox servers, the Front End Transport service on Mailbox servers, and on Edge Transport servers. By default, the Receive connectors that are required for inbound mail flow are created automatically when you install an Exchange Mailbox server, and when you subscribe an Edge Transport server to your Exchange organization.

A Receive connector is associated with the Mailbox server or Edge Transport server where it's created, and determines how that specific server listens for SMTP connections. On Mailbox servers, the Receive connector is stored in Active Directory as a child object of the server. On Edge Transport servers, the Receive connector is stored in Active Directory Lightweight Directory Services (AD LDS).

These are the important settings on Receive connectors:

- **Local adapter bindings**: Configure the combination of local IP addresses and TCP ports that the Receive connector uses to accept connections.
- **Remote network settings**: Configure the source IP addresses that the Receive connector listens to for connections.
- **Usage type**: Configure the default permission groups and smart host authentication mechanisms for the Receive connector.
- **Permission groups**: Configure who's allowed to use the Receive connector, and the permissions that they receive.

A Receive connector listens for inbound connections that match the configuration settings of the connector. Each Receive connector on the Exchange server uses a unique combination of local IP address bindings, TCP ports, and remote IP address ranges that define if and how connections from SMTP clients or servers are accepted.

Although the default Receive connectors are adequate in most cases, you can create custom Receive connectors for specific scenarios. For example:

- To apply special properties to an email source, for example, a larger maximum message size, more recipients per message or more simultaneous inbound connections.
- To accept encrypted mail by using a specific TLS certificate.

On Mailbox servers, you can create and manage Receive connectors in the Exchange admin center (EAC) or in the Exchange Management Shell. On Edge Transport servers, you can only use the Exchange Management Shell.

Receive connector changes in Exchange Server

These are the notable changes to Receive connectors in Exchange 2016 and Exchange 2019 compared to Exchange 2010:

- The *TlsCertificateName* parameter allows you to specify the certificate issuer and the certificate subject. This helps minimize the risk of fraudulent certificates.

- The *TransportRole* parameter allows you to distinguish between frontend (Client Access) and backend connectors on Mailbox servers.

Default Receive connectors created during setup

Several different Receive connectors are created by default when you install Exchange. By default, these connectors are enabled, and protocol logging is disabled for most of them. For more information about protocol logging on Receive connectors, see Protocol logging.

Default Receive connectors in the Front End Transport service on Mailbox servers

The primary function of Receive connectors in the Front End Transport service is to accept anonymous and authenticated SMTP connections into your Exchange organization. The **TransportRole** property value for these connectors is `FrontendTransport`. The Front End Transport service relays or *proxies* these connections to the Transport service for categorization and routing to the final destination.

The default Receive connectors that are created in the Front End Transport service on Mailbox servers are described in the following table.

NAME	DESCRIPTION	PROTOCOL LOGGING	TCP PORT	LOCAL IP ADDRESS BINDINGS	REMOTE IP ADDRESS RANGES	AUTHENTICATION MECHANISMS	PERMISSION GROUPS

Client Frontend *<ServerName>*	Accepts connections from authenticated SMTP clients.	None	587	All available IPv4 and IPv6 addresses (`0.0.0.0` and `[::]:`)	`{::-ffff:ffff:ffff` `ffff:ffff:ffff,` `0.0.0.0-255.25` (all IPv4 and IPv6 addresses)	`TLS`
						`BasicAuth`
						`BasicAuthRequireTLS`
						`Integrated`
							`ExchangeUsers`

Default Frontend *<ServerName>*	Accepts anonymous connections from external SMTP servers. This is the common messaging entry point into your Exchange organization.	Verbose	25	All available IPv4 and IPv6 addresses (`0.0.0.0` and `[::]:`)	`{::-ffff:ffff:ffff` `ffff:` `0.0.0.0-255.25` (all IPv4 and IPv6 addresses)	`TLS`
						`BasicAuth`
						`BasicAuthRequi`
						`ExchangeServer`
						`Integrated`
							`AnonymousUsers`
							`ExchangeLegacyServ`
							`ExchangeServers`

NAME	DESCRIPTION	PROTOCOL LOGGING	TCP PORT	LOCAL IP ADDRESS BINDINGS	REMOTE IP ADDRESS RANGES	AUTHENTICATION MECHANISMS	PERMISSION GROUPS

Outbound Proxy Frontend *<ServerName>*	Accepts authenticated connections from the Transport service on Mailbox servers. The connections are encrypted with the Exchange server's self-signed certificate. This connector is used only if the Send connector is configured to use outbound proxy. For more information, see Configure Send connectors to proxy outbound mail.	None	717	All available IPv4 and IPv6 addresses (`0.0.0.0` and `[::]:`)	`{::-ffff:ffff:ffff ffff:ffff:ffff, 0.0.0.0-255.25` (all IPv4 and IPv6 addresses)	`TLS`
						`BasicAuth`
						`BasicAuthRequireTLS`
						`ExchangeServer`
						`Integrated`
							`ExchangeServers`

Default Receive connectors in the Transport service on Mailbox servers

The primary function of Receive connectors in the Transport service is to accept authenticated and encrypted SMTP connections from other transport services on the local Mailbox server or remote Mailbox servers in your organization. The **TransportRole** property value on theses connectors is `HubTransport`. Clients don't directly connect to these connectors.

The default Receive connectors that are created in the Transport service on Mailbox servers are described in the following table.

NAME	DESCRIPTION	PROTOCOL LOGGING	TCP PORT	LOCAL IP ADDRESS BINDINGS	REMOTE IP ADDRESS RANGES	AUTHENTICATION MECHANISMS	PERMISSION GROUPS

NAME	DESCRIPTION	PROTOCOL LOGGING	TCP PORT	LOCAL IP ADDRESS BINDINGS	REMOTE IP ADDRESS RANGES	AUTHENTICATION MECHANISMS	PERMISSION GROUPS

Client Proxy *<ServerName>*	Accepts authenticated client connections that are proxied from the Front End Transport service.	None	465	All available IPv4 and IPv6 addresses (`0.0.0.0` and `[::]:`)	`{::- ffff:ffff:ffff ffff: 0.0.0.0-255.25` (all IPv4 and IPv6 addresses)	`TLS`
						`BasicAuth`
						`BasicAuthRequireTLS`
						`ExchangeServer`
						`Integrated`
							`ExchangeServers`
							`ExchangeUsers`

Default *<ServerName>*	Accepts authenticated connections from:
	• The Front End Transport service on the local or remote Mailbox servers
	• The Transport service on remote Mailbox servers
	• The Mailbox Transport service on the local or remote Mailbox servers
	• Edge Transport servers
	The connections are encrypted with the Exchange server's self-signed certificate.
		None	2525	All available IPv4 and IPv6 addresses (`0.0.0.0` and `[::]:`)	`{::- ffff:ffff:ffff ffff: 0.0.0.0-255.25` (all IPv4 and IPv6 addresses)	`TLS`
						`BasicAuth`
						`ExchangeServer`
						`Integrated`
							`ExchangeLegacyServ`
							`ExchangeServers`
							`ExchangeUsers`

Default Receive connectors in the Transport service on Edge Transport servers

The primary function of the Receive connector on Edge Transport servers is to accept mail from the Internet. Subscribing the Edge Transport server to your Exchange organization automatically configures the connector permissions and authentication mechanisms that are required for Internet mail flow to and from your organization. For more information, see Edge Transport servers.

The default Receive connector that's created in the Transport service on Edge Transport servers is described in the following table.

NAME	DESCRIPTION	PROTOCOL LOGGING	TCP PORT	LOCAL IP ADDRESS BINDINGS	REMOTE IP ADDRESS RANGES	AUTHENTICATION MECHANISMS	PERMISSION GROUPS

Default internal receive connector *<ServerName>*	Accepts anonymous connections from external SMTP servers.	None	25	All available IPv4 addresses (`0.0.0.0`)	`{0.0.0.0-255.255.255.25` (all IPv4 addresses)	`TLS`
						`ExchangeServer`
							`AnonymousUsers`
							`ExchangeServers`
							`Partners`

Implicit Receive connectors in the Mailbox Transport Delivery service on Mailbox servers

In addition to the Receive connectors are created during the installation of Exchange servers, there's a special *implicit Receive connector* in the Mailbox Transport Delivery service on Mailbox servers. This implicit Receive connector is automatically available, invisible, and requires no management. The primary function of this connector is to accept mail from the Transport service on the local Mailbox server or remote Mailbox servers in your organization.

The implicit Receive connector that exists in the Mailbox Transport Delivery service on Mailbox servers is described in the following table.

NAME	DESCRIPTION	PROTOCOL LOGGING	TCP PORT	LOCAL IP ADDRESS BINDINGS	REMOTE IP ADDRESS RANGES	AUTHENTICATION MECHANISMS	PERMISSION GROUPS
Mailbox delivery Receive connector	Accepts authenticated connections from the Transport service on the local or remote Mailbox servers.	None	475	All available IPv4 and IPv6 addresses (`0.0.0.0` and `[::]:`)	`{::-ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff, 0.0.0.0-255.255.255.255}` (all IPv4 and IPv6 addresses)	`ExchangeServer`	`ExchangeServers`

Receive connector local address bindings

Local address bindings restrict the Receive connector to listen for SMTP connections on a specific local IP address (network adapter) and TCP port. Typically, the combination of local IP address and TCP port is unique for every Receive connector on a server. However, multiple Receive connectors on a server can have the same local IP addresses and TCP ports if the remote IP address ranges are different. For more information, see the Receive connector remote addresses section.

By default, a Receive connector listens for connections on all available local IPv4 and IPv6 addresses (`0.0.0.0` and `[::]:`). If the server has multiple network adapters, you can configure Receive connectors to accept connections only from IP addresses that are configured for a specific network adapter. For example, on an Internet-facing Exchange server, you can have a Receive connector that's bound to the IP address of the external network adapter to listen for anonymous Internet connections. You can have a separate Receive connector that's bound to the IP address of the internal network adapter to listen for authenticated connections from internal Exchange servers.

> **NOTE**
>
> If you bind a Receive connector to a specific IP address, make sure that the address is configured on a local network adapter. If you specify an invalid local IP address, the Microsoft Exchange Transport service may fail to start when the server or service is restarted.

In the EAC, you use the **Network adapter bindings** field to configure the local address bindings in the new Receive connector wizard, or on the **Scoping** tab in the properties of existing Receive connectors. In the Exchange Management Shell, you use the *Bindings* parameter on the **New-ReceiveConnector** and **Set-ReceiveConnector** cmdlets. Depending on the usage type that you select, you might not be able to configure the local address bindings when you create the Receive connector, but you can modify them after you create the Receive connector. The affected usage types are identified in the Receive connector usage types section.

Receive connector remote addresses

Remote addresses define from where the Receive connector receives SMTP connections. By default, Receive connectors listen for connections from all IPv4 and IPv6 addresses (0.0.0.0-255.255.255.255 and ::-ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff). If you create a custom Receive connector to receive mail from a specific source, configure the connector to listen for connections only from the specific IP address or address ranges.

Multiple Receive connectors on the server can have overlapping remote IP address ranges as long as one range is completely overlapped by another. When remote IP address ranges overlap, the remote IP address range that has the most specific match to the connecting server's IP address is used.

For example, consider the following Receive connectors in the Front End Transport service on the server named Exchange01:

- **Connector name**: Client Frontend Exchange01
  - **Network adapter bindings**: All available IPv4 on port 25.
  - **Remote network settings**: 0.0.0.0-255.255.255.255
- **Connector name**: Custom Connector A
  - **Network adapter bindings**: All available IPv4 on port 25.
  - **Remote network settings**: 192.168.1.0-192.168.1.255
- **Connector name**: Custom Connector B
  - **Network adapter bindings**: All available IPv4 on port 25.
  - **Remote network settings**: 192.168.1.75

SMTP connections from 192.168.1.75 are accepted by Custom Connector B, because that connector has the most specific IP address match.

SMTP connections from 192.168.1.100 are accepted by Custom Connector A, because that connector has the most specific IP address match.

In the EAC, you use the **Remote network settings** field to configure the remote IP addresses in the new Receive connector wizard, or on the **Scoping** tab in the properties of existing Receive connectors. In the Exchange Management Shell, you use the *RemoteIPRanges* parameter on the **New-ReceiveConnector** and **Set-ReceiveConnector** cmdlets.

Receive connector usage types

The usage type determines the default security settings for the Receive connector. The usage type specifies who is authorized to use the connector, the permissions they get, and the authentication methods that are supported.

When you use the EAC to create Receive connectors, the wizard prompts you to select the **Type** value for the connector. When you use the **New-ReceiveConnector** cmdlet in the Exchange Management Shell, you use the *Usage* parameter with one of the available values (for example, `-Usage Custom`), or the designated switch for the usage type (for example, `-Custom`).

You can specify the connector usage type only when you create Receive connectors. After you create a connector, you can modify the available authentication mechanisms and permission groups in the EAC, or by using the **Set-ReceiveConnector** cmdlet in the Exchange Management Shell.

The available usage types are described in the following table.

USAGE TYPE	PERMISSION GROUPS ASSIGNED	AUTHENTICATION MECHANISMS AVAILABLE	COMMENTS
Client	**Exchange users** (`ExchangeUsers`)	**Transport Layer Security** (`TLS`) **Basic authentication** (`BasicAuth`) **Offer basic authentication only after starting TLS** (`BasicAuthRequireTLS`) **Integrated Windows authentication** (`Integrated`)	Used by POP3 and IMAP4 clients that need to submit email messages by using authenticated SMTP. When you create a Receive connector of this usage type in the EAC or in the Exchange Management Shell, you can't select the local IP address bindings or TCP port. By default, this usage type is bound to all local IPv4 and IPv6 addresses on TCP port 587. You can change these bindings after you create the connector. This usage type isn't available on Edge Transport servers.
Custom	None selected (`None`)	**Transport Layer Security** (`TLS`)	Used in cross-forest scenarios, for receiving mail from third-party messaging servers, and for external relay. After you create a Receive connector of this usage type, you need to add permissions groups in the EAC or in the Exchange Management Shell.
Internal	**Legacy Exchange servers** (`ExchangeLegacyServers`) **Exchange servers** (`ExchangeServers`)	**Transport Layer Security** (`TLS`) **Exchange Server authentication** (`ExchangeServers`)	Used in cross-forest scenarios, for receiving mail from previous versions of Exchange, for receiving mail from third-party messaging servers, or on Edge Transport servers to receive outbound mail from the internal Exchange organization. When you create a Receive connector of this usage type in the EAC or in the Exchange Management Shell, you can't select the local IP address bindings or TCP port. By default, the connector is bound to all local IPv4 and IPv6 addresses on TCP port 25. You can change these bindings after you create the connector. The `ExchangeLegacyServers` permission group isn't available on Edge Transport servers.

USAGE TYPE	PERMISSION GROUPS ASSIGNED	AUTHENTICATION MECHANISMS AVAILABLE	COMMENTS
Internet	**Anonymous users** (`AnonymousUsers`)	**Transport Layer Security** (`TLS`)	Used to receive mail from the Internet. When you create a Receive connector of this usage type in the EAC or in the Exchange Management Shell, you can't select the remote IP addresses. By default, the connector accepts remote connections from all IPv4 addresses (0.0.0.0-255.255.255.255). You can change these bindings after you create the connector.
Partner	**Partners** (`Partners`)	**Transport Layer Security** (`TLS`)	Used to configure secure communication with an external partner (mutual TLS authentication, also known as domain secure).

Receive connector authentication mechanisms

Authentication mechanisms specify the logon and encryption settings that are used for incoming SMTP connections. You can configure multiple authentication mechanisms for a Receive connector. In the EAC, authentication mechanisms are available in the **Security** tab in the properties of the Receive connector. In the Exchange Management Shell, permission groups are available in the *AuthMechanisms* parameter on the **New-ReceiveConnector** and **Set-ReceiveConnector** cmdlets.

The available authentication mechanisms are described in the following table.

AUTHENTICATION MECHANISM	DESCRIPTION
None selected (`None`)	No authentication.
Transport Layer Security (TLS) (`TLS`)	Advertise **STARTTLS** in the EHLO response. TLS encrypted connections require a server certificate that includes the name that the Receive connector advertises in the EHLO response. For more information, see Modify the SMTP banner on Receive connectors. Other Exchange servers in your organization trust the server's self-signed certificate, but clients and external servers typically use a trusted third-party certificate.
Basic authentication (`BasicAuth`)	Basic authentication (clear text).
Offer basic authentication only after starting TLS (`BasicAuthRequireTLS`)	Basic authentication that's encrypted with TLS.
Integrated Windows authentication (`Integrated`)	NTLM and Kerberos authentication.
Exchange Server authentication (`ExchangeServer`)	Generic Security Services application programming interface (GSSAPI) and Mutual GSSAPI authentication.

AUTHENTICATION MECHANISM	DESCRIPTION
Externally secured (`ExternalAuthoritative`)	The connection is presumed to be secured by using a security mechanism that's external to Exchange. The connection may be an Internet Protocol security (IPsec) association or a virtual private network (VPN). Alternatively, the servers may reside in a trusted, physically controlled network. This authentication mechanism requires the `ExchangeServers` permission group. This combination of authentication mechanism and security group permits the resolution of anonymous sender email addresses for messages that are received through the connector.

Receive connector permission groups

A *permission group* is a predefined set of permissions that's granted to well-known security principals and assigned to a Receive connector. Security principals include user accounts, computer accounts, and security groups (objects that are identifiable by a security identifier or SID that can have permissions assigned to them). Permission groups define who can use the Receive connector, and the permissions that they get. You can't create permission groups, nor can you modify the permission group members or the default permissions of the permission group.

In the EAC, permission groups are available in the **Security** tab in the properties of the Receive connector. In the Exchange Management Shell, permission groups are available in the *PermissionGroups* parameter in the **New-ReceiveConnector** and **Set-ReceiveConnector** cmdlets.

The available permission groups are described in the following table.

PERMISSION GROUP	ASSOCIATED SECURITY PRINCIPALS	PERMISSIONS GRANTED

Anonymous users (`Anonymous`)	`NT AUTHORITY\ANONYMOUS LOGON`	`ms-Exch-Accept-Headers-Routing`
		`ms-Exch-SMTP-Accept-Any-Sender`
		`ms-Exch-SMTP-Accept-Authoritative-Domain-Sender`
		`ms-Exch-SMTP-Submit`

Exchange users (`ExchangeUsers`)	`NT AUTHORITY\Authenticated Users`	`ms-Exch-Accept-Headers-Routing`
		`ms-Exch-Bypass-Anti-Spam`
		`ms-Exch-SMTP-Accept-Any-Recipient`
		`ms-Exch-SMTP-Submit`

Exchange servers (`ExchangeServers`)	`<Domain>\Exchange Servers`
	`MS Exchange\Edge Transport Servers`
	`MS Exchange\Hub Transport Servers`
	Note: These security principals also have other internal permissions assigned to them. For more information, see the end of the Receive connector permissions section.
		`ms-Exch-Accept-Headers-Forest`
		`ms-Exch-Accept-Headers-Organization`
		`ms-Exch-Accept-Headers-Routing`
		`ms-Exch-Bypass-Anti-Spam`
		`ms-Exch-Bypass-Message-Size-Limit`
		`ms-Exch-SMTP-Accept-Any-Recipient`
		`ms-Exch-SMTP-Accept-Any-Sender`
		`ms-Exch-SMTP-Accept-Authentication-Flag`
		`ms-Exch-SMTP-Accept-Authoritative-Domain-Sender`
		`ms-Exch-SMTP-Accept-Exch50`
		`ms-Exch-SMTP-Submit`

PERMISSION GROUP	ASSOCIATED SECURITY PRINCIPALS	PERMISSIONS GRANTED

Exchange servers (`ExchangeServers`)	`MS Exchange\Externally Secured Servers`	`ms-Exch-Accept-Headers-Routing`
		`ms-Exch-Bypass-Anti-Spam`
		`ms-Exch-Bypass-Message-Size-Limit`
		`s-Exch-SMTP-Accept-Any-Recipient`
		`ms-Exch-SMTP-Accept-Any-Sender`
		`ms-Exch-SMTP-Accept-Authentication-Flag`
		`ms-Exch-SMTP-Accept-Authoritative-Domain-Sender`
		`ms-Exch-SMTP-Accept-Exch50`
		`ms-Exch-SMTP-Submit`

Legacy Exchange servers (`ExchangeLegacyServers`)	`<Domain>\ExchangeLegacyInterop`	`ms-Exch-Accept-Headers-Routing`
		`ms-Exch-Bypass-Anti-Spam`
		`ms-Exch-Bypass-Message-Size-Limit`
		`ms-Exch-SMTP-Accept-Any-Recipient`
		`ms-Exch-SMTP-Accept-Any-Sender`
		`ms-Exch-SMTP-Accept-Authentication-Flag`
		`ms-Exch-SMTP-Accept-Authoritative-Domain-Sender`
		`ms-Exch-SMTP-Accept-Exch50`
		`ms-Exch-SMTP-Submit`

Partners (`Partner`)	`MS Exchange\Partner Servers`	`ms-Exch-Accept-Headers-Routing` `ms-Exch-SMTP-Submit`

The permissions are explained in the Receive connector permissions section later in this topic.

Receive connector permissions

Typically, you apply permissions to Receive connectors by using permission groups. However, you can configure granular permissions on a Receive connector by using the **Add-ADPermission** and **Remove-ADPermission** cmdlets.

Receive connector permissions are assigned to security principals by the permission groups for the connector. When an SMTP server or client establishes a connection to a Receive connector, the Receive connector permissions determine whether the connection is accepted, and how messages are processed.

The available Receive connector permissions are described in the following table.

RECEIVE CONNECTOR PERMISSION	DESCRIPTION
`ms-Exch-Accept-Headers-Forest`	Controls the preservation of Exchange forest headers in messages. Forest header names start with **X-MS-Exchange-Forest-**. If this permission isn't granted, all forest headers are removed from messages.
`ms-Exch-Accept-Headers-Organization`	Controls the preservation of Exchange organization headers in messages. Organization header names start with **X-MS-Exchange-Organization-**. If this permission isn't granted, all organization headers are removed from messages.
`ms-Exch-Accept-Headers-Routing`	Controls the preservation of **Received** and **Resent-*** headers in messages. If this permission isn't granted, all of these headers are removed from messages.
`ms-Exch-Bypass-Anti-Spam`	Allows SMTP clients or servers to bypass antispam filtering.

RECEIVE CONNECTOR PERMISSION	DESCRIPTION
`ms-Exch-Bypass-Message-Size-Limit`	Allows SMTP clients or servers to submit messages that exceed the maximum message size that's configured for the Receive connector.
`ms-Exch-SMTP-Accept-Any-Recipient`	Allows SMTP clients or servers to relay messages through the Receive connector. If this permission isn't granted, only messages that are sent to recipients in accepted domains that are configured for the Exchange organization are accepted by the Receive connector.
`ms-Exch-SMTP-Accept-Any-Sender`	Allows SMTP clients or servers to bypass the sender address spoofing check that normally requires the sender's email address to be in an accepted domain that's configured for Exchange organization.
`ms-Exch-SMTP-Accept-Authentication-Flag`	Controls whether messages from SMTP clients or servers are treated as authenticated. If this permission isn't granted, messages from theses sources are identified as external (unauthenticated). This setting is important for distribution groups that are configured to accept mail only from internal recipients (for example, the *RequireSenderAuthenticationEnabled* parameter value for the group is `$true`).
`ms-Exch-SMTP-Accept-Authoritative-Domain-Sender`	Allows access to the Receive connector by senders that have email addresses in authoritative domains that are configured for the Exchange organization.
`ms-Exch-SMTP-Accept-Exch50`	Allows SMTP clients or servers to submit **XEXCH50** commands on the Receive connector. The **X-EXCH50** binary large object (BLOB) was used by older versions of Exchange (Exchange 2003 and earlier) to store Exchange data in messages (for example, the spam confidence level or SCL).
`ms-Exch-SMTP-Submit`	This permission is required to submit messages to Receive connectors. If this permission isn't granted, the **MAIL FROM** and **AUTH** commands will fail.

Notes:

- In addition to the documented permissions, there are permissions that are assigned to all of the security principals in the **Exchange servers** (`ExchangeServers`) permission group except `MS Exchange\Externally Secured Servers` . These permissions are reserved for internal Microsoft use, and are presented here for reference purposes only.
  - `ms-Exch-SMTP-Accept-Xattr`
  - `ms-Exch-SMTP-Accept-XProxyFrom`
  - `ms-Exch-SMTP-Accept-XSessionParams`
  - `ms-Exch-SMTP-Accept-XShadow`
  - `ms-Exch-SMTP-Accept-XSysProbe`
  - `ms-Exch-SMTP-Send-XMessageContext-ADRecipientCache`
  - `ms-Exch-SMTP-Send-XMessageContext-ExtendedProperties`

  - `ms-Exch-SMTP-Send-XMessageContext-FastIndex`

- Permissions names that contain `ms-Exch-Accept-Headers-` are part of the *header firewall* feature. For more information, see Header firewall.

Receive connector permission procedures

To see the permissions that are assigned to security principals on a Receive connector, use the following syntax in the Exchange Management Shell:

```
Get-ADPermission -Identity <ReceiveConnector> [-User <SecurityPrincipal>] | where {($_.Deny -eq $false) -and
($_.IsInherited -eq $false)} | Format-Table User,ExtendedRights
```

For example, to see the permissions that are assigned to all security principals on the Receive connector named Client Frontend Mailbox01, run the following command:

```
Get-ADPermission -Identity "Client Frontend Mailbox01" | where {($_.Deny -eq $false) -and ($_.IsInherited -eq
$false)} | Format-Table User,ExtendedRights
```

To see the permissions that are assigned only to the security principal `NT AUTHORITY\Authenticated Users` on the Receive connector named Default Mailbox01, run the following command:

```
Get-ADPermission -Identity "Default Mailbox01" -User "NT AUTHORITY\Authenticated Users" | where {($_.Deny -eq
$false) -and ($_.IsInherited -eq $false)} | Format-Table User,ExtendedRights
```

To add permissions to a security principal on a Receive connector, use the following syntax:

```
Add-ADPermission -Identity <ReceiveConnector> -User <SecurityPrincipal> -ExtendedRights "<Permission1>","
<Permission2>"...
```

To remove permissions from a security principal on a Receive connector, use the following syntax:

```
Remove-ADPermission -Identity <ReceiveConnector> -User <SecurityPrincipal> -ExtendedRights "<Permission1>","
<Permission2>"...
```

Scenarios for custom Receive connectors in Exchange Server

8/3/2020 • 14 minutes to read • Edit Online

By default, Exchange Server comes with many different Receive connectors that are configured for most common mail flow scenarios. For more information about these connectors, see Default Receive connectors created during setup.

However, you might need to process messages from another messaging system that's not running Exchange. Or, if you have a network appliance that performs policy checks and then routes messages to your Exchange server, you'll need to manually configure a Receive connector.

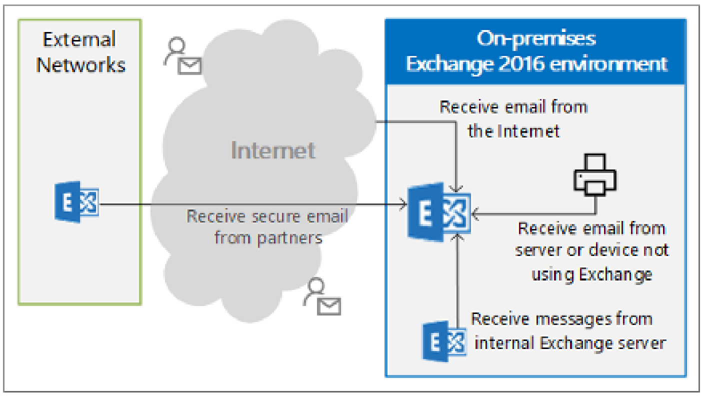
External Networks
Internet
Receive secure email from partners
On-premises Exchange 2016 environment
Receive email from the Internet
Receive email from server or device not using Exchange
Receive messages from internal Exchange server

If you need to create a custom Receive connector, consider these issues:

- You can create custom Receive connectors in the following services on Exchange servers:
  - **Mailbox servers**: The Transport (Hub) service and the Front End Transport service.
  - **Edge Transport servers**: The Transport service.
- Each Receive connector on an Exchange server requires a unique combination of network adapter bindings (the combination of **local IP address** and **TCP port**) and **remote network settings** (remote IP addresses).
  - A default Receive connector that listens on port 25 on all available local IP addresses from all remote IP addresses already exists on all Mailbox servers and Edge Transport servers.
  - If you create a custom Receive connector that listens on port 25 on all available local IP addresses, but the connector is restricted to a limited range of *remote* IP addresses, the new connector won't conflict with any of the default Receive connectors on the server. For a detailed explanation, see Receive connector remote addresses.

    If you can't restrict the remote IP addresses of the custom Receive connector, your only other option is to restrict the local IP address that the connector uses for port 25. You'll need to modify the local IP address of the conflicting default Receive connector, and then use a different local IP address when you create custom Receive connector.
- For Mailbox servers, you need to create custom Receive connectors that use port 25 in the Front End Transport service, not the Transport (Hub) service. Receive connectors in the Transport service on Mailbox servers accept authenticated and encrypted SMTP connections from other transport services on the local server or other Mailbox servers in your organization. Clients don't directly connect to these connectors. This is different than Exchange 2010, because you could only create Receive connectors on Hub Transport servers

(not Client Access servers).

Read more about Receive connectors in Exchange Server see, Receive connectors.

What do you need to know before you begin?

- Estimated time to complete each procedure: 10 minutes
- The Exchange admin center (EAC) procedures are only available on Mailbox servers. For more information about the EAC, see Exchange admin center in Exchange Server.
- The Exchange Management Shell procedures are available on Mailbox servers and Edge Transport servers. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Receive connectors" entry in the Mail flow permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Scenario 1: Receive email from the Internet

For this scenario, the Receive connector listens for anonymous SMTP connections on port 25 from all remote IP addresses. Typically, you don't need to manually configure a Receive connector to receive mail from the Internet. A Receive connector with these settings is automatically created by the installation of a Mailbox server or an Edge Transport server:

- The Receive connector named Default Frontend *<ServerName>* in the Front End Transport service on Mailbox servers.
- The Receive connector named Default internal receive connector *<ServerName>* on Edge Transport servers.

If one of these connectors exists, and you try to create a custom Receive connector on the server that also listens for anonymous SMTP connections on port 25 from all remote IP addresses, you'll get an error. You'll need to change the network adapter binding on the conflicting Receive connector to a specific local IP address. When you create the custom Internet Receive connector, you'll need to specify a different network adapter binding.

Use the EAC to create an Internet Receive connector on Mailbox servers

1. In the EAC, go to **Mail flow** > **Receive connectors**, and then click **Add** (+).
2. The **New receive connector** wizard opens. On the first page, configure these settings:
   - **Name**: Type something descriptive. For example, Internet Receive Connector.
   - **Role**: Select **Frontend Transport**.
   - **Type**: Select **Internet**.

   When you're finished, click **Next**.
3. On the last page of the wizard, do one of these steps in the **Network adapter bindings** section:

- If you're recreating an Internet Receive connector to replace the missing default Receive connector named Default Frontend *<ServerName>* on the Mailbox server, leave the default values of **IP addresses: (All available IPv4)** and **Port: 25** (when you click **Finish**, you won't receive an error message).

- If you're creating an Internet Receive connector while the default Receive connector named Default Frontend *<ServerName>* still exists on the Mailbox server, do these steps:

  a. Select the default entry **IP addresses: (All available IPv4)** and **Port: 25**, and then click **Edit** (✎).

  b. In the **Edit IP address** dialog that opens, configure these settings:

     - **Address**: Select **Specify an IPv4 address or an IPv6 address**, and type in a valid local IP address to use for the connector.

     - **Port**: Leave the default value **25** selected.

     When you're finished, click **Save**.

> **NOTE**
>
> After you've created the new Internet Receive connector on the Mailbox server, be sure to modify the local IP address settings in the properties of the default Receive connector named Default Frontend *<ServerName>*. You'll need to go to **Scoping** > **Network adapter bindings** in the properties of the connector, and then select a different local IP address to replace the default **IP addresses: (All available IPv4)** and **Port: 25** entry.

When you're finished, click **Finish**.

Use the Exchange Management Shell to create an Internet Receive connector

To create an Internet Receive connector, use this syntax:

```
New-ReceiveConnector -Name <UniqueName> [-TransportRole Frontend] -Internet -Bindings
<UniqueValidLocalIPAddress>
```

This example creates a new Receive connector named Internet Receive Connector on a Mailbox server that listens on port 25 on the local IP address 10.1.15 from all remote IP addresses:

```
New-ReceiveConnector -Name "Internet Receive Connector" -TransportRole Frontend -Internet -Bindings
10.10.1.1:25
```

Notes:

- To run this command on an Edge Transport server, omit the *TransportRole* parameter.

- If another Receive connector is configured to listen on port 25 using all available local IP addresses on the server, you'll need to use the *Bindings* parameter on the **Set-ReceiveConnector** cmdlet to specify a unique local IP address for the other connector after you create the new Internet Receive connector.

This example creates a new Receive connector named Internet Receive Connector that listens on port 25 from all remote IP addresses, but on all available local IP addresses. You can only run this command if the server has no other Receive connectors that are configured to listen on port 25 using all available local IP addresses.

```
New-ReceiveConnector -Name "Internet Receive Connector" -TransportRole Frontend -Internet -Bindings "0.0.0.0","
[::]:"
```

Note: To run this command on an Edge Transport server, omit the *TransportRole* parameter.

For detailed syntax and parameter information, see New-ReceiveConnector.

How do you know this worked?

To verify that you've successfully created a Receive connector to receive messages from the Internet, do any of these steps:

- In the EAC, go to **Mail flow** > **Receive connectors**, select the Receive connector, select **Edit** (✎), and verify the property values
- In the Exchange Management Shell, run this command on the server, and verify the property values:

```
Get-ReceiveConnector | where {$_.Bindings -like '*25' -AND $_.PermissionGroups -like '*AnonymousUsers*'}
| Format-List Identity,Bindings,RemoteIPRanges,PermissionGroups
```

- Enable protocol logging for the Receive connector. For more information, see Configure protocol logging.
- From an external client, send a test message to someone in your organization. You can also connect to the Receive connector by using Telnet. For more information, see Use Telnet to test SMTP communication on Exchange servers.

Scenario 2: Receive email from a partner

For this scenario, the Receive connector listens for TLS authenticated SMTP connections on port 25, but only from the specific IP addresses of the partner organization. No default Receive connector is suitable for this scenario; you need to create a custom Receive connector.

Note: Creating a dedicated Receive connector is only one step in TLS encrypting communication between your organization a trusted partner (for example, creating and installing certificates).

Use the EAC to create a Receive connector to encrypt messages from a partner on Mailbox servers

1. In the EAC, go to **Mail flow** > **Receive connectors**, and then click **Add** (+).
2. The **New receive connector** wizard opens. On the first page, configure these settings:
   - **Name**: Type something descriptive. For example, TLS Encrypted Messages from Fabrikam.com.
   - **Role**: Select **Frontend Transport**.
   - **Type**: Select **Partner**.

   When you're finished, click **Next**.
3. On the second page of the wizard, do one of these steps in the **Network adapter bindings** section:
   - Leave the default values of **IP addresses**: **(All available IPv4)** and **Port**: **25**.
   - If it's required for your scenario, you can restrict the Receive connector to a valid local IP address on the server:

     a. Select the default entry **IP addresses**: **(All available IPv4)** and **Port**: **25**, and then click **Edit** (✎).

     b. In the **Edit IP address** dialog that opens, configure these settings:

     - **Address**: Select **Specify an IPv4 address or an IPv6 address**, and type in a valid local IP address to use for the connector.
     - **Port**: Leave the default value **25** selected.

    When you're finished, click **Save**.

   When you're finished, click **Next**.

4. On the last page of the wizard, configure these settings in the **Remote network settings** section:

   a. Select the default entry **0.0.0.0-255.255.255.255**, and then click **Edit** (✏).

   b. In the **Edit IP address** dialog that opens, enter the IP address or IP address range of the remote partner organization.

      When you're finished, click **Save**.

   When you're finished, click **Finish**.

Use the Exchange Management Shell to create a Receive connector to encrypt messages from a partner

To create a Receive connector that uses TLS to encrypt messages from a partner, use this syntax:

```
New-ReceiveConnector -Name <UniqueName> [-TransportRole Frontend] -Partner  -Bindings <0.0.0.0:25 |
LocalIPAddress:25> -RemoteIPRanges <RemoteIPAddresses>
```

This example creates a Receive connector named Fabrikam.com TLS on a Mailbox server that only accepts messages from the IP addresses 17.17.17.1/24 using all available local IP addresses.

```
New-ReceiveConnector -Name "Fabrikam.com TLS" -TransportRole Frontend -Partner -RemoteIPRanges 17.17.17.1/24 -
Bindings 0.0.0.0:25
```

Note: To run this command on an Edge Transport server, omit the *TransportRole* parameter.

For detailed syntax and parameter information, see New-ReceiveConnector.

How do you know this worked?

To verify that you've successfully created a Receive connector to receive TLS encrypted messages from a partner, do any of these steps:

- In the EAC, go to **Mail flow** > **Receive connectors**, select the Receive connector, select **Edit** (✏), and verify the property values

- In the Exchange Management Shell, run this command on the server, and verify the property values:

  ```
  Get-ReceiveConnector | where {$_.Bindings -like '*25' -AND $_.PermissionGroups -like '*Partners*'} |
  Format-List Identity,Bindings,RemoteIPRanges,PermissionGroups
  ```

- Enable protocol logging for the Receive connector. For more information, see Configure protocol logging.

- Have someone in the partner organization send a test message to someone in your organization. Verify that the message is encrypted (you can verify that TLS is used by checking the message header).

Scenario 3: Receive messages from a server, service, or device that doesn't use Exchange

For this scenario, the Receive connector listens for connections on port 25, but only from the specific IP address of the service, or device. It's also likely that this scenario requires some type of authentication (consult the documentation for the service or device).

Use the EAC to create a Receive connector that only accepts messages from a specific service or device on Mailbox servers

1. In the EAC, go to **Mail flow** > **Receive connectors**, and then click **Add** (+).

2. The **New receive connector** wizard opens. On the first page, configure these settings:

   - **Name**: Type something descriptive. For example, Inbound mail from security appliance.

   - **Role**: Select **Frontend Transport**.

   - **Type**: Select **Custom**.

   When you're finished, click **Next**.

3. On the second page of the wizard, do one of these steps in the **Network adapter bindings** section:

   - Leave the default values of **IP addresses**: **(All available IPv4)** and **Port**: **25**.

   - If it's required for your scenario, you can restrict the Receive connector to a valid local IP address on the server:

     a. Select the default entry **IP addresses**: **(All available IPv4)** and **Port**: **25**, and then click **Edit** (✎).

     b. In the **Edit IP address** dialog that opens, configure these settings:

        - **Address**: Select **Specify an IPv4 address or an IPv6 address**, and type in a valid local IP address to use for the connector.

        - **Port**: Leave the default value **25** selected.

        When you're finished, click **Save**.

   When you're finished, click **Next**.

4. On the last page of the wizard, configure these settings in the **Remote network settings** section:

   a. Select the default entry **0.0.0.0-255.255.255.255**, and then click **Edit** (✎).

   b. In the **Edit IP address** dialog that opens, enter the IP address or IP address range of the service or device.

      When you're finished, click **Save**.

   When you're finished, click **Finish**.

5. Back at **Mail flow** > **Receive connectors**, select the connector you just created, and then click **Edit** (✎).

6. On the **Security** tab, configure the combination of authentication mechanisms and permission groups that are required for the service or device. For example:

   - Leave **Transport Layer Security (TLS)** selected, select **Basic authentication**, and then select the **Anonymous users** permission group.

   - Clear **Transport Layer Security (TLS)**, select **Basic authentication** and **Exchange server authentication**, and then select the **Exchange users** and **Legacy Exchange servers** permission group.

     For more information about permission groups, see Receive connector permission groups.

     Caution

     Be very careful using the authentication mechanism **Externally secured** with the permission group **Exchange servers**. This combination allows the remote IP addresses specified in the **Remote network settings** section on the **Scoping** tab to anonymously relay messages through the Exchange server. For more information, see Allow anonymous relay on Exchange servers.

> **WARNING**
>
> When using the authentication mechanism **Basic authentication** or **Offer basic authentication only after starting TLS** without the permission group **Anonymous users** as an authenticated relay connector, the routing of mail will always try to select the authenticated user or the organization's arbitration mailbox active mailbox server.

When you're finished, click **Save**.

Use the Exchange Management Shell to create a Receive connector that only accepts messages from a specific service or device

To create a Receive connector that only accepts messages from a specific service or device, use this syntax:

```
New-ReceiveConnector -Name <UniqueName> [-TransportRole Frontend] -Custom -Bindings <0.0.0.0:25 |
LocalIPAddress:25> -RemoteIPRanges <RemoteIPAddresses> -AuthMechanism <AuthMechanism1>,<AuthMechanism2>... -
PermissionGroups <PermissionGroup1>,<PermissionGroup2>...
```

This example creates a Receive connector named Inbound From Service on a Mailbox server:

- **Bindings**: All available local IP addresses.
- **Remote IP address ranges**: 192.168.5.1/24.
- **Authentication mechanisms**: Basic authentication.
- **Permission groups**: Anonymous users.

```
New-ReceiveConnector -Name "Inbound From Service" -TransportRole Frontend -Custom -Bindings 0.0.0.0:25 -
RemoteIPRanges 192.168.10.5 -AuthMechanism BasicAuth -PermissionGroups AnonymousUsers
```

Note: To run this command on an Edge Transport server, omit the *TransportRole* parameter.

For detailed syntax and parameter information, see New-ReceiveConnector.

How do you know this worked?

To verify that you've successfully created a Receive connector that only accepts messages from a specific service or device, do any of these steps:

- In the EAC, go to **Mail flow** > **Receive connectors**, select the Receive connector, select **Edit** (✎), and verify the property values.
- In the Exchange Management Shell, run this command on the server, and verify the property values:

  ```
  Get-ReceiveConnector | where {$_.Bindings -like '*25'} | Format-List
  Identity,RemoteIPRanges,PermissionGroups,AuthMechanism
  ```

- Enable protocol logging for the Receive connector. For more information, see Configure protocol logging.
- Send a test message or connect to the Receive connector by using Telnet from the server or device. For more information, see Use Telnet to test SMTP communication on Exchange servers.

Scenario 4: Receive messages from internal Exchange servers

You don't need to configure custom Receive connectors for internal mail flow between Mailbox servers. However, you might need to create a custom Receive connector on an unsubscribed Edge Transport server to receive messages from Mailbox servers. For this scenario, the Edge Transport server listens on port 25, but only from the IP address of the specified Mailbox servers.

Use the Exchange Management Shell to create a Receive connector that only accepts messages from an internal Exchange server

To create a Receive connector that only accepts messages from an internal Exchange server, use this syntax:

```
New-ReceiveConnector -Name <UniqueName> [-TransportRole Frontend] -Internal -RemoteIPRanges <RemoteIPAddress>
```

This example creates a Receive connector named Inbound From Organization on an unsubscribed Edge Transport server that listens for inbound messages from the internal Mailbox servers at IP addresses 10.1.2.10, 10.1.2.15, and 10.1.2.20.

```
New-ReceiveConnector -Name "Inbound From Organization" -Internal -RemoteIPRanges 10.1.2.10,10.1.2.15,10.1.2.20
```

Note: If your Edge Transport server uses different network adapters for internal and external networks, be sure to use the *Bindings* parameter on the **Set-ReceiveConnector** cmdlet after you create the connector to specify the correct local IP address for the connector.

For detailed syntax and parameter information, see New-ReceiveConnector.

How do you know this worked?

To verify that you've successfully created a Receive connector that only accepts messages from an internal Exchange server, do any of these steps:

- In the EAC, go to **Mail flow** > **Receive connectors**, select the Receive connector, select **Edit** (✎), and verify the property values.

- In the Exchange Management Shell, run this command on the server, and verify the property values:

  ```
  Get-ReceiveConnector | where {$_.Bindings -like '*25'} | Format-List
  Identity,RemoteIPRanges,PermissionGroups,AuthMechanism
  ```

- Enable protocol logging for the Receive connector. For more information, see Configure protocol logging.

- Send a test message or connect to the Receive connector by using Telnet from the remote Exchange server. For more information, see Use Telnet to test SMTP communication on Exchange servers.

Modify the SMTP banner on Receive connectors

8/3/2020 • 2 minutes to read • Edit Online

The *SMTP banner* is the initial SMTP connection response that a messaging server receives after it connects to an Exchange server. Specifically, the messaging server connects to a Receive connector that's configured on the Exchange server. For Exchange Mailbox servers, external messaging servers connect through Receive connectors that are configured in the Front End Transport service. The default Receive connector that's configured to accept anonymous SMTP connections is named Default Frontend *<ServerName>*. For Edge Transport servers, the default Receive connector in the Transport service named Default internal receive connector *<ServerName>*> is configured to accept anonymous SMTP connections. For more information, see How messages from external senders enter the transport pipeline and Default Receive connectors created during setup.

By default, the connection response looks like this:

```
220 <ServerName> Microsoft ESMTP MAIL service ready at <RegionalDay-Date-24HourTimeFormat>
<RegionalTimeZoneOffset>
```

Here are some reasons that you might want to modify the default SMTP banner:

- You don't want Exchange or the internal Exchange server name disclosed in the connection response to external messaging servers.
- You want the connection response to include your domain name to satisfy antispam or reverse DNS to SMTP banner checks.
- You want the connection response to include the name of the Receive connector to make it easier to troubleshoot connection problems.

What do you need to know before you begin?

- Estimated time to complete: 5 minutes
- You can only use PowerShell to perform this procedure. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- The replacement SMTP banner text string must always start with `220` (the default "Service ready" SMTP response code is 220).
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Receive connectors" entry in the Mail flow permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to modify the SMTP banner on a Receive connector

Use the following syntax:

```
Set-ReceiveConnector -Identity <ConnectorIdentity> -Banner "220 <Banner Text>"
```

This example changes the SMTP banner on the Receive connector named Default Frontend Mailbox01 to the value 220 contoso.com.

```
Set-ReceiveConnector -Identity "Default Frontend Mailbox01" -Banner "220 consoso.com"
```

This example removes the custom SMTP banner, which returns the SMTP banner to the default value.

```
Set-ReceiveConnector -Identity "Default Frontend Mailbox01" -Banner $null
```

How do you know this worked?

To verify that you have successfully modified the SMTP banner on a Receive connector, do these steps:

1. Open a Telnet client on a computer that can access the Receive connector, and run the following command:

   ```
   open <Connector FQDN or IP address><TCPPort>
   ```

2. Verify the that response contains the SMTP banner you configured.

Note that this procedure only works on Receive connectors that allow anonymous or Basic authentication. For more information, see Use Telnet to test SMTP communication on Exchange servers.

Allow anonymous relay on Exchange servers

8/3/2020 • 11 minutes to read • Edit Online

Open relay is a very bad thing for messaging servers on the Internet. Messaging servers that are accidentally or intentionally configured as open relays allow mail from any source to be transparently re-routed through the open relay server. This behavior masks the original source of the messages, and makes it look like the mail originated from the open relay server. Open relay servers are eagerly sought out and used by spammers, so you never want your messaging servers to be configured for open relay.

On the other hand, *anonymous relay* is a common requirement for many businesses that have internal web servers, database servers, monitoring applications, or other network devices that generate email messages, but are incapable of actually sending those messages.

In Exchange Server, you can create a dedicated Receive connector in the Front End Transport service on a Mailbox server that allows anonymous relay from a specific list of internal network hosts. Here are some key considerations for the anonymous relay Receive connector:

- You need to create a dedicated Receive connector to specify the network hosts that are allowed to anonymously relay messages, so you can exclude anyone or anything else from using the connector. Don't attempt to add anonymous relay capability to the default Receive connectors that are created by Exchange. Restricting access to the Receive connector is critical, because you don't want to configure the server as an open relay.

- You need to create the dedicated Receive connector in the Front End Transport service, not in the Transport service. In Exchange Server, the Front End Transport service and the Transport service are always located together on Mailbox servers. The Front End Transport service has a default Receive connector named Default Frontend *<ServerName>* that's configured to listen for inbound SMTP connections from any source on TCP port 25. You can create another Receive connector in the Front End Transport service that also listens for incoming SMTP connections on TCP port 25, but you need to specify the IP addresses that are allowed to use the connector. The dedicated Receive connector will always be used for incoming connections from those specific network hosts (the Receive connector that's configured with the most specific match to the connecting server's IP address wins).

  In contrast, the Transport service has a Default receive connector named Default *<ServerName>* that's also configured to listed for inbound SMTP connections from any source, but this connector listens on TCP port 2525 so that it doesn't conflict with the Receive connector in the Front End Transport service. Furthermore, only other transport services and Exchange servers in your organization are expected to use this Receive connector, so the authentication and encryption methods are set accordingly.

  For more information, see Mail flow and the transport pipeline and Default Receive connectors created during setup.

- After you create the dedicated Receive connector, you need to modify its permissions to allow anonymous relay only by the specified network hosts as identified by their IP addresses. At a minimum, the network hosts need the following permissions on the Receive connector to anonymously relay messages:

  - `ms-Exch-Accept-Headers-Routing`
  - `ms-Exch-SMTP-Accept-Any-Recipient`
  - `ms-Exch-SMTP-Accept-Any-Sender`
  - `ms-Exch-SMTP-Accept-Authoritative-Domain-Sender`

- `ms-Exch-SMTP-Submit`

For more information about permissions on Receive connectors, see Receive connector permission groups and Receive connector permissions.

There are two different methods that you can use to configure the permissions that are required for anonymous relay on a Receive connector. These methods are described in the following table.

METHOD	PERMISSIONS GRANTED	PROS	CONS
Add the **Anonymous users** (`Anonymous`) permission group to the Receive connector and add the `Ms-Exch-SMTP-Accept-Any-Recipient` permission to the `NT AUTHORITY\ANONYMOUS LOGON` security principal on the Receive connector.	Connections use the `NT AUTHORITY\ANONYMOUS LOGON` security principal with the following permissions: • `ms-Exch-Accept-Headers-Routing` • `ms-Exch-SMTP-Accept-Any-Recipient` • `ms-Exch-SMTP-Accept-Any-Sender` • `ms-Exch-SMTP-Accept-Authoritative-Domain-Sender` • `ms-Exch-SMTP-Submit`	Grants the minimum required permissions to allow anonymous relay.	More difficult to configure (must use the Exchange Management Shell). The network hosts are considered anonymous senders. Messages don't bypass antispam or message size limit checks, and the sender's email address can't be resolved to the corresponding display name (if any) in the global address list.
Add the **Exchange servers** (`ExchangeServers`) permission group and the **Externally secured** (`ExternalAuthoritative`) authentication mechanism to the Receive connector.	Connections use the `MS Exchange\Externally Secured Servers` security principal with the following permissions: • `ms-Exch-Accept-Headers-Routing` • `ms-Exch-Bypass-Anti-Spam` • `ms-Exch-Bypass-Message-Size-Limit` • `ms-Exch-SMTP-Accept-Any-Recipient` • `ms-Exch-SMTP-Accept-Any-Sender` • `ms-Exch-SMTP-Accept-Authentication-Flag` • `ms-Exch-SMTP-Accept-Authoritative-Domain-Sender` • `ms-Exch-SMTP-Accept-Exch50` • `ms-Exch-SMTP-Submit`	Easier to configure (can do everything in the Exchange admin center). The network hosts are considered authenticated senders. Messages bypass antispam and message size limit checks, and the sender's email address can be resolved to a corresponding display name in the global address list.	Grants the permissions to submit messages as if they originated from internal senders within your Exchange organization. The network hosts are considered completely trustworthy, regardless of the volume, size, or content of the messages that they send.

Ultimately, you need to decide on the approach that best fits the needs of your organization. We'll show you how to configure both methods. Just remember that it's one method or the other, and not both at the same time.

What do you need to know before you begin?

- Estimated time to complete this task: 10 minutes.
- Some of these procedures require the Exchange Management Shell. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Receive connectors" entry in the Mail flow permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Step 1: Create a dedicated Receive connector for anonymous relay

You can create the Receive connector in the EAC or in the Exchange Management Shell.

Use the EAC to create a dedicated Receive connector for anonymous relay

1. In the EAC, navigate to **Mail flow** > **Receive connectors**, and then click **Add** +. This starts the **New Receive connector** wizard.
2. On the first page, enter the following information:
   - **Name**: Enter a descriptive name for the Receive connector, for example, Anonymous Relay.
   - **Role**: Select **Frontend Transport**.
   - **Type**: Select **Custom**.

   When you're finished, click **Next**.
3. On the next page, in the **Network adapter bindings** section, do one of the following:
   - If the Exchange server has one network adapter, and doesn't segregate internal and external traffic by using different subnets, accept the existing **(All available IPv4)** entry on port 25.
   - If the Exchange server has an internal network adapter and an external network adapter, and segregates internal and external network traffic by using different subnets, you can further enhance security for the connector by limiting the use of the connector to requests that originate on the internal network adapter. To do this:
     a. Select the existing **(All available IPv4)** entry, click **Remove** −, and then click **Add** +.
     b. In the resulting **Network Adapter Bindings** dialog, select **Specify an IPv4 address or an IPv6 address**, and enter a valid and available IP address that's configured on the internal network adapter, and then click **Save**.

   When you're finished, click **Next**.
4. On the next page, in the **Remote network settings** section, do the following steps:
   a. Select the existing **0.0.0.0-255.255.255.255** entry, and then click **Remove** −, and then click **Add** +.
   b. In the resulting **Remote Address Settings** dialog, enter an IP address or IP address range that

identifies the network hosts that are allowed use this connector, and then click **Save**. You can repeat this step to add multiple IP addresses or IP address ranges. Err on the side of being too specific instead of too general to clearly identify the network hosts that are allowed to use this connector.

When you're finished, click **Finish**.

Use the Exchange Management Shell to create a dedicated Receive connector for anonymous relay

To create the same Receive connector in the Exchange Management Shell, use the following syntax:

```
New-ReceiveConnector -Name <ConnectorName> -TransportRole FrontendTransport -Custom -Bindings
<LocalIPAddresses>:25 -RemoteIpRanges <RemoteIPAddresses>
```

This example creates a new Receive connector with the following configuration options:

- **Name**: Anonymous Relay
- **Transport role**: `FrontEndTransport`
- **Usage type**: Custom
- **Bindings**: `0.0.0.0:25` (listen for inbound messages on all IP addresses that are configured on all network adapters in the Exchange server on TCP port 25.)
- **Remote IP addresses that are allowed to use this connector**: 192.168.5.10 and 192.168.5.11

```
New-ReceiveConnector -Name "Anonymous Relay" -TransportRole FrontendTransport -Custom -Bindings 0.0.0.0:25 -
RemoteIpRanges 192.168.5.10,192.168.5.11
```

Notes:

- The *Bindings* parameter is required when you specify the Custom usage type.
- The *RemoteIpRanges* parameter accepts an individual IP address, an IP address range (for example, `192.168.5.10-192.168.5.20`), or Classless InterDomain Routing (CIDR) (for example, `192.168.5.1/24`). You can specify multiple values separated by commas.

Step 2: Configure the permissions for anonymous relay on the dedicated Receive connector

As described in the introduction, there are two different methods you can use to configure the required permissions on the Receive connector:

- Configure the connections as anonymous.
- Configure the connections as externally secured.

Choose one method or the other. The examples use the Receive connector named Anonymous Relay that you created in Step 1.

Configure the connections as anonymous

Run the following commands in the Exchange Management Shell:

1.

```
Set-ReceiveConnector "Anonymous Relay" -PermissionGroups AnonymousUsers
```

2.

```
Get-ReceiveConnector "Anonymous Relay" | Add-ADPermission -User "NT AUTHORITY\ANONYMOUS LOGON" -ExtendedRights
"Ms-Exch-SMTP-Accept-Any-Recipient"
```

Configure the connections as externally secured

1. In the EAC, navigate to **Mail flow** > **Receive connectors**, select the Anonymous Relay connector, and then click **Edit** ✎.

2. In the properties of the connector, click **Security** and make the following selections:

   - **Authentication**: Deselect **Transport Layer Security (TLS)** and select **Externally secured (for example, with IPsec)**.

   - **Permission groups**: Select **Exchange servers**.

   When you're finished, click **Save**.

To perform these same steps in the Exchange Management Shell, run the following command:

```
Set-ReceiveConnector "Anonymous Relay" -AuthMechanism ExternalAuthoritative -PermissionGroups ExchangeServers
```

How do you know this worked?

To verify that you've successfully configured anonymous relay, do the following steps:

- Verify the configuration of the dedicated Receive connector.

  ```
  Get-ReceiveConnector "Anonymous Relay" | Format-List Enabled,TransportRole,Bindings,RemoteIPRanges
  ```

- Verify the permissions on the dedicated Receive connector.

  ```
  Get-ADPermission "Anonymous Relay" -User "NT AUTHORITY\ANONYMOUS LOGON" | where {($_.Deny -eq $false) -
  and ($_.IsInherited -eq $false)} | Format-Table User,ExtendedRights
  ```

  Or

  ```
  Get-ADPermission "Anonymous Relay" -User "MS Exchange\Externally Secured Servers" | where {($_.Deny -eq
  $false) -and ($_.IsInherited -eq $false)} | Format-Table User,ExtendedRights
  ```

- Use Telnet to test if one or more of the specified network hosts can connect to the dedicated Receive connector, and can anonymously relay mail through the connector. By default, the Telnet Client isn't installed in most client or server versions of Microsoft Windows. To install it, see Install Telnet Client.

  For more information, see Use Telnet to test SMTP communication on Exchange servers.

  If the network host is a device that doesn't have Telnet, you could temporarily add the IP address of a computer to the Receive connector, and then remove the IP address from the Receive connector when you're finished testing.

  For the test, the you'll need the following values:

  - **Destination**: This is the IP address or FQDN that you use to connect to the dedicated Receive connector. This is likely the IP address of the Mailbox server where the Receive connector is defined.

  This relates to the **Network adapter bindings** property (or the *Bindings* parameter) value that you configured on the connector. You'll need to use the valid value for your environment. In this example, we'll use 10.1.1.1.

- **Sender's email address**: You'll probably configure the servers or devices that are anonymously relaying mail to use a sending email address that's in an authoritative domain for your organization. In this example, we'll use chris@contoso.com.

- **Recipient's email address**: Use a valid email address. In this example, we'll use kate@fabrikam.com.

- **Message subject**: Test

- **Message body**: This is a test message

1. Open a Command Prompt window, type telnet, and then press Enter.

2. Type set localecho, and then press Enter.

3. Type OPEN 10.1.1.1 25, and then press Enter.

4. Type EHLO, and then press Enter.

5. Type MAIL FROM:chris@contoso.com, and then press Enter.

6. Type RCPT TO:kate@fabrikam.com, and then press Enter.

   - If you receive the response `250 2.1.5 Recipient OK`, the Receive connector allows anonymous relay from the network host. Continue to the next step to finish sending the test message.

   - If you receive the response `550 5.7.1 Unable to relay`, the Receive connector doesn't allow anonymous relay from the network host. If this happens, do the following:

     - Verify that you're connecting to the correct IP address or FQDN for the dedicated Receive connector.

     - Verify that the computer where you're running Telnet is allowed to use the Receive connector.

     - Verify the permissions on the Receive connector.

7. Type DATA, and then press Enter.

   You should receive a response that looks like this:

   ```
   354 Start mail input; end with <CLRF>.<CLRF>
   ```

8. Type Subject: Test, and then press Enter.

9. Press Enter again.

10. Type This is a test message, and then press Enter.

11. Press Enter, type a period (.), and then press Enter.

    You should receive a response that looks like this:

    ```
    250 2.6.0 <GUID> Queued mail for delivery
    ```

12. To disconnect from the SMTP server, type QUIT, and then press Enter.

    You should receive a response that looks like this:

    ```
    221 2.0.0 Service closing transmission channel
    ```

13. To close the Telnet session, type quit, and then press Enter.

- If anonymous relay works intermittently, you may need to modify the default message rate and throttling limits on the Receive connector. For more information, see Message throttling on Receive connectors.

Send connectors

8/3/2020 • 12 minutes to read • Edit Online

Exchange uses Send connectors for outbound SMTP connections from source Exchange servers to destination email servers. The Send connector that's used to route messages to a recipient is selected during the routing resolution phase of message categorization. For more information, see Mail routing.

You can create Send connectors in the Transport service on Mailbox servers and on Edge Transport servers. Send connectors are stored in Active Directory and are (by default) visible to all Mailbox servers in the organization.

> **IMPORTANT**
> By default, no Send connectors exist for external mail flow when you install Exchange. To enable outbound internet mail flow, you need to create a Send connector, or subscribe an Edge Transport server to your Exchange organization. For more information, see Create a Send connector to send mail to the Internet and Edge Transport servers.

You don't need to configure Send connectors to send mail between Exchange servers in the same Active Directory forest. Implicit and invisible Send connectors that are fully aware of the Exchange server topology are available for sending mail to internal Exchange servers. These connectors are described in the Implicit Send connectors section.

These are the important settings on Send connectors:

- **Usage type**
- **Network settings**: Configure how the Send connector routes mail: by using DNS or by automatically forward all mail to a smart host.
- **Address spaces**: Configure the destination domains that the Send connector is responsible for.
- **Scope**: Configures the visibility of the Send connector to other Exchange servers in the organization.
- **Source servers**: Configure the Exchange servers where the Send connector is hosted. Mail that needs to be delivered by using the Send connector is routed to one of the source servers.

On Mailbox servers, you can create and manage Send connectors in the Exchange admin center or in the Exchange Management Shell. On Edge Transport servers, you can only use the Exchange Management Shell.

Send connector changes in Exchange Server

These are the notable changes to Send connectors in Exchange 2016 or Exchange 2019 compared to Exchange 2010:

- You can configure Send connectors to redirect or *proxy* outbound mail through the Front End Transport service. For more information, see Configure Send connectors to proxy outbound mail.
- The *IsCoexistenceConnector* parameter is no longer available.
- The *LinkedReceiveConnector* parameter is no longer available.
- The default maximum message size is increased to 35 MB (approximately 25 MB due to Base64 encoding). For more information, see Message size and recipient limits in Exchange Server.
- The *TlsCertificateName* parameter allows you to specify the certificate issuer and the certificate subject. This helps minimize the risk of fraudulent certificates.

Implicit Send connectors

Although no Send connectors are created during the installation of Exchange servers, a special *implicit Send connector* named the intra-organization Send connector is present. This implicit Send connector is automatically available, invisible, and requires no management. The intra-organization Send connector exists in the transport services to send mail, either internally between services on the local Exchange server, or to services on remote Exchange servers in the organization. For example:

- Front End Transport service to the Transport service.
- Transport service to the Transport service on other servers.
- Transport service to subscribed Edge Transport servers.
- Transport service to the Mailbox Transport Delivery service.
- Mailbox Transport Submission service to the Transport service.

For more information, see Mail flow and the transport pipeline.

Send connector usage types

For Send connectors, the usage type is basically a descriptive label that identifies what the Send connector is used for. All usage type values receive the same permissions.

You can specify the connector usage type only when you create Send connectors. When you use the EAC, you must select a **Type** value. But when you use the **New-SendConnector** cmdlet in the Exchange Management Shell, the usage type isn't required (either by using `-Usage <UsageType>` or `-<UsageType>`).

Specifying a usage type does configure a default maximum message size, which you can change after you create the connector.

The available usage type values are described in the following table.

USAGE TYPE	MAXIMUM MESSAGE SIZE	COMMENTS
Custom	35 MB	None
Internal	unlimited	When you create a Send connector of this usage type in the EAC, you can't select **MX record associated with recipient domain**. After you create the connector, you can go to the **Delivery** tab in the properties of the Send connector and select **MX record associated with recipient domain**. This same restriction doesn't exist in the Exchange Management Shell. You can use the *Internal* switch and set the *DNSRoutingEnabled* to `$true` on the **New-SendConnector** cmdlet.
Internet	35 MB	None

USAGE TYPE	MAXIMUM MESSAGE SIZE	COMMENTS
Partner	35 MB	When you create a Send connector of this usage type in the EAC, you can't select **Route mail through smart hosts** or a smart host authentication mechanism. After you create the connector, you can go to the **Delivery** tab in the properties of the Send connector and select **Route mail through smart hosts** and the smart host authentication mechanism. This same restriction doesn't exist in the Exchange Management Shell. You can use the *Partner* switch and set the *DNSRoutingEnabled* to `$false` and use the *SmartHosts* and *SmartHostAuthMechanism* parameters on the **New-SendConnector** cmdlet.

Send connector network settings

Every Send connector needs to be configured with one of these options:

- Use DNS to route mail.
- Use one or more smart hosts to route mail.

Use DNS to route mail

When you select DNS resolution to deliver mail, the source Exchange server for the Send connector must be able to resolve the MX records for the address spaces that are configured on the connector. Depending on the nature of the connector, and how many network adapters are in the server, the Send connector could require access to an internal DNS server, or an external (public) DNS server. You can configure the server to use specific DNS servers for internal and external DNS lookups:

- In the EAC at **Servers** > **Server** > select the server and click **Edit** ✎ > **DNS lookups** tab.
- In the Exchange Management Shell, you use the *ExternalDNS** and *InternalDNS** parameters on the **Set-TransportService** cmdlet.

If you've already configured the Exchange server with separate DNS settings to use for internal and external DNS lookups, and the Send connector routes mail to an external address space, you need to configure the Send connector to use the external DNS server:

- In the EAC, select **Use the external DNS lookup setting on servers with transport roles** (in the new Send connector wizard, or on the **Delivery** tab in the properties of existing connectors).
- In the Exchange Management Shell, use the *UseExternalDNSServersEnabled* parameter on the **New-SendConnector** and **Set-SendConnector** cmdlets.

Use smart hosts to route mail

When you route mail through a smart host, the Send connector forwards mail to the smart host, and the smart host is responsible for routing mail to next hop on its way to the ultimate destination. A common use for smart host routing is to send outgoing mail through an antispam service or device.

You identify one or more smart hosts to use for the Send connector by an individual IP address (for example 10.1.1.1), a fully qualified domain name (FQDN) (for example spamservice.contoso.com), or combinations of both types of values. If you use an FQDN, the source Exchange server for the Send connector must be able to resolve the FQDN (which could be an MX record or an A record) by using DNS.

An important part of smart host routing is the authentication mechanism that the smart hosts uses. The available authentication mechanisms are described in the following table.

AUTHENTICATION MECHANISM	DESCRIPTION
None (`None`)	No authentication. For example, when access to the smart host is restricted by the source IP address.
Basic authentication (`BasicAuth`)	Basic authentication. Requires a username and password. The username and password are sent in clear text.
Offer basic authentication only after starting TLS (`BasicAuthRequireTLS`)	Basic authentication that's encrypted with TLS. This requires a server certificate on the smart host that contains the exact FQDN of the smart host that's defined on the Send connector. The Send connector attempts to establish the TLS session by sending the **STARTTLS** command to the smart host, and only performs Basic authentication after the TLS session is established. A client certificate is also required to support mutual TLS authentication.
Exchange Server authentication (`ExchangeServer`)	Generic Security Services application programming interface (GSSAPI) and Mutual GSSAPI authentication.
Externally secured (`ExternalAuthoritative`)	The connection is presumed to be secured by using a security mechanism that's external to Exchange. The connection may be an Internet Protocol security (IPsec) association or a virtual private network (VPN). Alternatively, the servers may reside in a trusted, physically controlled network.

Send connector address spaces

The address space specifies the destination domains that are serviced by the Send connector. Mail is routed through a Send connector based on the domain of the recipient's email address.

The available SMTP address space values are described in the following table.

ADDRESS SPACE	EXPLANATION
`*`	The Send connector routes mail to recipients in all domains.
Domain (for example, `contoso.com`)	The Send connector routes mail to recipients in the specified domain, but not in any subdomains.
Domain and subdomains (for example, `*.contoso.com`)	The Send connector routes mail to recipients in the specified domain, and in all subdomains.
`--`	The Send connector routes mail to recipients in all accepted domains in the Exchange organization. This value is only available on Send connectors on Edge Transport servers that send mail to the internal Exchange organization.

An address space also has **Type** and **Cost** values that you can configure.

On Edge Transport servers, the **Type** value must be `SMTP`. On Mailbox servers, you can also use non-SMTP address space types like `X400` or any other text string. X.400 addresses need to be RFC 1685 compliant (for example, `o=MySite;p=MyOrg;a=adatum;c=us`), but other **Type** values accept any text value for the address space. If you specify a non-SMTP address space type, the Send connector must use smart host routing, and SMTP is used to send messages to the smart host. Delivery Agent connectors and Foreign connectors send non-SMTP messages to non-SMTP servers without using SMTP. For more information, see Delivery Agents and Delivery Agent Connectors and Foreign Connectors.

The **Cost** value on the address space is used for mail flow optimization and fault tolerance when you have the same address spaces configured on multiple Send connectors on different source servers. A lower priority value indicates a preferred Send connector.

The Send connector that's used to route messages to a recipient is selected during the routing resolution phase of message categorization. The Send connector whose address space most closely matches the recipient's email address, and whose priority value is lowest is selected.

For example, suppose the recipient is julia@marketing.contoso.com. If a Send connector is configured for *.contoso.com, the message is routed through that connector. If no Send connector is configured for *.contoso.com, the message is routed through the connector that's configured for *. If multiple Send connectors in the same Active Directory site are configured for *.contoso.com, the connector with the lower priority value is selected.

Send connector scope

The source servers for a Send connector determine the destination Exchange server for mail that needs to be routed through the Send connector. The Send connector scope controls the visibility of the connector within the Exchange organization.

By default, Send connectors are visible to all the Exchange servers in the entire Active Directory forest, and are used in routing decisions. However, you can limit the scope of a Send connector so that it's only visible to other Exchange servers in the same Active Directory site. The Send connector is invisible to Exchange servers in other Active Directory sites, and isn't used in their routing decisions. A Send connector that's restricted in this way is said to be *scoped*.

To configure scoped Send connectors in the EAC, you select **Scoped send connector** in the **Address space** section of the new Send connector wizard, or on the **Scoping** tab in the properties of existing Send connectors. In the Exchange Management Shell, you use the *IsScopedConnector* parameter on the **New-SendConnector** and **Set-SendConnector** cmdlets.

Send connector permissions

When the Send connector establishes a connection with the destination email server, the Send connector permissions determine the types of headers that can be sent in messages. If a message includes headers that aren't allowed by the permissions, those headers are removed from messages.

Permissions are assigned to Send connectors by well-known security principals. Security principals include user accounts, computer accounts, and security groups (objects that are identifiable by a security identifier or SID that can have permissions assigned to them). By default, the same security principals with the same permissions are assigned on all Send connectors, regardless of the usage type that you selected when you created the connector. To modify the default permissions for a Send connector, you need to use the **Add-ADPermission** and **Remove-ADPermission** cmdlets in the Exchange Management Shell.

The available Send connector permissions are described in the following table.

PERMISSION	ASSIGNED TO	DESCRIPTION

`ms-Exch-Send-Headers-Forest`	`<Domain>\Exchange Servers`
	`MS Exchange\Edge Transport Servers`
	`MS Exchange\Hub Transport Servers`
		Controls the preservation of Exchange forest headers in messages. Forest header names start with **X-MS-Exchange-Forest-**. If this permission isn't granted, all forest headers are removed from messages.

`ms-Exch-Send-Headers-Organization`	`<Domain>\Exchange Servers`
	`MS Exchange\Edge Transport Servers`
	`MS Exchange\Hub Transport Servers`
		Controls the preservation of Exchange organization headers in messages. Organization header names start with **X-MS-Exchange-Organization-**. If this permission isn't granted, all organization headers are removed from messages.

`ms-Exch-Send-Headers-Routing`	`NT AUTHORITY\ANONYMOUS LOGON`
	`<Domain>\Exchange Servers`
	`MS Exchange\Edge Transport Servers`
	`MS Exchange\Externally Secured Servers`
	`MS Exchange\Hub Transport Servers`
	`MS Exchange\Legacy Exchange Servers`
	`MS Exchange\Partner Servers`
		Controls the preservation of **RECEIVED** headers in messages. If this permission isn't granted, all received headers are removed from messages.

`ms-Exch-SMTP-Send-Exch50`	`<Domain>\Exchange Servers`
	`MS Exchange\Edge Transport Servers`
	`MS Exchange\Externally Secured Servers`
	`MS Exchange\Hub Transport Servers`
	`MS Exchange\Legacy Exchange Servers`
		Allows the source Exchange server to submit **XEXCH50** commands on the Send connector. The **X-EXCH50** binary large object (BLOB) was used by older versions of Exchange (Exchange 2003 and earlier) to store Exchange data in messages (for example, the spam confidence level or SCL).
		If this permission isn't granted, and messages contain the **X-EXCH50** BLOB, the Exchange server sends the message without the **X-EXCH50** BLOB.

`ms-Exch-SMTP-Send-XShadow`	`<Domain>\Exchange Servers`
	`MS Exchange\Edge Transport Servers`
	`MS Exchange\Hub Transport Servers`
		This permission is reserved for internal Microsoft use, and is presented here for reference purposes only.

Note: Permissions names that contain `ms-Exch-Send-Headers-` are part of the *header firewall* feature. For more information, see Header firewall.

Send connector permission procedures

To see the permissions that are assigned to security principals on a Send connector, use the following syntax in the Exchange Management Shell:

```
Get-ADPermission -Identity <SendConnector> [-User <SecurityPrincipal>] | where {($_.Deny -eq $false) -and
($_.IsInherited -eq $false)} | Format-Table User,ExtendedRights
```

For example, to see the permissions that are assigned to all security principals on the Send connector named To Fabrikam.com, run the following command:

```
Get-ADPermission -Identity "To Fabrikam.com" | where {($_.Deny -eq $false) -and ($_.IsInherited -eq $false)} |
Format-Table User,ExtendedRights
```

To see the permissions that are assigned only to the security principal `NT AUTHORITY\ANONYMOUS LOGON` on the Send connector named To Fabrikam, run the following command:

```
Get-ADPermission -Identity "To Fabrikam.com" -User "NT AUTHORITY\ANONYMOUS LOGON" | where {($_.Deny -eq
$false) -and ($_.IsInherited -eq $false)} | Format-Table User,ExtendedRights
```

To add permissions to a security principal on a Send connector, use the following syntax:

```
Add-ADPermission -Identity <SendConnector> -User <SecurityPrincipal> -ExtendedRights "<Permission1>","
<Permission2>"...
```

To remove permissions from a security principal on a Send connector, use the following syntax:

```
Remove-ADPermission -Identity <SendConnector> -User <SecurityPrincipal> -ExtendedRights "<Permission1>","
<Permission2>"...
```

Create a Send connector in Exchange Server to send mail to the internet

8/3/2020 • 4 minutes to read • Edit Online

When install your first Exchange Server 2016 or Exchange 2019 server, the server isn't able to send mail outside of your Exchange organization. To send mail outside your Exchange organization, you need to create a Send connector.

What do you need to know before you begin?

- Estimated time to complete: 5 minutes
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Send connectors" entry in the Mail flow permissions topic.
- See Deploy a new installation of Exchange Server if you're beginning your installation. After the installation, you can use the steps in this topic to create an internet Send connector.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Create a Send connector to send mail to the internet

Until you create a Send connector, mail can't flow from your Exchange to the internet. The exception is if you install an Edge Transport in your perimeter network and subscribe the Edge Transport to your Exchange organization. For more information, see Edge Transport servers.

See also Send connectors for more information about connectors and why you would may or may not want to use them instead of Edge Transport Servers.

Use the EAC to create an internet Send connector

1. In the EAC, navigate to **Mail flow** > **Send connectors**, and then click **Add** +. This starts the **New Send connector** wizard.

2. On the first page, enter the following information:

   - **Name**: Enter a descriptive name for the Send connector (for example, To internet).
   - **Type**: Select **Internet**.

   When you're finished, click **Next**.

3. On the next page, verify that **MX record associated with recipient domain** is selected. This means the connector uses DNS on the internet to route mail, as opposed to routing all outbound mail to a smart host. For information about creating a Send connector that uses smart host routing, see Create a Send connector to route outbound mail through a smart host.

When you're finished, click **Next**.

4. On the next page, enter the following information:

   - In the **Address space** section, click **Add** +. In the **Add domain** dialog box that appears, in **Fully Qualified Domain Name (FQDN)**, enter an asterisk (*), and then click **Save**. This value indicates that the Send connector applies to messages addressed to all external domains.

   - The **Scoped send connector** setting is important if your organization has Exchange servers installed in multiple Active Directory sites:

     - If you don't select **Scoped send connector**, the connector is usable by all transport servers (Exchange 2013 or later Mailbox servers and Exchange 2010 Hub Transport servers) in the entire Active Directory forest. This is the default value.

     - If you select **Scoped send connector**, the connector is only usable by other transport servers in the same Active Directory site.

   When you're finished, click **Next**.

5. On the next page, in the **Source server** section, click **Add** +. In the **Select a Server** dialog box that appears, select one or more Mailbox servers that you want to use to send mail to the internet. If you have multiple Mailbox servers in your environment, select the ones that can route mail to the internet. If you have only one Mailbox server, select that one. After you've selected at least one Mailbox server, click **Add**, click **OK**, and then click **Finish**.

After you create the Send connector, it appears in the Send connector list. To configure the Send connector to proxy outbound mail through the Front End Transport service, see Configure Send connectors to proxy outbound mail.

Use the Exchange Management Shell to create an internet Send connector

1. Open the Exchange Management Shell. For more information, see Open the Exchange Management Shell.

2. Use the following syntax:

   ```
   New-SendConnector -Name <Name> -AddressSpaces * -Internet [-SourceTransportServer <fqdn1>,<fqdn2>...]
   ```

   This example creates the internet Send connector named "To internet" with the following properties:

   - The usage type is Internet.

   - The Send connector uses DNS routing. We aren't using the *DNSRoutingEnabled* parameter, and the default value is `$true`.

   - The Send connector is for all external domains (*).

   - The local Exchange server is the source server. We aren't using the *SourceTransportServer* parameter, and the default value is the local Exchange server.

   - The Send connector isn't scoped to the local Active Directory site. We aren't using the *IsScopedConnector* parameter, and the default value is `$false`.

   ```
   New-SendConnector -Name "To internet" -AddressSpaces * -Internet
   ```

   For information about other options, see New-SendConnector.

NOTE

To configure the Send connector to proxy outbound mail through the Front End Transport service, add `-FrontEndProxyEnabled $true` to the command. For more information, see Configure Send connectors to proxy outbound mail.

How do you know this worked?

To verify that you have successfully created a Send Connector that sends mail to the internet, create and send a message from an internal mailbox to an outside recipient, and verify the recipient receives the message.

You can also turn on protocol logging for the Send connector, and view the information in the log. For more information, see Protocol logging.

Create a Send connector to route outbound mail through a smart host

8/3/2020 • 5 minutes to read • Edit Online

Instead of routing all outbound messages directly to the Internet, you may need to route your organization's outbound mail through a third-party smart host. For example, your organization may have an appliance that scans outbound mail for spam and malware.

What do you need to know before you begin?

- Estimated time to complete: 5 minutes
- See Deploy a new installation of Exchange Server if you're beginning your installation. After the installation you can use the steps in this topic to create your outbound connector.
- The smart host described in this topic needs to use SMTP to transmit messages. If it doesn't, you need to use a Delivery Agent connector or a Foreign connector. For more information, seeDelivery Agents and Delivery Agent Connectors and Foreign Connectors.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Send connectors" entry in the Mail flow permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to create a Send connector that uses smart host routing

1. In the EAC, navigate to **Mail flow** > **Send connectors**, and then click **Add** +. This starts the **New Send connector** wizard.
2. On the first page, enter the following information:

   - **Name**: Enter a descriptive name for the Send connector, for example, Smart host to Internet.
   - **Type**: Select a descriptive value. For example, **Internet** or **Custom**. For more information about Send connector usage types, see Send connector usage types.

   When you're finished, click **Next**.
3. On the next page, select **Route mail through smart hosts**, and then click **Add** +. In the **Add smart host** dialog box that appears, identify the smart host by using one of the following values:

   - **IP address**: For example, 192.168.3.2.
   - **Fully qualified domain name (FQDN)**: For example, securitydevice01.contoso.com. Note that the Exchange source servers for the Send connector must be able to resolve the smart host in DNS by using this FQDN.

When you're finished, click **Save**.

4. You can enter multiple smart hosts by repeating Step 3. When you're finished, click **Next**.

5. On the next page, in the **Route mail through smart hosts** section, select the authentication method that's required by the smart host. Valid values are:

AUTHENTICATION MECHANISM	DESCRIPTION
None	No authentication. For example, when access to the smart host is restricted by the source IP address.
Basic authentication	Basic authentication. Requires a username and password. The username and password are sent in clear text.
Offer basic authentication only after starting TLS	Basic authentication that's encrypted with TLS. This requires a server certificate on the smart host that contains the exact FQDN of the smart host that's defined on the Send connector.
Exchange Server authentication	Generic Security Services application programming interface (GSSAPI) and Mutual GSSAPI authentication.
Externally secured	The connection is presumed to be secured by using a security mechanism that's external to Exchange. The connection may be an Internet Protocol security (IPsec) association or a virtual private network (VPN). Alternatively, the servers may reside in a trusted, physically controlled network.

When you're finished, click **Next**.

6. On the next page, in the **Address space** section, click **Add** ➕. In the **Add domain** dialog box that appears, enter the following information:

   - **Type**: Verify SMTP is entered.

   - **Fully Qualified Domain Name (FQDN)**: Enter an asterisk (*) to indicate the Send connector applies to messages addressed to all external domains. Alternatively, you can enter a specific external domain (for example, contoso.com), or a domain and all subdomains (for example, *.contoso.com).

   - **Cost**: Verify 1 is entered. A lower value indicates a more preferred route for the domains you specified.

   When you're finished, click **Save**.

7. Back on the previous page, the **Scoped send connector** setting is important if your organization has Exchange servers installed in multiple Active Directory sites:

   - If you don't select **Scoped send connector**, the connector is usable by all transport servers (Exchange 2013 or later Mailbox servers and Exchange 2010 Hub Transport servers) in the entire Active Directory forest. This is the default value.

   - If you select **Scoped send connector**, the connector is only usable by other transport servers in the same Active Directory site.

   When you're finished, click **Next**.

8. On the next page, in the **Source server** section, click **Add** ➕. In the **Select a Server** dialog box that appears, select one or more Mailbox servers that you want to use to send outbound mail to the smart host.

If you have multiple Mailbox servers in your environment, select the ones that can route mail to the smart host. If you have only one Mailbox server, select that one. After you've selected at least one Mailbox server, click **Add**, click **OK**, and then click **Finish**.

After you create the Send connector, it appears in the Send connector list.

Use the Exchange Management Shell to create a Send connector that uses smart host routing

1. Open the Exchange Management Shell. For more information, see Open the Exchange Management Shell.

2. Use the following syntax:

```
New-SendConnector -Name <Name> -AddressSpaces * -Custom -DnsRoutingEnabled $false -SmartHosts
<SmartHost1>[,<SmartHost2>...] [-SourceTransportServer <fqdn1>,<fqdn2>...]
```

This example creates the Internet Send connector named "Smart host to Internet" with the following properties:

- The usage type is Custom.

- The Send connector uses smart host routing (the *DNSRoutingEnabled* parameter is set to the value `$false`). The smart host's IP address is 192.168.3.2, and the authentication method is None, because the smart host is configured to listen for connections only from a restricted list of source servers.

- The Send connector is for all external domains (*). The value `*` is equivalent to the value `"SMTP:*;1"`, where the address space type is `SMTP`, and the address space cost value is `1`.

- The local Exchange server is the source server. We aren't using the *SourceTransportServer* parameter, and the default value is the local Exchange server.

- The Send connector isn't scoped to the local Active Directory site. We aren't using the *IsScopedConnector* parameter, and the default value is `$false`. The Send connector is useable by all Exchange transport servers in the Active Directory forest.

```
New-SendConnector -Name "Smart host to Internet" -AddressSpaces * -Custom -DNSRoutingEnabled $false -
SmartHosts 192.168.3.2 -SmartHostAuthMechanism None
```

For information about other options, see New-SendConnector.

How do you know this worked?

To verify that you have successfully created a Send connector to route outbound email through a smart host, send a message from a user in your organization to an external domain that's serviced by the Send connector.

You can also turn on protocol logging for the Send connector, and view the information in the log. For more information, see Protocol logging.

Configure Send connectors to proxy outbound mail

8/3/2020 • 2 minutes to read • Edit Online

When you create Send connectors, outbound mail flows through the Send connector in the Transport service on the Mailbox server or servers you specify, as shown in the following diagram.

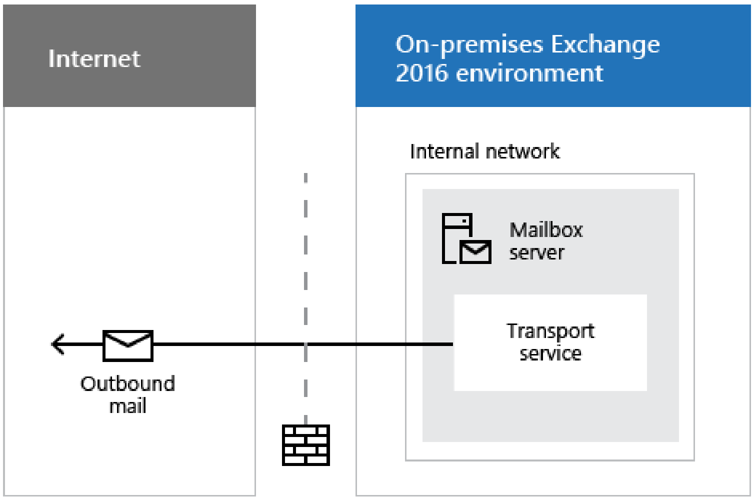

However, you can configure a Send connector to relay or *proxy* outbound mail through the Front End Transport service on the Mailbox server, as shown in the following diagram.

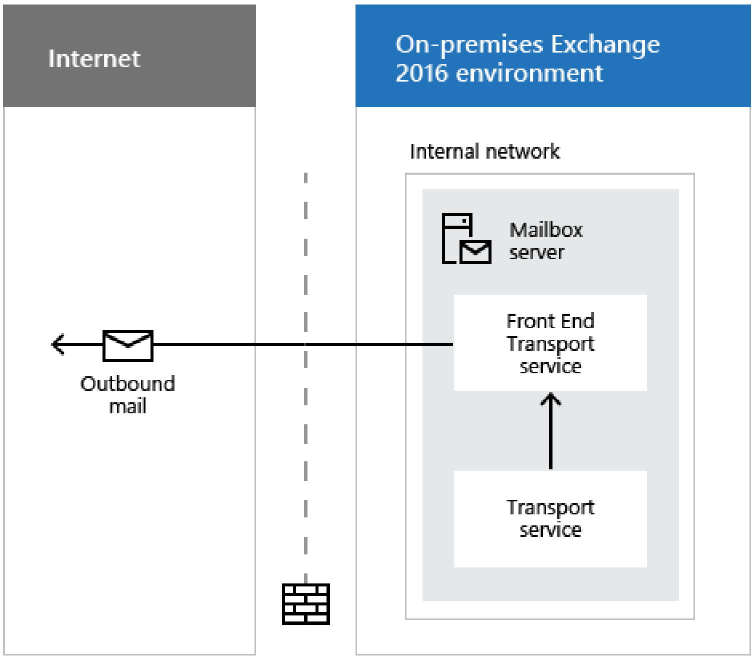

By default, all inbound mail enters your Exchange organization through the Front End Transport service, and the Front End Transport service proxies inbound mail to the Transport service. For more information, see Mail flow and the transport pipeline.

When you configure a Send connector to proxy outbound mail through the Front End Transport service, the Receive connector named "Outbound Proxy Frontend *<Mailbox server name>*" in the Front End Transport service listens for these outbound messages from the Transport service, and then the Front End Transport service sends the messages to the internet.

What do you need to know before you begin?

- Estimated time to complete: less than 5 minutes
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Send connectors" entry in the Mail flow permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Configure Send connectors to proxy outbound mail through the Front End Transport service

Use the EAC to configure Send connectors to proxy outbound mail

In the Exchange admin center (EAC), you can only configure existing Send connectors to proxy outbound mail.

1. In the EAC, navigate to **Mail flow** > **Send connectors**, select the Send connector, and then click **Edit** ✎.
2. On the **General** tab, in the **Connector status** section, select **Proxy through client access server**, and then click **Save**.

Use PowerShell to configure Send Connectors to proxy outbound mail

In the Exchange Management Shell, you can configure new or existing Send connectors to proxy outbound mail.

For information about how to open the Exchange Management Shell, see Open the Exchange Management Shell.

- To configure a new Send connector to proxy outbound mail, add `-FrontEndProxyEnabled $true` to the **New-SendConnector** command.
- To configure an existing Send connector to proxy outbound mail, run the following command:

  ```
  Set-SendConnector <Send connector identity>  -FrontEndProxyEnabled $true
  ```

  This example configures the existing Send connector named "Contoso.com Outbound" to proxy outbound mail.

  ```
  Set-SendConnector "Contoso.com Outbound" -FrontendProxyEnabled $true
  ```

How do you know this worked?

To verify that a Send connector is configured for outbound proxy, perform either of the following procedures:

- In the EAC, navigate to **Mail flow** > **Send connectors**, select the Send connector, and then click **Edit** ✎. On the **General** tab, in the **Connector status** section, verify **Proxy through client access server** is selected.
- In the Exchange Management Shell, run the following command:

```
Get-SendConnector | Format-Table -Auto Name,FrontEndProxyEnabled
```

Verify the **FrontEndProxyEnabled** value is `True` for the Send connector.

Protocol logging

8/3/2020 • 5 minutes to read • Edit Online

Protocol logging records the SMTP conversations that occur between messaging servers and between Exchange services in the transport pipeline as part of message delivery. You can use protocol logging to diagnose mail flow problems. The SMTP conversations that can be recorded by protocol logging occur in the following locations:

- Send connectors and Receive connectors in the Transport service on Mailbox servers.
- Send connectors and Receive connectors in the Transport service on Edge Transport servers.
- Receive connectors in the Front End Transport service on Mailbox servers.
- The implicit and invisible intra-organization Send connector in the Transport service on Mailbox servers.
- The implicit and invisible intra-organization Send connector in the Front End Transport service on Mailbox servers.
- The implicit and invisible intra-organization Send connector in the Mailbox Transport Submission service on Mailbox servers.
- The implicit and invisible Mailbox Delivery Receive connector in the Mailbox Transport Delivery service on Mailbox servers.

By default, protocol logging is enabled on the following connectors:

- The default Receive connector named Default Frontend *<ServerName>* in the Front End Transport service on Mailbox servers.
- The implicit and invisible Send connector in the Front End Transport service on Mailbox servers.

By default, protocol logging is disabled on all other connectors. You need to enable or disable protocol logging on each individual connector. You configure other protocol logging options for all Receive connectors or all Send connectors that exist in each individual transport service on the Exchange server. All Receive connectors in a transport service share the same protocol log files and protocol log options. These files and options are separate from the Send connector protocol log files and protocol log options in the same transport service on the Exchange server.

By default, Exchange uses circular logging to limit the protocol log based on file size and file age to help control the hard disk space that's used by the log files. To configure protocol logging, see Configure protocol logging.

Structure of the protocol log files

By default, the protocol log files exist in the following locations:

- Front End Transport service on Mailbox servers:
  - **Receive connectors**: `%ExchangeInstallPath%TransportRoles\Logs\FrontEnd\ProtocolLog\SmtpReceive`
  - **Send connectors**: `%ExchangeInstallPath%TransportRoles\Logs\FrontEnd\ProtocolLog\SmtpSend`
- Transport service on Mailbox servers:
  - **Receive connectors**: `%ExchangeInstallPath%TransportRoles\Logs\Hub\ProtocolLog\SmtpReceive`
  - **Send connectors**: `%ExchangeInstallPath%TransportRoles\Logs\Hub\ProtocolLog\SmtpSend`

- Mailbox Transport Delivery service on Mailbox servers (**Receive connectors**): `%ExchangeInstallPath%TransportRoles\Logs\Mailbox\ProtocolLog\SmtpReceive\Delivery`
- Mailbox Transport Submission service on Mailbox servers (**Send connectors**): `%ExchangeInstallPath%TransportRoles\Logs\Mailbox\ProtocolLog\SmtpSend\Submission`

  Note: Protocol logging for side effect messages that are submitted after messages are delivered to mailboxes occurs in `%ExchangeInstallPath%TransportRoles\Logs\Mailbox\ProtocolLog\SmtpSend\Delivery`. For example, a message that's delivered to a mailbox triggers an Inbox rule that redirects the message to another recipient.

- Transport service on Edge Transport servers:
  - **Receive connectors**: `%ExchangeInstallPath%TransportRoles\Logs\Edge\ProtocolLog\SmtpReceive`
  - **Send connectors**: `%ExchangeInstallPath%TransportRoles\Logs\Edge\ProtocolLog\SmtpSend`

The naming convention for log files is `SENDyyyymmdd-nnnn.log` for Send connectors and `RECVyyyymmdd-nnnn.log` for Receive connectors. The placeholders represent the following information:

- *yyyymmdd* is the coordinated universal time (UTC) date when the log file was created. *yyyy* = year, *mm* = month, and *dd* = day.
- *nnnn* is an instance number that starts at the value 1 every day.

Information is written to the log file until the file reaches its maximum size. Then, a new log file that has an incremented instance number is opened (the first log file is -1, the next is -2, and so on). Circular logging deletes the oldest log files when either of the following conditions is true:

- A log file reaches its maximum age.
- The protocol log folder reaches its maximum size.

The protocol log files are text files that contain data in the comma-separated value file (CSV) format. Each protocol log file has a header that contains the following information:

- **#Software**: The value is `Microsoft Exchange Server`.
- **#Version**: Version number of the Exchange server that created the message tracking log file. The value uses the format `15.01.nnnn.nnn`.
- **#Log-Type**: The value is either `SMTP Receive Protocol Log` or `SMTP Send Protocol Log`.
- **#Date**: UTC date-time when the log file was created. The UTC date-time is represented in the ISO 8601 date-time format: *yyyy-mm-dd*T*hh:mm:ss.fff*Z, where *yyyy* = year, *mm* = month, *dd* = day, T indicates the beginning of the time component, *hh* = hour, *mm* = minute, *ss* = second, *fff* = fractions of a second, and Z signifies Zulu, which is another way to denote UTC.
- **#Fields**: Comma-delimited field names that are used in the protocol log files.

Fields in the protocol log

The protocol log stores each SMTP protocol event on a single line in the log. The information stored on each line is organized by fields, and these fields are separated by commas. The fields that are used in the protocol log are described in the following table.

FIELD NAME	DESCRIPTION

FIELD NAME	DESCRIPTION
date-time	UTC date-time of the protocol event. The UTC date-time is represented in the ISO 8601 date-time format: *yyyy-mm-dd*T*hh:mm:ss.fffZ*, where *yyyy* = year, *mm* = month, *dd* = day, T indicates the beginning of the time component, *hh* = hour, *mm* = minute, *ss* = second, *fff* = fractions of a second, and Z signifies Zulu, which is another way to denote UTC.
connector-id	Distinguished name (DN) of the connector that's associated with the SMTP event.
session-id	GUID value that's unique for each SMTP session, but is the same for every event that's associated with that SMTP session.
sequence-number	Counter that starts at 0 and is incremented for each event in the same SMTP session.
local-endpoint	Local endpoint of an SMTP session. This consists of an IP address and TCP port number formatted as *<IP address>*:*<port>*.
remote-endpoint	Remote endpoint of an SMTP session. This consists of an IP address and TCP port number formatted as *<IP address>*:*<port>*.
event	Single character that represents the protocol event. Valid values are: `+` : Connect `-` : Disconnect `>` : Send `<` : Receive `*` : Information
data	Text information associated with the SMTP event.
context	Additional contextual information that may be associated with the SMTP event.

One SMTP conversation that represents sending or receiving a single email message generates multiple SMTP events. Each event is recorded on a separate line in the protocol log. An Exchange server has many SMTP conversations going on at any given time. This creates protocol log entries from different SMTP conversations that are mixed together. You can use the **session-id** and **sequence-number** fields to sort the protocol log entries by each individual SMTP conversation.

Configure protocol logging

8/3/2020 • 9 minutes to read • Edit Online

Protocol logging records the SMTP conversations that occur between messaging servers and between Exchange services in the transport pipeline as part of message delivery.

The following options are available for the protocol logs of all Send connectors and Receive connectors on the Exchange server:

- Specify the location of the protocol log files. The default locations are:
  - Front End Transport service on Mailbox servers:
  - **Receive connectors**: `%ExchangeInstallPath%TransportRoles\Logs\FrontEnd\ProtocolLog\SmtpReceive`
  - **Send connectors**: `%ExchangeInstallPath%TransportRoles\Logs\FrontEnd\ProtocolLog\SmtpSend`
  - Transport service on Mailbox servers:
  - **Receive connectors**: `%ExchangeInstallPath%TransportRoles\Logs\Hub\ProtocolLog\SmtpReceive`
  - **Send connectors**: `%ExchangeInstallPath%TransportRoles\Logs\Hub\ProtocolLog\SmtpSend`
  - Mailbox Transport Delivery service on Mailbox servers (**Receive connectors**): `%ExchangeInstallPath%TransportRoles\Logs\Mailbox\ProtocolLog\SmtpReceive\Delivery`
  - Mailbox Transport Submission service on Mailbox servers (**Send connectors**): `%ExchangeInstallPath%TransportRoles\Logs\Mailbox\ProtocolLog\SmtpSend\Submission`

    Note: Protocol logging for side effect messages that are submitted after messages are delivered to mailboxes occurs in `%ExchangeInstallPath%TransportRoles\Logs\Mailbox\ProtocolLog\SmtpSend\Delivery`. For example, a message that's delivered to a mailbox triggers an Inbox rule that redirects the message to another recipient.
  - Transport service on Edge Transport servers:
  - **Receive connectors**: `%ExchangeInstallPath%TransportRoles\Logs\Edge\ProtocolLog\SmtpReceive`
  - **Send connectors**: `%ExchangeInstallPath%TransportRoles\Logs\Edge\ProtocolLog\SmtpSend`
- Specify a maximum size for the protocol log files. The default size is 10 megabytes (MB).
- Specify a maximum size for the protocol log files. The default size is 250 MB.
- Specify a maximum age for the protocol log files. The default age is 30 days.

What do you need to know before you begin?

- Estimated time to complete: 5 minutes
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Transport Service", "Front End Transport service", "Mailbox Transport service", "Receive connectors" and "Send connectors" entries in the Mail flow permissions topic.
- You can use the Exchange admin center (EAC) to enable or disable protocol logging for Receive connectors and Send connectors on Mailbox servers. You can also use the EAC to configure the protocol log paths for

the Transport service only. For all other protocol logging options, you need to use the Exchange Management Shell. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.

- You enable or disable protocol logging on each individual connector. You configure other protocol logging options for all Receive connectors or all Send connectors that affect each individual transport service on the Exchange server. All Receive connectors in a transport service share the same protocol log files and protocol log options. These files and options are separate from the Send connector protocol log files and protocol log options in the same transport service.

  Caution

  Don't perform this procedure on an Edge Transport server that has been subscribed to the Exchange organization by using EdgeSync. Instead, make the changes in the Transport service on the Mailbox server. The changes are then replicated to the Edge Transport server the next time EdgeSync synchronization occurs.

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to configure protocol logging

Use the EAC to enable or disable protocol logging on a connector

Use this procedure to enable or disable protocol logging on a Send connector or a Receive connector in the Transport service on Mailbox servers, or a Receive connector in the Front End Transport service on Mailbox servers.

1. Open the EAC and navigate to one of the following locations:

   - **Mail flow** > **Send connectors**.

   - **Mail flow** > **Receive connectors**.

2. Select the connector you want to configure, and then click **Edit** ✎.

3. On the **General** tab in the **Protocol logging level** section, select one of the following options:

   - **None**: Protocol logging disabled on the connector.

   - **Verbose**: Protocol logging is enabled on the connector.

4. When you're finished, click **Save**.

Use the EAC to configure the location of the protocol logs on an Exchange server

Use this procedure to configure the location of the protocol logs for all Send connectors or all Receive connectors in the Transport service on Mailbox servers.

1. Open the EAC and navigate to **Servers** > **Servers**.

2. Select the Mailbox server you want to configure, and then click **Edit** ✎.

3. On the server properties page, click **Transport logs**. In the **Protocol log** section, change the following settings:

   - **Send protocol log path**

- **Receive protocol log path**

  Specify a location on the local Exchange server. If the folder doesn't exist, it's created when you click **Save**.

4. When you're finished, click **Save**.

How do you know this worked?

To verify that you have successfully used the EAC to configure protocol logging, browse to the location that you specified for the Send connector or the Receive connector protocol logs. If you enabled protocol logging, verify that a log file exists, and that the file is being updated for the connector. If you disabled protocol logging, verify that the latest log file is no longer being updated for the connector.

Use the Exchange Management Shell to enable or disable protocol logging on a connector

Use the Exchange Management Shell to enable or disable protocol logging on a Send connector or a Receive connector

Use this procedure to enable or disable protocol logging on:

- A Send connector or a Receive connector in the Transport service on Mailbox servers.
- A Receive connector in the Front End Transport service on Mailbox servers.
- A Send connector or a Receive connector in the Transport service on Edge Transport servers.

To enable or disable protocol logging on a Send connector or a Receive connector, use the following syntax in the Exchange Management Shell:

```
<Set-SendConnector | Set-ReceiveConnector> <ConnectorIdentity> -ProtocolLoggingLevel <Verbose | None>
```

This example enables protocol logging for the Receive connector named Connection from Contoso.com on the server named Mailbox01.

```
Set-ReceiveConnector "Mailbox01\Connection from Contoso.com" -ProtocolLoggingLevel Verbose
```

This example disables protocol logging for the Send connector named Connection to Internet.

```
Set-ReceiveConnector "Connection to Internet" -ProtocolLoggingLevel None
```

Use the Exchange Management Shell to enable or disable protocol logging on the intra-organization Send connector

Use this procedure to enable or disable protocol logging on the implicit and invisible intra-organization Send connector that exists in the Transport service, the Front End Transport service, and the Mailbox Transport Submission service on Mailbox servers. For more information about these connectors, see Implicit Send connectors.

Protocol logging for the intra-organization Send connector occurs in the Send connector protocol logs for the specified transport service. Note that the Transport service setting controls protocol logging for the intra-organization Send connector in the Transport service and in the Mailbox Transport Submission service.

To enable or disable protocol logging on the intra-organization Send connector, use the following syntax in the Exchange Management Shell:

```
<Set-TransportService | Set-FrontEndTransportService> <ServerIdentity> -IntraOrgConnectorProtocolLoggingLevel
<Verbose | None>
```

This example enables protocol logging on the intra-organization Send connector in the Transport service and in the Mailbox Transport Submission service on the server named Mailbox01.

```
Set-TransportService Mailbox01 -IntraOrgConnectorProtocolLoggingLevel Verbose
```

This example disables protocol logging on the intra-organization Send connector in the Front End Transport service on the same server.

```
Set-FrontEndTransportService Mailbox01 -IntraOrgConnectorProtocolLoggingLevel None
```

Use the Exchange Management Shell to enable or disable protocol logging on the mailbox delivery Receive connector

Use this procedure to enable or disable protocol logging on the implicit and invisible mailbox delivery Receive connector that exists in the Mailbox Transport Delivery service. Protocol logging for this connector occurs in the Receive connector protocol logs for the Mailbox Transport Delivery service. For more information about this connector, see Implicit Receive connectors in the Mailbox Transport Delivery service on Mailbox servers.

To enable or disable protocol logging on the mailbox delivery Receive connector, use the following syntax in the Exchange Management Shell:

```
Set-MailboxTransportService <ServerIdentity> -MailboxDeliveryConnectorProtocolLoggingLevel <Verbose | None>
```

This example enables protocol logging on the mailbox delivery Receive connector on the server named Mailbox01.

```
Set-MailboxTransportService Mailbox01 -MailboxDeliveryConnectorProtocolLoggingLevel Verbose
```

This example disables protocol logging on the mailbox delivery Receive connector on the same server.

```
Set-MailboxTransportService Mailbox01 -MailboxDeliveryConnectorProtocolLoggingLevel None
```

How do you know this worked?

To verify that you have successfully used the Exchange Management Shell to enable or disable protocol logging on a connector, perform the following steps:

1. Run the following command in the Exchange Management Shell to verify whether protocol logging is enabled or disabled for all connectors on the Exchange server:

   ```
   Write-Host "Send Connectors:" -ForegroundColor yellow; Get-SendConnector | Format-List
   Name,ProtocolLoggingLevel; Write-Host "Receive Connectors:" -ForegroundColor yellow; Get-
   ReceiveConnector | Format-List Name,TransportRole,ProtocolLoggingLevel; Write-Host "Mailbox Transport
   Delivery service:" -ForegroundColor yellow; Get-MailboxTransportService | Format-List
   *ProtocolLoggingLevel; Write-Host "Front End Transport service:" -ForegroundColor yellow; Get-
   FrontEndTransportService | Format-List *ProtocolLoggingLevel; Write-Host "Transport service and Mailbox
   Transport Submission service:" -ForegroundColor yellow; Get-TransportService | Format-List
   *ProtocolLoggingLevel
   ```

2. Browse to the location of the protocol log. If you enabled protocol logging, verify that a log file exists, and

that the file is being updated for the connector. If you disabled protocol logging, verify that the latest log file is no longer being updated for the connector.

Use the Exchange Management Shell to configure the protocol log settings on an Exchange server

Use this procedure to configure the protocol log settings for all Send connectors or Receive connectors in a transport service on a Mailbox server, and in the Transport service on an Edge Transport server.

To configure the protocol log settings in the Exchange Management Shell, use the following syntax:

```
<Set-FrontEndTransportService | Set-MailboxTransportService | Set-TransportService> <ServerIdentity> -
ReceiveProtocolLogPath <LocalFilePath> -ReceiveProtocolLogMaxFileSize <Size> -
ReceiveProtocolLogMaxDirectorySize <Size> -ReceiveProtocolLogMaxAge <dd.hh:mm:ss> -SendProtocolLogPath
<LocalFilePath> -SendProtocolLogMaxFileSize <Size> -SendProtocolLogMaxDirectorySize <Size> -
SendProtocolLogMaxAge <dd.hh:mm:ss>
```

This example sets the following protocol log settings in the Transport service on the server named Mailbox01:

- Sets the location of protocol log for all Receive connectors to D:\Hub SMTP Receive Log and the location for all Send connectors to D:\Hub SMTP Send Log. If the folder doesn't exist, it's created for you.
- Sets the maximum size of a connector protocol log file for Receive connectors and Send connectors to 20 MB.
- Sets the maximum size of the connector protocol log folder for Receive connectors and Send connectors to 400 MB.
- Sets the maximum age of a protocol log file for Receive connectors and Send connectors to 45 days.

```
Set-TransportService Mailbox01 -ReceiveProtocolLogPath "D:\Hub SMTP Receive Log" -
ReceiveProtocolLogMaxFileSize 20MB -ReceiveProtocolLogMaxDirectorySize 400MB -ReceiveProtocolLogMaxAge
45.00:00:00 -SendProtocolLogPath "D:\Hub SMTP Send Log" -SendProtocolLogMaxFileSize 20MB -
SendProtocolLogMaxDirectorySize 400MB -SendProtocolLogMaxAge 45.00:00:00
```

Notes:

- Setting the *SendProtocolLogPath* or *ReceiveProtocolLogPath* parameters to the value `$null` effectively disables protocol logging for all Send connectors or Receive connectors on the server. However, setting the value to `$null` generates event log errors when protocol logging is enabled for any Send connector or Receive connector on the server, including the intra-organization Send connector or the mailbox delivery Receive connector.
- Setting the *ReceiveProtocolLogMaxAge* or *SendProtocolLogMaxAge* parameters to the value `00:00:00` prevents the automatic removal of protocol log files because of their age.

How do you know this worked?

To verify that you have successfully used the Exchange Management Shell to configure the protocol logging settings on an Exchange server, perform the following steps:

1. Run the following command in the Exchange Management Shell and verify the protocol log settings on the Exchange server:

```
Write-Host "Front End Transport service:" -ForegroundColor yellow; Get-FrontEndTransportService |
Format-List ReceiveProtocolLog*,SendProtocolLog*; Write-Host "Mailbox Transport Submission and Mailbox
Transport Delivery services:" -ForegroundColor yellow; Get-MailboxTransportService | Format-List
ReceiveProtocolLog*,SendProtocolLog*; Write-Host "Transport service:" -ForegroundColor yellow; Get-
TransportService | Format-List ReceiveProtocolLog*,SendProtocolLog*
```

2. Open the location of the protocol log in Windows Explorer or File Explorer to verify that the log files exist, that data is being written to the files, and that the files are being recycled based on the maximum file size and maximum directory size values that you configured.

Mail routing in Exchange Server

8/3/2020 • 19 minutes to read • Edit Online

The primary task of the Transport service that exists on Mailbox servers in your Exchange organization is to route messages received from users and external sources to their ultimate destinations. Routing decisions are made during message categorization. The categorizer is a component of the Transport service that processes all incoming messages and determines what to do with the message based on information about their destinations.

Routing in Exchange 2016 and Exchange 2019 is virtually unchanged from Exchange 2013. These are the notable changes to routing compared to Exchange 2010:

- Routing is fully aware of database availability groups (DAGs), and is able to use DAG membership in routing decisions, even when the DAG members are in different Active Directory sites. For Mailbox servers that don't belong to DAGs and for interoperability with previous versions of Exchange, Active Directory site membership is still used in routing decisions.
- The Transport service never communicates directly with a mailbox database. Instead, the Transport service communicates with the Mailbox Transport service locally or on remote Mailbox servers. Only the Mailbox Transport service communicates with the local mailbox database. When the Mailbox server is a member of a DAG, only the Mailbox Transport service on the Mailbox server that holds the active copy of the mailbox database accepts messages for the destination recipient.
- Remote procedure calls (RPCs) are used only by the Mailbox Transport service to send messages to or receive messages from the local mailbox database. When the Mailbox server is a member of a DAG, the Mailbox Transport service only uses RPCs to communicate locally with the active copies of the mailbox databases. In other words, RPC is never used for cross-server or cross-service communication. Instead, the Mailbox Transport service and the Transport service always communicate using SMTP.
- Exchange now uses more precise queuing for remote destinations. Instead of using one queue for all destinations in a remote Active Directory site, Exchange now queues messages for specific destinations within the Active Directory site, such as individual Send connectors.
- Linked connectors are no longer available. A linked connector was a Receive connector that was linked to a Send connector. All messages received by the Receive connector were automatically forwarded to the Send connector.

Routing components

When a message is received by the Transport service on a Mailbox server, the message must be categorized. The first phase of message categorization is recipient resolution. After the recipient has been resolved, the ultimate destination can be determined. The next phase, routing, determines how to best reach that destination. Routing in Exchange is generalized for increased flexibility and decreased complexity by using the concepts of *routing destinations* and *delivery groups*.

Routing destinations

The ultimate destination for a message is called a *routing destination*. Regardless of the complexity of an Exchange organization, there are surprisingly few routing destinations. They are:

- **A mailbox database**: This is the routing destination for any recipient with a mailbox in the Exchange organization. In Exchange 2013 or later, public folders are a type of mailbox, so routing messages to public folder recipients is the same as routing messages to mailbox recipients.

- **A connector**: A Send connector is used as a routing destination for SMTP messages based on the configuration of the Send connector (address spaces, scoped or not, etc.). Similarly, a Delivery Agent connector or Foreign connector is used as a routing destination for non-SMTP messages.
- **A distribution group expansion server**: This is the routing destination when a distribution group has a designated expansion server (a server that's responsible for expanding the membership list of the group). A distribution group expansion server is an Exchange 2013 or later Mailbox server or an Exchange 2010 Hub Transport server.

Note that these routing destinations existed in previous versions of Exchange, but they weren't called routing destinations.

Delivery groups

A collection of one or more transport servers is responsible for delivering mail to each routing destination. This collection of transport servers is called a *delivery group*. The term *transport servers* is used because the servers could be a mixture of Exchange 2013 or later Mailbox servers (the Transport service) or Exchange 2010 Hub Transport servers. The relationship between routing destinations and delivery groups is explained in the following table:

ROUTING DESTINATION	DELIVERY GROUP
Exchange 2013 or later mailbox databases	Exchange 2013 or later Mailbox servers.
Exchange 2010 mailbox databases in Exchange 2016 organizations	Only Exchange 2010 Hub Transport servers.
Connectors	Exchange 2013 or later Mailbox servers or Exchange 2010 Hub Transport servers.
Distribution group expansion servers	Exchange 2013 or later Mailbox servers or Exchange 2010 Hub Transport servers.

How the message is routed depends on the relationship between the source delivery group and the destination delivery group:

- If the source and destination delivery group are the same, no routing decisions are required. The routing destination is the next hop for the message.
- If the source delivery group is outside the destination delivery group, routing decisions are required. The message is relayed along the least-cost routing path to the destination delivery group. Depending on the size and complexity of the Exchange environment, the message might be relayed through many transport servers to reach the destination delivery group for delivery to the routing destination.

The different types of delivery groups that exist in Exchange 2016 are summarized in the following table.

DELIVERY GROUP TYPE	DELIVERY GROUP	ROUTING DESTINATION	COMMENTS

DELIVERY GROUP TYPE	DELIVERY GROUP	ROUTING DESTINATION	COMMENTS
Routable DAG	• Exchange 2019 Mailbox servers that belong to the Exchange 2019 DAG. • Exchange 2016 Mailbox servers that belong to the Exchange 2016 DAG. • Exchange 2013 Mailbox servers that belong to the Exchange 2013 DAG.	Mailbox databases in the DAG	After the message arrives at a Mailbox server in the DAG, the Transport service routes the message to the Mailbox Transport Delivery service on the DAG member that holds the active copy of the destination mailbox database. The Mailbox Transport Delivery service then delivers the message to the local mailbox database. Although a DAG might contain Mailbox servers located in different Active Directory sites, the DAG defines the delivery group, not the Active Directory site.

DELIVERY GROUP TYPE	DELIVERY GROUP	ROUTING DESTINATION	COMMENTS
Mailbox delivery group (Exchange 2013 or later)	Exchange 2013 or later Mailbox servers in the Active Directory site.	Mailbox databases on Exchange 2013 or later servers in the Active Directory site that don't belong to a DAG.	Mailbox databases located on servers that don't belong to a DAG are serviced by the Transport service on Mailbox servers in the same Active Directory site. After the message arrives on an Mailbox server in the Active Directory site, the Transport service uses SMTP to transfer the message to the Mailbox Transport Delivery service on the Mailbox server that holds the mailbox database. The Mailbox Transport Delivery service then delivers the message to the local mailbox database using RPC. In other words, the following mail delivery paths are supported between the different versions of Exchange: • Exchange 2019 Transport service to Exchange 2016 Mailbox Transport Delivery service to Exchange 2016 mailbox database. • Exchange 2019 Transport service to Exchange 2013 Mailbox Transport Delivery service to Exchange 2013 mailbox database. • Exchange 2016 Transport service to Exchange 2019 Mailbox Transport Delivery service to Exchange 2019 mailbox database. • Exchange 2016 Transport service to Exchange 2013 Mailbox Transport Delivery service to Exchange 2013 mailbox database. • Exchange 2013 Transport service to Exchange 2019 Mailbox Transport Delivery service to Exchange 2019 mailbox database. • Exchange 2013 Transport service to Exchange 2016 Mailbox Transport Delivery service to Exchange 2016 mailbox database.

DELIVERY GROUP TYPE	DELIVERY GROUP	ROUTING DESTINATION	COMMENTS
Mailbox delivery group (Exchange 2010)	Exchange 2010 Hub Transport servers in the Active Directory site.	Mailbox databases on Exchange 2010 Mailbox servers in the Active Directory site.	Mailbox databases located on Exchange 2010 Mailbox servers are serviced by the Exchange 2010 Hub Transport servers in the same Active Directory site. After the message arrives at a random Exchange 2010 Hub Transport server in the Active Directory site, the store driver on the Hub Transport server uses RPC to write the message to the mailbox database.
Connector source server	A mixture of any Exchange 2013 or later Mailbox servers or Exchange 2010 Hub Transport servers that are defined as source transport servers for the connector.	A Send connector, Delivery Agent connector, or Foreign connector.	If the connector is scoped (that is, restricted to transport servers in the same Active Directory site), then only other transport servers in that site are aware of the connector, and can use the connector to route mail. If the connector isn't scoped, then all transport servers in the entire Active Directory forest are aware of the connector, and can use the connector to route mail.
Server list	The Exchange 2013 or later Mailbox server or Exchange 2010 Hub Transport server that's defined as the expansion server for the distribution group.	The distribution group expansion server.	none

DELIVERY GROUP TYPE	DELIVERY GROUP	ROUTING DESTINATION	COMMENTS
AD site	Any mixture of Exchange 2013 or later Mailbox servers or Exchange 2010 Hub Transport servers that exist in: • Active Directory sites that are configured as hub sites. • Active Directory sites that have subscribed Edge Transport servers.	None. The message must travel through the Active Directory site on the way to the actual routing destination.	This delivery group type is the only routing scenario in Exchange 2013 or later where *delayed fan-out* is still used. Delayed fan-out attempts to reduce the number of message transmissions when multiple routing destinations share part of the least-cost routing path. Hub sites are used only if the Active Directory site exists along the least-cost routing path for the message. br/> For Edge Transport servers, the Transport service on any Mailbox server in the subscribed Active Directory site is able to send messages to the Edge Transport server, regardless of whether that server participates in EdgeSync synchronization. For more information, see Edge Transport servers.

> **NOTE**
>
> Delivery group membership isn't mutually exclusive. For example, a Mailbox server that's a member of a DAG can also be the source transport server of a Send connector. The Mailbox server belongs to the routable DAG delivery group for the mailbox databases in the DAG, and the connector source server delivery group for the Send connector.

Queues

From the perspective of the sending transport server, each message delivery queue represents the destination for a particular message. When the Transport service selects the destination for a message, the destination is stamped on the recipient as the **NextHopSolutionKey** attribute. If a single message is sent to more than one recipient, each recipient has the **NextHopSolutionKey** attribute. The receiving transport server also performs message categorization and queues the message for delivery. After a message is queued, you can examine the delivery type for a particular queue to determine whether a message will be relayed again when it reaches the next hop destination. Every unique value of the **NextHopSolutionKey** attribute corresponds to a separate message delivery queue.

For more information, see NextHopSolutionKey.

Routing messages

When a message needs to be delivered to a remote delivery group, a routing path must be determined for the message. Exchange uses the following logic to select the routing path for a message. This logic is basically unchanged from Exchange 2010:

1. Calculate the least-cost routing path by adding the cost of the IP site links that must be traversed to reach the destination. If the destination is a connector, the cost assigned to the address space is added to the cost to reach the selected connector. If multiple routing paths are possible, the routing path with the lowest

aggregate cost is used.

Note: Size limits on connectors are a factor here. Connectors that are configured with message sizes limits smaller than the size of the message are eliminated from consideration. For more information, see Connector selection in external message routing.

2. If more than one routing path has the same aggregate cost, the number of hops in each path is evaluated and the routing path with the least number of hops is used.

3. If more than one routing path is still available, the name assigned to the Active Directory sites before the destination is considered. The routing path where the Active Directory site nearest the destination is lowest in alphanumeric order is used. If the site nearest the destination is the same for all routing paths being evaluated, an earlier site name is considered.

In Exchange 2010, each message recipient is associated with only one Active Directory site, and there is only one least cost routing from the source Active Directory site to the destination site. In Exchange 2013 or later, a delivery group might span multiple Active Directory sites, and there might be multiple least-cost routing paths to those sites. Exchange designates a single Active Directory site in the destination delivery group as the *primary site*. The primary site is closest Active Directory site based on the routing logic described earlier. To successfully route messages between delivery groups, Exchange takes the following issues into consideration:

- **The presence of one or more hub sites along the least-cost routing path**: If the least-cost routing path to the primary site contains any hub sites, the message must be routed through the hub sites. The closest hub site along the least-cost routing path is selected as a new delivery group of the type **AD site**, which includes all transport servers in the hub site. After the message traverses the hub site, routing of the message along the least-cost routing path continues. If the primary site happens to be a hub site, the primary site is still considered a hub site for the following reasons:

  - If the destination delivery group spans multiple Active Directory sites, the source server should only attempt to connect to the servers in the hub site.

  - The servers in the hub site that belong to the target delivery group are preferred.

    As in previous version of Exchange, hub sites that aren't in the least-cost routing path to the primary site are ignored.

- **The target Exchange server to select in the destination routing group**: When the destination delivery group spans multiple Active Directory sites, the routing path to specific servers within the delivery group might have different costs. Servers located in the closest Active Directory site are selected as the target servers for the delivery group based on the least-cost routing path, and the Active Directory site those servers are in is selected as the primary site.

- **Fallback options when connection attempts to all servers in the destination routing group fail**: If the destination delivery group spans multiple Active Directory sites, the first fallback option is all other servers in the destination delivery group in other Active Directory sites that aren't selected as target servers. Server selection is based on the least-cost routing path to the other Active Directory sites. If the destination delivery group has any servers in the local Active Directory site, there are no other fallback options because the message is already as close to the target routing destination as possible. If the destination delivery group has servers in remote Active Directory sites, the option is to try to connect to all other servers in the primary site. If that fails, a backoff path in the least-cost routing path to the primary site is used. Exchange tries to deliver the message as close to the destination as possible by backing off, hop by hop, along the least-cost routing path until a connection is made.

Routing messages between Active Directory sites

The way that Exchange routes messages between Active Directory sites is virtually the same as Exchange 2010. For more information, see Route Mail Between Active Directory Sites.

Routing in the Front End Transport service on Mailbox servers

The Front End Transport service acts as a stateless proxy for all inbound and (optionally) outbound external SMTP traffic for the Exchange organization. For outgoing messages, the Transport service communicates with the Front End Transport service only when it's specifically configured to do so. For more information, see Configure Send connectors to proxy outbound mail.

For incoming messages, the Front End Transport service must quickly find a single, healthy Transport service to receive the message transmission, regardless of the number or type of recipients. Failure to do so results in the email service being perceived as unavailable by the sending server. Like the Transport service, the Front End Transport service loads routing tables based on information from Active Directory, and uses delivery groups to determine how to route messages. However, the routing tables used by the Front End Transport service have the following unique characteristics:

- The Front End Transport service is never considered a member of a delivery group, even when the Mailbox server and the Client access server are installed on the same physical server (which is always the case in Exchange 2016 or later). This forces the Front End Transport service to communicate only with the Transport service.
- The routing tables don't contain any Send connector routes.
- The routing tables contain a special list of Mailbox servers in the local Active Directory site for fast fail-over purposes.

Routing in the Front End Transport service resolves message recipients to mailbox databases. The list of Mailbox servers used by the Front End Transport service is based on the mailbox databases of the message recipients. Note that it's possible that none of the recipients have mailboxes, for example, if the recipient is a distribution group or a mail user. For each mailbox database, the Front End Transport service looks up the delivery group and the associated routing information. The delivery groups used by the Front End Transport service are:

- Routable DAG
- Mailbox delivery group
- AD site

Depending on the number and type of recipients, the Front End Transport service performs one of the following actions:

- For messages with a single mailbox recipient, select a Mailbox server in the target delivery group, and give preference to the Mailbox server based on the proximity of the Active Directory site. Routing the message to the recipient might involve routing the message through a hub site.
- For messages with multiple mailbox recipients, use the first 20 recipients to select a Mailbox server in the closest delivery group, based on the proximity of the Active Directory site. Note that message bifurcation doesn't occur in Front End Transport, so only one Mailbox server is ultimately selected, regardless of number of recipients in a message.
- If the message has no mailbox recipients, select a random Mailbox server in the local Active Directory site.

Routing in the Mailbox Transport service on Mailbox servers

The Mailbox Transport service consists of two separate services: the Mailbox Transport Submission service and Mailbox Transport Delivery service. The Mailbox Transport Delivery service receives SMTP messages from the Transport service, and connects to the local mailbox database by using RPC to deliver the message. The Mailbox Transport Submission service connects to the local mailbox database by using RPC to retrieve messages, and submits the messages over SMTP to the Transport service. The Mailbox Transport service is stateless, and doesn't use message delivery queues.

Like the Transport service, the Mailbox Transport service loads routing tables based on information from Active

Directory, and uses delivery groups to determine how to route messages. However, there are routing aspects that are unique to the Mailbox Transport service:

- Because the Transport service and the Mailbox Transport service exist on the same Mailbox server, the Mailbox Transport service always belongs to the same delivery group as the Mailbox server. This delivery group is referred to as the *local delivery group*.
- The Mailbox Transport Submission service doesn't automatically send messages to the Transport service on the local Mailbox server or on other Mailbox servers in its own local delivery group. The Mailbox Transport Submission service has access to the same routing topology information as the Transport service, so the Mailbox Transport submission service can send messages to the Transport service on Mailbox servers outside the delivery group. The Mailbox servers in the local delivery group are used as fallback options, and for delivery to non-mailbox recipients.
- The Mailbox Transport service only communicates with the Transport service on Mailbox servers.
- The Mailbox Transport service only communicates with local mailbox databases. The Mailbox Transport service never communicates with mailbox databases on other Mailbox servers.

When a user sends a message from their mailbox, the Mailbox Transport Submission service resolves the message recipients to mailbox databases. The list of Mailbox servers used by the Mailbox Transport Submission service is based on the mailbox databases of the message recipients. Note that it's possible that none of the recipients have mailboxes, for example, if the recipient is a distribution group or a mail user. For each mailbox database, the Mailbox Transport Submission service looks up the delivery group and the associated routing information. The delivery groups used by the Mailbox Transport Submission service are:

- Routable DAG
- Mailbox delivery group
- AD site

Depending on the number and type of recipients, the Mailbox Transport Submission service performs one of the following actions:

- For messages with a single mailbox recipient, select a Mailbox server in the target delivery group, and give preference to the Mailbox server based on the proximity of the Active Directory site. Routing the message to the recipient might involve routing the message through a hub site.
- For messages with multiple mailbox recipients, use the first 20 recipients to select a Mailbox server in the closest delivery group, based on the proximity of the Active Directory site.
- If the message has no mailbox recipients, select a Mailbox server in the local delivery group.

When the Mailbox Transport Delivery service receives a message from the Transport service, it accepts or rejects the message for delivery to a local mailbox database. The Mailbox Transport Delivery service can deliver the message if the recipient resides in an active copy of a local mailbox database. But, if the recipient doesn't reside in an active copy of a local mailbox database, the Mailbox Transport Delivery service can't deliver the message, and must provide a non-delivery response to the Transport service. For example, if the active copy of the mailbox database recently moved to another server, the Transport service might erroneously transmit a message to a Mailbox server that now holds an inactive copy of the mailbox database. The non-delivery responses that the Mailbox Transport Delivery service returns to the Transport service include:

- Retry delivery
- Generate an NDR (also known as a non-delivery report, delivery status notification, DSN, or bounce message)
- Reroute the message

Routing in the Transport service on Edge Transport servers

The Transport service on Edge Transport servers provides SMTP relay and smart host services for all Internet mail flow. Messages that come and go from the Internet are stored in message delivery queues on the Edge Transport server. The queues correspond to external domains or Send connectors. For more information, see NextHopSolutionKey.

Typically, when you install an Edge Transport server in your perimeter network, you subscribe the Edge Transport server to an Active Directory site. The Active Directory site contains the Mailbox servers that relay messages to and from the Edge Transport server. The Edge Subscription process creates an Active Directory site membership affiliation for the Edge Transport server. The site affiliation enables the Mailbox servers in the Active Directory site to relay messages to the Edge Transport server without having to configure explicit Send connectors.

In organizations that have Exchange servers in multiple Active Directory sites, outbound mail from internal recipients to external recipients is first routed to the subscribed Active Directory site. Transport servers in the target Active Directory site are the delivery group. The routing destination is the intra-organization Send connector in the Transport service on any of the Mailbox servers in the subscribed Active Directory site. The *intra-organization Send connector* is special Send connector that exists in the Transport service on every Mailbox server. This Send connector is implicitly created, invisible, requires no management, and is used to relay messages between Exchange servers.

For more information about how mail is routed to and from Edge Transport servers, see Mail flow and the transport pipeline.

Connector selection in external message routing

8/3/2020 • 4 minutes to read • Edit Online

Like previous versions of Exchange, Exchange Server 2016 and Exchange Server 2019 use connectors to deliver messages to external recipients (recipients that don't exist in the Exchange organization). Exchange uses Send connectors to route messages to external SMTP domains. If the external recipient isn't on an SMTP messaging system, Exchange uses Delivery Agent connectors or Foreign connectors.

For more information about the different types of connectors, see Connectors. For more information about how Exchange makes routing decisions, see Mail routing.

Connector considerations in message routing

The settings that are configured on connectors might eliminate an otherwise available connector from routing consideration. These settings are described in the following table:

CONNECTOR SETTING	COMMENTS
State (enabled or disabled)	Only enabled connectors are used in routing decisions. If a connector is disabled, it's not considered when routing messages.
Address space	The address spaces defines the destination domains or other address spaces that are serviced by the connector. When Exchange selects a connector for routing a message, it only considers connectors that have a matching address space. If more than one connector matches the destination address space, the connector with the more precise address match is selected. For example, suppose the recipient is julia@marketing.contoso.com, and separate Send connectors are configured for *, *.contoso.com and marketing.contoso.com. The order of connector preference based solely on the address space is: 1: marketing.contoso.com 2: *.contoso.com 3: *
Address space type	By default, the address space type on a new Send connector is SMTP. If you specify a non-SMTP address space, the messages are still sent to the destination (a smart host) by using SMTP. You need to create a Delivery Agent connector or a Foreign connector to route non-SMTP messages to non-SMTP messaging servers without using SMTP.
Address space cost	You use the cost value on the address space for mail flow optimization and fault tolerance when the same address space is configured on multiple connectors. A lower cost value indicates a preferred connector.

CONNECTOR SETTING	COMMENTS
Source server	At least one Mailbox server or Edge Transport server must be configured to host the connector. You can configure multiple source servers to provide load balancing and fault tolerance for the address spaces that are defined on the connector.
Scope	The connector's scope controls its visibility within the Exchange organization. By default, connectors are visible to all the Exchange servers in the entire Active Directory forest. However, you can limit the scope of a connector so that it's only visible to other Exchange servers in the local Active Directory site. The connector is invisible to Exchange servers in other Active Directory sites, and isn't used in their routing decisions. A connector that's restricted in this way is said to be *scoped*.
Message size limits	A message size restriction on a connector can eliminate the connector from selection if the message is larger than the maximum size that's allowed on the connector. For more information about message size limits on connectors, see Connector limits.

Selecting the connector for an external recipient

For messages that are sent to external recipients, Exchange must select the best connector to route the message through. The decisions that are required to select this connector are described in the following list:

1. Exchange eliminates all connectors that have a message size limit that's smaller than the size of the message.

2. Exchange narrows the list of remaining connectors to those that satisfy all of the following criteria:

   - The connector is scoped to another Exchange server in the local Active Directory site, or isn't scoped at all (is available to all Exchange in the Active Directory forest).

   - The connector is enabled.

   - The connector is configured with an address space that matches the recipient's email address.

3. From the resulting list of connectors, Exchange selects the connector with the most specific address space match. If multiple connectors have the same address space specificity, Exchange uses the following criteria to select a connector:

   a. **Aggregate cost**: This is the sum of the cost that's assigned to all the IP site links between the source Active Directory site and the Active Directory site that contains the source servers for the connector, and the cost that's assigned to the address space on the connector (IP site link costs + connector cost). The connector with the lowest aggregate cost is selected. If multiple connectors have the same aggregate cost, the selection process continues to the next step.

   b. **Hop count**: The source server for the connector that can be reached in the least number of hops is selected. Typically, this means the general order of preference is:

      a. The local Exchange server.

      b. An Exchange server in the same Active Directory site.

      c. An Exchange server in a remote Active Directory site.

If multiple connectors have the same hop count, the selection process continues to the next step.

c. **Connector name**: If more than one routing path has the same aggregate cost and hop count, the connector with the name that has the lowest alphanumeric value is selected.

Handling messages that can't be routed

If no connector satisfies all of the selection criteria, one of the following actions occurs:

- If there is no matching connector for an SMTP address space, the recipient is marked as unreachable and the message is routed to the Unreachable queue. For more information about the Unreachable queue, see Types of queues.
- If there is no matching connector for a non-SMTP address space, a non-delivery report (also known as an NDR or bounce message) is returned to the sender.
- If the message size exceeds the connector size restriction for all connectors, an NDR is returned to the sender.

Recipient resolution in Exchange Server

8/3/2020 • 21 minutes to read • Edit Online

Recipient resolution is when the Exchange server expands and resolves all recipients in a message. Recipient resolution is responsible for these actions:

- Matching the recipient to the corresponding Active Directory object.
- Expanding distribution groups into a list of individual recipients.
- Applying message limits and any alternative recipients to each recipient.

Recipient resolution in Exchange 2016 or Exchange 2019 is basically unchanged from Exchange 2013. Recipient resolution is done by the categorizer in the Transport service on Mailbox servers. Each message is processed by the categorizer after the message is put in the Submission queue, but before the message is put in a delivery queue. The component of the categorizer that's responsible for recipient resolution is frequently called the *resolver*. For more information about the categorizer and the Submission queue, see Understanding the Transport service on Mailbox servers.

This topic explains the various stages and components of recipient resolution that occur on the Exchange server.

Top-level resolution

Top-level resolution is the first stage of recipient resolution that associates each recipient in an incoming message to a matching recipient object in Active Directory. The categorizer creates a list that contains the sender and the initial, unexpanded recipient email addresses from the message, and uses that list to query Active Directory. When a match is found in the email address properties for mail-enabled Active Directory objects, the categorizer caches the properties of the objects, and enforces any sender message restrictions.

Recipient email addresses

Top-level resolution begins with a message and the initial, unexpanded list of recipients from the *message envelope*. The message envelope contains the SMTP commands that are used to transmit messages between messaging servers:

- The sender's email address is contained in the **MAIL FROM** command.
- Each recipient's email address is contained in a separate **RCPT TO** command.

The envelope sender and envelope recipients are typically created from the sender and recipients in the `To:`, `From:`, `Cc:`, and `Bcc:` header fields in the message header. However, these header fields are easily forged and may not match the actual sender or recipient email addresses that were used to transmit the message.

Encapsulated email addresses

Standard SMTP email addresses follow the specifications in RFC 5321 and RFC 5322 (for example chris@contoso.com). However, an email address can also be a non-SMTP address that's encapsulated in an SMTP address. Exchange supports the Internet Mail Connector Encapsulated Address (IMCEA) encapsulation method that replaces characters that would be invalid in an SMTP email address with valid characters.

IMCEA addresses use the syntax `IMCEA<Type>-<address>@<domain>`:

- *<Type>* identifies the type of non-SMTP address (for example `EX`, `X400`, or `FAX`). Although `SMTP` and `X500` are theoretically valid values for *<Type>*, Exchange recipient resolution rejects any IMCEA-encoded addresses that use either of these types.

- *<address>* is the encoded version of original address:
  - Alphanumeric characters, the equal sign (=) and the hyphen (-) aren't replaced.
  - Forward slashes (/) are replaced by underscores (_).
  - Other US-ASCII characters are replaced by a plus sign (+) and the two digits of the ASCII value in hexadecimal (for example, the space character has the encoded value `+20`).
- *<domain>* is SMTP domain that's used to encapsulate the non-SMTP address (for example, contoso.com).

IMCEA addresses are returned to their original values (unencapsulated) only when the domain matches the default accepted domain in the Exchange organization. For more information about the default accepted domain, see Default domain.

The maximum length for an SMTP email address in Exchange is 571 characters. This limit includes:

- 315 characters for the name part of the address
- 255 characters for the domain name
- The at sign (@) character that separates the name from the domain name

Exchange doesn't support IMCEA-encoded messages where the name part of the address exceeds 315 characters, even if the complete email address is less than 571 characters.

Address resolution

For each message, the sender's email address and all recipient addresses are added to a list that's used to query Active Directory. Any encapsulated addresses are unencapsulated before they're added to the list. The Active Directory query is performed on up to 20 email addresses at a time. If the query encounters transient errors, the message is returned to the Submission queue and deferred for the time that's specified by the *ResolverRetryInterval* key in the `%ExchangeInstallPath%Bin\EdgeTransport.exe.config` XML application configuration file that's associated with the Transport service. The default value is 30 minutes.

The Active Directory recipient objects that are used by Exchange are described in the following table. For more information about Exchange recipient types, see Recipients.

ACTIVE DIRECTORY RECIPIENT TYPE	DESCRIPTION
DistributionGroup	Any mail-enabled group object. The distribution group object types are: • **MailUniversalDistributionGroup**: A universal distribution group object • **MailUniversalSecurityGroup**: A universal security group (USG) object that has an email address
DynamicDistributionGroup	An object that has the Active Directory class **msExchDynamicDistributionList**. For more information, see Manage dynamic distribution groups.
Mailbox	An object that has an email address and a defined *Database* parameter.
MailUser	A user object that has an email address without a defined *Database* parameter. For more information, see Manage mail users.

ACTIVE DIRECTORY RECIPIENT TYPE	DESCRIPTION
MailContact	A contact object that has an email address. Typically, a mail contact is used for recipients outside the Exchange organization. A mail contact is also used in cross-forest Exchange environments. For more information, see Manage mail contacts.
MailPublicFolder	A public folder object that has an email address. For more information, see Public folders.
MicrosoftExchangeRecipient	An object that has the Active Directory class **msExchExchangeServerRecipient**. For more information about the Exchange recipient object, see Recipients.
SystemMailbox	A user object that has an email address and is located in the Microsoft Exchange System Objects container. There should be one system mailbox for each mailbox database in the Exchange organization.

The Active Directory query classifies an object with missing or malformed critical properties as an invalid object (for example, a dynamic distribution group object without an email address). Messages sent to recipients that are classified as invalid objects generate a non-delivery report (also known as an NDR or bounce message).

For each email address, the categorizer does a single initial query for all possible recipient properties (for example, the recipient identifiers, recipient type, message limits, email addresses, and alternative recipients). The applicable recipient properties are cached for later use. Recipient resolution classifies recipients based on similarities in how the recipients are resolved, and the similarity of the applicable recipient properties.

The LDAP filter that's used for address resolution depends on the recipient's email address:

- For the **EX** email address type, the LDAP filter is based on the recipient's **legacyExchangeDN** attribute (higher priority) or the recipient's **proxyAddresses** attribute (lower priority).
- For all other email addresses types, the recipient **proxyAddresses** attribute is used as the LDAP filter.

If the email address doesn't match the recipient's primary SMTP address, the categorizer rewrites the email address in the message to match the primary SMTP address. The original email address is saved in the `ORCPT=` entry in the **RCPT TO** command in the message envelope.

Sender message restrictions

The size of a message can change because of content conversion, encoding, and agent processing. When a message enters the Exchange organization, the original size of the message is recorded in the **X-MS-Exchange-Organization-OriginalSize:** header field in the message header. The lower value of the current message size or the original message size is used to enforce sender message size limits. If the original message size header field doesn't exist, it's created using the current size of the message. If the message is too large, it's returned to the sender in an NDR, and additional message processing is stopped.

The sender recipient limit is only enforced in the Transport service on the first Mailbox server that processes the message. The original, unexpanded message envelope recipient count is compared to the sender recipient limit.

The message sender and all recipients are marked as resolved by stamping an extended property in the message. This extended property allows the message to bypass top-level resolution if the message goes through recipient resolution again (for example, because the Exchange Transport service restarted.

Expansion

Expansion occurs after top-level resolution. Expansion completely expands nested levels of recipients into individual recipients. Expansion may require multiple trips through the expansion process to expand all recipients. Not all recipients have to be expanded. However, all recipients go through the expansion process to enforce recipient message restrictions for all kinds of recipients.

The types of recipients that require expansion are described in this list:

- **Distribution groups and dynamic distribution groups**: Distribution groups are expanded based on the **memberOf** Active Directory property. Dynamic distribution groups are expanded by using the Active Directory query definition. If the *ExpansionServer* parameter is set on the group in the Exchange Management Shell, the group is routed to the specified server for expansion.

  Note: When you specify an expansion server for a group, the group becomes dependent on the availability of the expansion server (messages can't be delivered to the group if the expansion server is unavailable). Therefore, consider implementing high availability solutions for expansion servers.

- **Alternative recipients**: You can configure mailboxes and mail-enabled public folders to forward messages to other recipients:

  - **Mailboxes**: You can configure forwarding to another recipient in the Exchange organization, or to an external email address. For more information, see Configure email forwarding for a mailbox.

  - **Mail-enabled public folders**: You can configure forwarding to another recipient in the Exchange organization.

    You can configure the mailbox or mail-enabled public folder to only send messages to the forwarding address, or to the forwarding address and the original recipient.

- **Contact chains**: A *contact chain* is a mail user or mail contact where the external email address is set to the email address of another recipient in the Exchange organization.

Recipient loop detection

As groups, alternative recipients, and contacts chains are expanded, the categorizer checks for *recipient loops*. A recipient loop is a configuration problem that causes message delivery to the same recipients in an endless circle. The different types of recipient loops are described in this list:

- **Harmless recipient loop**: These are the two scenarios when harmless recipient can loops occur:

  - When two groups contain one another as members.

  - When mailboxes or mail-enabled public folders are set to deliver and forward to one another (the message is delivered to the original recipient and forwarded).

    When the categorizer detects a harmless recipient loop, the message is delivered to the recipient, but no additional attempts are made to deliver the message to the same recipient.

- **Broken recipient loop**: When mailboxes or mail-enabled public folders are set to forward to one another (the messages are only forwarded).

  A broken recipient loop can't result in successful message delivery. When the categorizer detects a broken recipient loop, expansion activity for the current recipient stops, and an NDR is generated.

Recipient loop detection doesn't prevent duplicate message delivery. For example, consider this scenario:

- Distribution Group A has Distribution Group B and Distribution Group C as members.

- Distribution Group C is also a member of Distribution Group B.

In this scenario, Distribution Group C will experience duplicate message delivery.

Delivery report redirection for groups

When a group is expanded, the message type is checked to see if it's a delivery report message. If the message is a delivery report, the redirection settings of the group are checked to see if redirection of the delivery report is required. You may want to suppress delivery reports for the group because a delivery report might disclose unwanted information about the membership of the group.

The delivery report redirection settings that are available in the Exchange Management Shell for distribution groups and dynamic distribution groups are described in this list:

- **ReportToManagerEnabled parameter**: Enables or disables sending delivery reports to the group manager. Valid values are `$true` or `$false`. The default value is `$false`. For a distribution group, the manager is controlled by the *ManagedBy* parameter on the **Set-Group** (distribution groups), or **Set-DynamicDistributionGroup** (dynamic distribution groups) cmdlets.

- **ReportToOriginatorEnabled parameter**: Enables or disables sending delivery reports to the message sender for messages that are sent to the group. Valid values are `$true` or `$false`. The default value is `$true`.

  Note: The values of *ReportToManagerEnabled* parameter and *ReportToOriginatorEnabled* can't both be `$true`. If one parameter is set to `$true`, the other must be set to `$false`. The values of both parameters can be `$false`, which suppresses the redirection of all delivery report messages for the group.

The different types of delivery report messages that can be affected by delivery report redirection for groups are described in this list:

- **Delivery receipt (DR)**: Confirms that a message was delivered to its intended recipient.

- **Delivery status notification (DSN)**: Describes the result of an attempt to deliver a message that didn't result in the message being returned to the sender in an non-delivery report (NDR). For more information about DSN messages, see DSNs and NDRs in Exchange Server.

- **Message disposition notification (MDN)**: Describes the status of a message after it has been successfully delivered to a recipient. Read notifications (RNs) and non-read notification (NRNs) are both examples of MDN messages. MDN messages are defined in RFC 2298 and are controlled by the **Disposition-Notification-To:** header field in the message header. MDN settings that use this header field are compatible with many different kinds of messaging servers. MDN settings can also be defined by using MAPI properties in Outlook and Exchange.

- **Non-delivery report (NDR)**: Indicates to the message sender that the message couldn't be delivered to the specified recipients. The message is returned to the sender in the NDR.

- **Non-read notification (NRN)**: Indicates that a message was deleted before it was read.

- **Out of office (OOF)**: Indicates that the recipient won't respond to email messages. The acronym OOF dates back to the original Microsoft messaging system where the corresponding notification was named "out of facility."

- **Read notification (RN)**: Indicates that a message was read.

- **Recall Report**: Indicates the status of a recall request for a specific recipient (the sender tried to recall a sent message by using Outlook).

These are the settings that cause delivery report messages to be deleted when they're sent to a group:

- Report redirection isn't configured for the group, or report redirection is set to the message sender.

- Report redirection is set to the group manager, and the delivery report message isn't an NDR.

If report redirection is set to the group manager, and the delivery report message is an NDR, the message is delivered to the group manager.

The affect of group delivery report redirection settings on regular messages that contain report requests are described in this list:

- If report redirection is set to the message sender, the report request settings aren't modified.
- If report redirection isn't configured for the group, all report requests are suppressed. The `NOTIFY=NEVER` entry is added to **RCPT TO** for each recipient in the message envelope.
- If report redirection is set to the group manager, NDRs are sent to the group manager, but all other report requests are suppressed.

Message restrictions on recipients

The expansion process also enforces any message restrictions that are configured on recipients. These restrictions may be configured individually for each recipient or organizationally for all servers in the Exchange organization.

For more information on message size limits, see Recipient limits and Organizational limits.

RESTRICTION	EAC CONFIGURATION	EXCHANGE MANAGEMENT SHELL CONFIGURATION	DESCRIPTION
Maximum size of a message received	Organization: **Mail flow** > **Receive connectors** > **More options** ••• > **Organization transport settings** > **Limits** tab > **Maximum receive message size (MB)** Mailboxes: **Recipients** > **Mailboxes** > select the mailbox > **Edit** ✏ > **Mailbox features** > **Mail flow** section > **Message size restrictions** section > **View details** > **Received messages** section > **Maximum message size (KB)** Mail users: **Recipients** > **Contacts** > select the mail user > **Edit** ✏ > **Mail flow settings** > **Message size restrictions** section > **View details** > **Received messages** section > **Maximum message size (KB)**	Organization cmdlet: **Set-TransportConfig** Recipient cmdlets: **Set-DistributionGroup**, **Set-DynamicDistributionGroup**, **Set-Mailbox**, **Set-MailContact**, **Set-MailPublicFolder**, and **Set-MailUser** Parameter: *MaxReceiveSize*	Specifies the maximum size of a message, which includes the message header, the message body, and any attachments. Whenever Exchange checks the message size, the lower value of the current message size or the original message size (the **X-MS-Exchange-Organization-OriginalSize:** message header) is used. The size of the message can change because of content conversion, encoding, and transport agent processing. **Note**: The specified maximum message size is inflated by approximately 33% to account for Base64 encoding (for example, the specified value 64 MB results in a realistic maximum message size of approximately 48 MB). For more information, see Message size and recipient limits in Exchange Server.

RESTRICTION	EAC CONFIGURATION	EXCHANGE MANAGEMENT SHELL CONFIGURATION	DESCRIPTION
The recipient can only accept messages from internal senders, and must reject messages from external senders	Mailboxes: **Recipients** > **Mailboxes** > select the mailbox > **Edit** ✏ > **Mailbox features** > **Mail flow** section > **Message delivery restrictions** section > **View details** > **Accept messages from** section > check or uncheck **Require that all senders are authenticated** Remote mailboxes: **Recipients** > **Mailboxes** > select the Microsoft 365 or Office 365 mailbox > **Edit** ✏ > **Mail flow settings** > **Message delivery restrictions** section > **View details** > **Accept messages from** section > check or uncheck **Require that all senders are authenticated** Mail users: **Recipients** > **Contacts** > select the mail user > **Edit** ✏ > **Mail flow settings** > **Message delivery restrictions** section > **View details** > **Accept messages from** section > check or uncheck **Require that all senders are authenticated** Groups: **Recipients** > **Groups** > select the group > **Edit** ✏ > **Delivery management** > select **Only senders inside my organization** or **Senders inside and outside of my organization**	Cmdlets: **New-DistributionGroup**, **Set-DistributionGroup**, **Set-DynamicDistributionGroup**, **Set-Mailbox**, **Set-MailContact**, **Set-MailPublicFolder**, **Set-MailUser**, and **Set-RemoteMailbox** Parameter: *RequireSenderAuthenticationEnabled*	You can configure the recipient to only accept messages from authenticated (internal) senders, or to accept messages from authenticated and unauthenticated (external) senders.

RESTRICTION	EAC CONFIGURATION	EXCHANGE MANAGEMENT SHELL CONFIGURATION	DESCRIPTION
Senders who are allowed or aren't allowed to send messages to the recipient	Mailboxes: **Recipients** > **Mailboxes** > select the mailbox > **Edit** ✏ > **Mailbox features** > **Mail flow** section > **Message delivery restrictions** section > **View details** > **Accept messages from** section: **All senders** or **Only senders in the following list** or **Reject messages from** section: **No senders** or **Senders in the following list** Remote mailboxes: **Recipients** > **Mailboxes** > select the Microsoft 365 or Office 365 mailbox > **Edit** ✏ > **Mail flow settings** > **Message delivery restrictions** section > **View details** > **Accept messages from** section: **All senders** or **Only senders in the following list** or **Reject messages from** section: **No senders** or **Senders in the following list** Mail users: **Recipients** > **Contacts** > select the mail user > **Edit** ✏ > **Mail flow settings** > **Message delivery restrictions** section > **View details** > **Accept messages from** section: **All senders** or **Only senders in the following list** or **Reject messages from** section: **No senders** or **Senders in the following list** Groups: **Recipients** > **Groups** > select the group > **Edit** ✏ > **Delivery management** > click **Add** + or **Remove** − to specify users or group members who can send to the group (messages from others senders are rejected).	Cmdlets: **Set-DistributionGroup**, **Set-DynamicDistributionGroup**, **Set-Mailbox**, **Set-MailContact**, **Set-MailPublicFolder**, **Set-MailUser**, and **Set-RemoteMailbox** Accept parameters: *AcceptMessagesOnlyFromSendersOrMembers* (or *AcceptMessagesOnlyFrom* for individual recipients only and *AcceptMessagesOnlyFromDLMembers* for group members only) Reject parameters: *RejectMessagesFromSendersOrMembers* (or *RejectMessagesOnlyFrom* for individual recipients only and *RejectMessagesOnlyFromDLMembers* for group members only)	The categorizer checks the recipient permission in two passes. The first pass determines whether the sender is present in the accept or reject lists. If the sender isn't found in either list, the distribution groups in those parameters are fully expanded. This complete group expansion might take some time, so we recommend that you minimize the depth of nested groups in the accept or reject lists.

Certain types of messages that are sent by authenticated senders are exempt from restrictions. The following list describes the messages that are exempt from recipient restrictions:

- **Messages sent by the Microsoft Exchange recipient**: These messages include DSNs and NDRs , journal reports, quota messages, and other system-generated messages that are sent to internal message senders. For more information about the Microsoft Exchange recipient, see Recipients.
- **Messages sent by the external postmaster address**: These messages include DSNs and NDRs, and

other system-generated messages that are sent to external message senders. For more information about the external postmaster address, see Managing the External Postmaster Address.

Exchange blocks certain types of messages that are sent to external domains (for example, internal OOF messages, automatic replies, and meeting forward notifications). You configure these settings in remote domains (the default remote domain, or remote domains for specific external domains). For more information, see Managing Remote Domains.

Bifurcation and controlling recipient expansion

Because the complete list of message recipients is expanded and resolved by recipient resolution, there are occasions when different copies of the same message need to be created:

- **Recipients require different message settings**: Creating a new version of the message that has slightly different properties than the original is called *bifurcation*. For example, Exchange might need to bifurcate a message when read receipts are enabled for some recipients and blocked for others.

- **Limit the number of envelope recipients in a single message**: Expanding large group can generate thousands of individual recipients. Instead of creating a single copy of the message that has thousands of envelope recipients, Exchange creates multiple copies of the same message that have a limited number of recipients in the message envelope.

Bifurcation

Recipient resolution bifurcates a message if the following conditions are true:

- When the message sender in **MAIL FROM** in the message envelope is updated (for example, when the *ReportToManagerEnabled* parameter on a group has the value `$true`).

- When auto-response messages (for example, DSNs and NDRs, OOF messages, and recall reports) need to be suppressed.

- When alternative recipients are expanded.

- When a **Resent-From:** header field is added to the message header. Resent header fields are informational header fields that can be used to determine whether a message has been forwarded by a user. Resent header fields are used so that the message appears to the recipient as if it was sent directly by the original sender. The recipient can view the message header to discover who forwarded the message. Resent header fields are defined in section 3.6.6 of RFC 5322.

- When the expansion history of the group needs to be transmitted.

Controlling recipient expansion

When the number of expanded recipients is too large, the categorizer splits the message into multiple copies to reduce the system resources that are used during message expansion. The maximum number of envelope recipients in a message is controlled by the *ExpansionSizeLimit* key in the `%ExchangeInstallPath%Bin\EdgeTransport.exe.config` application configuration file. The default value is 1000.

Caution

We recommend that you don't modify the value of the *ExpansionSizeLimit* key on an Exchange transport server in a production environment.

Recipient resolution diagnostics

Exchange provides reporting and diagnostic information for recipient resolution in performance counters, message tracking log entries, and recipient resolution logging. These sources can help you identify and diagnose problems with recipient resolution.

Recipient resolution performance counters

The performance counters that are available for recipient resolution are described in this table.

COUNTER NAME	DISPLAY NAME	DESCRIPTION
AmbiguousRecipientsTotal	Ambiguous Recipients	The total number of ambiguous recipients that were detected during recipient resolution. Ambiguous recipients are different recipients that have matching **legacyExchangeDN** Active Directory attributes or matching **proxyAddresses** Active Directory attributes.
AmbiguousSendersTotal	Ambiguous Senders	This is the number of ambiguous senders that were detected during recipient resolution. Ambiguous senders are different senders that have matching **legacyExchangeDN** Active Directory attributes or matching **proxyAddresses** Active Directory attributes.
FailedRecipientsTotal	Failed Recipients	The number of failed recipients that were detected during recipient resolution.
LoopRecipientsTotal	Loop Recipients	The number of recipients that failed recipient resolution because of recipient loops.
MessagesChippedTotal	Messages Chipped	The total number of copies of the same message that were created during recipient resolution to control the number of envelope recipients in a single message. This process is referred to as *chipping*.
MessagesCreatedTotal	Messages Created	The number of messages that were created during recipient resolution.
MessagesRetriedTotal	Messages Retried	The number of messages that were scheduled for retry during recipient resolution.
UnresolvedOrgRecipientsTotal	Unresolved Org Recipients	The number of unresolved recipients from an authoritative domain that were detected during recipient resolution.
UnresolvedOrgSendersTotal	Unresolved Org Senders	The number of unresolved senders from an authoritative domain that were detected during recipient resolution.

Recipient resolution events in the message tracking log

The recipient resolution events that are written in the message tracking log are described in this table.

MESSAGE TRACKING EVENT	DESCRIPTION
EXPAND	A distribution group was expanded.
REDIRECT	A message was redirected to the forwarding address that's configured on the mailbox or mail-enabled public folder.
RESOLVE	A recipient's email address was changed to the primary SMTP email address of the corresponding Active Directory recipient object (in other words, the message was sent to a proxy address of the recipient).
TRANSFER	Message bifurcation or chipping occurred (for example, due to content conversion, message recipient limits, or transport agents).

For more information about message tracking, see Message tracking.

Recipient resolution logging

Recipient resolution logging is controlled by the *ResolverLogLevel* key in the `%ExchangeInstallPath%Bin\EdgeTransport.exe.config` application configuration file. Valid values for this key are:

- `Disabled`: No recipient resolution data is logged. This is the default value.
- `Enabled`: Only message envelope data is logged.
- `FullContent`: Message envelope data and message header data is logged

The log files are stored at `%ExchangeInstallPath%Logging\Resolver`.

> **NOTE**
>
> Any customized Exchange or Internet Information Server (IIS) settings that you made in Exchange XML application configuration files on the Exchange server (for example, web.config files or the EdgeTransport.exe.config file) **will be overwritten** when you install an Exchange CU. Be sure save this information so you can easily re-apply the settings after the install. After you install the Exchange CU, you need to re-configure these settings.

Transport high availability in Exchange Server

8/3/2020 • 3 minutes to read • Edit Online

In Exchange Server, transport high availability is responsible for keeping redundant copies of messages before and after the messages are successfully delivered. These features were introduced in Exchange 2013 as improvements to the transport high availability features in Exchange 2010 (for example, shadow redundancy and the transport dumpster) to help ensure messages aren't lost in transit.

Key features that improve transport high availability in Exchange 2013, Exchange 2016, and Exchange 2019 over Exchange 2010 include:

- Shadow redundancy creates a redundant copy of the message on another server before the message is accepted or acknowledged. The sending server's support or lack of support for shadow redundancy is irrelevant.
- Shadow redundancy recognizes both database availability groups (DAGs) and Active Directory sites as transport high availability boundaries. This reduces the number of servers that can hold redundant copies of messages, and eliminates unnecessary redundant message maintenance traffic across DAGs or Active Directory sites.

  For more information, see Shadow redundancy in Exchange Server.
- The transport dumpster has been improved and is now named *Safety Net*. Safety Net stores messages that were successfully processed by the Transport service on Mailbox servers. Safety Net works best for Mailbox servers in a DAG, but Safety Net also works for multiple Mailbox servers in the same Active Directory site that don't belong to a DAG.
- Safety Net itself is now made redundant on another server. This is important to avoid a single point of failure, because the Transport service and the mailbox databases are both located on the Mailbox server.

  For more information, see Safety Net in Exchange Server.

This diagram provides a high-level overview of how transport high availability works in Exchange Server.

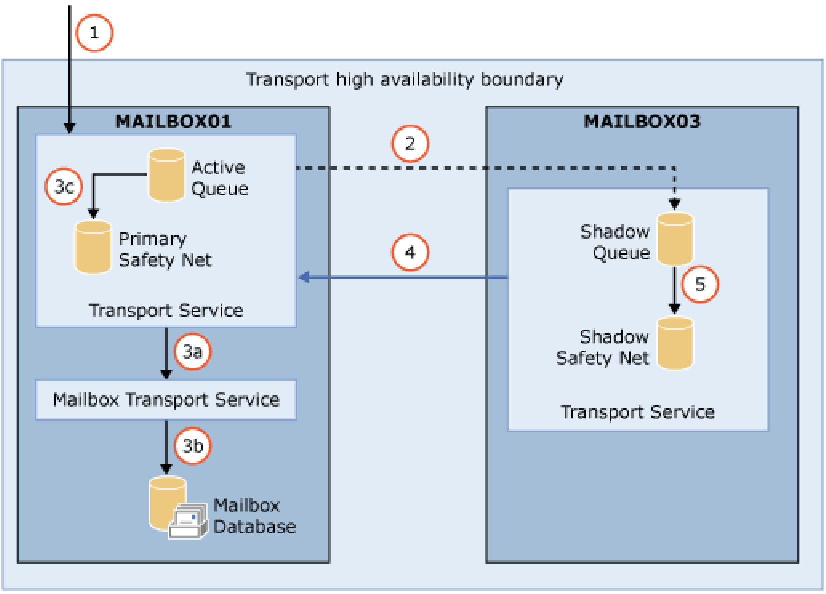

1. An Exchange Mailbox server named Mailbox01 receives a message from an SMTP server that's outside the transport high availability boundary. The *transport high availability boundary* is a DAG or an Active Directory site in non-DAG environments. The message could come from:

- An internal third-party messaging server.
- An Internet messaging server that's proxied through the Front End Transport service on a Mailbox server.
- Another Exchange server in your organization.

2. Before acknowledging receipt of the message, Mailbox01 initiates a new SMTP session to another Exchange Mailbox server named Mailbox03 that's within the Transport high availability boundary, and Mailbox03 makes a shadow copy of the message. In DAG environments, a shadow server in a remote Active Directory site is preferred. Mailbox01 is the primary server holding the primary message, and Mailbox03 is the shadow server holding the shadow message.

3. The Transport service on Mailbox01 processes the primary message.

   a. In this example, the recipient's mailbox is located on Mailbox01, so the Transport service transmits the message to the local Mailbox Transport service.

   b. The Mailbox Transport service delivers the message to the local mailbox database.

   c. Mailbox01 queues a discard status for Mailbox03 that indicates the primary message was successfully processed, and Mailbox01 moves a copy of the primary message into the local Primary Safety Net. Note that the message moves between queues within the same queue database.

4. Mailbox03 periodically polls Mailbox01 for the discard status of the primary message.

5. When Mailbox03 determines Mailbox01 successfully processed the primary message, Mailbox03 moves the shadow message into the local Shadow Safety Net. Note that the message moves between queues within the same queue database.

The message is retained in Primary Safety Net and Shadow Safety Net until the message expires based on a configurable timeout value. If a mailbox database failover occurs before the message expires, the Primary Safety Net on Mailbox01 resubmits the message. If the Mailbox01 isn't available, the Shadow Safety Net on Mailbox03 takes over and resubmits the message.

Message redundancy in the Front End Transport service on Mailbox servers

The Front End Transport service on a Mailbox server (part of the Client Access services) has no message queues. It's a stateless proxy that accepts incoming SMTP connections, and proxies them to the Transport service on a Mailbox server. The Front End Transport service keeps the SMTP session with the sending server open while:

- The primary message is transmitted to the Transport service on a Mailbox server.

  and

- A shadow copy of the message is made by the Transport service on a different Mailbox server within the transport high availability boundary (DAG or Active Directory site).

Only after both the primary message and shadow message are successfully created, the end of data SMTP command is sent back to the sending SMTP server through the Front End Transport service.

Shadow redundancy in Exchange Server

8/3/2020 • 19 minutes to read • Edit Online

Shadow redundancy was introduced in Exchange 2010 to provide redundant copies of messages before they're delivered to mailboxes. In Exchange 2010, shadow redundancy delayed deleting a message from the queue database on a Hub Transport server until the server verified that the next hop in the message delivery path had completed delivery. If the next hop failed before reporting successful delivery back to the Hub Transport server, the server resubmitted the message to that next hop. Exchange 2010 Hub Transport servers used the **XSHADOW** verb to advertise their shadow redundancy support. If a source messaging server didn't support shadow redundancy, Exchange 2010 used delayed acknowledgment based on a configured time interval on the Receive connector to make a redundant copy of the message.

Exchange 2016 and Exchange 2019 have the same improvements that were made to shadow redundancy in Exchange 2013: the Transport service on a Mailbox server now makes a redundant copy of any message it receives before acknowledging the receipt of the message to the sending server. Maintaining redundant copies of messages in transit is more than a best effort that may or may not work, because now shadow redundancy doesn't depend the sending server's supported features (support or lack of support for shadow redundancy doesn't matter). This helps to ensure that all messages in the transport pipeline are made redundant while they're in transit. If Exchange determines the original message was lost in transit, the redundant copy of the message is redelivered.

For more information about transport high availability features in Exchange Server, see Transport high availability in Exchange Server. For more information about message redundancy after a message has been successfully delivered, see Safety Net in Exchange Server.

Shadow redundancy components

This table describes the components of shadow redundancy in the Transport service on Mailbox servers. These terms are used throughout the topic.

TERM	DESCRIPTION
Transport high availability boundary	A database availability group (DAG) in DAG environments, or an Active Directory site in non-DAG environments. For DAGs that span multiple Active Directory sites, the DAG itself is still the boundary (not the site). When a message arrives on a Mailbox server in the transport high availability boundary, Exchange tries to maintain two redundant copies of the message on Mailbox servers within the boundary. When a message leaves the transport high availability boundary, Exchange stops maintaining redundant copies of the message.
Primary message	The message submitted into the transport pipeline for delivery.
Shadow message	The redundant copy of the message that the shadow server retains until it confirms the primary message was successfully processed by the primary server.

TERM	DESCRIPTION
Primary server	The Mailbox server that's currently processing the primary message.
Shadow server	The Mailbox server that holds the shadow message for the primary server. A Mailbox server may be the primary server for some messages and the shadow server for other messages simultaneously.
Shadow queue	The delivery queue where the shadow server stores shadow messages. For messages with multiple recipients, each next hop for the primary message requires separate shadow queues.
Discard status	The information that the Mailbox server maintains for shadow messages to indicate the primary message has been successfully processed.
Discard notification	The response a shadow server receives from a primary server indicating a shadow message is ready to be discarded.
Safety Net	The improved version of the transport dumpster in Exchange 2013 or later. Messages that are successfully processed or delivered to a mailbox recipient by the Transport service on a Mailbox server are moved into Safety Net. For more information, see Safety Net in Exchange Server.
Shadow Redundancy Manager	The transport component that manages shadow redundancy.
Heartbeat	The process that allows primary servers and shadow servers to verify the availability of each other.

Requirements for shadow redundancy

Although it may seem obvious, shadow redundancy requires multiple Mailbox servers:

- If the Mailbox server isn't a member of a DAG, the other Mailbox servers must be in the local Active Directory site.
- If the Mailbox server is a member of a DAG, the other Mailbox servers must belong to the same DAG. The other DAG members can be in the local Active Directory site, or in a remote site. By default, if the DAG spans multiple Active Directory sites, shadow redundancy prefers creating a redundant copy of the message in a remote site for site resiliency.

These are the situations where shadow redundancy can't protect messages in transit:

- In single Exchange server environments.
- In under-provisioned DAGs.
- During the simultaneous failure of two or more Mailbox servers involved in the shadow redundancy of a message.

Shadow redundancy is enabled by default

By default, shadow redundancy is enabled globally in the Transport service on all Mailbox servers. This table describes the parameters that enable shadow redundancy.

PARAMETER	DEFAULT VALUE	DESCRIPTION
ShadowRedundancyEnabled on **Set-TransportConfig**	`$true`	`$true` : Shadow redundancy is enabled on all Mailbox servers in the organization. `$false` : Shadow redundancy is disabled on all Mailbox servers in the organization.
RejectMessageOnShadowFailure on **Set-TransportConfig**	`$false`	`$false` : When a shadow copy of the message can't be created, the primary message is accepted anyway by Mailbox servers in the organization. These messages aren't redundantly persisted while they're in transit. `$true` : No message is accepted or acknowledged by any Mailbox server in the organization until a shadow copy of the message is successfully created. If a shadow copy of the message can't be created, the primary message is rejected with a transient error, but the sending server can transmit the message again. The SMTP response code is `451 4.4.0 Message failed to be made redundant` . All messages in the organization are redundantly persisted while they're in transit. **Note**: Use `$true` only if you have multiple Mailbox servers in the same DAG or Active Directory site so a shadow copy of the message can be created. This parameter is only meaningful when *ShadowRedundancyEnabled* is `$true` .

How shadow messages are created

The main goal of shadow redundancy is to always have two copies of a message within a transport high availability boundary while the message is in transit. Where and when the redundant copy of the message is created depends on where the message came from, and where the message is going. There are three determining factors for creating shadow messages:

- Messages received from outside a transport high availability boundary (the DAG, or an Active Directory site in non-DAG environments).
- Messages sent outside a transport high availability boundary.
- Messages received from the Mailbox Transport Submission service from a Mailbox server within the transport high availability boundary.

Shadow redundancy never tracks shadow messages across a transport high availability boundary. When a message crosses the transport high availability boundary, shadow redundancy begins or restarts. This reduces shadow message maintenance traffic and prevents shadow message resubmissions across the transport high availability boundary. Exchange 2010 Hub Transport servers are a special case, and are discussed later in this topic.

Messages received from outside a transport high availability boundary

When the Transport service on a Mailbox server receives a message from outside the transport high availability boundary, the Mailbox server isn't concerned about the support or lack of support for shadow redundancy by the sending server. As long as shadow redundancy is enabled, the Mailbox server that receives the message makes a redundant copy of the message on another Mailbox server within the transport high availability boundary before acknowledging receipt of the message back to the sending server. Here's an example of how the process works:

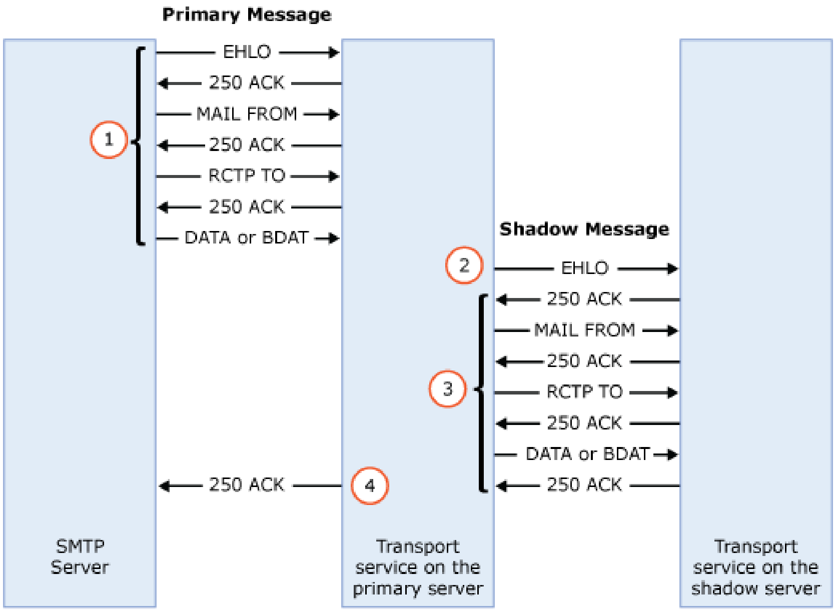

1. A messaging server transmits a message to the Transport service on a Mailbox server. The Mailbox server is the primary server, and the message is the primary message.

2. While the original SMTP session with the messaging server is still active, the Transport service on primary server opens a new, simultaneous SMTP session with the Transport service on a different Mailbox server in the organization to create a redundant copy of the message.

   - If the primary server is a member of a DAG, the primary server connects to a different Mailbox server in the same DAG. If the DAG spans multiple Active Directory sites, a Mailbox server in a different Active Directory site is preferred by default (the default value of the *ShadowMessagePreferenceSetting* parameter on the **Set-TransportConfig** cmdlet is `PreferRemote`, but you can change it to `RemoteOnly` or `LocalOnly`).

   - If the primary server isn't a member of a DAG, the primary server connects to a different Mailbox server in the same Active Directory site (regardless of the value of the *ShadowMessagePreferenceSetting* parameter).

3. The primary server transmits a copy of the message to the Transport service on another Mailbox server, and 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.

4. After the primary server receives acknowledgment from the shadow server, the primary server acknowledges the receipt of the primary message to the original messaging server in the original SMTP session, and the SMTP session is closed.

Messages sent outside a transport high availability boundary

When a Mailbox server transmits a message outside the transport high availability boundary, and the messaging server on the other side acknowledges successful receipt of the message, and the Mailbox server moves the message into Safety Net. No resubmission of the message from Safety Net can occur after the primary message has been successfully transmitted across the transport high availability boundary. For more information about Safety Net, see Safety Net in Exchange Server.

Messages transmitted within a transport high availability boundary

Message routing is optimized so that when the ultimate destination is in a DAG or Active Directory site, multiple hops between servers within the destination DAG or site aren't typically required. After the message is accepted by the Transport service on a Mailbox server in the destination DAG or Active Directory, the next hop for the message is typically the ultimate destination itself (for example, the Mailbox server that holds the active copy of the destination mailbox). Shadow redundancy's goal of keeping two copies of a message in transit is fulfilled when one shadow copy of the message exists *anywhere* within the DAG or Active Directory site. Typically, only failover scenarios in a DAG that require the **Redirect-Message** cmdlet to drain the active message queues on a Mailbox server would require multiple hops within the same transport high availability boundary.

Shadow redundancy with Exchange 2010 Hub Transport servers in the same Active Directory site in Exchange 2016 organizations

When an Exchange 2010 Hub Transport server transmits a message to an Exchange 2016 Mailbox server in the same Active Directory site, the Exchange 2010 Hub Transport server advertises support for shadow redundancy using the XSHADOW command, but the Mailbox server doesn't advertise support for shadow redundancy. This prevents the Exchange 2010 Hub Transport server from creating a shadow copy of the message on an Exchange 2016 Mailbox server.

When the Transport service on an Exchange 2016 Mailbox server transmits a message to an Exchange 2010 Hub Transport in the same Active Directory site, the Exchange 2016 Mailbox server shadows the message for the Exchange 2010 Hub Transport server. After the Exchange 2016 Mailbox server receives acknowledgment from the Exchange 2010 Hub Transport server that the message was successfully received, the Exchange 2016 Mailbox server moves the successfully processed message into Safety Net. However, the successfully processed messages stored in Safety Net by Exchange 2016 Mailbox are never resubmitted to the Exchange 2010 Hub Transport servers.

SMTP timeouts

During the attempt to make a redundant copy of the message, the SMTP connection between servers (the sending server and the primary server, or the primary server and the shadow server) could timeout. Receive connectors and Send connectors both have a *ConnectionInactivityTimeOut* parameter for when data is actually being transmitted on the connector. Receive connectors also have an absolute *ConnectionTimeOut* parameter.

If any of the SMTP sessions time out before the shadow copy of the message is successfully created and acknowledged, the result is controlled by the *RejectMessageOnShadowFailure* parameter on the **Set-TransportConfig** cmdlet. By default, the value of this parameter is `$false`, which means the primary message is accepted without a shadow copy being created. If the value of this parameter is `$true` the primary message is rejected with the transient error `451 4.4.0`.

If the shadow copy of a message is successfully created, but the SMTP session between the sending server and the primary server times out, the primary server accepts and processes the primary message. The sending server will re-deliver the unacknowledged message, but duplicate message detection will prevent Exchange mailbox users from seeing the duplicate messages. When the sending server resubmits the message, the primary server will create another shadow copy of the message. There's no relationship between the shadow messages created during message resubmissions by the sending server.

The following table describes the parameters that control the creation of shadow messages

SOURCE	DEFAULT VALUE	DESCRIPTION
ShadowMessagePreferenceSetting on **Set-TransportConfig**	`PreferRemote`	This parameter is only used when the primary server that's trying to make a shadow copy of the message is a member of a DAG that spans multiple Active Directory sites. • `PreferRemote`: Try to make a shadow copy of the message on a DAG member in a different Active Directory site based on the number of attempts specified by the *MaxRetriesForRemoteSiteShadow* parameter. If the operation fails, try make a shadow copy of the message on a DAG member in the local Active Directory site based on the number of attempts specified by the *MaxRetriesForLocalSiteShadow* parameter. • `LocalOnly`: Try to make a shadow copy of the message only on a DAG member in the local Active Directory site based on the number of attempts specified by the *MaxRetriesForLocalSiteShadow* parameter. • `RemoteOnly`: Try to make shadow copy of the message only on a DAG member in a different Active Directory site based on the number of attempts specified by the *MaxRetriesForRemoteSiteShadow* parameter.
MaxRetriesForRemoteSiteShadow on **Set-TransportConfig**	4	This parameter specifies the maximum number of attempts to create a shadow copy of the message on another server in the DAG when the value of the *ShadowMessagePreferenceSetting* parameter is `PreferRemote` (the default value) or `RemoteOnly`. This parameter is only used when the Mailbox server is a member of a DAG that spans multiple Active Directory sites. If a shadow copy of the message isn't successfully created after the specified number of attempts, the result depends of the value of the *RejectMessageOnShadowFailure* parameter: • `$true`: The primary message is rejected with a transient error. • `$false`: The primary message is accepted anyway, but isn't redundantly persisted.

SOURCE	DEFAULT VALUE	DESCRIPTION
MaxRetriesForLocalSiteShadow on **Set-TransportConfig**	2	This parameter specifies the maximum number of attempts to create a shadow copy of the message on another Mailbox server in the local Active Directory site when: • The Mailbox server is a member of a DAG that spans multiple Active Directory sites, and the value of the *ShadowMessagePreferenceSetting* parameter is `PreferRemote` (the default value) or `LocalOnly`. • The Mailbox server is a member of a DAG that's in one Active Directory site. • The Mailbox server isn't a member of a DAG. If a shadow copy of the message isn't successfully created after the specified number of attempts, the result depends of the value of the *RejectMessageOnShadowFailure* parameter: • `$true`: The primary message is rejected with a transient error. ò `$false`: The primary message is accepted anyway, but isn't redundantly persisted.
ConnectionInactivityTimeout on **Set-ReceiveConnector**	5 minutes for Receive connectors in the Transport service on Mailbox servers	This parameter specifies the maximum time that an open SMTP connection with the source messaging server can remain idle before the connection is closed. The value of this parameter must be greater than the value of the *ConnectionTimeout* parameter.
ConnectionTimeout on **Set-ReceiveConnector**	10 minutes for Receive connectors in the Transport service on Mailbox servers	This parameter specifies the maximum time that an SMTP connection with the source messaging server can remain open, even if the server is transmitting data. The value of this parameter must be greater than the value of the *ConnectionInactivityTimeout* parameter.
ConnectionInactivityTimeOut on **Set-SendConnector**	10 minutes	This parameter specifies the maximum time that an open SMTP connection with a destination messaging server can remain idle before the connection is closed.

How shadow messages are maintained

After a shadow message is successfully created, the work of shadow redundancy has only just begun. The primary server and the shadow server need to stay in contact with each other to track the progress of the message.

When the primary server successfully transmits the message to the next hop, and the next hop acknowledges receipt of the message, the primary server updates the *discard status* of the message as delivery complete. The discard status is basically a message that contains of list of messages that are being monitored. A successfully delivered message doesn't need to be kept in a shadow queue, so once the shadow server knows the primary server has successfully transmitted the message to the next hop, the shadow server moves the shadow message from the shadow queue into Safety Net.

The shadow server determines the discard status of the shadow messages in its shadow queues by querying the primary server. If the shadow server opens an SMTP session with the primary server for any reason (including the transmission of other unrelated messages), the shadow server issues the **XQDISCARD** command to determine the discard status of the primary messages. Otherwise, the shadow server will automatically open an SMTP session with the primary server after a preconfigured time interval (the *ShadowHeartbeatFrequentcy* parameter on the **Set-TransportConfig** cmdlet; the default value is 2 minutes).

After the shadow server opens an SMTP session with the primary server, the primary server responds with the *discard notifications* for messages that apply to the querying shadow server. Discard notifications are stored on disk (not in memory) so, if the Microsoft Exchange Transport service restarts, the discard notifications persist. After the service starts, the primary server still knows about the messages it successfully processed, and that information is available to the shadow server.

The SMTP communication between the shadow server and the primary server is used as the *heartbeat* that determines the availability of the servers. If the shadow server can't open an SMTP session with the primary server after a preconfigured time interval (the *ShadowResubmitTimeSpan* parameter on the **Set-TransportConfig** cmdlet; the default value is 3 hours) the shadow server promotes itself as the primary server, promotes the shadow messages as primary messages, and transmits the messages to the next hop. But, whenever the shadow server detects that the queue database ID of the primary server has changed, the shadow server also promotes itself as the primary server, promotes the shadow messages as primary messages, and transmits the messages to the next hop. This could happen well before the *ShadowResubmitTimeSpan* parameter value has passed.

Shadow Redundancy Manager is the core component on a Mailbox server that's responsible for managing shadow redundancy. Shadow Redundancy Manager is responsible for maintaining the following information for all the primary messages that a server is currently processing:

- The shadow server for each primary message that's being processed.
- The discard status to be sent to shadow servers.

Shadow Redundancy Manager is responsible for the following actions for all the shadow messages that a shadow server has in its shadow queues:

- Maintaining the list of primary servers for each shadow message.
- Comparing the original database ID and the current database ID of the queue database where the primary copy of the message is stored.
- Checking the availability of each primary server for which a shadow message is queued.
- Processing discard notifications from primary servers.
- Removing the shadow messages from the shadow queues after all expected discard notifications are received.
- Deciding when the shadow server should take ownership of shadow messages, becoming a primary server.
- Tracking message bifurcations and other side-effect messages like delivery status notifications (also known as DSNs, non-delivery reports, NDRs, or bounce messages) and journal reports to verify that the redundant copy of the message isn't released until all forks of the message are fully processed.

This table describes the parameters that control how shadow messages are maintained.

PARAMETER	DEFAULT VALUE	DESCRIPTION
ShadowHeartbeatFrequency on **Set-TransportConfig**	2 minutes	The maximum amount of time a shadow server waits before opening an SMTP connection to the primary server to check the discard status of messages.
ShadowResubmitTimeSpan on **Set-TransportConfig**	3 hours	How long a server waits before deciding that a primary server has failed and assumes ownership of shadow messages in the shadow queue for the primary server that's unreachable. Note that a shadow server can also promote itself as the primary server before the value of this parameter when the queue database of the primary server is found to have a different database ID.
ShadowMessageAutoDiscardInterval on **Set-TransportConfig**	2 days	How long a server retains discard events for successfully delivered messages. A primary server queues discard events until it's queried by the shadow server. However, if the shadow server doesn't query the primary server for the duration specified in this parameter, the primary server deletes the queued discard events.
SafetyNetHoldTime on **Set-TransportConfig**	2 days	How long successfully processed messages are retained in Safety Net. Unacknowledged shadow messages eventually expire from Safety Net after the sum of the *SafetyNetHoldTime* and *MessageExpirationTimeout* parameter values on the **Set-TransportService** cmdlet.
MessageExpirationTimeout on **Set-TransportService**	2 days	How long a message can remain in a queue before it expires.

Message processing after an outage

This table summarizes how shadow redundancy minimizes message loss due to server outages. For clarity, the server that had an outage is named Mailbox01.

RECOVERY SCENARIO	ACTIONS TAKEN

RECOVERY SCENARIO	ACTIONS TAKEN
Mailbox01 comes back online with a new queue database before the value of the *ShadowResubmitTimeSpan* parameter has passed (by default, 3 hours). This scenario can occur when the queue database is unrecoverable due to data corruption or hardware failure.	When the new queue database ID is detected on Mailbox01, each server that has shadow messages queued for Mailbox01 will assume ownership of those messages and resubmit them. The messages are then delivered to their destinations. The maximum delay for message submission after the new queue database is detected is the value of the *ShadowHeartbeatFrequency* parameter (by default, 2 minutes).
Mailbox01 comes back online with the same database after the value of the *ShadowResubmitTimeSpan* parameter has passed (by default, 3 hours). This scenario can occur after a network card failure, or time-consuming maintenance on the server.	After Mailbox01 comes back online, it will deliver the messages in its queues, which have already been delivered by the servers that hold shadow copies of messages for Mailbox01. This will result in duplicate delivery of these messages. Exchange mailbox users won't see duplicate messages due to duplicate message detection. However, recipients on other messaging systems might see duplicate copies of messages. The maximum delay for message submission is the value of the *ShadowResubmitTimeSpan* parameter.

Safety Net in Exchange Server

8/3/2020 • 11 minutes to read • Edit Online

In Exchange 2010, the *transport dumpster* helped protect against data loss by maintaining a queue of successfully delivered messages that hadn't replicated to the passive mailbox database copies in the database availability group (DAG). When a mailbox database or server failure required the promotion of an out-of-date copy of the mailbox database, the messages in the transport dumpster were automatically resubmitted to the new active copy of the mailbox database.

The transport dumpster was improved in Exchange 2013 and is now called *Safety Net*. Exchange 2016 and Exchange 2019 have these same improvements.

Here's how Safety Net is similar to the transport dumpster in Exchange 2010:

- Safety Net is a queue that's associated with the Transport service on a Mailbox server. This queue stores copies of messages that were successfully processed by the server.

- You can specify how long Safety Net stores copies of the successfully processed messages before they expire and are automatically deleted. The default is 2 days.

Here's how Safety Net is improved from the transport dumpster in Exchange 2010:

- **Safety Net doesn't require a DAG**: For Mailbox servers that don't belong to a DAG, Safety Net stores copies of the delivered messages on other Mailbox servers in the local Active Directory site.

- **Safety Net itself isn't a single point of failure**: Redundancy is provided by using a *Primary Safety Net* and a *Shadow Safety Net*. If the Primary Safety Net is unavailable for more than 12 hours, resubmit requests become shadow resubmit requests, and messages are re-delivered from the Shadow Safety Net.

- **Safety Net takes over some responsibility from shadow redundancy in DAG environments**: Shadow redundancy doesn't need to keep another copy of the delivered message in a shadow queue while it waits for the delivered message to replicate to the passive copies of mailbox database. The copy of the delivered message is already stored in Safety Net, so the message can be resubmitted from Safety Net if necessary.

- **Safety Net tries to guarantee message redundancy**: Safety Net is more than just a best effort for message redundancy, so you can't specify a maximum size limit for Safety Net. You can only specify how long Safety Net stores messages before they're automatically deleted.

For more information about transport high availability features in Exchange Server, see Transport high availability in Exchange Server. For more information about message redundancy for messages in transit, see Shadow redundancy in Exchange Server.

How Safety Net works

Shadow redundancy keeps a redundant copy of the message while the message is in transit. Safety Net keeps a redundant copy of a message after the message is successfully processed. So, Safety Net begins where shadow redundancy ends. concepts in shadow redundancy, including the transport high availability boundary, primary messages, primary servers, shadow messages and shadow servers also apply to Safety Net. For more information, see Shadow redundancy in Exchange Server.

The Primary Safety Net exists on the Mailbox server that held the primary message before the message was successfully processed by the Transport service. This could mean the message was delivered to the Mailbox

Transport Delivery service on the destination Mailbox server. Or, the message could have been relayed through the Mailbox server in an Active Directory site that's designated as a hub site on the way to the destination DAG or Active Directory site. After the primary server processes the primary message, the message is moved from the active delivery queue into the Primary Safety Net on the same server.

The Shadow Safety Net exists on the Mailbox server that held the shadow message. After the shadow server determines the primary server has successfully processed the primary message, the shadow server moves the shadow message from the shadow queue into the Shadow Safety Net on the same server. Although it may seem obvious, the existence of the Shadow Safety Net requires shadow redundancy to be enabled (it's is enabled by default).

This table describes the parameters that are used by Safety Net.

PARAMETER	DEFAULT VALUE	DESCRIPTION
SafetyNetHoldTime on **Set-TransportConfig**	2 days	The length of time successfully processed primary messages are stored in Primary Safety Net, and acknowledged shadow messages are stored in Shadow Safety Net. You can also specify this value in the Exchange admin center (EAC) at **Mail flow** > **Receive connectors** > **More options** ••• > **Organization transport settings** > **Safety Net** > **Safety Net hold time**. Unacknowledged shadow messages eventually expire from Shadow Safety Net after the sum of *SafetyNetHoldTime* and *MessageExpirationTimeout* parameter values. To avoid data loss during Safety Net resubmits, the value of this parameter must be greater than or equal to the value of *ReplayLagTime* on **Set-MailboxDatabaseCopy** for the lagged copy of the mailbox database.
ReplayLagTime on **Set-MailboxDatabaseCopy**	Not configured	The amount of time that the Microsoft Exchange Replication service should wait before replaying log files that have been copied to the passive database copy. Setting this parameter to a value greater than 0 creates a lagged copy of the mailbox database. The maximum value is 14 days. To avoid data loss during Safety Net resubmits, the value of this parameter for the lagged copy of the mailbox database must be less than or equal to the value of *SafetyNetHoldTime* on **Set-TransportConfig**.
MessageExpirationTimeout on **Set-TransportService**	2 days	How long a message can remain in a queue before it expires.

PARAMETER	DEFAULT VALUE	DESCRIPTION
ShadowRedundancyEnabled on **Set-TransportConfig**	`$true`	`$true` : Shadow redundancy is enabled on all Mailbox servers in the organization. `$false` : Shadow redundancy is disabled on all transport servers in the organization. Redundancy for Safety Net requires shadow redundancy to be enabled.

Safety Net maximum supported sizes

In Microsoft Exchange Server 2019 and 2016, the maximum supported database size for the transport Safety Net JET database is 2 TB.

When a Hub-and-spoke topology is used, the transport Safety Net JET database can grow beyond 2 TB. To stay within the supported limit of 2 TB, follow these guidelines:

- Hub servers that are used for message relay can't be configured to deliver messages to mailboxes.
- Disable Safety Net on hub servers that are used for message relay. To do this, follow these steps:

  1. In a Command prompt window, open the EdgeTransport.exe.config file in **Notepad** by running the following command on the server:

     ```
     Notepad %ExchangeInstallPath%Bin\EdgeTransport.exe.config
     ```

  2. Add the following key in the **appSettings** section.

     ```
     <add key="SafetyNetHoldTimeInterval" value="0.00:00:15" />
     ```

     When you're finished, save and close the EdgeTransport.exe.config file.

  3. Restart the Exchange Transport service by running the following command:

     ```
     net stop MSExchangeTransport && net start MSExchangeTransport
     ```

Message resubmission from Safety Net

The Active Manager component of the Microsoft Exchange Replication service (MRS) manages DAGs and mailbox database copies. Message resubmissions from Safety Net require no manual actions, and are initiated by the Active Manager. For more information about Active Manager, see Active Manager.

There are two basic Safety Net message resubmission scenarios:

- After the automatic or manual failover of a mailbox database in a DAG.
- After you activate a lagged copy of a mailbox database.

A *lagged mailbox database copy* or *lagged copy* is a passive copy of a mailbox database where updates to the database are intentionally delayed to protect against logical corruption of the mailbox database. For more information, see Manage mailbox database copies.

The only significant difference between the two scenarios is how far back in time to go to resubmit messages from Safety Net. Typically, for database failover in a DAG, the new active copy of the mailbox database is anywhere from several minutes to several hours behind the old active copy. A lagged copy of a mailbox database is typically several days behind the old active copy.

The main requirement for successful message resubmission from Safety Net for a lagged copy is: the length of time messages are stored in Safety Net must be greater than or equal to the lag time of the lagged copy. In other words, the value of *SafetyNetHoldTime* on **Set-TransportConfig** must be greater than or equal to the value of the *ReplayLagTime* on **Set-MailboxDatabaseCopy** for the lagged copy.

Message resubmission from Shadow Safety Net

Message resubmission from Shadow Safety Net (like message resubmission from Primary Safety Net), is fully automated, and requires no manual intervention. This scenario describes the interaction of Primary Safety Net and Shadow Safety Net during message resubmission:

1. Active Manager requests message resubmission from Safety Net for a mailbox database for the specified time interval (for example, 5:00 to 9:00). However, the Mailbox server that holds the Primary Safety Net has crashed due to a hardware failure. Active Manager unsuccessfully tries to contact the Primary Safety Net for the next 12 hours.

2. After 12 hours, Active Manager sends a broadcast message to the Transport service on all Mailbox servers in the transport high availability boundary (the DAG or Active Directory site in non-DAG environments) looking for other Safety Nets that contain messages for the target mailbox database for the specified time interval. The Shadow Safety Net responds and resubmits messages for the mailbox database for the time interval 5:00 to 9:00.

When a Shadow Safety Net responds, it only resubmits the messages for the required mailbox database during the required time interval. This restriction by mailbox database and time interval helps reduce these potential issues:

- Resubmitting messages from Safety Net could result in duplicate deliveries. This isn't an issue for mailboxes in the Exchange organization, because duplicate message detection prevents mailbox users from seeing the duplicate messages. But, duplicate message delivery to external recipients could result in duplicate copies of messages that the recipient would see.

- Shadow messages resubmitted from Shadow Safety Net require full categorization and processing through the Transport service on the Mailbox server. Resubmission of a large number of shadow messages can be expensive in terms of Mailbox server system resources.

These are some important considerations for the shadow messages that are stored in Shadow Safety Net:

- Shadow Safety Net doesn't know where the primary server transmitted the primary message.

- The shadow messages in Shadow Safety Net only contain original message envelope recipients, not the actual recipients where the primary message was delivered (for example, the message envelope recipient might be a distribution group that requires expansion).

- The messages in Shadow Safety net don't contain any message updates that occurred after the primary server processed the message (for example, message encoding or content conversion).

This scenario describes what happens if the Primary Safety Net is offline during part of the requested resubmit interval:

1. The queue database on the Mailbox server that holds the Primary Safety Net is corrupt, and a new queue database is created at 7:00. All of the primary messages stored in the Primary Safety Net from 1:00 to 7:00 are lost, but the server is able to store copies of successfully delivered messages in Safety Net starting at

7:00.

2. Active Manager requests a resubmission of messages from Safety Net for a mailbox database for the time interval 1:00 to 9:00.

3. The Primary Safety Net resubmits messages for the time interval 7:00 to 9:00.

4. Because the Primary Safety Net doesn't have the required messages for 1:00 to 7:00. the Primary Safety Net sends a broadcast message to the Transport service on all Mailbox servers in the transport high availability boundary looking for other Safety Nets that contain the required messages. The Shadow Safety Net generates a second resubmit request on behalf of the Primary Safety Net to resubmit the shadow messages for the target mailbox database for the time interval 1:00 to 7:00.

These are some other issues to consider when messages are resubmitted from Safety Net:

- **All delivery status notifications (also known as DSNs, non-delivery reports, NDRs or bounce messages) are suppressed for Safety Net message resubmissions**: For example, if the primary message resulted in an NDR, the NDR for the resubmitted message won't be delivered.

- **Users removed from a distribution group may not receive a resubmitted message when the Shadow Safety Net resubmits the message**: For example, a message is sent to a group containing User A and User B, and both recipients receive the message. User B is subsequently removed from the group. Later, a resubmit request from Primary Safety Net is made for the mailbox database that holds User B's mailbox. However, the Primary Safety Net is unavailable for more than 12 hours, so the Shadow Safety Net server responds and resubmits the affected message. During message resubmission when the distribution group is expanded, User B is no longer a member of the group, and won't receive a copy of the resubmitted message.

- **New Users added to a distribution group may receive an old resubmitted message when the Shadow Safety Net resubmits the message**: For example, a message is sent to a group containing User A and User B, and both recipients receive the message. User C is subsequently added to the group. Later, a resubmit request from Primary Safety Net is made for the mailbox database that holds User C's mailbox. However, the Primary Safety Net server is unavailable for more than 12 hours, so the Shadow Safety Net server responds and resubmits the affected messages. During message resubmission when the distribution group is expanded, User C is now a member of the group, and will receive a copy of the resubmitted message.

- **Deploying Safety Net in Hub and Spoke Topology**: Safety Net is designed to protect message delivery on Exchange Servers hosting end-user mailboxes. Customers who have deployed a hub and spoke routing topology should disable Safety Net on transport servers in hub sites to avoid a large growth in the size of the transport database in hub locations.

Transport logs in Exchange Server

8/3/2020 • 3 minutes to read • Edit Online

Transport logs provide information about what's happening in the transport pipeline. For more information about the transport pipeline, see Mail flow and the transport pipeline.

The transport logs in Exchange Server are described in the following sections.

Agent logging

Agent logging records the actions that are performed on messages by specific antispam transport agents on the Exchange server. For more information, see these topics:

- Antispam Agent Logging
- Configure Antispam Agent Logging
- Enable antispam functionality on Mailbox servers

Enabled by default?: Yes

Default location of log files: Note that the folder isn't created until an agent attempts to write information to the log.

- **Mailbox servers**:
  - **Front End Transport service**: `%ExchangeInstallPath%TransportRoles\Logs\FrontEnd\AgentLog`
  - **Transport service**: `%ExchangeInstallPath%TransportRoles\Logs\Hub\AgentLog`
- **Transport service on Edge Transport servers**: `%ExchangeInstallPath%TransportRoles\Logs\Edge\AgentLog`

Connectivity logging

Connectivity logging records outbound message transmission activity by the transport services on the Exchange server. For more information, see these topics:

- Connectivity logging in Exchange Server
- Configure connectivity logging in Exchange Server

Enabled by default?: Yes

Default location of log files:

- **Mailbox servers**:
  - **Front End Transport service**: `%ExchangeInstallPath%TransportRoles\Logs\FrontEnd\Connectivity`
  - **Transport service**: `%ExchangeInstallPath%TransportRoles\Logs\Hub\Connectivity`
  - **Mailbox Transport Delivery service**: `%ExchangeInstallPath%TransportRoles\Logs\Mailbox\Connectivity\Delivery`
  - **Mailbox Transport Submission service**:
- **Transport service on Edge Transport servers**:

`%ExchangeInstallPath%TransportRoles\Logs\Edge\Connectivity`

Message tracking and delivery reports for administrators

Message tracking is a detailed record of all message activity as mail flows through the transport pipeline on an Exchange server. For more information, see these topics:

- Message tracking
- Configure message tracking
- Search message tracking logs

Delivery reports for administrators is a targeted search of the message tracking log for messages that were sent to or from a specified mailbox. For more information, see these topics:

- Delivery reports for administrators
- Track messages with delivery reports

Enabled by default?: Yes

Default location of log files:

- **Mailbox servers**: `%ExchangeInstallPath%TransportRoles\Logs\MessageTracking`:
  - `MSGTRK` files for the Transport service.
  - `MSGTRMD` files for the Mailbox Transport Delivery service.
  - `MSGTRMS` files for the Mailbox Transport Submission service.
- **Transport service on Edge Transport servers**: `%ExchangeInstallPath%TransportRoles\Logs\MessageTracking`

Pipeline tracing

Pipeline tracing records snapshots of messages before and after the message is affected by transport agents in the transport pipeline. For more information, see these topics:

- Pipeline Tracing
- Configure Pipeline Tracing

Enabled by default?: No

Default location of log files: Note that the folder isn't created until pipeline tracing is enabled.

- **Mailbox servers**:
  - **Transport service**: `%ExchangeInstallPath%TransportRoles\Logs\Hub\PipelineTracing`
  - **Mailbox Transport service**: `%ExchangeInstallPath%TransportRoles\Logs\Mailbox\PipelineTracing`
- **Transport service on Edge Transport servers**: `%ExchangeInstallPath%TransportRoles\Logs\Edge\PipelineTracking`

Protocol logging

Protocol logging records the SMTP conversations that occur on Send connectors and Receive connectors during message delivery. For more information, see these topics:

- Protocol logging
- Configure protocol logging

Enabled by default?: Only on these connectors:

- The default Receive connector named Default Frontend *<ServerName>* in the Front End Transport service on Mailbox servers.
- The implicit and invisible Send connector in the Front End Transport service on Mailbox servers.

For more information about these connectors, see Default Receive connectors created during setup and Implicit Send connectors.

Default location of log files:

- **Mailbox servers**:
  - **Front End Transport service**:
  - **Receive connectors**: `%ExchangeInstallPath%TransportRoles\Logs\FrontEnd\ProtocolLog\SmtpReceive`
  - **Send connectors**: `%ExchangeInstallPath%TransportRoles\Logs\FrontEnd\ProtocolLog\SmtpSend`
  - **Transport service**:
  - **Receive connectors**: `%ExchangeInstallPath%TransportRoles\Logs\Hub\ProtocolLog\SmtpReceive`
  - **Send connectors**: `%ExchangeInstallPath%TransportRoles\Logs\Hub\ProtocolLog\SmtpSend`
  - **Mailbox Transport Delivery service (Receive Connectors)**: `%ExchangeInstallPath%TransportRoles\Logs\Mailbox\ProtocolLog\SmtpReceive\Delivery`
  - **Mailbox Transport Submission service**:
  - **Send connectors**: `%ExchangeInstallPath%TransportRoles\Logs\Mailbox\ProtocolLog\SmtpSend\Submission`
  - **Side effect messages**: `%ExchangeInstallPath%TransportRoles\Logs\Mailbox\ProtocolLog\SmtpSend\Delivery`
- **Transport service on Edge Transport servers**:
  - **Receive connectors**: `%ExchangeInstallPath%TransportRoles\Logs\Edge\ProtocolLog\SmtpReceive`
  - **Send connectors**: `%ExchangeInstallPath%TransportRoles\Logs\Edge\ProtocolLog\SmtpSend`

Routing table logging

Note: The Routing Log Viewer is no longer available in the Exchange Toolbox.

Routing table logging periodically records snapshots of the routing table that Exchange servers uses to deliver messages. For more information, see these topics:

- Understanding Routing Table Logging
- Configure Routing Table Logging

Enabled by default?: Yes

Default location of log files:

- **Mailbox servers**:

  - **Front End Transport service**: `%ExchangeInstallPath%TransportRoles\Logs\FrontEnd\Routing`
  - **Transport service**: `%ExchangeInstallPath%TransportRoles\Logs\Hub\Routing`
  - **Mailbox Transport service:**: `%ExchangeInstallPath%TransportRoles\Logs\Mailbox\Routing`:
  - `MDRoutingConfig` files for the Mailbox Transport Delivery service.
  - `MSRoutingConfig` files for the Mailbox Transport Submission service.
- **Transport service on Edge Transport servers**: `%ExchangeInstallPath%TransportRoles\Logs\Edge\Routing`

Message tracking

8/3/2020 • 15 minutes to read • Edit Online

The message tracking log is a detailed record of all activity as mail flows through the transport pipeline on Mailbox servers and Edge Transport servers. You can use message tracking for message forensics, mail flow analysis, reporting, and troubleshooting.

By default, Exchange uses circular logging to limit the message tracking log based on file size and file age to help control the hard disk space that's used by the log files. To configure the message tracking log, see Configure message tracking.

Search the message tracking log

Message tracking logs contain vast amounts of data as messages move through a Mailbox server or Edge Transport server. When it comes to searching the message tracking logs, you have options:

- **Get-MessageTrackingLog**: Administrators can use this Exchange Management Shell cmdlet to search the message tracking log for information about messages using a wide range of filter criteria. For more information, see Search message tracking logs.
- **Delivery reports for administrators**: Administrators can use the **Delivery reports** tab in the Exchange admin center or the underlying **Search-MessageTrackingReport** and **Get-MessageTrackingReport** cmdlets in the Exchange Management Shell to search the message tracking logs for information about messages sent by or received by a specific mailbox in the organization. For more information, see Delivery reports for administrators.

Structure of the message tracking log files

By default, the message tracking log files exist in `%ExchangeInstallPath%TransportRoles\Logs\MessageTracking`. The folder contains log files that have different names, but they all follow the naming convention `MSGTRKServiceyyyymmdd-nnnn.log`. The different log file names are described in the following table.

FILE NAME	SERVERS	DESCRIPTION
`MSGTRK`	Mailbox servers and Edge Transport servers	Log files for the Transport service.
`MSGTRKMA`	Mailbox servers	Log files for the approvals and rejections in moderated transport. For more information, see Manage message approval.
`MSGTRKMD`	Mailbox servers	Log files for messages delivered to mailboxes by the Mailbox Transport Delivery service.
`MSGTRKMS`	Mailbox servers	Log files for messages sent from mailboxes by the Mailbox Transport Submission service.

The other placeholders in the log file names represent the following information:

- *yyyymmdd* is the coordinated universal time (UTC) date when the log file was created. *yyyy* = year, *mm* = month, and *dd* = day.

- *nnnn* is an instance number that starts at the value 1 every day for each log.

Information is written to the log file until the file reaches its maximum size. Then, a new log file that has an incremented instance number is opened (the first log file is -1, the next is -2, and so on). Circular logging deletes the oldest log files for a service when either of the following conditions are true:

- A log file reaches its maximum age.

- The message tracking log folder reaches its maximum size.

  Notes:

  - The maximum size of the message tracking log folder is calculated as the total size of all log files that have the same name prefix. Other files that do not follow the name prefix convention are not counted in the total folder size calculation. Renaming old log files or copying other files into the message tracking log folder could cause the folder to exceed its specified maximum size.

  - On Mailbox servers, the maximum size of the message tracking log folder is three times the specified value. Although the message tracking log files are generated by the four different services and have four different name prefixes, the amount and frequency of data written to the moderated transport log (`MSGTRKMA`) is negligible compared to the other three logs.

The message tracking log files are text files that contain data in the comma-separated value (CSV) format. Each message tracking log file has a header that contains the following information:

- **#Software**: The value is `Microsoft Exchange Server`.

- **#Version**: Version number of the Exchange server that created the message tracking log file. The value uses the format `15.01.nnnn.nnn`.

- **#Log-Type**: The value is `Message Tracking Log`.

- **#Date**: The UTC date-time when the log file was created. The UTC date-time is represented in the ISO 8601 date-time format: *yyyy-mm-dd*T*hh:mm:ss.fff*Z, where *yyyy* = year, *mm* = month, *dd* = day, T indicates the beginning of the time component, *hh* = hour, *mm* = minute, *ss* = second, *fff* = fractions of a second, and Z signifies Zulu, which is another way to denote UTC.

- **#Fields**: Comma-delimited field names that are used in the message tracking log files.

Fields in the message tracking log files

The message tracking log stores each message event on a single line in the log. The message event information is organized by fields, and these fields are separated by commas. The field name is generally descriptive enough to determine the type of information that it contains. However, some fields may be blank, or the type of information in the field may change based on the message event type and the service that recorded the event. General descriptions of the fields that are used to classify each message tracking event are explained in the following table.

FIELD NAME	DESCRIPTION

FIELD NAME	DESCRIPTION
date-time	The UTC date-time of the message tracking event. The UTC date-time is represented in the ISO 8601 date-time format: *yyyy-mm-dd*T*hh:mm:ss.fff*Z, where *yyyy* = year, *mm* = month, *dd* = day, T indicates the beginning of the time component, *hh* = hour, *mm* = minute, *ss* = second, *fff* = fractions of a second, and Z signifies Zulu, which is another way to denote UTC.
client-ip	The IPv4 or IPv6 address of the messaging server or messaging client that submitted the message.
client-hostname	The host name or FQDN of the messaging server or messaging client that submitted the message.
server-ip	The IPv4 or IPv6 address of the source or destination server.
server-hostname	The host name or FQDN of the destination server.

source-context	Extra information associated with the **source** field. For example:
	`CatContentConversion`
	`250 2.0.0 OK;ClientSubmitTime:<UTC>`

connector-id	The name of the Send connector or Receive connector that accepted the message. For example, *ServerName*\ConnectorName or *ConnectorName*.
source	The Exchange transport component that's responsible for the event. These values are described in the Source values in the message tracking log section later in this topic.
event-id	The message event type. These values are described in the Event types in the message tracking log section later in this topic.
internal-message-id	A message identifier that's assigned by the Exchange server that's currently processing the message. The **internal-message-id** of a message is different in the message tracking log of every Exchange server that's involved in the transmission of the message. An example value is `73014444033`.
message-id	The value of the **Message-Id:** header field in the message header. If the **Message-Id:** header field doesn't exist or is blank, Exchange assigns an arbitrary value. This value is constant for the lifetime of the message. For messages created in Exchange, the value is in the format `<GUID@ServerFQDN>`, including the angle brackets (`< >`). For example, `<4867a3d78a50438bad95c0f6d072fca5@mailbox01.contoso.com>`. Other messaging systems may use different syntax or values.
network-message-id	A unique message ID value that persists across copies of the message that may be created due to bifurcation or distribution group expansion. An example value is `1341ac7b13fb42ab4d4408cf7f55890f`.

FIELD NAME	DESCRIPTION
recipient-address	The email addresses of the message's recipients. Multiple email addresses are separated by the semicolon character (;).

recipient-status	The recipient status for each recipient separated by the semicolon character (;). The status values are presented for the recipients in the same order as the values in the **recipient-address** field. Example status values include:
	`To`, `Cc` or `Bcc`
	`250 2.1.5 Recipient OK`
	`550 4.4.7 QUEUE.Expired;<ErrorText>`

total-bytes	The total size of the message in bytes, including all attachments.
recipient-count	The total number of recipients in the message.
related-recipient-address	This field is used with **EXPAND**, **REDIRECT**, and **RESOLVE** events to display other recipient email addresses that are associated with the message.
reference	This field contains additional information for specific types of events. For example: **DSN**: Contains the report link, which is the **Message-Id** value of the associated delivery status notification (also known as a DSN, bounce message, non-delivery report, or NDR) if a DSN is generated subsequent to this event. If this is a DSN message, the **Reference** field contains the **Message-Id** value of the original message that the DSN was generated for. **EXPAND**: Contains the **related-recipient-address** value of the related messages. **RECEIVE**: May contain the **Message-Id** value of the related message if the message was generated by other processes, for example, journaling or inbox rules. **SEND**: Contains the **Internal-Message-Id** value of any DSN messages. **THROTTLE**: Contains the reason why the message was throttled. **TRANSFER**: Contains the **Internal-Message-Id** value of the message that's being forked. **Message generated by inbox rules**: Contains the **Internal-Message-Id** value of the inbound message that caused the inbox rule to generate the outbound message. **Forked messages**: Might contain the **Internal-Message-Id** value. For other types of events, this field is usually blank.
message-subject	The message's subject found in the **Subject:** header field. The tracking of message subjects is controlled by the *MessageTrackingLogSubjectLoggingEnabled* parameter on the **Set-TransportService** cmdlet. By default, message subject tracking is enabled.
sender-address	The email address specified in the **Sender:** header field, or the **From:** header field if the **Sender:** field doesn't exist.

FIELD NAME	DESCRIPTION
return-path	The return email address specified by the **MAIL FROM** command that sent the message. Although this field is never empty, it can have the null sender address value represented as `<>`.
message-info	Additional information about the message. For example: The message origination date-time in UTC for **DELIVER** and **SEND** events. The origination date-time is the time when the message first entered the Exchange organization. The UTC date-time is represented in the ISO 8601 date-time format: *yyyy-mm-dd*T*hh:mm:ss.fff*Z, where *yyyy* = year, *mm* = month, *dd* = day, T indicates the beginning of the time component, *hh* = hour, *mm* = minute, *ss* = second, *fff* = fractions of a second, and Z signifies Zulu, which is another way to denote UTC. Authentication errors. For example, you may see the value `11a` and the type of authentication that was used when the authentication error occurred.
directionality	The direction of the message. Example values include `Incoming`, `Undefined`, and `Originating`.
tenant-id	This field isn't used in on-premises Exchange organizations.
original-client-ip	The IPv4 or IPv6 address of the original client.
original-server-ip	The IPv4 or IPv6 address of the original server.
custom-data	This field contains data related to specific event types. For example, the Transport Rule agent uses this field to record the GUID of the mail flow rule (also known as a transport rule) or DLP policy that acted on the message. For more information, see View DLP policy detection reports.
transport-traffic-type	In on-premises Exchange, this field is blank or has the value `Email`.
log-id	A unique identifier for a row in the in the message tracking log. This field isn't important in on-premises Exchange organizations.
schema-version	Version number of the Exchange server that created the entry in the message tracking log. The value uses the format `15.01.nnnn.nnn`.

Event types in the message tracking log

Various event types in the **event-id** field are used to classify the message events in the message tracking log. Some message events appear in only one type of message tracking log file, and some message events appear in all types of message tracking log files. The events types that are used to classify each message event are explained in the following table.

EVENT NAME	DESCRIPTION
AGENTINFO	This event is used by transport agents to log custom data.
BADMAIL	A message submitted by the Pickup directory or the Replay directory that can't be delivered or returned.
CLIENTSUBMISSION	A message was submitted from the Outbox of a mailbox.
DEFER	Message delivery was delayed.
DELIVER	A message was delivered to a local mailbox.
DELIVERFAIL	An agent tried to deliver the message to a folder that doesn't exist in the mailbox.
DROP	A message was dropped without a delivery status notification (also known as a DSN, bounce message, non-delivery report, or NDR). For example: • Completed moderation approval request messages. • Spam messages that were silently dropped without an NDR.
DSN	A delivery status notification (DSN) was generated.
DUPLICATEDELIVER	A duplicate message was delivered to the recipient. Duplication may occur if a recipient is a member of multiple nested distribution groups. Duplicate messages are detected and removed by the information store.
DUPLICATEEXPAND	During the expansion of the distribution group, a duplicate recipient was detected.
DUPLICATEREDIRECT	An alternate recipient for the message was already a recipient.
EXPAND	A distribution group was expanded.
FAIL	Message delivery failed. Sources include **SMTP**, **DNS**, **QUEUE**, and **ROUTING**.
HADISCARD	A shadow message was discarded after the primary copy was delivered to the next hop. For more information, see Shadow redundancy in Exchange Server.
HARECEIVE	A shadow message was received by the server in the local database availability group (DAG) or Active Directory site.
HAREDIRECT	A shadow message was created.
HAREDIRECTFAIL	A shadow message failed to be created. The details are stored in the **source-context** field.
INITMESSAGECREATED	A message was sent to a moderated recipient, so the message was sent to the arbitration mailbox for approval. For more information, see Manage message approval.

EVENT NAME	DESCRIPTION
LOAD	A message was successfully loaded at boot.
MODERATIONEXPIRE	A moderator for a moderated recipient never approved or rejected the message, so the message expired. For more information about moderated recipients, see Manage message approval.
MODERATORAPPROVE	A moderator for a moderated recipient approved the message, so the message was delivered to the moderated recipient.
MODERATORREJECT	A moderator for a moderated recipient rejected the message, so the message wasn't delivered to the moderated recipient.
MODERATORSALLNDR	All approval requests sent to all moderators of a moderated recipient were undeliverable, and resulted in non-delivery reports (also known as NDRs or bounce messages).
NOTIFYMAPI	A message was detected in the Outbox of a mailbox on the local server.
NOTIFYSHADOW	A message was detected in the Outbox of a mailbox on the local server, and a shadow copy of the message needs to be created.
POISONMESSAGE	A message was put in the poison message queue or removed from the poison message queue.
PROCESS	The message was successfully processed.
PROCESSMEETINGMESSAGE	A meeting message was processed by the Mailbox Transport Delivery service.
RECEIVE	A message was received by the SMTP receive component of the transport service or from the Pickup or Replay directories (source: `SMTP`), or a message was submitted from a mailbox to the Mailbox Transport Submission service (source: `STOREDRIVER`).
REDIRECT	A message was redirected to an alternative recipient after an Active Directory lookup.
RESOLVE	A message's recipients were resolved to a different email address after an Active Directory lookup.
RESUBMIT	A message was automatically resubmitted from Safety Net. For more information, see Safety Net in Exchange Server.
RESUBMITDEFER	A message resubmitted from Safety Net was deferred.
RESUBMITFAIL	A message resubmitted from Safety Net failed.
SEND	A message was sent by SMTP between transport services.

EVENT NAME	DESCRIPTION
SUBMIT	The Mailbox Transport Submission service successfully transmitted the message to the Transport service. For **SUBMIT** events, the **source-context** property contains the following details: **MDB**: The mailbox database GUID. **Mailbox**: The mailbox GUID. **Event**: The event sequence number. **MessageClass**: The type of message. For example, `IPM.Note`. **CreationTime**: Date-time of the message submission. **ClientType**: For example, `User`, `OWA`, or `ActiveSync`.
SUBMITDEFER	The message transmission from the Mailbox Transport Submission service to the Transport service was deferred.
SUBMITFAIL	The message transmission from the Mailbox Transport Submission service to the Transport service failed.
SUPPRESSED	The message transmission was suppressed.
THROTTLE	The message was throttled.
TRANSFER	Recipients were moved to a forked message because of content conversion, message recipient limits, or agents. Sources include **ROUTING** or **QUEUE**.

Source values in the message tracking log

The values in the **source** field in the message tracking log indicate the transport component that's responsible for the message tracking event. The following table describes the values of the **source** field.

SOURCE VALUE	DESCRIPTION
ADMIN	The event source was human intervention. For example, an administrator used Queue Viewer to delete a message, or submitted message files using the Replay directory.
AGENT	The event source was a transport agent.
APPROVAL	The event source was the approval framework that's used with moderated recipients. For more information, see Manage message approval.
BOOTLOADER	The event source was unprocessed messages that exist on the server at boot time. This is related to the **LOAD** event type.
DNS	The event source was DNS.
DSN	The event source was a delivery status notification (also known as a DSN, bounce message, non-delivery report, or NDR).
GATEWAY	The event source was a Foreign connector. For more information, see Foreign Connectors.

SOURCE VALUE	DESCRIPTION
MAILBOXRULE	The event source was an Inbox rule. For more information, see Inbox rules.
MEETINGMESSAGEPROCESSOR	The event source was the meeting message processor, which updates calendars based on meeting updates.
ORAR	The event source was an Originator Requested Alternate Recipient (ORAR). You can enable or disable support for ORAR on Receive connectors using the *OrarEnabled* parameter on the **New-ReceiveConnector** or **Set-ReceiveConnector** cmdlets.
PICKUP	The event source was the Pickup directory. For more information, see Pickup Directory and Replay Directory.
POISONMESSAGE	The event source was the poison message identifier. For more information about poison messages and the poison message queue, see Queues and messages in queues
PUBLICFOLDER	The event source was a mail-enabled public folder.
QUEUE	The event source was a queue.
REDUNDANCY	The event source was Shadow Redundancy. For more information, see Shadow redundancy in Exchange Server.
ROUTING	The event source was the routing resolution component of the categorizer in the Transport service.
SAFETYNET	The event source was Safety Net. For more information, see Safety Net in Exchange Server.
SMTP	The message was submitted by the SMTP send or SMTP receive component of the transport service.
STOREDRIVER	The event source was a MAPI submission from a mailbox on the local server.

Example entries in the message tracking log

An uneventful message sent between two users generates several entries in the message tracking log. You can see the results using the **Get-MessageTrackingLog** cmdlet. For more information, see Search message tracking logs.

This is an example of the message tracking log entries created when the user chris@contoso.com successfully sends a test message to the user michelle@contoso.com. Both users have mailboxes on the same server.

```
EventId    Source      Sender            Recipients             MessageSubject
-------    ------      ------            ----------             --------------
NOTIFYMAPI STOREDRIVER                   {}
RECEIVE    STOREDRIVER chris@contoso.com {michelle@contoso.com} test
SUBMIT     STOREDRIVER chris@contoso.com {michelle@contoso.com} test
HAREDIRECT SMTP        chris@contoso.com {michelle@contoso.com} test
RECEIVE    SMTP        chris@contoso.com {michelle@contoso.com} test
AGENTINFO  AGENT       chris@contoso.com {michelle@contoso.com} test
SEND       SMTP        chris@contoso.com {michelle@contoso.com} test
DELIVER    STOREDRIVER chris@contoso.com {michelle@contoso.com} test
```

Security concerns for the message tracking log

No message content is stored in the message tracking log. By default, the subject line of an email message is stored in the message tracking log. You might need to disable subject logging to comply with increased security or privacy requirements. For instructions on how to disable subject logging, see Configure message tracking.

Configure message tracking

8/3/2020 • 4 minutes to read • Edit Online

Message tracking records the message activity as mail flows through the transport pipeline on Mailbox servers and Edge Transport servers. You can use message tracking logs for message forensics, mail flow analysis, reporting, and troubleshooting.

You use the **Set-TransportService** cmdlet in the Exchange Management Shell on Mailbox servers and Edge Transport servers for all message tracking configuration tasks. For example:

- Enable or disable message tracking. The default is enabled.
- Specify the location of the message tracking log files. The default location is `%ExchangeInstallPath%TransportRoles\Logs\MessageTracking`.
- Specify a maximum size for the individual message tracking log files. The default is 10 MB.
- Specify a maximum size for the directory that contains the message tracking log files: The default is 1000 MB.
- Specify maximum age for the message tracking log files: The default is 30 days.
- Enable or disable message subject logging in the message tracking logs. The default is enabled.

> **NOTE**
>
> On Mailbox servers, you can also use the Exchange admin center (EAC) to enable or disable message tracking, and to specify the location of the message tracking log files.

What do you need to know before you begin?

- Estimated time to complete: 5 minutes
- To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Transport service" entries in the Mail flow permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to configure message tracking on Mailbox servers

1. Open the EAC and navigate to **Servers** > **Servers** > select the Mailbox server that you want to configure > and click **Edit** ✎.
2. On the server properties page, click **Transport Logs**. In the **Message tracking log** section, change any of

the following settings:

- **Enable message tracking log**: To disable message tracking on the server, clear the check box. To enable message tracking on the server, select the check box.

- **Message tracking log path**: The value you specify must be on the local Exchange server. If the folder doesn't exist, it's created for you when you click **Save**.

3. When you're finished, click **Save**.

Use the Exchange Management Shell to configure message tracking

As previously explained, you can use the **Set-TransportService** cmdlet to perform all message tracking configuration tasks on Mailbox servers and Edge Transport servers. To configure message tracking in the Exchange Management Shell, use the following syntax:

```
Set-TransportService [<ServerIdentity>] -MessageTrackingLogEnabled <$true | $false> -MessageTrackingLogMaxAge
<dd.hh:mm:ss> -MessageTrackingLogMaxDirectorySize <Size> -MessageTrackingLogMaxFileSize <Size> -
MessageTrackingLogPath <LocalFilePath> -MessageTrackingLogSubjectLoggingEnabled <$true | $false>
```

Note that you don't need to specify the Exchange server when you run the command on the server that you want to configure.

This example configures the following message tracking log settings on the server named Mailbox01:

- Sets the location of the message tracking log files to D:\Message Tracking Log. Note that if the folder doesn't exist, it's created for you.

- Sets the maximum size of a message tracking log file to 20 MB.

- Sets the maximum size of the message tracking log directory to 1.5 GB.

- Sets the maximum age of a message tracking log file to 45 days.

```
Set-TransportService Mailbox01 -MessageTrackingLogPath "D:\Message Tracking Log" -
MessageTrackingLogMaxFileSize 20MB -MessageTrackingLogMaxDirectorySize 1.5GB -MessageTrackingLogMaxAge
45.00:00:00
```

NOTE

• Setting the *MessageTrackingLogPath* parameter to the value `$null`, effectively disables message tracking. However, if the value of the *MessageTrackingLogEnabled* parameter is `$true`, event log errors are generated.

• Setting the *MessageTrackingLogMaxAge* parameter to the value `00:00:00` prevents the automatic removal of message tracking log files because of their age.

• The maximum size of the message tracking log directory is three times the value of the *MessageTrackingLogMaxDirectorySize* parameter. Although the message tracking log files that are generated by the four different services have four different name prefixes, the amount and frequency of data written to the moderated transport log (**MSGTRKMA**) is negligible compared to the other three logs. For more information, see Structure of the message tracking log files.

This example disables message subject logging in the message tracking log on the server named Mailbox01:

```
Set-TransportService Mailbox01 -MessageTrackingLogSubjectLoggingEnabled $false
```

This example disables message tracking on the Mailbox server named Mailbox01:

```
Set-TransportService Mailbox01 -MessageTrackingLogEnabled $false
```

How do you know this worked?

To verify that you have successfully configured message tracking, run the following command in the Exchange Management Shell:

```
Get-TransportService [<ServerIdentity>] | Format-List MessageTrackingLog*
```

You can also open the location of the message tracking log in Windows Explorer or File Explorer to verify that the log files exist, that data is being written to the files, and that they're being recycled based on the maximum file size and maximum directory size values that you configured.

Search message tracking logs

8/3/2020 • 4 minutes to read • Edit Online

Message tracking records the message activity as mail flows through the transport pipeline on Mailbox servers and Edge Transport servers. You can use the **Get-MessageTrackingLog** cmdlet in the Exchange Management Shell to search for entries in the message tracking log by using specific search criteria. For example:

- Find out what happened to a message that was sent by a user to a specific recipient.
- Find out if a mail flow rule (also known as a transport rule) acted on a message.
- Find out if a message sent from an Internet sender made it into your Exchange organization.
- Find all messages sent by a specified user during a specified time period.

What do you need to know before you begin?

- Estimated time to complete: 10 minutes
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Message tracking" entry in the Mail flow permissions topic.
- Searching the message tracking logs requires that the Microsoft Exchange Transport Log Search service is running. If you disable or stop this service, you can't search the message tracking logs or run delivery reports. However, stopping this service does not affect other features in Exchange.
- The field names displayed in the results from the **Get-MessageTrackingLog** cmdlet are similar to the actual field names found in the message tracking log files. The biggest differences are:
  - Dashes are removed from the field names. For example, **internal-message-id** is displayed as `InternalMessageId`.
  - The **date-time** field is displayed as `Timestamp`.
  - The **recipient-address** field is displayed as `Recipients`.
  - The **sender-address** field is displayed as `Sender`.
- The **date-time** field in the message tracking log stores information in Coordinated Universal Time (UTC). However, you need to enter your date-time search criteria for the *Start* or *End* parameters in the regional date-time format of the computer that you're using to perform the search.
- You can't copy the message tracking log files from another Exchange server and then search them by using the **Get-MessageTrackingLog** cmdlet. Also, if you manually save an existing message tracking log file, the change in the file's date-time stamp breaks the query logic that Exchange uses to search the message tracking logs.
- In Exchange 2016, the **Get-MessageTrackingLog** cmdlet is able to search the message tracking logs on Exchange 2013 Mailbox servers and Exchange 2010 Hub Transport servers in the same Active Directory site. In Exchange 2019, the **Get-MessageTrackingLog** cmdlet is able to search the message tracking logs on Exchange 2016 and Exchange 2013 Mailbox servers in the same Active Directory site.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to search the message tracking logs

To search the message tracking log entries for specific events, use the following syntax.

```
Get-MessageTrackingLog [-Server <ServerIdentity>] [-ResultSize <Integer> | Unlimited] [-Start <DateTime>] [-
End <DateTime>] [-EventId <EventId>] [-InternalMessageId <InternalMessageId>] [-MessageId <MessageId>] [-
MessageSubject <Subject>] [-Recipients <RecipientAddress1,RecipientAddress2...>] [-Reference <Reference>] [-
Sender <SenderAddress>]
```

To view the 1000 most recent message tracking log entries on the server, run the following command:

```
Get-MessageTrackingLog
```

This example searches the message tracking logs on the local server for all entries from 3/28/2015 8:00 AM to 3/28/2015 5:00 PM for all **FAIL** events where the message sender was pat@contoso.com.

```
Get-MessageTrackingLog -ResultSize Unlimited -Start "3/28/2015 8:00AM" -End "3/28/2015 5:00PM" -EventId "Fail"
-Sender "pat@contoso.com"
```

Use the Exchange Management Shell to control the output of a message tracking log search

Use the following syntax.

```
Get-MessageTrackingLog <SearchFilters> | <Format-Table | Format-List> [<FieldNames>] [<OutputFileOptions>]
```

This example searches the message tracking logs using the following search criteria:

- Return results for the first 1,000 **Send** events.
- Display the results in the list format.
- Display only those field names that begin with `Send` or `Recipient`.
- Write the output to a new file named `D:\Send Search.txt`

```
Get-MessageTrackingLog -EventId Send | Format-List Send*,Recipient* | Set-Content -Path "D:\Send Search.txt"
```

Use the Exchange Management Shell to search the message tracking logs for message entries on multiple servers

Typically, the value in the **MessageID:** header field remains constant as the message travels throughout the Exchange organization. This property is named **InternetMessageId** in queue viewing utilities, and **MessageId** in the message tracking log viewing utilities. After you have determined the **MessageID:** value of a specific message, you can search for information about that message in the message tracking logs on every Mailbox server in your Exchange organization.

To search all message tracking log entries for a specific message across all Mailbox servers and Exchange 2010 Hub Transport servers, use the following syntax.

```
$Servers = Get-ExchangeServer;  $Servers | where {$_.isHubTransportServer -eq $true -or $_.isMailboxServer -eq
$true} | Get-MessageTrackingLog -MessageId <MessageID>  | Select-Object <CommaSeparatedFieldNames>  | Sort-
Object -Property <FieldName>
```

This example searches the message tracking logs on all Mailbox servers and Exchange 2010 Hub Transport server by using the following search criteria:

- Find any entries related to a message that has a **MessageID:** value of `<ba18339e-8151-4ff3-aeea-87ccf5fc9796@mailbox01.contoso.com>`. Note that you can omit the angle bracket characters (`<` `>`). If you don't, you need to enclose the entire **MessageID:** value in quotation marks.
- For each entry, display the fields **date-time**, **server-hostname**, **client-hostname**, **source**, **event-id**, and **recipient-address**.
- Sort the results by the **date-time** field.

```
$Servers = Get-ExchangeServer; $Servers | where {$_.isHubTransportServer -eq $true -or $_.isMailboxServer -eq
$true} | Get-MessageTrackingLog -MessageId ba18339e-8151-4ff3-aeea-87ccf5fc9796@mailbox01.contoso.com |
Select-Object Timestamp,ServerHostname,ClientHostname,Source,EventId,Recipients | Sort-Object -Property
Timestamp
```

Use the EAC to search the message tracking logs

You can use the Delivery Reports for administrators feature in the Exchange admin center (EAC) to search the message tracking logs for information about messages sent by or received by a specific mailbox in your organization. For more information, see Track messages with delivery reports.

Delivery reports for administrators

8/3/2020 • 2 minutes to read • Edit Online

With delivery reports for administrators, you can track delivery information about messages sent by or received from any specific mailbox in your organization. Specifically, delivery reports for administrators uses the Exchange admin center (EAC) to perform a targeted search of the message tracking logs. The search is always scoped to a specific mailbox. You can search for messages sent by the mailbox, or sent to the mailbox, and you can filter the search results by the message subject.

The content of the message body isn't returned in a delivery report, but the subject line is displayed in the results. If you want to search the mailboxes in your organization for specific email messages based on message content, see In-Place eDiscovery in Exchange Server.

You may find delivery report searches useful in the following situations:

- A manager gives a poor review for a trainee because the trainee didn't turn in an assignment on time. The trainee insists he sent a message with the assignment attached. The manager asks you to verify the status of the message.
- A security bulletin has been sent to users asking that they reply immediately, but no one has replied. Are they ignoring the message or did they just not receive it?
- Users complain that no one is receiving their messages. They check delivery status for their mail but can't figure out what is going on. This may be because a rule is being applied to messages at the organization level.

After you create a delivery report search, the resulting delivery report will show the following information: Who the message was sent from and to, the subject line, and when the message was sent. The delivery report also shows message delivery status and reasons why delivery may be delayed or failed.

More about delivery reports

- Here's how administrators in on-premises Exchange organizations create delivery reports: Track messages with delivery reports.
- A more powerful option for administrators in on-premises Exchange organizations is to use the Exchange Management Shell to query the message tracking logs directly. For more information, see Search message tracking logs.
- Exchange 2016 or Exchange 2019 delivery reports can track messages across Exchange 2019, Exchange 2016, and Exchange 2013 servers in the same Active Directory site.

Track messages with delivery reports

8/3/2020 • 4 minutes to read • Edit Online

Delivery Reports is a message tracking tool in the Exchange admin center (EAC) that you can use to search for delivery status on email messages that were sent to or from users in your organization's address book. You can track delivery information about messages sent by or received from any specific mailbox in your organization. The message's content isn't returned in the delivery report, but the subject line is displayed in the results. You can track messages for up to 14 days after they were sent or received.

> **NOTE**
>
> Delivery Reports tracks messages that were sent by people using Microsoft Outlook or Outlook on the web. It doesn't track messages sent from POP3 or IMAP4 email clients, such as Windows Live Mail or Mozilla Thunderbird.

What do you need to know before you begin?

- Estimated time to complete each procedure: Time to complete will vary based on the scope of your search.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Message tracking" entry in the Mail flow permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.
- Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to track messages

1. In the EAC, navigate to **Mail Flow** > **Delivery Reports**.

2. Enter the following information:

   - **Mailbox to search**: Click **Browse** to select the mailbox from the address book and then click **OK**. Selecting the mailbox to search is required.

   - Select one of the following:

   - **Search for messages sent to**: Use this option to search for messages that were sent to specific users from the mailbox you selected in **Mailbox to search**. Click **Select users** and then pick users from the address book by selecting a user from the list and clicking **Add**. You can select more than one user here. When you're finished selecting users, click **OK** to return to the **Delivery Reports** page. If you select this option, you can also leave the field blank to find messages sent to anyone.

   - **Search for messages received from**: Use this option to search for messages that were sent by specific user to the mailbox you selected in **Mailbox to search**. Again, select the user from the address book and click **OK** to return to the **Delivery Reports** page. If you select this option, you have to specify a sender.

   - **Search for these words in the subject line**: Enter subject line information here, or leave it blank.

3. When you're finished, click **Search**. If you want to start over, click **Clear**.

Use the EAC to review a delivery report

To view delivery information, select a message in the **Search results** pane and click **Details** ✎.

The delivery report shows delivery status and detailed delivery information for the message you have selected from the **Search results** pane. At the top of the report, you'll see the following fields:

- **Subject**: The subject line of the message appears as the heading of the report.
- **From**: Alias, display name, or email address of the person who sent the message.
- **To**: Alias, display name, or email address for each recipient of the message.
- **Sent**: Date and time the message was sent.

Summary to date section

This section appears in the delivery report if a message was sent to more than one recipient. The top of this section tells you the total number of recipients that the message was sent to and gives brief delivery information for each recipient.

- **Summary to date**: Displays total number of recipients, and if there are messages **Pending**, **Delivered**, or **Unsuccessful**. Click the hyperlinks to sort by status.
- **Search box**: The search box is useful if you sent the message to a group of more than 30 recipients. In the search box, type an email address that you want to get delivery information about and click the magnifying glass 🔍.
- **To**: Shows the email address of the recipient.
- **Status**: This column displays the status of the message for each recipient.

Detailed report information

This section contains detailed delivery information for a message sent to the recipient you select in the **Summary to date** section.

- **Delivery Report for**: The email address of the selected recipient is shown here.
- **Submitted**: Date and time that the message was submitted for delivery by the system.

Depending on the delivery status of the message, you may see a variety of status states, including:

- **Delivered**: Indicates successful delivery.
- **Deferred**: Indicates that a message is delayed.
- **Pending**: If message delivery is pending because a message meets the criteria for an organization-wide rule or policy or because it's subject to message approval, the status message explains what action a rule is performing or that the message must be approved by a moderator before delivery.
- **Moderator**: The status indicates whether the message was approved or rejected by the moderator.
- **Groups Expanded**: If a message was sent to a group, the individual users are shown in the **Summary to date** section so you can see the delivery status for each recipient. If you need to remove or add a user to a group during a delivery report investigation, you can modify a group by clicking **Edit Groups**.
- **Failed**: Shows the date, time, and reason for a message delivery failure. For example, an organization-wide rule may be blocking message delivery or the message couldn't be delivered.

When you're done reviewing the report, click **Close**. Delivery reports aren't saved, but you can re-run a report at any time. Remember there is a two-week search window.

How do you know this worked?

If your search was successful, messages that fit the search criteria are listed in the **Search results** pane. To view the delivery information for a specific message, select it and then click **Details** ✎. If no messages are displayed in the **Search results** pane, change the search criteria and then re-run the search.

Connectivity logging in Exchange Server

8/3/2020 • 4 minutes to read • Edit Online

Connectivity logging records the outbound connection activity that's used to transmit messages on Exchange servers. In Exchange Server, the following services transmit messages, so they have connectivity logs:

- The Transport service on Mailbox servers and Edge Transport servers.
- The Front End Transport service on Mailbox servers.
- The Mailbox Transport Submission service on Mailbox servers.
- The Mailbox Transport Delivery service on Mailbox servers.

For more information about these transport services, and where they can transmit messages, see Mail flow and the transport pipeline.

Connectivity logging doesn't track the transmission of individual messages. Instead, it tracks the number and size of messages that were transmitted over a connection, DNS resolution information for the destination, and informational messages that are related to the connection.

By default, connectivity logging is enabled, and Exchange uses circular logging to limit the connectivity log files based on size and age to help control the hard disk space that's used. To configure connectivity logging, see Configure connectivity logging in Exchange Server.

Note: If you're interested in a detailed record of the entire SMTP protocol conversation from start to finish, see Protocol logging.

Structure of the connectivity log files

By default, the connectivity log files exist in these locations:

- **Mailbox servers**:
  - **Transport service**: `%ExchangeInstallPath%TransportRoles\Logs\Hub\Connectivity`
  - **Front End Transport service**: `%ExchangeInstallPath%TransportRoles\Logs\FrontEnd\Connectivity`
  - **Mailbox Transport Delivery service**: `%ExchangeInstallPath%TransportRoles\Logs\Mailbox\Connectivity\Delivery`
  - **Mailbox Transport Submission service**: `%ExchangeInstallPath%TransportRoles\Logs\Mailbox\Connectivity\Submission`
- **Edge Transport servers**

The naming convention for the connectivity log files is `CONNECTLOGyyyymmdd-nnnn.log`. The placeholders represent the following information:

- *yyyymmdd* is the Coordinated Universal Time (UTC) when the log file was created. *yyyy* = year, *mm* = month, and *dd* = day.
- *nnnn* is an instance number that starts at the value of 1 for each day.

Information is written to the log file until the file reaches its maximum size. Then, a new log file that has an incremented instance number is opened (the first log file is -1, the next is -2, and so on). Circular logging deletes the oldest log files when either of the following conditions are true:

- A log file reaches its maximum age.
- The connectivity log folder reaches its maximum size.

The connectivity log files are text files that contain data in the comma-separated value file (CSV) format. Each connectivity log file has a header that contains the following information:

- **#Software**: The value is `Microsoft Exchange Server`.
- **#Version**: The value is `15.0.0.0`.
- **#Log-Type**: The value is `Transport Connectivity Log`.
- **#Date**: The UTC date-time when the log file was created. The UTC date-time is represented in the ISO 8601 date-time format: *yyyy-mm-dd*T*hh:mm:ss.fff*Z, where *yyyy* = year, *mm* = month, *dd* = day, T indicates the beginning of the time component, *hh* = hour, *mm* = minute, *ss* = second, *fff* = fractions of a second, and Z signifies Zulu, which is another way to denote UTC.
- **#Fields**: Comma delimited field names that are used in the connectivity log files. These values are described in the next section.

Fields in the connectivity log files

Connectivity logging stores each outbound connection event on a single line in the log. The information on each line is organized by fields, and these fields are separated by commas. The following table describes the fields that are used to classify each outgoing connection event.

FIELD NAME	DESCRIPTION
date-time	UTC date-time of the connection event. The UTC date-time is represented in the ISO 8601 date-time format: *yyyy-mm-dd*T*hh:mm:ss.fff*Z, where *yyyy* = year, *mm* = month, *dd* = day, T indicates the beginning of the time component, *hh* = hour, *mm* = minute, *ss* = second, *fff* = fractions of a second, and Z signifies Zulu, which is another way to denote UTC.
session	A GUID value. The value is the same for every event that's associated with the session, but different for each session.
source	One of these values: **SMTP** for SMTP connections. **MapiDelivery** for connections from the local mailbox database by the Mailbox Transport Delivery service. **MapiSubmission** for connections from the local mailbox database by the Mailbox Transport Submission service.

FIELD NAME	DESCRIPTION
destination	These are some examples of values you'll see here: **In the Transport service**: • The FQDN of the destination messaging server • `shadowredundancy` (on Mailbox servers only) **In the Front End Transport service**: • `internalproxy` • `client proxy` **In the Mailbox Transport Delivery service**: The GUID of the destination mailbox database. **In the Mailbox Transport Submission service**: • The GUID of the destination mailbox database. • `mailboxtransportsubmissioninternalproxy`
direction	Single character that represents the start, middle, or end of the connection. The values you'll see here are: `+`: Connect `-`: Disconnect `>`: Send
description	Text information that's associated with the connection event. For example: Number and size of messages that were transmitted. DNS MX resource record resolution information for destination domains. DNS resolution information for destination Mailbox servers. Connection establishment messages. Connection failure messages.

The transport services connect to and transmit messages to multiple destinations simultaneously. Entries in the log file from different connection events are interlaced (they typically aren't grouped together as one uninterrupted series of connection events). However you can use the fields (in particular, the unique **session** field value for a connection) to organize and arrange the log entries for each separate connection from start to finish.

Configure connectivity logging in Exchange Server

8/3/2020 • 4 minutes to read • Edit Online

Connectivity logging records outbound connection activity (source, destination, number and size of messages, and connection information) for the transport services on Exchange servers. For more information about connectivity logging, see Connectivity logging in Exchange Server.

What do you need to know before you begin?

- Estimated time to complete: 15 minutes
- You can use the Exchange admin center (EAC) to enable or disable connectivity logging and set the log path for the Transport service on Mailbox servers only. For all other connectivity logging options in the other transport services, you need to use the Exchange Management Shell. For more information about the EAC, see Exchange admin center in Exchange Server. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- The folder for connectivity logging needs the following permissions:
  - Network Service: Full Control
  - System: Full Control
  - Administrators: Full Control

  If the folder doesn't exist, but the parent folder has these permissions, the new folder is created automatically.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Transport service", "Front End Transport service", and "Mailbox Transport service" entries in the Mail flow permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to configure connectivity logging in the Transport service on Mailbox servers

1. In the EAC, go to **Servers** > **Servers**.
2. Select the Mailbox server you want to configure, and then click **Edit** ✎.
3. On the server properties page that opens, click **Transport Logs**.
4. In the **Connectivity log** section, change any of these settings:
   - **Enable connectivity log**: To disable connectivity logging for the Transport service on the server, clear the check box. To enable connectivity logging for the Transport service on the server, select the

check box.

- **Connectivity log path**: The value you specify must be on the local Exchange server. If the folder doesn't exist, it will be created for if the parent folder has the required permissions.

When you're finished, click **Save**.

Use the Exchange Management Shell to configure connectivity logging

On Mailbox servers, connectivity logging is available on the following transport services:

- The Transport service (use the **Set-TransportService** cmdlet).
- The Front End Transport service (use the **Set-FrontEndTransportService** cmdlet).
- The Mailbox Transport Delivery and Mailbox Transport Submission services (use the **Set-MailboxTransportService** cmdlet to configure both).

On Edge Transport servers, connectivity logging is available on the Transport service (use the **Set-TransportService** cmdlet).

To configure connectivity logging, use the following syntax:

```
<Set-TransportService | Set-MailboxTransportService | Set-FrontEndTransportService> -Identity <ServerIdentity>
-ConnectivityLogEnabled <$true | $false> -ConnectivityLogMaxAge <dd.hh:mm:ss> -ConnectivityLogMaxDirectorySize
<Size> -ConnectivityLogMaxFileSize <Size> -ConnectivityLogPath <LocalFilePath>
```

This example sets the following connectivity log settings in the Transport service on the Mailbox server named Mailbox01:

- **Location of the connectivity log**: D:\Connectivity Log\Hub. Note that if the folder doesn't exist, it will be created for you if the parent folder has the required permissions.
- **Maximum size of a connectivity log file**: Sets the maximum size of a connectivity log file to 20 MB.
- **Maximum size of the connectivity log folder**: Sets the maximum size of the connectivity log directory to 1.5 GB.
- **Maximum age of a connectivity log file**: Sets the maximum age of a connectivity log file to 45 days.

```
Set-TransportService -Identity Mailbox01 -ConnectivityLogPath "D:\Connectivity Log\Hub" -
ConnectivityLogMaxFileSize 20MB -ConnectivityLogMaxDirectorySize 1.5GB -ConnectivityLogMaxAge 45.00:00:00
```

For detailed syntax and parameter information, see Set-TransportService, Set-FrontendTransportService, and Set-MailboxTransportService.

Notes:

- Setting the *ConnectivityLogPath* parameter to the value `$null`, effectively disables connectivity logging. However, this value generates event log errors if the value of the *ConnectivityLogEnabled* parameter is also `$true`.
- When you use the *ConnectivityLogPath* parameter on the **Set-MailboxTransportService** cmdlet, two subfolders are automatically created in the folder you specify:
  - `Delivery` for the Mailbox Transport Delivery service.
  - `Submission` for the Mailbox Transport Submission service.

- Setting the *ConnectivityLogMaxAge* parameter to the value `00:00:00` prevents the automatic removal of connectivity log files because of their age.

How do you know this worked?

To verify that you've successfully configured connectivity logging, use these steps:

1. Run the following command in the Exchange Management Shell to verify the connectivity log settings on the Exchange servers:

   ```
   Write-Host "Front End Transport service:" -ForegroundColor yellow; Get-FrontEndTransportService |
   Format-List Name,ConnectivityLog*; Write-Host "Mailbox Transport Submission and Mailbox Transport
   Delivery services:" -ForegroundColor yellow; Get-MailboxTransportService | Format-List
   Name,ConnectivityLog*; Write-Host "Transport service:" -ForegroundColor yellow; Get-TransportService |
   Format-List Name,ConnectivityLog*
   ```

2. Open the location of the connectivity log in Windows Explorer or File Explorer to verify that the log files exist, that data is being written to the files, and that the files are being recycled based on the maximum file size and maximum directory size values that you configured. If you disabled connectivity logging, verify that the log files aren't being updated.

Queues and messages in queues in Exchange Server

8/3/2020 • 26 minutes to read • Edit Online

A *queue* is a temporary holding location for messages that are waiting to enter the next stage of processing or delivery to a destination. Each queue represents a logical set of messages that the Exchange server processes in a specific order. In Exchange 2016 and Exchange 2019, queues hold messages before, during, and after delivery. Queues exist in the Transport service on Mailbox servers and on Edge Transport servers. Mailbox servers and Edge Transport servers are called *transport servers* throughout this topic.

Like all previous versions of Exchange, a single Extensible Storage Engine (ESE) database is used for queue storage.

You can manage queues and messages in queues by using the Exchange Management Shell and Queue Viewer in the Exchange Toolbox. You can use these interfaces to view the status and contents of queues and detailed message properties. You can also perform actions that modify queues or the messages in queues. For more information, see Procedures for queues and Procedures for messages in queues.

Types of queues

The following types of queues are used in Exchange 2016 and Exchage 2019, which are the same as Exchange 2013:

QUEUE	SERVER ROLE	DESCRIPTION
Delivery queues	Mailbox servers and Edge Transport servers	Holds messages that are being delivered to all internal and external destinations. Delivery queues are dynamically created when they're required, and are automatically deleted when the queue is empty and the expiration time has passed. The queue expiration time is controlled by the *QueueMaxIdleTime* parameter on the **Set-TransportService** cmdlet. The default value is three minutes. On Edge Transport servers, there's a queue for every unique destination SMTP domain or smart host. On Mailbox servers, there's a queue for every unique destination as indicated by the **NextHopSolutionKey** property. For more information, see the NextHopSolutionKey section later in this topic. All messages are transmitted between Exchange 2016 and Exchange 2013 servers by using SMTP. Non-SMTP destinations also use delivery queues if the destination is serviced by a Delivery Agent connector. For more information, see Delivery Agents and Delivery Agent Connectors.

QUEUE	SERVER ROLE	DESCRIPTION
Poison message queue	Mailbox servers and Edge Transport servers	Isolates messages that contain errors and are determined to be harmful to Exchange after a server or service failure. The messages may be genuinely harmful in their content and format, or the messages might have been the victims of a poorly written transport agent or a software bug that crashed the Exchange server while it was processing the otherwise valid messages. The poison message queue is typically empty. If the poison message queue contains no messages, then it doesn't appear in the queue management tools. Messages in the poison message queue are never automatically resumed or expired. Messages remain in the poison message queue until they're manually resumed or removed by an administrator. Every Mailbox server or Edge Transport server has only one poison message queue.
Shadow queues	Mailbox servers	Shadow queues hold redundant copies of messages while the messages are in transit. For more information, see Shadow redundancy in Exchange Server.
Submission queue	Mailbox servers and Edge Transport servers	Holds messages that have been accepted by the Transport service, but haven't been processed. Messages in the Submission queue are either waiting to be processed, or are actively being processed. On Mailbox servers, messages are received by a Receive connector, the Pickup or Replay directories, or the Mailbox Transport Submission service. On Edge Transport servers, messages are typically received by a Receive connector, but the Pickup and Replay directories are also available. The categorizer retrieves messages from this queue and, among other things, determines the location of the recipient and the route to that location. After categorization, the message is moved to a delivery queue or to the Unreachable queue. For more information about the categorizer and the transport pipeline, see Mail flow and the transport pipeline. Every Mailbox server or Edge Transport server has only one Submission queue.
Unreachable queue	Mailbox servers and Edge Transport servers	Contains messages that can't be routed to their destinations. Typically, an unreachable destination is caused by configuration changes that have modified the routing path for delivery. Regardless of destination, all messages that have unreachable recipients reside in this queue. Every Mailbox server or Edge Transport server has only one Unreachable queue.

Queue database files

All the different queues are stored in a single ESE database. By default, this queue database is located on the transport server at `%ExchangeInstallPath%TransportRoles\data\Queue`.

Like any ESE database, the queue database uses log files to accept, track, and maintain data. To enhance performance, all message transactions are written first to log files and memory, and then to the database file. The checkpoint file tracks the transaction log entries that have been committed to the database. During an ordinary shutdown of the Microsoft Exchange Transport service, uncommitted database changes that are found in the transaction logs are committed to the database.

Circular logging is used for the queue database. This means that transaction logs that are older than the current checkpoint are immediately and automatically deleted. Therefore, the transaction logs can't be replayed for queue database recovery from backup.

The following table lists the files that constitute the queue database.

FILE	DESCRIPTION
Mail.que	This queue database file stores all the queued messages.
Tmp.edb	This temporary database file is used to verify the queue database schema on startup.
Trn*.log	Transaction logs record all changes to the queue database. Changes to the database are first written to the transaction log and then committed to the database. Trn.log is the current active transaction log file. Trntmp.log is the next provisioned transaction log file that's created in advance. If the existing Trn.log transaction log file reaches its maximum size, Trn.log is renamed to Trn *nnnn*.log, where *nnnn* is a sequence number. Trntmp.log is then renamed Trn.log and becomes the current active transaction log file.
Trn.chk	This checkpoint file tracks the transaction log entries that have been committed to the database. This file is always in the same location as the mail.que file.
Trnres00001.jrs Trnres00002.jrs	These reserve transaction log files act as placeholders. They're only used when the hard disk that contains the transaction log runs out of space to stop the queue database cleanly.

Exchange uses *generation tables* for storage and clean-up of messages in the queue database. Instead of processing and deleting individual message records from one large table, the queue database stores messages in time-based tables, and only deletes the entire table after all the messages in the table have been successfully processed. For example, consider the following example:

- All messages queued from 1:00 PM to 2:00 PM, regardless of the queue or destination, are stored in the `1p-2p_msgs` table.
- At 2:00 PM, new messages are stored in the `2p-3p_msgs` table.
- At 4:00 PM, a new table named `4p-5p_msgs` is created. The entire `1p-2p_msgs` table is deleted, but only if all messages in the table have been successfully processed.

This approach of deleting entire messages tables instead of individual messages helps improves the I/O performance of the drive that holds the queue database.

Options for configuring the queue database

You configure the queue database by adding or modifying keys in the `%ExchangeInstallPath%Bin\EdgeTransport.exe.config` XML application configuration file. This file is associated with the Microsoft Exchange Transport service. Changes you make to the EdgeTransport.exe.config file take effect after you restart the Microsoft Exchange Transport service.

> **NOTE**
>
> Any customized per-server Exchange or Internet Information Server settings you make in exExchangeNoVersion XML application configuration files (for example, web.config files or the EdgeTransport.exe.config file) will be overwritten when you install an exExchangeNoVersion Cumulative Update (CU). Make sure that you save this information so that you can easily re-configure your server after the install. You must re-configure these settings after you install an exExchangeNoVersion CU.

The `<appSettings>` section of the EdgeTransport.exe.config file is where you can add new keys or modify existing keys. If a specific key doesn't exist, you can add it manually to change its value.

The keys for the queue database that are available in the EdgeTransport.exe.config file are described in the following table.

Message queue database keys that are available in the EdgeTransport.exe.config file

KEY	DEFAULT VALUE	DESCRIPTION
QueueDatabaseBatchSize	40	Specifies the number of database I/O operations that can be grouped together before they're executed. By default, this key doesn't exist in the EdgeTransport.exe.config file.
QueueDatabaseBatchTimeout	100	Specifies the maximum time in milliseconds that the database will wait for multiple database I/O operations to group before it executes them. The database I/O operations are executed without waiting for any more if the following conditions are true: • The number of database I/O operations that's specified by the *QueueDatabaseBatchSize* key hasn't been reached. • The time specified by the *QueueDatabaseBatchTimeout* key has passed. By default, this key doesn't exist in the EdgeTransport.exe.config file.
QueueDatabaseMaxConnections	4	Specifies the number of ESE database connections that can be open.
QueueDatabaseLoggingBufferSize	5MB	Specifies the memory that's used to cache the transaction records before they're written to the transaction log file.
QueueDatabaseLoggingFileSize	5MB	Specifies the maximum size of a transaction log file. When the maximum log file size is reached, a new log file is opened.
QueueDatabaseLoggingPath	`%ExchangeInstallPath%TransportRoles\data\Queue`	Specifies the default directory for the queue database log files. For instructions on how to change the location of the queue database, see Change the location of the queue database.

KEY	DEFAULT VALUE	DESCRIPTION
QueueDatabaseMaxBackgroundCleanupTasks	32	Specifies the maximum number of background cleanup work items that can be queued to the database engine thread pool at any time.
QueueDatabaseOnlineDefragEnabled	True	Enables or disables scheduled online defragmentation of the mail queue database. By default, this key doesn't exist in the EdgeTransport.exe.config file.
QueueDatabaseOnlineDefragSchedule	`1:00:00` or 1:00 A.M.	Specifies the time of day in 24 hour format to start the online defragmentation of the mail queue database. To specify a value, enter the value as a time span: *hh:mm:ss*, where *h* = hours, *m* = minutes, and *s* = seconds.
QueueDatabaseOnlineDefragTimeToRun	`3:00:00` or 3 hours	Specifies the length of time the online defragmentation task is allowed to run. Even if the defragmentation task doesn't finish in the time specified, the queue database is left in a consistent state. To specify a value, enter the value as a time span: *hh:mm:ss*, where *h* = hours, *m* = minutes, and *s* = seconds.
QueueDatabasePath	`%ExchangeInstallPath%TransportRoles\data\Queue`	Specifies the default directory for the queue database files. For instructions on how to change the location of the queue database, see Change the location of the queue database.

Queue properties

A queue has many properties that describe the purpose and status of the queue. Some queue properties are applied to the queue when the queue is created, and don't change. Other properties contain status, size, time, or other indicators that are updated frequently.

NextHopSolutionKey

The routing component of the categorizer in the Microsoft Exchange Transport service selects the destination for a message, and this destination is used to create the delivery queue. The destination is stamped on every recipient as the **NextHopSolutionKey** property. Every unique value of the **NextHopSolutionKey** property corresponds to a separate delivery queue.

The **NextHopSolutionKey** property contains the following fields:

- **DeliveryType**: Represents the results of the categorization of the message, and how the Transport service intends to transmit the message to the next hop, which could be the ultimate destination of the message, or an intermediate hop along the way. The Transport service uses a predefined list of values for **DeliveryType**.

  Based on the value of **DeliveryType**, the **NextHopCategory** property is added to the queue.

  - The value `External` indicates the next hop for the queue is outside the Exchange organization.
  - The value `Internal` indicates the next hop for the queue is inside the Exchange organization.

    Note that a message for an external recipient may require one or more internal hops before the message is delivered externally.

- **NextHopDomain**: Uses specific values based on the value of the **DeliveryType** field. For delivery queues, the value of this field is effectively the name of the queue.

The value of **NextHopDomain** isn't always a domain name. For example, the value could be the name of the target Active Directory site or database availability group (DAG). Think of this field as the *next hop name*.

- **NextHopConnector**: Uses specific values based on the value of the **DeliveryType** field. The value is always expressed as a GUID. If this field isn't used, the value is a GUID with all zeroes.

  The value of **NextHopConnector** isn't always the GUID of a connector. For example, the value could be the GUID of the target Active Directory site or DAG. Think of this field as the *next hop GUID*.

The values of **DeliveryType**, **NextHopCategory**, **NextHopDomain** and **NextHopConnector** are described in the following table.

DELIVERY TYPE IN QUEUE VIEWER	DELIVERYTYPE IN THE EXCHANGE MANAGEMENT SHELL	DESCRIPTION	NEXTHOPCATEGORY	NEXTHOPDOMAIN	NEXTHOPCONNECTOR
Delivery Agent	`DeliveryAgent`	The queue holds messages for delivery to recipients in a non-SMTP address space that's serviced by a delivery agent and a Delivery Agent connector. The connector has the local Mailbox server configured as a source server. For more information, see Delivery Agents and Delivery Agent Connectors.	External	This value is the destination address space that's configured on the Delivery Agent connector. For example, `MOBILE`.	This value is the GUID of the Delivery Agent connector. For example, `4520e633-d83d-411a-bbe4-6a84648674ee`.
DnsConnectorDelivery	`DnsConnectorDelivery`	The queue holds messages for delivery to recipients in an SMTP domain. The Send connector that services the domain has the local transport server configured as source server, and the Send connector is configured to use DNS routing.	External	This value is the destination address space that's configured on the Send connector. For example, `contoso.com`.	This value is the GUID of the Send connector. For example, `4520e633-d83d-411a-bbe4-6a84648674ee`.
Heartbeat	`Heartbeat`	This value is reserved for internal Microsoft use. For more information about heartbeat, see Shadow redundancy in Exchange Server.	n/a	n/a	n/a

DELIVERY TYPE IN QUEUE VIEWER	DELIVERYTYPE IN THE EXCHANGE MANAGEMENT SHELL	DESCRIPTION	NEXTHOPCATEGORY	NEXTHOPDOMAIN	NEXTHOPCONNECTOR
MapiDelivery	`MapiDelivery`	**Note**: This value isn't used by Exchange 2013 or later. It's included for backwards compatibility with Exchange 2010. The queue holds messages for delivery by an Exchange 2010 Hub Transport server to a mailbox on an Exchange 2010 Mailbox server in the local Active Directory site.	n/a	n/a	n/a
NonSmtpGatewayDelivery	`NonSmtpGatewayDelivery`	The queue holds messages for delivery to recipients in a non-SMTP address space that's serviced by a Foreign connector. The connector has the local Mailbox server configured as a source server. For more information, see Foreign Connectors.	External	This value is the destination address space that's configured on the Foreign connector. For example, `FAX`.	This value is the GUID of the Foreign connector. For example, `4520e633-d83d-411a-bbe4-6a84648674ee`.
Shadow Redundancy	`ShadowRedundancy`	The queue holds messages in a shadow queue. A shadow queue holds redundant copies messages in transit in case the primary messages aren't successfully delivered. For more information, see Shadow redundancy in Exchange Server.	Internal	This value is the FQDN of the primary transport server for which the shadow queue is holding redundant copies of the primary messages. For example, `mailbox01.contoso.com`.	This value is `00000000-0000-0000-0000-000000000000`.

DELIVERY TYPE IN QUEUE VIEWER	DELIVERYTYPE IN THE EXCHANGE MANAGEMENT SHELL	DESCRIPTION	NEXTHOPCATEGORY	NEXTHOPDOMAIN	NEXTHOPCONNECTOR
SmartHostConnectorDelivery	`SmartHostConnectorDelivery`	The queue holds messages for delivery to recipients in an SMTP domain. The Send connector that services the domain has the local transport server configured as source server, and the Send connector is configured to use smart host routing.	External	This value is the list of smart hosts that are configured on the Send connector. Smart hosts can be configured as FQDNs, IP addresses or both. The values can be one of the following: **FQDN**: The syntax is `<FQDN1,FQDN2,...>` . For example, `smarthost01.contoso.com` or `smarthost01.contoso.com,smarthost02.fabrikam` . **IP address**: The syntax is `<[IPAddress1],[IPAddress2],...>` . For example, `[10.10.10.100]` or `[10.10.10.100],[10.10.10.101]` . **FQDN and IP address**: The syntax is `<[IPAddress1],FQDN1,...>` , and depends on how the smart hosts are listed on the Send connector. For example, `[172.17.17.7],relay.tailspintoys.com` or `mail.contoso.com,[192.168.1.50]` .	This value is the GUID of the Send connector. For example, `4520e633-d83d-411a-bbe4-6a84648674ee` .
SMTP Delivery to Ex Online	`SmtpDeliveryToExo`	This value isn't used in on-premises Exchange.	n/a	n/a	n/a

DELIVERY TYPE IN QUEUE VIEWER	DELIVERYTYPE IN THE EXCHANGE MANAGEMENT SHELL	DESCRIPTION	NEXTHOPCATEGORY	NEXTHOPDOMAIN	NEXTHOPCONNECTOR
SMTP Delivery to Mailbox	`SmtpDeliveryToMailbox`	The queue holds messages for delivery to Exchange 2013 or later mailbox recipients. The destination mailbox database is in one of the following locations: • The local Exchange 2013 or later Mailbox server. • An Exchange 2019 Mailbox server in the same Exchange 2019 DAG. • An Exchange 2016 Mailbox server in the same Exchange 2016 DAG. An Exchange 2013 Mailbox server in the same Exchange 2013 DAG. • An Exchange 2013 or later Mailbox server in the same Active Directory site in non-DAG environments.	Internal	This value is the name of the destination mailbox database. For example, `Mailbox Database 0471695037` .	This value is the GUID of the target mailbox database. For example, `6dcb5a1e-0a88-4fc9-b8f9-634c34b1a123` .

DELIVERY TYPE IN QUEUE VIEWER	DELIVERYTYPE IN THE EXCHANGE MANAGEMENT SHELL	DESCRIPTION	NEXTHOPCATEGORY	NEXTHOPDOMAIN	NEXTHOPCONNECTOR
SMTP Relay to Send Connector Source Servers	`SmtpRelayToConnectorSourceServers`	The queue holds messages for delivery to an SMTP or non-SMTP address space that's serviced by a Send connector, Delivery Agent connector, or Foreign connector. The connector has a remote transport server configured as a source server. The remote transport server could be an Exchange 2013 or later Mailbox server or an Exchange 2010 Hub Transport server. The remote transport server could be located in the local Active Directory site, or in a remote Active Directory site.	Internal	This value is the name of the destination Send connector, Delivery Agent connector, or Foreign connector. For example, `Contoso.com Send Connector` .	This value is the GUID of the destination Send connector, Delivery Agent connector, or Foreign connector. For example, `4520e633-d83d-411a-bbe4-6a84648674ee` .
SMTP Relay to Database Availability Group	`SmtpRelayToDag`	The queue holds messages for delivery to Exchange 2013 or later mailbox recipients, where the destination mailbox database is located in a remote DAG. The remote DAG could be located in the local Active Directory site, or in a remote Active Directory site.	Internal	This value is the name of the destination DAG. For example, `DAG1` .	This value is the GUID of the destination DAG. For example, `6dcb5a1e-0a88-4fc9-b8f9-634c34b1a123`

DELIVERY TYPE IN QUEUE VIEWER	DELIVERYTYPE IN THE EXCHANGE MANAGEMENT SHELL	DESCRIPTION	NEXTHOPCATEGORY	NEXTHOPDOMAIN	NEXTHOPCONNECTOR
SMTP Relay to Mailbox Delivery Group	`SmtpRelayToMailboxDeliveryGroup`	The queue holds messages for delivery to legacy mailbox recipients, where the destination mailbox is on an Exchange 2010 Mailbox server. The message is related to an Exchange 2010 Hub Transport server. The destination Exchange 2010 Hub Transport server could be in the local Active Directory site, or a remote Active Directory site.	Internal	The queue name uses the syntax: `Site:<ADSiteName>;Version:<ExchangeVersion>` , where *<ADSiteName>* is the name of the destination Active Directory site, and *<ExchangeVersion>* is the version of Exchange 2010 on the Mailbox server.	This value is blank.
SMTP Relay to Remote Active Directory Site	`SmtpRelayToRemoteActiveDirectorySite`	The queue holds messages for delivery to a remote destination, and the routing topology requires the message to be routed through a specific Active Directory site. The site is an intermediate hop on the way to the final destination. This situation occurs under the following circumstances: The message needs to be routed through a hub site. The message requires delivery through a Send connector that's configured on an Edge Transport server that's subscribed to a remote Active Directory site.	Internal	This value is the target Active Directory site name. For example, `NorthAmericaSite` .	This value is the GUID of the target Active Directory site. For example, `bfd6c3df-5b65-8bfb-53f1f2c0d55c` .
SMTP Relay to specified remote forest	`SmtpRelayToRemoteForest`	This value isn't used in on-premises Exchange	n/a	n/a	n/a

DELIVERY TYPE IN QUEUE VIEWER	DELIVERYTYPE IN THE EXCHANGE MANAGEMENT SHELL	DESCRIPTION	NEXTHOPCATEGORY	NEXTHOPDOMAIN	NEXTHOPCONNECTOR
SMTP Relay to Specified Exchange Servers	`SmtpRelayToServers`	The queue holds messages for delivery to a distribution group that's configured for a specific expansion server. The expansion server could be an Exchange 2013 or later Mailbox server or an Exchange 2010 Hub Transport server. The expansion server could be located in the local Active Directory site, or in a remote Active Directory site.	Internal	This value is the FQDN of the target expansion server. For example, `mailbox01.contoso.com`.	This value is `0000000-0000-0000-0000-000000000000`.
SmtpRelayToTiRg	`SmtpRelayToTiRg`	**Note**: This value isn't used by Exchange 2013 or later. It's included for backwards compatibility with Exchange 2010. The queue holds messages for delivery by an Exchange 2010 Hub Transport server to an Exchange 2003 routing group.	n/a	n/a	n/a
Smtp Relay in Active Directory Site	`SmtpRelayWithinAdSite`	**Note**: This value isn't used by Exchange 2013 or later. It's included for backwards compatibility with Exchange 2010. The queue holds messages for delivery by an Exchange 2010 Hub Transport server to another Hub Transport server in the same Active Directory site.	n/a	n/a	n/a

DELIVERY TYPE IN QUEUE VIEWER	DELIVERYTYPE IN THE EXCHANGE MANAGEMENT SHELL	DESCRIPTION	NEXTHOPCATEGORY	NEXTHOPDOMAIN	NEXTHOPCONNECTOR
SMTP Relay in Active Directory Site to Edge Transport Server	`SmtpRelayWithinAdSiteToEdge`	The queue holds messages for delivery to an external SMTP domain that's serviced by a Send connector that's configured on an Edge Transport server. The Edge Transport server is subscribed to the local Active Directory site.	Internal	This value is the name of the Send connector that sends outbound Internet mail from the Edge Transport server to the Internet. This Send connector is automatically created by the Edge subscription, and is named EdgeSync - *<ADSiteName>* to Internet.	This value is the GUID of the Send connector. For example, `4520e633-d83d-411a-bbe4-6a84648674ee` .
Undefined	`Undefined`	This value is used only on the Submission queue and the poison message queue.	Internal	For the Submission queue, this value is `Submisssion` . For the poison message queue, this value is `Poison Message` .	This value is `00000000-0000-0000-0000-000000000000` .
Unreachable	`Unreachable`	This value is used only on the Unreachable queue.	Internal	This value is `Unreachable Domain` .	This value is `00000000-0000-0000-0000-000000000000` .

IncomingRate, OutgoingRate, and Velocity

Exchange measures the rate of messages entering and leaving a queue and stores these values in queue properties. You can use these rates as an indicator of queue and transport server health. The properties are described in the following table:

PROPERTY	DESCRIPTION
IncomingRate	The rate that messages are entering the queue. The rate is the number of messages per second averaged over the last minute.
OutgoingRate	The rate that messages are leaving the queue. The rate is the number of messages per second averaged over the last minute.
Velocity	The drain rate of the queue, calculated by subtracting the value of **IncomingRate** from the value of **OutgoingRate**. If the value is greater than 0, messages are leaving the queue faster than they are entering the queue. If the value equals 0, messages are leaving the queue as fast as they are entering the queue. This is also the value you'll see when the queue is inactive. If the value is less than 0, messages are entering the queue faster than they are leaving the queue. The **Velocity** value is displayed in the results of **Get-Queue**.

At a basic level, a positive value of **Velocity** indicates a healthy queue that's efficiently draining, and a negative value of **Velocity** indicates a queue that isn't efficiently draining. However, you also need to consider the values of **IncomingRate**,

OutgoingRate, and **MessageCount**, as well as the magnitude of **Velocity**.

For example, consider a queue that has the following property values.

- **Velocity**: -50
- **MessageCount**: 1000
- **OutgoingRate**: 10
- **IncomingRate**: 60

Based on the property values for this queue, the negative value for **Velocity** clearly indicates that the queue isn't draining properly.

Now consider a queue that has the following property values.

- **Velocity**: -0.85
- **MessageCount**: 2
- **OutgoingRate**: 0.15
- **IncomingRate**: 1

Although the value for **Velocity** is negative, it's very close to zero, and the values of the other properties are also very small. Therefore, a negative **Velocity** value for this queue doesn't indicate a problem with the queue.

Queue status

The current status of a queue is stored in the **Status** property of the queue. A queue can have one of the status values that's described in the following table:

QUEUE STATUS	DESCRIPTION
Active	The queue is actively transmitting messages.
Connecting	The queue is in the process of connecting to the next hop.
Ready	The queue recently transmitted messages, but the queue is now empty.
Retry	The last automatic or manual connection attempt failed, and the queue is waiting to retry the connection.
Suspended	The queue has been manually suspended by an administrator to prevent message delivery. New messages can enter the queue, and messages that are in the act of being transmitted to the next hop will finish delivery and leave the queue. Otherwise, messages won't leave the queue until the queue is manually resumed by an administrator. **Notes:** You can suspend the following queues: Delivery queues that have any status. The Unreachable queue. When you suspend this queue, messages are no longer automatically resubmitted to the categorizer when configuration updates are detected. To automatically resubmit these messages, you need to manually resume the queue. The Submission queue. When you suspend this queue, messages aren't picked up by the categorizer until the queue is resumed. Suspending a queue doesn't change the status of the messages in the queue.

Other queue properties

There are other queue properties that are self-explanatory. You can use most of the queue properties as filter options. By specifying filter criteria, you can quickly locate queues and take action on them. For a complete description of the filterable queue properties, see Queue properties.

An important queue property that's also worth mentioning here is the **MessageCount** property that shows how many messages are in a queue. This property is an important indicator of queue health. For example, a delivery queue that contains a large number of messages that continues to grow and never decreases could indicate a routing or transport pipeline issue that requires your attention.

Message properties

A message in a queue has many properties. Many of the properties reflect the information that was used to create the message. Some of the messages status and information properties are heavily influenced by corresponding properties on the queue. However, an individual message may have a different value than the corresponding property of the queue. Other properties contain status, time, or other indicators that are updated frequently.

Message status

The current status of a message is stored in the **Status** property of the message. A message can have one of the status values that's described in the following table:

MESSAGE STATUS	DESCRIPTION
Active	If the message is in a delivery queue, the message is being delivered to its destination. If the message is in the Submission queue, the message is being processed by the categorizer.
Locked	This value is reserved for internal Microsoft use, and isn't used in on-premises Exchange organizations.
PendingRemove	The message was deleted by the administrator, but the message was already in the act of being transmitted to the next hop. The message will be deleted if the delivery ends in an error that causes the message to reenter the queue. Otherwise, delivery will continue.
PendingSuspend	The message was suspended by the administrator, but the message was already in the act of being transmitted to the next hop. The message will be suspended if the delivery ends in an error that causes the message to reenter the queue. Otherwise, delivery will continue.
Ready	The message is waiting in the queue and is ready to be processed.
Retry	The last automatic or manual connection attempt fail for the queue that holds the message. The message is waiting for the next automatic queue connection retry.
Suspended	The message was manually suspended by an administrator. Any messages in the poison message queue are in a permanently suspended state.

Other message properties

There are other message properties that are self-explanatory. You can use most of the message properties as filter options. By specifying filter criteria, you can quickly locate messages and take action on them. For a complete description of the filterable message properties, see Properties of messages in queues.

Manage queues and messages in queues

Queue Viewer and the historical queue and message management cmdlets in the Exchange Management Shell are restricted to a single Exchange server. You can view or operate on individual queues or messages, or multiple queues or messages, but only on a specific server.

The **Get-QueueDigest** cmdlet was introduced in Exchange 2013 to provide a high-level, aggregate view of the state of queues on all servers within a specific scope. The scope could be a DAG, an Active Directory site, a list of servers, or the entire Active Directory forest. Note that queues on a subscribed Edge Transport server in the perimeter network aren't included in the results. Also, **Get-QueueDigest** is available on Edge Transport servers, but the results are restricted to queues on the Edge Transport server.

> **NOTE**
>
> By default, the **Get-QueueDigest** cmdlet displays delivery queues that contain ten or more messages, and the results are between one and two minutes old. For instructions on how to change these default values, see Configure Get-QueueDigest.

The following table describes the management tasks you can perform on queues or messages in queues.

TASK	DESCRIPTION	TOOL TO USE	INSTRUCTIONS
View and filter queues on a server	Displays one or more queues on a transport server. You can use the results to take action on the queues.	Queue Viewer or the **Get-Queue** cmdlet.	Procedures for queues
View and filter queues on specific servers in specific DAGs, specific Active Directory sites, or in the whole Active Directory forest.	Displays a summary list of queues.	**Get-QueueDigest** cmdlet	Procedures for queues
Suspend queues	Temporarily prevent delivery of messages that are currently in the queue. The queue continues to accept new messages, but no messages leave the queue.	Queue Viewer or the **Suspend-Queue** cmdlet.	Procedures for queues
Resume queues	Reverses the effect of the suspend queue action, and enables delivery of queued messages to resume.	Queue Viewer or the **Resume-Queue** cmdlet.	Procedures for queues
Retry queues	Immediately tries to connect to the next hop. Without manual intervention, when the connection to the next hop fails, the connection is attempted a specific number of times after a specific time interval between each attempt. Whether the connection attempt is manual or automatic, any connection attempt resets the next retry time. For more information, see Message retry, resubmit, and expiration intervals.	Queue Viewer or the **Retry-Queue** cmdlet.	Procedures for queues

TASK	DESCRIPTION	TOOL TO USE	INSTRUCTIONS
Resubmit messages in queues	Causes messages in the queue to be resubmitted to the Submission queue and to go back through the categorization process.	**Retry-Queue** with the *Resubmit* parameter Note that you can use Queue Viewer to resubmit messages, but only from the poison message queue. To resubmit a poison message, you first need to resume the message in Queue Viewer, or by using the **Resume-Message** cmdlet.	Procedures for queues
Suspend messages in queues	Temporarily prevents delivery of a message. You can use the suspend message action to prevent delivery of a message to all the recipients in a specific queue or to all recipients in all queues.	Queue Viewer or the **Suspend-Message** cmdlet.	Procedures for messages in queues
Resume messages in queues	Reverses the effect of the suspend message action, and enables the delivery of queued messages to resume. You can resume the delivery of a message to all recipients in a specific queue, or to all recipients in all queues.	Queue Viewer or the **Resume-Message** cmdlet.	Procedures for messages in queues
Remove messages from queues	Permanently prevents the delivery of a message. You can prevent the delivery of a message to any recipients in a specific queue, or to all recipients in all queues. Optionally, you can send a non-delivery report (also known as an NDR, delivery status notification, DSN or bounce message) to the sender when the message is removed.	Queue Viewer or the **Remove-Message** cmdlet.	Procedures for messages in queues
Export messages from queues	Copies a message to the location that you specify. The messages aren't deleted from the queue, but a copy of the message is saved as a file in the specified location. This enables administrators or officials in an organization to later examine the messages. Before you export a message, you need to temporarily suspend the message.	**Export-Message** cmdlet only.	Export messages from queues

Procedures for queues

8/3/2020 • 12 minutes to read • Edit Online

In Exchange Server, you can use the Queue Viewer in the Exchange Toolbox or the Exchange Management Shell to manage queues. For more information about queues, see Queues and messages in queues.

This topic describes how to perform the following procedures on queues:

- **View queues**
- **Retry queues**: When an Exchange server can't connect to the next hop, the queue is put into a status of Retry, and the server periodically tries to connect and deliver the messages. When you manually retry a queue, you override the scheduled retry time by forcing an immediate connection attempt.
- **Resubmit queues**: Resubmitting a queue is similar to retrying a queue, except the messages are sent back to the Submission queue for the categorizer to process, instead of immediately trying to connect to the next hop. This is useful if changes to your network infrastructure are preventing the messages in the queue from being delivered.
- **Suspend queues**: New messages can enter the queue, and messages that are in the act of being transmitted to the next hop will leave the queue, but otherwise, messages won't leave the queue until the queue is manually resumed.
- **Resume queues**: Restart outgoing message delivery for a queue that has a status of Suspended. When you resume a queue, the status of messages in the queue doesn't change (for example, messages that have a status of Suspended remain suspended and won't leave the queue).

For procedures on messages in queues, see Procedures for messages in queues.

What do you need to know before you begin?

- Estimated time to complete each procedure: 5 minutes
- To find and open the Exchange Toolbox, use one of the following procedures:
  - **Windows 10**: Click **Start** > **All Apps** > **Microsoft Exchange Server <Version>** > **Exchange Toolbox**.
  - **Windows Server 2012 R2 or Windows 8.1**: On the Start screen, open the Apps view by clicking the down arrow near the lower-left corner or swiping up from the middle of the screen. The **Exchange Toolbox** shortcut is in a group named **Microsoft Exchange Server <Version>**.
  - **Windows Server 2012**: Use any of the following methods:
    - On the Start screen, click an empty area, and type Exchange Toolbox.
    - On the desktop or the Start screen, press Windows key + Q. In the Search charm, type Exchange Toolbox.
    - On the desktop or the Start screen, move your cursor to the upper-right corner, or swipe left from the right edge of the screen to show the charms. Click the Search charm, and type Exchange Toolbox.

    When the shortcut appears in the results, you can select it.
- To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see

Open the Exchange Management Shell.

- For more information about using filters and identity values in the Exchange Management Shell, see Find queues and messages in queues in the Exchange Management Shell.

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Queues" entry in the Mail flow permissions topic.

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

View queues

Use Queue Viewer to view queues

1. In the **Exchange Toolbox**, in the **Mail flow tools** section, double-click **Queue Viewer** to open the tool in a new window.

2. In Queue Viewer, click the **Queues** tab. A list of all queues on the server to which you're connected is displayed.

3. You can use the **Export List** link in the action pane to export the list of queues. For more information, see How to Export Lists from the Exchange Management Consoles.

Use the Exchange Management Shell to view queues

To view queues, use the following syntax.

```
Get-Queue [-Filter <Filter> -Server <ServerIdentity> -Include <Internal | External | Empty | DeliveryType> -
Exclude <Internal | External | Empty | DeliveryType>]
```

This example displays basic information about all non-empty queues on the server named Mailbox01.

```
Get-Queue -Server Mailbox01 -Exclude Empty
```

This example displays detailed information for all queues on the local Exchange server that contain more than 100 messages.

```
Get-Queue -Filter "MessageCount -gt 100" | Format-List
```

For more information, see Get-Queue and Find queues and messages in queues in the Exchange Management Shell.

Use the Exchange Management Shell to view queue summary information on multiple Exchange servers

The **Get-QueueDigest** cmdlet provides a high-level, aggregate view of the state of queues on all servers within a specific scope (for example, a DAG, an Active Directory site, a list of servers, or the entire Active Directory forest).

By default, the **Get-QueueDigest** cmdlet displays delivery queues that contain ten or more messages, and the results are between one and two minutes old. For instructions on how to change these default values, see

Configure Get-QueueDigest.

Notes:

- Queues on a subscribed Edge Transport server aren't included in the results of **Get-QueueDigest**.
- **Get-QueueDigest** is available on Edge Transport servers, but the results are restricted to local queues on the server.

To view summary information about queues on multiple Exchange servers, run the following command:

```
Get-QueueDigest <-Server <ServerIdentity1,ServerIdentity2...> | -Dag <DagIdentity1,DagIdentity2...> | -Site
<ADSiteIdentity1,ADSiteIdentity2...> | -Forest> [-Filter <Filter>]
```

This example displays summary information about the queues on all Exchange 2013 or later Mailbox servers in the Active Directory site named FirstSite where the message count is greater than 100.

```
Get-QueueDigest -Site FirstSite -Filter "MessageCount -gt 100"
```

This example displays summary information about the queues on all Mailbox servers in the database availability group (DAG) named DAG01 where the queue status has the value **Retry**.

```
Get-QueueDigest -Dag DAG01 -Filter "Status -eq 'Retry'"
```

For more information, see Get-QueueDigest.

Retry queues

When you retry a delivery queue, you force an immediate connection attempt and override the next scheduled retry time. For more information about the schedule retry time for queues, see Message retry, resubmit, and expiration intervals.

Notes:

- The queue must be in a status of Retry for this action to have any effect.
- If the connection isn't successful, the retry interval timer is reset.

Use Queue Viewer to retry a queue

1. In the **Exchange Toolbox**, in the **Mail flow tools** section, double-click **Queue Viewer** to open the tool in a new window.
2. In Queue Viewer, click the **Queues** tab. A list of all queues on the server that you're connected to is displayed.
3. Click **Create Filter**, and enter your filter expression as follows:
   a. Select **Status** from the queue property drop-down list.
   b. Select **Equals** from the comparison operator drop-down list.
   c. Select **Retry** from the value drop-down list.
   d. Click **Apply Filter**. All queues that currently have a **Retry** status are displayed.
   e. Select one or more queues from the list. Right-click, and then select **Retry Queue**. If the connection attempt is successful, the queue status changes to **Active**. If no connection can be made, the queue remains in a status of **Retry** and the next retry time is updated.

Use the Exchange Management Shell to retry a queue

To retry queues, use the following syntax.

```
Retry-Queue <-Identity QueueIdentity | -Filter QueueFilter [-Server ServerIdentity]>
```

This example retries all queues on the local server with the status of Retry.

```
Retry-Queue -Filter "Status -eq 'Retry'"
```

This example retries the queue named contoso.com on the server named Mailbox01.

```
Retry-Queue -Identity Mailbox01\contoso.com
```

How do you know this worked?

To verify that you have successfully retried a queue, use either of the following procedures:

- In Queue Viewer, verify the values of the **Status**, **Next Retry Time**, and **Last Error** properties.

- In the Exchange Management Shell, replace *<QueueIdentity>* with the identity of the queue, and use the following syntax to verify the property values:

  ```
  Get-Queue -Identity <QueueIdentity> | Format-Table -Auto Identity,Status,LastRetryTime,NextRetryTime
  ```

Resubmit queues

Resubmitting a queue sends all messages in the queue back to the Submission queue for the categorizer to process. For more information about the categorizer, see Mail flow and the transport pipeline.

Notes:

- You can't use Queue Viewer to resubmit queues. You can only use the Exchange Management Shell.

- You can resubmit the following queues:

  - A delivery queue that has the status of Retry.

  - The Unreachable queue.

    Any messages in the queue that have the status value of Suspended aren't resubmitted.

- You can't resubmit the poison message queue, but you can resubmit individual messages in the queue. For more information, see the Resubmit messages in the poison message queue section later in this topic.

- Instead of resubmitting the queue, you can export the messages to .eml files and resubmit them by using the Replay directory on any Exchange server. For more information, see Export messages from queues

Use the Exchange Management Shell to resubmit queues

To resubmit queues, use the following syntax:

```
Retry-Queue <-Identity QueueIdentity | -Filter "Status -eq 'Retry'" -Server ServerIdentity> -Resubmit $true
```

This example resubmits all messages located in any delivery queues with the status of Retry on the server named Mailbox01.

```
Retry-Queue -Filter "Status -eq 'Retry'" -Server Mailbox01 -Resubmit $true
```

This example resubmits all messages located in the Unreachable queue on the server Mailbox01.

```
Retry-Queue -Identity Mailbox01\Unreachable -Resubmit $true
```

For more information, see Retry-Queue.

How do you know this worked?

To verify that you have successfully resubmitted a queue, use either of the following procedures:

- In Queue Viewer, verify the properties of the queue.
- In the Exchange Management Shell, replace *<QueueIdentity>* with the identity of the queue, and run the following command to verify the property values:

  ```
  Get-Queue -Identity <QueueIdentity>
  ```

Resubmit messages in the poison message queue

A special case for resubmitting messages is the poison message queue. You can't resubmit the poison message queue like other queues, but you can resubmit individual messages in the poison message queue.

Notes:

- Messages in the poison message queue might be genuinely harmful, or they might be valid messages that are the victims of an poorly written transport agent or a software bug. If you're unsure of the safety of the messages in the poison message queue, you should export the messages to files so you can examine them. For more information, see Export messages from queues.
- The procedure to resubmit messages from the poison message queue is the same as resuming suspended messages from other queues. You can use Queue Viewer or the Exchange Management Shell. For more information about resuming messages, see Resume messages in queues.
- The poison message queue is only visible when the queue contains messages.

Use Queue Viewer to resubmit messages in the poison message queue

1. In the **Exchange Toolbox**, in the **Mail flow tools** section, double-click **Queue Viewer** to open the tool in a new window.
2. In Queue Viewer, click the **Queues** tab. A list of all queues on the server that you're connected to is displayed.
3. Select the poison message queue. In the action pane, select **View Messages**.
4. Select one or more messages from the list, right-click, and select **Resume**.

Use the Exchange Management Shell to resubmit messages in the poison message queue

To resubmit a message from the poison message queue, perform the following steps.

1. Find the identity of the message by running the following command on the local server.

   ```
   Get-Message -Queue Poison | Format-Table Identity
   ```

2. Use the identity of the message from the previous step in the following command.

```
Resume-Message <PoisonMessageIdentity>
```

This example resumes a message from the poison message queue that has the message Identity value of 222.

```
Resume-Message 222
```

For more information, see Resume-Message.

How do you know this worked?

To verify that you have successfully resubmitted a message from the poison message queue, use either of the following procedures to verify that the message is no longer in the queue:

- In Queue Viewer, view the poison message queue where you attempted to resubmit the message.
- In the Exchange Management Shell, run the following command:

```
Get-Message -Queue Poison
```

If the message you resubmitted was the only message in the poison message queue, and the queue is no longer visible, that's also an indication of a successful message resubmission.

Suspend queues

You can suspend a queue to stop mail flow, and then suspend one or more messages in the queue. For more information, see Suspend messages in queues.

Notes:

- You can suspend the following queues:
  - A delivery queue that has any status.
  - The Unreachable queue. Until you manually resume this queue, messages are no longer automatically resubmitted to the categorizer when configuration updates are detected.
  - The Submission queue. Until you manually resume this queue, messages aren't picked up by the categorizer.
- Suspending a queue doesn't change the status of the messages in the queue to Suspended.

Use Queue Viewer to suspend a queue

1. In the **Exchange Toolbox**, in the **Mail flow tools** section, double-click **Queue Viewer** to open the tool in a new window.
2. In Queue Viewer, click the **Queues** tab. A list of all queues on the server that you're connected to is displayed. You can create a filter to display only queues that meet specific criteria.
3. Select one or more queues, right-click, and then select **Suspend**.

Use the Exchange Management Shell to suspend a queue

To suspend a queue, use the following syntax:

```
Suspend-Queue <-Identity QueueIdentity | -Filter "QueueFilter" [-Server ServerIdentity]>
```

This example suspends all queues on the local server that have a message count equal to or greater than 1,000 and that have a status of Retry.

```
Suspend-Queue -Filter "MessageCount -ge 1000 -and Status -eq 'Retry'"
```

This example suspends the queue named contoso.com on the server named Mailbox01.

```
Suspend-Queue -Identity Mailbox01\contoso.com
```

For more information, see Suspend-Queue.

How do you know this worked?

To verify that you have successfully suspended a queue, use either of the following procedures:

- In Queue Viewer, verify the queue has the **Status** value of Retry.

- In the Exchange Management Shell, replace *<QueueIdentity>* with the identity of the queue, and run the following command to verify the **Status** property value:

  ```
  Get-Queue -Identity <QueueIdentity>
  ```

Resume queues

By resuming a queue, you restart outgoing message delivery from a queue that has a status of Suspended.

Notes:

- You can only resume queues that have been suspended.

- Resuming a queue doesn't change the status of messages in the queue. For example, messages that have a status of Suspended remain suspended and don't leave the queue after you resume the queue.

Use Queue Viewer to resume queues

1. In the **Exchange Toolbox**, in the **Mail flow tools** section, double-click **Queue Viewer** to open the tool in a new window.

2. In Queue Viewer, click the **Queues** tab. A list of all queues on the server that you're connected to is displayed.

3. Click **Create Filter**, and enter your filter expression as follows:

   a. Select **Status** from the queue property drop-down list.

   b. Select **Equals** from the comparison operator drop-down list.

   c. Select **Suspended** from the value drop-down list.

4. Click **Apply Filter**. All queues on the server that are currently suspended are displayed.

5. Select one or more queues from the list, right-click, and then select **Resume**.

Use the Exchange Management Shell to resume queues

To resume queues, use the following syntax:

```
Resume-Queue <-Identity QueueIdentity | -Filter "QueueFilter" [-Server ServerIdentity]>
```

This example resumes all queues on the local server that have a status of Suspended.

```
Resume-Queue -Filter "Status -eq 'Suspended'"
```

This example resumes the suspended delivery queue named contoso.com on the server named Mailbox01.

```
Resume-Queue -Identity Mailbox01\contoso.com
```

For more information, see Resume-Queue.

How do you know this worked?

To verify that you have successfully resumed a queue, use either of the following procedures:

- In Queue Viewer, verify the queue doesn't have the **Status** value Suspended (for example, Active, Connecting, or Ready).
- In the Exchange Management Shell, replace *<QueueIdentity>* with the identity of the queue, and run the following command to verify the **Status** property value:

  ```
  Get-Queue -Identity <QueueIdentity>
  ```

Queue properties in Exchange Server

8/3/2020 • 6 minutes to read • Edit Online

Filtering queues by one or more queue properties in Exchange Server allows you to quickly find and take action on those queues. The following scenarios are examples of how you might use queue filtering to manage mail flow:

- You receive a message from System Center Operations Manager that indicates a queue length has exceeded the established threshold. You want to investigate whether a server-wide mail flow problem exists.

  You create a filter to view all the queues on a server whose message count exceeds what you consider to be typical. If a mail flow problem is indicated, you can select all the queues in the results and suspend the queues while you continue to investigate.

- You suspend several queues to investigate the cause of mail flow problems. You determine that the problem was caused by an incorrect connector configuration that is now fixed.

  You can create a filter to view all the queues that have a status of Suspended, and then select all the queues in the filter results and resume the queues.

You can create queue filters in Queue Viewer in the Exchange Toolbox, or by using the *Filter* parameter on the queue management cmdlets. Note that the queue management cmdlets support more filterable properties than Queue Viewer.

For more information about Queue Viewer, see Queue Viewer. For more information about the queue management cmdlets, see Procedures for queues and Find queues and messages in queues in the Exchange Management Shell.

Queue properties to use as filters

The following table describes the queue properties that you can use as filters in Queue Viewer and the Exchange Management Shell.

QUEUE VIEWER	EXCHANGE MANAGEMENT SHELL	COMPARISON OPERATORS	DESCRIPTION
n/a	`DeferredMessageCount`	Equals (`-eq`) Does not equal (`-ne`) Greater than (`-gt`) Greater than or equal to (`-ge`) Less than (`-lt`) Less than or equal to (`-le`)	The number of messages returned to the Submission queue because of transient errors that were encountered during recipient resolution. For more information about deferred messages, see Recipient resolution in Exchange Server.

QUEUE VIEWER	EXCHANGE MANAGEMENT SHELL	COMPARISON OPERATORS	DESCRIPTION
n/a	`DDeferredMessageCountsPerPriority`	Equals (`-eq`) Does Not Equal (`-ne`) Contains (`-like`)	An array that shows the number of deferred messages in the queue by priority (importance) value. The **MessageCountsPerPriority** property shows what each number means. For example, the value `{1, 5, 10, 0}` indicates the queue contains 1 deferred High priority message, 5 deferred Normal priority messages, 10 deferred Low priority messages, and no deferred messages that have the priority value None.
Delivery Type	`DeliveryType`	**Equals** (`-eq`) **Does Not Equal** (`-ne`)	The results of the categorization of the message, and how the Transport service intends to transmit the message to the next hop. For a list of the available **DeliveryType** values, see NextHopSolutionKey.
n/a	`FirstRetryTime`	Equals (`-eq`) Does not equal (`-ne`) Greater than (`-gt`) Greater than or equal to (`-ge`) Less than (`-lt`) Less than or equal to (`-le`)	The date/time of the first connection attempt for a queue that has a status of `Retry`. For more information, see Message retry, resubmit, and expiration intervals.
n/a	`Identity`	n/a	The identity of the queue in the form of *<Server>**<Queue>*. For more information see Queue identity.
n/a	`IncomingRate`	Equals (`-eq`) Does not equal (`-ne`) Greater than (`-gt`) Greater than or equal to (`-ge`) Less than (`-lt`) Less than or equal to (`-le`)	A calculated number that indicates how quickly messages are entering the queue. For more information, see IncomingRate, OutgoingRate, and Velocity.
Last Error	`LastError`	**Equals** (`-eq`) **Does Not Equal** (`-ne`) **Contains** (`-contains`) **Is Present** **Is Not Present**	The last error that was recorded for the queue. For more information about SMTP error codes, see DSNs and NDRs in Exchange Server.

QUEUE VIEWER	EXCHANGE MANAGEMENT SHELL	COMPARISON OPERATORS	DESCRIPTION
Last Retry Time	`LastRetryTime`	**Greater Than (`-gt`)** **Greater Than or Equals (`-ge`)** **Less Than (`-lt`)** **Less Than or Equals (`-le`)** **Is Present** **Is Not Present**	The date/time of the last connection attempt for a queue that has a status of `Retry`. For more information, see Message retry, resubmit, and expiration intervals.
n/a	`LockedMessageCount`	n/a	This property is reserved for internal Microsoft use, and isn't used in on-premises Exchange organizations.
Message Count	`MessageCount`	**Equals (`-eq`)** **Does Not Equal (`-ne`)** **Greater Than (`-gt`)** **Greater Than or Equals (`-ge`)** **Less Than (`-lt`)** **Less Than or Equals (`-le`)**	The number of messages in the queue.
n/a	`MessageCountsPerPriority`	Equals (`-eq`) Does Not Equal (`-ne`) Contains (`-like`)	An array that shows the number of messages in the queue by priority (importance) value. The **MessageCountsPerPriority** property shows what each number means. For example, the value `{1, 100, 10, 0}` indicates the queue contains 1 High priority message, 100 Normal priority messages, 10 Low priority messages, and no messages that have the priority value None. For more information about priority queuing, see Priority Queuing.
n/a	`NextHopCategory`	Equals (`-eq`) Does Not Equal (`-ne`)	The value `Internal` or `External` for the next hop based on the value of the **DeliveryType** property. For more information, see NextHopSolutionKey.
n/a	`NextHopConnector`	Equals (`-eq`) Does Not Equal (`-ne`) Contains (`-like`)	The GUID of the next hop based on the value of the **DeliveryType** property. For more information, see NextHopSolutionKey.

QUEUE VIEWER	EXCHANGE MANAGEMENT SHELL	COMPARISON OPERATORS	DESCRIPTION
Next Hop Domain	`NextHopDomain`	**Equals** (`-eq`) **Does Not Equal** (`-ne`) **Contains** (`-like`)	The name of next hop based on the value of the **DeliveryType** property. For more information, see NextHopSolutionKey.
Next Retry Time	`NextRetryTime`	**Greater Than** (`-gt`) **Greater Than or Equals** (`-ge`) **Less Than** (`-lt`) **Less Than or Equals** (`-le`) **Is Present** **Is Not Present**	The date/time of the next connection attempt for a queue that has a status of `Retry` . For more information, see Message retry, resubmit, and expiration intervals.
n/a	`OutboundIPPool`	n/a	This property is reserved for internal Microsoft use, and isn't used in on-premises Exchange organizations.
n/a	`OutgoingRate`	Equals (`-eq`) Does not equal (`-ne`) Greater than (`-gt`) Greater than or equal to (`-ge`) Less than (`-lt`) Less than or equal to (`-le`)	A calculated number that indicates how quickly messages are leaving the queue. For more information, see IncomingRate, OutgoingRate, and Velocity.
n/a	`OverrideSource`	n/a	This property is reserved for internal Microsoft use, and isn't used in on-premises Exchange organizations.
n/a	`PriorityDescriptions`	n/a	The value descriptions in the **DeferredMessageCountsPerPriority** and **MessageCountsPerPriority** properties. The value of this property is `{High, Normal, Low, None}` . Because the value of this property is always the same, it won't make a good filter.
n/a	`RetryCount`	Equals (`-eq`) Does not equal (`-ne`) Greater than (`-gt`) Greater than or equal to (`-ge`) Less than (`-lt`) Less than or equal to (`-le`)	The number connection attempts for a queue that has a status of `Retry` . For more information, see Message retry, resubmit, and expiration intervals.

QUEUE VIEWER	EXCHANGE MANAGEMENT SHELL	COMPARISON OPERATORS	DESCRIPTION
n/a	`RiskLevel`	n/a	This property is reserved for internal Microsoft use, and isn't used in on-premises Exchange organizations.
Status	`Status`	**Equals** (`eq`) **Does Not Equal** (`-ne`)	The current queue status. A queue can have one of the following status values: Active, Connecting, Suspended, Ready, or Retry. For more information, see Queue status.
n/a	`TlsDomain`	Equals (`-eq`) Does Not Equal (`-ne`) Contains (`-like`)	The FQDN of the destination domain if the domain is configured for Domain Security (mutual TLS authentication).
n/a	`Velocity`	Equals (`-eq`) Does not equal (`-ne`) Greater than (`-gt`) Greater than or equal to (`-ge`) Less than (`-lt`) Less than or equal to (`-le`)	A calculated number that indicates how effectively the queue is draining. For more information, see IncomingRate, OutgoingRate, and Velocity

Export messages from queues

8/3/2020 • 4 minutes to read • Edit Online

On Mailbox servers and Edge Transport servers in Exchange Server, you can export the messages in a queue to files. The exported messages aren't removed from the queue. Copies of the messages are made in the specified location as a plain text files. You can view the message files in Notepad or Outlook, and you can resubmit the message files by using the Replay directory on any other Mailbox server or Edge Transport server inside or outside your Exchange organization.

What do you need to know before you begin?

- Estimated time to complete each procedure: 5 minutes

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Queues" entry in Mail flow permissions topic.

- To export messages from a delivery queue, the Submission queue, or the Unreachable queue, the messages need to be in the Suspended state. For active, healthy queues, you first suspend the queue so you can then suspend the messages. Messages in the poison message queue are already in the Suspended state. For more information, see Suspend queues and Suspend messages in queues.

- You can't use Queue Viewer in the Exchange Toolbox to export messages. However, you can use Queue Viewer to locate, identify, and suspend the messages before you export them using the Exchange Management Shell. For more information about Queue Viewer, see Queue Viewer. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.

- When you export messages from a queue, you don't remove the messages from the queue. If you resubmit the exported messages by using the Replay directory, you should remove the messages from the queue to avoid duplicate message delivery. For more information, see Remove messages from queues.

- Verify the following information about the target location for the exported message files:

  - The target folder needs to exist before you export any messages, and won't be created for you. If you don't specify the complete path, the files are written to the current Exchange Management Shell working directory.

  - The path can be local to the Exchange server, or it can be a UNC path to a share on a remote server (\server\share).

  - Your account needs to have the **Write** permission in the target folder.

- We use the message's **InternetMessageID** property value for the exported message file names to help ensure uniqueness. The procedures include steps to remove angled brackets (> and <), because they aren't allowed in file names. Also, we use the .eml file name extension so you can easily open the files in Outlook or resubmit the files by using the Replay directory.

- For more information about identity and filters for queues and messages in queues, see the following topics:

  - Find queues and messages in queues in the Exchange Management Shell

  - Queue properties

  - Properties of messages in queues

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to export a specific message from a queue

To export a specific message from a queue, use the following syntax:

```
Export-Message -Identity <MessageIdentity> | AssembleMessage -Path <FilePath>\<FileName>.eml
```

This example takes the following actions on the server named Mailbox01:

1. Suspends the contoso.com delivery queue.

2. Suspends the message in the queue that has the **InternalMessageID** value 1234.

3. Exports a copy of the message to the file D:\contoso Export\export.eml.

```
Suspend-Queue Mailbox01\contoso.com
```

```
Suspend-Message -Identity Mailbox01\contoso.com\1234
```

```
Export-Message -Identity Mailbox01\contoso.com\1234 | AssembleMessage -Path "D:\Contoso Export\export.eml"
```

Use the Exchange Management Shell to export all messages from a queue

To export all messages from a queue, and use the **InternetMessageID** value of each message as the file name, use the following syntax:

```
Get-Message -Queue <QueueIdentity> -ResultSize Unlimited | ForEach-Object {$Temp=
<Path>+$_.InternetMessageID+".eml"; $Temp=$Temp.Replace("<","_"); $Temp=$Temp.Replace(">","_"); Export-Message
$_.Identity | AssembleMessage -Path $Temp}
```

This example takes the following actions on the server named Mailbox01:

1. Suspends the contoso.com delivery queue.

2. Suspends all messages in the queue.

3. Exports copies of the messages to the local folder named D:\Contoso Export.

```
Suspend-Queue Mailbox01\contoso.com
```

```
Get-Queue Mailbox01\contoso.com | Get-Message -ResultSize Unlimited | Suspend-Message
```

```
Get-Message -Queue Mailbox01\Contoso.com -ResultSize Unlimited | ForEach-Object {$Temp="D:\Contoso
Export\"+$_.InternetMessageID+".eml"; $Temp=$Temp.Replace("<","_"); $Temp=$Temp.Replace(">","_"); Export-
Message $_.Identity | AssembleMessage -Path $Temp}
```

Use the Exchange Management Shell to export specific messages from all queues on a server

To export specific messages from all queues on a server, and use the **InternetMessageID** value of each message as the file name, use the following syntax:

```
Get-Message -Filter "<MessageFilter>" [-Server <ServerIdentity>] -ResultSize Unlimited | ForEach-Object
{$Temp=<Path>+$_.InternetMessageID+".eml"; $Temp=$Temp.Replace("<","_"); $Temp=$Temp.Replace(">","_"); Export-
Message $_.Identity | AssembleMessage -Path $Temp}
```

This example takes the following actions on the server named Mailbox01:

1. Suspends all queues on the server.
2. Suspends all messages in all queues on the server from senders in the fabrikam.com domain.
3. Exports copies of the messages to the local folder named D:\Fabrikam Export.

```
Suspend-Queue -Server Mailbox01
```

```
Suspend-Message -Filter "FromAddress -like '*@fabrikam.com'" -Server Mailbox01
```

```
Get-Message -Filter "FromAddress -like '*@fabrikam.com'" -Server Mailbox01 -ResultSize Unlimited | ForEach-
Object {$Temp="D:\Fabrikam Export\"+$_.InternetMessageID+".eml"; $Temp=$Temp.Replace("<","_");
$Temp=$Temp.Replace(">","_"); Export-Message $_.Identity | AssembleMessage -Path $Temp}
```

Use the Exchange Management Shell to export all messages from all queues on a server

To export all messages from all queues on a server, and use the **InternetMessageID** value of each message as the file name, use the following syntax:

```
Get-Message [-Server <ServerIdentity>] -ResultSize Unlimited | ForEach-Object {$Temp=
<Path>+$_.InternetMessageID+".eml"; $Temp=$Temp.Replace("<","_"); $Temp=$Temp.Replace(">","_"); Export-Message
$_.Identity | AssembleMessage -Path $Temp}
```

This example takes the following actions on the server named Mailbox01:

1. Suspends all queues on the server.
2. Suspends all messages in all queues on the server.
3. Exports copies of the messages to the local folder named D:\Mailbox01 Export.

```
Suspend-Queue -Server Mailbox01
```

```
Get-Queue -Server Mailbox01 | Get-Message -ResultSize Unlimited | Suspend-Message
```

```
Get-Message -Server Mailbox01 -ResultSize Unlimited | ForEach-Object {$Temp="D:\Mailbox01
Export\"+$_.InternetMessageID+".eml"; $Temp=$Temp.Replace("<","_"); $Temp=$Temp.Replace(">","_"); Export-
Message $_.Identity | AssembleMessage -Path $Temp}
```

Procedures for messages in queues

8/3/2020 • 9 minutes to read • Edit Online

In Exchange Server, you can use the Queue Viewer in the Exchange Toolbox or the Exchange Management Shell to manage messages in queues. For more information about messages in queues, see Message properties.

This topic describes how to perform the following procedures on messages in queues:

- **Remove messages**: You can remove messages from queues with our without a non-delivery report to the sender (also known as an NDR, delivery status notification, DSN, or bounce message).
- **Suspend messages**: When you suspend a message, you prevent delivery of the message. The message won't leave the queue until you resume the message.
- **Resume messages**: You can resume a message that currently has a status of Suspended. By resuming a message, you enable delivery of the message.
- **Redirect messages**: You can drain messages from all the delivery queues on a Mailbox server, and transfer those messages to another Mailbox server.

For information about exporting messages from queues, see Export messages from queues.

What do you need to know before you begin?

- Estimated time to complete each procedure: 5 minutes
- To find and open the Exchange Toolbox, use one of the following procedures:
  - **Windows 10**: Click **Start** > **All Apps** > **Microsoft Exchange Server <Version>** > **Exchange Toolbox**.
  - **Windows Server 2012 R2 or Windows 8.1**: On the Start screen, open the Apps view by clicking the down arrow near the lower-left corner or swiping up from the middle of the screen. The **Exchange Toolbox** shortcut is in a group named **Microsoft Exchange Server <Version>**.
  - **Windows Server 2012**: Use any of the following methods:
  - On the Start screen, click an empty area, and type Exchange Toolbox.
  - On the desktop or the Start screen, press Windows key + Q. In the Search charm, type Exchange Toolbox.
  - On the desktop or the Start screen, move your cursor to the upper-right corner, or swipe left from the right edge of the screen to show the charms. Click the Search charm, and type Exchange Toolbox.

    When the shortcut appears in the results, you can select it.
- To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- For more information about using filters and identity values in the Exchange Management Shell, see Find queues and messages in queues in the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Queues" entry in the Mail flow permissions topic.

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Remove messages from queues

Notes:

- A message that's being sent to multiple recipients might be located in more than one queue. To remove a message from more than one queue in a single operation, you need to use a filter. For more information, see Properties of messages in queues and Message filtering parameters.

- You can't remove messages from the Submission queue.

Use Queue Viewer to remove messages from queues

1. In the **Exchange Toolbox**, in the **Mail flow tools** section, double-click **Queue Viewer** to open the tool in a new window.

2. In Queue Viewer, click the **Messages** tab. A list of all messages on the server that you're connected to is displayed. To adjust the action to a single queue, click the **Queues** tab, double-click the queue name, and then click the *Server\Queue* tab that appears.

3. Select one or more messages from the list, right-click, and then select **Remove Messages (with NDR)** or **Remove Messages (without NDR)**. A dialog box appears that confirms the selected action and displays, **Do you want to continue?**. Click **Yes**.

4. To remove all messages from a particular queue, click the **Queues** tab. Select a queue, right-click, and then select **Remove Messages (with NDR)** or **Remove Messages (without NDR)**. A dialog box appears that confirms the selected action and displays, **Do you want to continue?**. Click **Yes**.

   > **NOTE**
   >
   > If you're working with a filtered list, the displayed page may not include all items in the filter. In this case, a prompt appears that displays: **This action will affect all items on this page. To expand the scope of this action to include all items in this filter, check the following box before you click OK.**

Use the Exchange Management Shell to remove messages

To remove messages from queues, use the following syntax.

```
Remove-Message <-Identity MessageIdentity | -Filter "MessageFilter"> -WithNDR <$true | $false>
```

This example removes messages in the queues that have a subject of "Win Big" without sending an NDR.

```
Remove-Message -Filter "Subject -eq 'Win Big'" -WithNDR $false
```

This example removes the message with the message ID 3 from the Unreachable queue on server named Mailbox01 and sends an NDR.

```
Remove-Message -Identity Mailbox01\Unreachable\3 -WithNDR $true
```

For more information, see Remove-Message

How do you know this worked?

To verify that you have successfully removed messages from queues, use either of the following procedures:

- In Queue Viewer, select the queue or create a filter to verify the messages no longer exist.

- In the Exchange Management Shell, replace *MessageFilter* with the filter that you used, or *<QueueIdentity>* with the identity of the queue, and run either of the following commands to verify the messages no longer exist:

  ```
  Get-Message -Filter "MessageFilter"
  ```

  Or

  ```
  Get-Message -Queue <QueueIdentity>
  ```

  For more information, see Get-Message.

Suspend messages in queues

Notes:

- A message that's being sent to multiple recipients might be located in more than one queue. To suspend a message in more than one queue in a single operation, you need to use a filter. For more information, see Properties of messages in queues and Message filtering parameters.

- If you suspend a message that's in the act of being transmitted to the next hop, delivery of the message will continue, and the message status will be **PendingSuspend**. If delivery fails, the message will re-enter the queue, and then the message will be suspended.

Use Queue Viewer to suspend messages

1. In the **Exchange Toolbox**, in the **Mail flow tools** section, double-click **Queue Viewer** to open the tool in a new window.

2. In Queue Viewer, click the **Messages** tab. A list of all messages on the server that you're connected to is displayed. To limit the view to a single queue, click the **Queues** tab, double-click the queue name, and then click the *Server**Queue* tab that appears.

3. Select one or more messages, right-click, and then select **Suspend**.

Use the Exchange Management Shell to suspend messages

To suspend messages, use the following syntax:

```
Suspend-Message <-Identity MessageIdentity | -Filter "MessageFilter">
```

This example suspends the message with the message ID 3 in the Unreachable queue on server named Mailbox01.

```
Suspend-Message -Identity Mailbox01\Unreachable\3
```

This example suspends all messages in all queues on the local server that are from any sender in the domain contoso.com.

```
Suspend-Message -Filter "FromAddress -like '*contoso.com'"
```

This example suspends all messages in the delivery queue for contoso.com on the server named Mailbox01.

```
Get-Queue Mailbox01\contoso.com | Get-Message | Suspend-Message
```

This example suspends all messages in all queues on the local server.

```
Get-Queue | Get-Message | Suspend-Message
```

For more information, see Suspend-Message.

How do you know this worked?

To verify that you have successfully suspended messages in queues, use either of the following procedures:

- In Queue Viewer, select the queue or create a filter to verify messages are suspended.
- In the Exchange Management Shell, replace *MessageFilter* with the filter that you used, or *<QueueIdentity>* with the identity of the queue, and run either of the following commands to verify that the messages are suspended:

  ```
  Get-Message -Filter "MessageFilter"
  ```

  Or

  ```
  Get-Message -Queue <QueueIdentity>
  ```

  For more information, see Get-Message.

Resume messages in queues

Notes:

- You can only resume messages that have a status of Suspended.
- The status of the queue that holds the messages affects the delivery of the message. For example, if you resume suspended messages in a queue that has a status of Suspended, the messages can't be delivered until you resume the queue. For more information about resuming queues, see Resume queues.

Use Queue Viewer to resume messages

1. In the **Exchange Toolbox**, in the **Mail flow tools** section, double-click **Queue Viewer** to open the tool in a new window.
2. In Queue Viewer, click the **Messages** tab. A list of all messages on the server that you're connected to is displayed. To adjust the action to focus on a single queue, click the **Queues** tab, double-click the queue name, and then click the *Server\Queue* tab that appears.
3. Click **Create Filter**, and enter your filter expression as follows:
   a. Select **Status** from the message property drop-down list.

b. Select **Equals** from the comparison operator drop-down list.

c. Select **Suspended** from the value drop-down list.

4. Click **Apply Filter**. All messages that have a status of Suspended are displayed.

5. Select one or more messages from the list, right-click, and select **Resume**.

Use the Exchange Management Shell to resume messages

To resume messages, use the following syntax:

```
Resume-Message <-Identity MessageIdentity | -Filter "MessageFilter">
```

This example resumes all messages being sent from any sender in the contoso.com domain.

```
Resume-Message -Filter "FromAddress -like '*contoso.com'"
```

This example resumes the message with the message ID 3 in the Unreachable queue on server named Mailbox01.

```
Resume-Message -Identity Mailbox01\Unreachable\3
```

How do you know this worked?

To verify that you have successfully resumed messages in queues, use either of the following procedures:

- In Queue Viewer, select the queue or create a filter to verify the that messages are no longer suspended.

- In the Exchange Management Shell, replace *MessageFilter* with the filter that you used, or *<QueueIdentity>* with the identity of the queue, and run either of the following commands to verify that the messages are no longer suspended:

  ```
  Get-Message -Filter "MessageFilter"
  ```

  Or

  ```
  Get-Message -Queue <QueueIdentity>
  ```

  For more information, see Get-Message.

If you can't find the messages in any queues on the server, this likely indicates the message was successfully delivered to the next hop.

Redirect messages in queues

Redirecting messages drains all active messages from delivery queues on the source Mailbox server and routes them to the target Mailbox server. The messages are queued for delivery and routed to the next hop.

Notes:

- Only active messages are redirected.

- Shadow queues and messages in the poison message queue aren't redirected.

- The source Mailbox server doesn't accept new messages while messages are being redirected.

- You can only use the Exchange Management Shell to redirect messages.

Use the Exchange Management Shell to redirect messages

To redirect messages, use the following syntax:

```
Redirect-Message -Server <ServerIdentity> -Target <ServerFQDN>
```

This example redirects messages from all delivery queues on the server named Mailbox01 to the server named Mailbox02.contoso.com.

```
Redirect-Message -Server Mailbox01 -Target Mailbox02.contoso.com
```

For more information, see Redirect-Message.

How do you know this worked?

To verify that you have successfully redirected messages in queues, use either of the following procedures:

- In Queue Viewer, verify that the **Message Count** value on delivery queues on the source server is empty or decreasing.

- In the Exchange Management Shell, run the following command to verify that the **MessageCount** property value for the delivery queues on the source server is decreasing or empty.

  ```
  Get-Queue
  ```

Properties of messages in queues

8/3/2020 • 6 minutes to read • Edit Online

Filtering messages in queues by one or more message properties in Exchange Server allows you to quickly locate messages and take action on them. When an email message is sent to multiple recipients, the message might be located in multiple queues on the server. When you filter messages in queues by message properties, you can locate messages across all queues. The following scenarios are examples of how you might use message filtering to manage mail flow:

- The Submission queue on the Mailbox server or Edge Transport server that receives email from the Internet has a high volume of messages that are queued for delivery. Many of the messages have the same subject. Therefore, you suspect that spam is being sent to your organization. You can create a filter to view all the messages that meet the subject criteria. If you determine that the messages are spam, you can select them all and delete them from the delivery queue without sending an NDR.
- A user reports that mail flow is slow. You examine the queues and see that many messages with random subjects appear to be coming from a single domain. You can create a filter to view all the queued messages from that domain. If you determine that the messages are spam, you can select them all and delete them from the queues without sending an NDR.

You can create message filters in Queue Viewer in the Exchange Toolbox, or by using the *Filter* parameter on the message management cmdlets. Note that the message management cmdlets support more filterable properties than Queue Viewer.

For more information about Queue Viewer, see Queue Viewer. For more information about the message management cmdlets, see Procedures for messages in queues and Find queues and messages in queues in the Exchange Management Shell.

Message properties to use as filters

The following table describes the message properties that you can use as filters in Queue Viewer and the Exchange Management Shell.

QUEUE VIEWER	EXCHANGE MANAGEMENT SHELL	COMPARISON OPERATORS	DESCRIPTION
n/a	`AccountForest`	n/a	his property is reserved for internal Microsoft use, and isn't used in on-premises Exchange organizations. In on-premises Exchange, this property is the forest root domain where the mailbox resides (for example, contoso.com).
n/a	`ComponentLatency`	n/a	This property is reserved for internal Microsoft use, and isn't used in on-premises Exchange organizations.
Date Received	`DateReceived`	**Greater Than** (`-gt`) **Greater Than or Equals** (`-ge`) **Less Than** (`-lt`) **Less Than or Equals** (`-le`)	The date/time when the message was placed in the queue.

QUEUE VIEWER	EXCHANGE MANAGEMENT SHELL	COMPARISON OPERATORS	DESCRIPTION

n/a	`DeferReason`	Equals (`-eq`)
		Does not equal (`-ne`)
		Contains (`-like`)
			Indicates why the message was deferred. If the message wasn't deferred, this property has the value `None` . A deferred message is returned to the Submission queue because of transient errors that were encountered during recipient resolution. For more information about deferred messages, see Recipient resolution in Exchange Server. The possible values are: `AD Transient Failure During Content Conversion`
			)
			`AD Transient Failure During Resolve`
			`Agent`
			`Ambiguous Recipient`
			`Config Update`
			`Loop Detected`
			`Marked As Retry Delivery If Rejected`
			`Recipient does not have a mailbox database`
			`Recipient Thread Limit Exceeded`
			`Rerouted By Store Driver`
			`Storage Transient Failure During Content Conversion`
			`Target Site Inbound Mail Disabled`
			`Transient Accepted Domains Load Failure`
			`Transient Attribution Failure`
			`Transient Failure`

n/a	`Directionality`	Equals (`-eq`) Does Not Equal (`-ne`)	Valid values are `Incoming` , `Originating` , and `Undefined` .
Expiration Time	`ExpirationTime`	**Greater Than** (`-gt`) **Greater Than or Equals** (`-ge`) **Less Than** (`-lt`) **Less Than or Equals** (`-le`)	The date/time when the message will expire and be deleted from the queue if the message can't be delivered.
n/a	`ExternalDirectoryOrganizationId`	n/a	This property is reserved for internal Microsoft use, and isn't used in on-premises Exchange organizations. In on-premises Exchange, the value is `00000000-0000-0000-0000-000000000000` .
From Address	`FromAddress`	**Equals** (`-eq`) **Does Not Equal** (`-ne`) **Contains** (`-contains`)	The SMTP address of the sender.
n/a	`Identity`	n/a	The identity of the message in the form of *<Server>*\ *<Queue>*\ *<MessageInteger>*. For more information see Message identity.

QUEUE VIEWER	EXCHANGE MANAGEMENT SHELL	COMPARISON OPERATORS	DESCRIPTION
Internet Message ID	`InternetMessageId`	**Equals** (`-eq`) **Does Not Equal** (`-ne`) **Contains** (`-contains`)	The value of the **Message-Id:** header field in the message header. This value is constant for the lifetime of the message. For messages created in Exchange, the value is in the format `<GUID@ServerFQDN>` , including the angle brackets (< >). For example, `<4867a3d78a50438bad95c0f6d072fca5@mailb` .
Last Error	`LastError`	**Equals** (`-eq`) **Does Not Equal** (`-ne`) **Contains** (`-contains`) **Is Present** **Is Not Present**	The last error that was recorded for a message. For example, `A matching connector cannot be found to route the external recipient` .
n/a	`LockReason`	n/a	This property is reserved for internal Microsoft use, and isn't used in on-premises Exchange organizations.
n/a	`MessageLatency`	Equals (`-eq`) Does not equal (`-ne`) Greater than (`-gt`) Greater than or equal to (`-ge`) Less than (`-lt`) Less than or equal to (`-le`	The amount of time that elapsed between when the message first entered the Submission queue on the server, and when the message was placed in the queue. The value uses the syntax *hh:mm:ss.ff*, where *hh* = hour, *mm* = minute, *ss* = second, and *ff* = fractions of a second.
Message Source Name	`MessageSourceName`	**Equals** (`-eq`) **Does Not Equal** (`-ne`) **Contains** (`-contains`)	The name of the transport component that submitted the message to the queue. For example, if the message came in through a Receive connector, the value is: `SMTP:` *<ConnectorName>*. If the message is a delivery status notification (DSN), the value is `DSN` .
n/a	`OriginalFromAddress`	Equals (`-eq`) Does not equal (`-ne`) Contains (`-like`)	The original sender's email address for any new side effect messages that are created during categorization (for example, journal rules, NDRs, or mail flow rules rules, also known as transport rules).
n/a	`Priority`	Equals (`-eq`) Does not equal (`-ne`)	The priority (importance) of the message that's assigned by the user in Microsoft Outlook or Outlook on the web. Valid values are `Low` , `Normal` , and `High` . For more information, see Priority Queuing.
Queue ID	`Queue`	**Equals** (`-eq`) **Does Not Equal** (`-ne`) **Contains** (`-contains`)	The queue that holds the message. The queue identity uses the syntax *<Server>**<Queue>*. For more information, see Queue identity.

QUEUE VIEWER	EXCHANGE MANAGEMENT SHELL	COMPARISON OPERATORS	DESCRIPTION
n/a	`Recipients`	Contains (`-like`)	An array that contains details about the recipient and the Send connector that will be used, or any errors that were encountered. For example: `{chris@contoso.com;2;2;A matching connector cannot be found to route the external recipient;16;<No Matching Connector>;0}`
n/a	`RetryCount`	Equals (`-eq`) Does not equal (`-ne`) Greater than (`-gt`) Greater than or equal to (`-ge`) Less than (`-lt`) Less than or equal to (`-le`	The number of times that delivery of the message to the destination was tried, either automatically or manually.
SCL	`SCL`	Equals (`-eq`) Does Not Equal (`-ne`) Greater Than (`-gt`) Greater Than or Equals (`-ge`) Less Than (`-lt`) Less Than or Equals (`-le`	The spam confidence level (SCL) rating of the message. Valid SCL entries are integers 0 through 9, or -1 for internal (authenticated) messages. For more information, see Exchange spam confidence level (SCL) thresholds.
Size (KB)	`Size`	Equals (`-eq`) Does Not Equal (`-ne`) Greater Than (`-gt`) Greater Than or Equals (`-ge`) Less Than (`-lt`) Less Than or Equals (`-le`	The size of the message. In Queue Viewer, you need to specify the message size in kilobytes (KB), but in the Exchange Management Shell, you can also specify other sizes, for example, bytes (B) or megabytes (MB).
Source IP	`SourceIP`	Equals (`-eq`) Does Not Equal (`-ne`)	The IPv4 or IPv6 address of the server that submitted the message to the Exchange server that holds the message in the queue. The address could be the IP address of a remote SMTP server, or the IP address of the local Exchange server.
Status	`Status`	Equals (`-eq`) Does Not Equal (`-ne`)	The current message status. Valid values are: Active Locked Pending Remove (`PendingRemove`) Pending Suspend (`PendingSuspend`) Ready Retry Suspended For more information, see Message status.

QUEUE VIEWER	EXCHANGE MANAGEMENT SHELL	COMPARISON OPERATORS	DESCRIPTION
Subject	`Subject`	**Equals** (`-eq`) **Does Not Equal** (`-ne`) **Contains** (`-contains`) **Is Present** **Is Not Present**	The subject of the message (from the **Subject:** header field).
n/a	`TrafficType`	n/a	This property is reserved for internal Microsoft use, and isn't used in on-premises Exchange organizations. In on-premises Exchange, this property is blank or has the value `Email` .
n/a	`TrafficSubType`	n/a	This property is reserved for internal Microsoft use, and isn't used in on-premises Exchange organizations.

Queue Viewer

8/3/2020 • 2 minutes to read • Edit Online

Queue Viewer is part of the Exchange Toolbox that's installed on Mailbox servers and Edge Transport servers in Exchange Server 2016 and Exchange Server 2019. Queue Viewer is a Microsoft Management Console (MMC) snap-in that you can use to view information about and take action on queues and messages in queues. Queue Viewer is useful for troubleshooting mail flow issues and identifying spam.

Queue Viewer is located in the **Mail flow tools** section of the Exchange Toolbox.

To find and open the Exchange Toolbox, use one of the following procedures:

- **Windows 10**: Click **Start** > **All Apps** > **Microsoft Exchange Server <Version>** > **Exchange Toolbox**.
- **Windows Server 2012 R2 or Windows 8.1**: On the Start screen, open the Apps view by clicking the down arrow near the lower-left corner or swiping up from the middle of the screen. The **Exchange Toolbox** shortcut is in a group named **Microsoft Exchange Server <Version>**.
- **Windows Server 2012**: Use any of the following methods:
  - On the Start screen, click an empty area, and type Exchange Toolbox.
  - On the desktop or the Start screen, press Windows key + Q. In the Search charm, type Exchange Toolbox.
  - On the desktop or the Start screen, move your cursor to the upper-right corner, or swipe left from the right edge of the screen to show the charms. Click the Search charm, and type Exchange Toolbox.

  When the shortcut appears in the results, you can select it.

For more information about queues and messages in queues, see Queues and messages in queues.

Topics that contain Queue Viewer procedures

The topics in the following table contain procedures that use Queue Viewer:

TOPIC	DESCRIPTION
Connect to a Server in Queue Viewer	By default, Queue Viewer opens the queue database on the server where you opened Queue Viewer. However, you can connect to a different server.
Set Queue Viewer Options	You can configure the queue and message refresh intervals, and the number of items that are displayed on each page.
View queued message properties in Queue Viewer	Explains how to use Queue Viewer to view messages, and explains the message properties.
Export Lists from Queue Viewer	You can use the **Export List** link in the action pane to export the list of queues or a list of messages for troubleshooting and diagnostics.

TOPIC	DESCRIPTION
Queue properties	Describes the queue properties, and shows the properties that are available in Queue View versus the Exchange Management Shell.
Properties of messages in queues	Describes the message properties, and shows the properties that are available in Queue View versus the Exchange Management Shell.
Procedures for queues	Explains how to view, retry, resubmit, suspend, and resume queues.
Procedures for messages in queues	Explains how to remove, suspend, resume, and redirect messages in queues.

View queued message properties in Queue Viewer

8/3/2020 • 6 minutes to read • Edit Online

You can use the Queue Viewer in the Exchange Toolbox to view queues and the properties of messages in queues. In Exchange Server 2016 and Exchange Server 2019, Queue Viewer is available on Mailbox servers and Edge Transport servers.

For more information about queues, see Queues and messages in queues. For more information about Queue Viewer, see Queue Viewer.

What do you need to know before you begin?

- Estimated time to complete: 5 minutes
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Queues" entry in the Mail flow permissions topic.
- To find and open the Exchange Toolbox, use one of the following procedures:
  - **Windows 10**: Click **Start** > **All Apps** > **Microsoft Exchange Server <Version>** > **Exchange Toolbox**.
  - **Windows Server 2012 R2 or Windows 8.1**: On the Start screen, open the Apps view by clicking the down arrow near the lower-left corner or swiping up from the middle of the screen. The **Exchange Toolbox** shortcut is in a group named **Microsoft Exchange Server <Version>**.
  - **Windows Server 2012**: Use any of the following methods:
    - On the Start screen, click an empty area, and type Exchange Toolbox.
    - On the desktop or the Start screen, press Windows key + Q. In the Search charm, type Exchange Toolbox.
    - On the desktop or the Start screen, move your cursor to the upper-right corner, or swipe left from the right edge of the screen to show the charms. Click the Search charm, and type Exchange Toolbox.

    When the shortcut appears in the results, you can select it.
- You can also use the **Get-Message** cmdlet in the Exchange Management Shell to view additional message properties that aren't visible in Queue Viewer. For more information, see Properties of messages in queues and Find queues and messages in queues in the Exchange Management Shell.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use Queue Viewer to view the properties of a message

1. In the **Exchange Toolbox**, in the **Mail flow tools** section, double-click **Queue Viewer** to open the tool in

a new window.

2. In Queue Viewer, select the **Messages** tab to see the list of messages that are currently queued for delivery in your organization. The list of messages displays the following information:

   - **From Address**: The sender's email address.

   - **Status**: A message can have one of the following status values:

   - **Active**: If the message is in a delivery queue, the message is being delivered to its destination. If the message is in the Submission queue, the message is being processed by the categorizer.

   - **Pending Remove**: The message was deleted by the administrator, but was already being delivered. The message will be deleted if the delivery ends in an error that causes the message to re-enter the queue. Otherwise, delivery will continue.

   - **Pending Suspend**: The message was suspended by the administrator, but was already being delivered. The message will be suspended if the delivery ends in an error that causes the message to re-enter the queue. Otherwise, delivery will continue.

   - **Ready**: The message is waiting in the queue, and is ready to be processed.

   - **Retry**: The queue's last connection attempt failed. The message is waiting for the next queue retry.

   - **Suspended**: The message was suspended by the administrator. For more information, see Suspend messages in queues.

   - **Size (KB)**: The size of the message rounded up to the nearest kilobyte (KB).

   - **SCL**: The spam confidence level (SCL) rating of the message. Valid SCL entries are integers 0 through 9, or -1 for internal (authenticated) messages. For more information, see Exchange spam confidence level (SCL) thresholds.

   - **Queue ID**: The queue that holds the message. The queue identity uses the syntax *<Server>*\ *<Queue>*, where *<Queue>* is one of the following values:

   - **Persistent queue name**

   - **Poison**: Isolates messages that contain errors and are determined to be harmful to Exchange after a server or service failure. The messages may be genuinely harmful in their content and format, or the messages might have been the victims of a poorly written transport agent or a software bug that crashed the Exchange server while it was processing the otherwise valid messages.

   - **Submission**: Holds messages that have been accepted by the Transport service, but haven't been processed. Messages in the Submission queue are either waiting to be processed, or are actively being processed.

   - **Unreachable**: Contains messages that can't be routed to their destinations. Typically, an unreachable destination is caused by configuration changes that have modified the routing path for delivery. Regardless of destination, all messages that have unreachable recipients reside in this queue.

   - **Delivery queue name**: The value of the **NextHopDomain** property of the queue, which is effectively the name of the queue. For example, a domain name, Active Directory site name, or database availability group (DAG) name. For more information, see NextHopSolutionKey.

   - **Message Source Name**: The Exchange component that submitted the message to the queue. For example, `SMTP:Default <ServerName>`.

   - **Subject**: The subject of the message.

   - **Last Error**: The last error that was encountered for the message.

For example, if you didn't create a Send connector to deliver Internet mail, messages that are addressed to external recipients will go to the Unreachable queue, and the **Last Error** value for the message will be: `A matching connector cannot be found to route the external recipient`. For more information about creating a Send connector, see Create a Send connector to send mail to the Internet.

For more information about SMTP error codes, see DSNs and NDRs in Exchange Server.

3. When you right-click a message and select **Properties**, additional details are available on the **General** and **Recipient Information** tabs.

   - The **General** tab displays the same **Subject**, **From Address**, **Status**, **Size**, **Message Source Name**, **SCL**, and **Last Error** values that are shown in the list of messages. The following additional properties are also displayed on the **General** tab:

   - **Identity**: The identity of the message. The message identity uses the syntax *<Server>*\ *<Queue>*\ *<MessageInteger>*, where *<Queue>* is the identity of the queue as described in the **Queue ID** property, and *<MessageInteger>* is the unique integer value of the message that's displayed in the **Identity** property of the **Get-Message** cmdlet.

   - **Internet Message ID**: The value of the **Message-Id:** header field in the message header. This value is constant for the lifetime of the message. For messages created in Exchange, the value is in the format `<GUID@ServerFQDN>`, including the angle brackets (< >). For example, `<4867a3d78a50438bad95c0f6d072fca5@mailbox01.contoso.com>`.

   - **Source IP**: The IPv4 or IPv6 address of the internal Exchange server or external messaging server that submitted the message.

   - **Date Received**: The date-time when the message entered the queue.

   - **Expiration Time**: The date-time when the message will expire and will be deleted from the queue if the message can't be delivered.

   - **Recipients**: The recipients in the message, with any corresponding error messages. To see the status and error messages for each recipient, go to the **Recipient Information** tab.

   - The **Recipient Information** tab displays the **Address**, **Status**, and **Last Error** values for each recipient in the message. The **Status** value for a recipient can be **Complete**, **Ready**, or **Retry**.

Find queues and messages in queues in the Exchange Management Shell

8/3/2020 • 17 minutes to read • Edit Online

As in previous versions of Exchange, you can use the Exchange Management Shell in Exchange Server to view information about queues and messages, and use that information to take action on queues and messages. Typically, an active Exchange contains a large number of queues and messages to be delivered, so it's important to understand how to identify the queues or messages that you want to manage.

Note that you can also use Queue Viewer in the Exchange Toolbox to manage queues and messages in queues. However, the queue and message viewing cmdlets in the Exchange Management Shell support more filterable properties and filter options than Queue Viewer. For more information about using Queue Viewer, see Queue Viewer.

Also remember that queues exist on Mailbox servers and Edge Transport servers (the Transport service). For more information about queues and messages in queues, see Queues and messages in queues.

Queue filtering parameters

The following table summarizes the filtering parameters that are available on the queue management cmdlets.

CMDLET	FILTERING PARAMETERS	COMMENTS
Get-Queue	*Exclude* *Filter* *Identity* *Include* *Server*	You can use the *Include* and *Exclude* parameters with the other filtering parameters in the same command. You can't use the *Identity* and *Filter* parameters in the same command. The *Server* parameter specifies the server where you want to run the command. You can't use the *Server* and *Identity* parameters in the same command, but you can use the *Server* parameter with the other filtering parameters in the same command.
Resume-Queue Retry-Queue Suspend-Queue	*Identity* *Filter* *Server*	You can't use the *Identity* parameter with the other filtering parameters in the same command. The *Server* parameter specifies the server where you want to run the command. You can use the *Server* and *Filter* parameters in the same command.
Get-QueueDigest	*Dag* *Filter* *Forest* *Server* *Site*	You need to use one of the *Dag*, *Site*, *Server*, or *Forest* parameters, but you can't use any of them together in the same command. You can use the *Filter* parameter with any of the other filtering parameters.

Queue identity

The *Identity* parameter uses the basic syntax *<Server>\ <Queue>*. Typically, this value uniquely identifies the queue, so you can't use other filtering parameters with the *Identity* parameter. The exception is the **Get-Queue** cmdlet, where you can use the *Include* and *Exclude* parameters with the *Identity* parameter.

The following table explains the *Identity* parameter syntax on the queue management cmdlets.

IDENTITY PARAMETER VALUE	DESCRIPTION
`<Server>\<PersistentQueueName>` or `<PersistentQueueName>`	A persistent queue on the specified or local server. `<PersistentQueueName>` is `Submission`, `Unreachable`, or `Poison`. For more information about persistent queues, see Types of queues.
`<Server>\<NextHopDomain>` or `<NextHopDomain>`	A delivery queue on the specified or local server. `<NextHopDomain>` is the name of the queue from the value of the **NextHopDomain** property of the queue. For example, the address space of a Send connector, the name of an Active Directory site, or the name of a DAG. For more information, see NextHopSolutionKey.
`<Server>\<QueueInteger>` or `<QueueInteger>`	A delivery queue on the specified or local server. `<QueueInteger>` is the unique integer value that's assigned to a delivery queue or a shadow queue in the queue database. However, you need to run the **Get-Queue** cmdlet to find this value in the **Identity** or **QueueIdentity** properties.
`<Server>\Shadow\<QueueInteger>` or `Shadow\<QueueInteger>`	A shadow queue on the specified or local server. For more information about shadow queues and shadow redundancy, see Shadow redundancy in Exchange Server.

`<Server>\*` or `*`	All queues on the specified or local server.
	Note: *Identity* is a positional parameter, which means you can specify the value without specifying the `-Identity` qualifier. For example, the following commands produce the same result:
	`Get-Queue -Identity *`
	`Get-Queue *`
	`Get-Queue`

Filter parameter on queue cmdlets

You can use the *Filter* parameter on all of the queue management cmdlets to identify one or more queues based on the properties of the queues. The *Filter* parameter creates an OPath filter with comparison operators to restrict the command to queues that meet the filter criteria. You can use the logical operator `-and` to specify multiple conditions for the match. Here's a generic example of the syntax:

```
Get-Queue -Filter "<Property1> -<ComparisonOperator> '<Value1>' -and <Property2> -<ComparisonOperator>
'<Value2>'..."
```

For a complete list of queue properties you can use with the *Filter* parameter, see Queue properties.

For a list of comparison operators you can use with the *Filter* parameter, see the Comparison operators to use when filtering queues or messages section in this topic.

For examples of procedures that use the *Filter* parameter to view and manage queues, see Procedures for queues.

Include and Exclude parameters on Get-Queue

You can use the *Include* and *Exclude* parameters on the **Get-Queue** cmdlet by themselves, with each othe , or with the other filtering parameters to fine-tune your results. For example, you can:

- Exclude empty queues.
- Exclude queues to external destinations.
- Include queues that have a specific value of **DeliveryType**.

The *Include* and *Exclude* parameters use the following queue properties to filter queues:

VALUE	DESCRIPTION	EXAMPLE
`DeliveryType`	Includes or excludes queues based on the **DeliveryType** property that defines how the message will be transmitted to the next hop. The valid values are described in NextHopSolutionKey. You can specify multiple values separated by commas.	Returns all delivery queues on the local server where the next hop is a Send connector that's hosted on the local server and is configured for smart host routing. `Get-Queue -Include SmartHostConnectorDelivery`
`Empty`	Includes or excludes empty queues. Empty queues have the value `0` in the **MessageCount** property.	Returns all queues on the local server that contain messages. `Get-Queue -Exclude Empty`
`External`	Includes or excludes queues that have the value `External` in the **NextHopCategory** property. External queues always have one of the following values for **DeliveryType**: • `DeliveryAgent` • `DnsConnectorDelivery` • `NonSmtpGatewayDelivery` • `SmartHostConnectorDelivery` For more information, see NextHopSolutionKey.	Returns all internal queues on the local server. `Get-Queue -Exclude External`
`Internal`	This value includes or excludes queues that have the value `Internal` in the **NextHopCategory** property. Note that a message for an external recipient may require multiple internal hops before it reaches a gateway server where it's delivered externally.	Returns all internal queues on the local server. `Get-Queue -Include Internal`

Note that you can duplicate the functionality of the *Include* and *Exclude* parameters by using the *Filter* parameter. For example, the following commands produce the same result:

- `Get-Queue -Exclude Empty`
- `Get-Queue -Filter "MessageCount -gt 0"`

However, as you can see, the syntax of the *Include* and *Exclude* parameters is simpler and easier to remember.

Get-QueueDigest

The **Get-QueueDigest** cmdlet allows you to view information about some or all of the queues in your organization by using a single command. Specifically, the **Get-QueueDigest** cmdlet allows you to view information about queues based on their location on servers, in DAGs, in Active Directory sites, or in the whole

Active Directory forest.

Note that queues on a subscribed Edge Transport server aren't included in the results. Also, **Get-QueueDigest** is available on an Edge Transport server, but the results are restricted to local queues on the Edge Transport server.

NOTE

By default, the **Get-QueueDigest** cmdlet displays delivery queues that contain ten or more messages, and the results are between one and two minutes old. For instructions on how to change these default values, see Configure Get-QueueDigest.

The following table describes the filtering and sorting parameters that are available on the **Get-QueueDigest** cmdlet.

PARAMETER	DESCRIPTION
Dag, *Server*, or *Site*	These parameters are mutually exclusive (can't be used in the same command), and set the scope for the cmdlet. You need to specify one of these parameters or the *Forest* switch. Typically, you would use the name of the server, DAG or Active Directory site, but you can use any value that uniquely identifies the server, DAG, or site. You can specify multiple servers, DAGs, or sites separated by commas.
Forest	This switch is required if you aren't using the *Dag*, *Server*, or *Site* parameters. You don't specify a value with this switch. By using this switch, you get queues from all Exchange Mailbox servers in the local Active Directory forest. You can't use this switch to view queues in remote Active Directory forests.
DetailsLevel	`Normal` is the default value. The following properties are returned in the results: • **QueueIdentity** • **ServerIdentity** • **MessageCount** `Verbose` returns the following additional properties in the results: • **DeferredMessageCount** • **LockedMessageCount*** • **IncomingRate** • **OutgoingRate** • **Velocity** • **NextHopDomain** • **NextHopCategory** • **NextHopConnector** • **DeliveryType** • **Status** • **RiskLevel*** • **OutboundIPPool*** • **LastError** • **TlsDomain** `None` omits the queue name from the **Details** column in the results. * These properties are reserved for internal Microsoft use, and aren't used in on-premises Exchange organizations. For more information about all properties in this list, see Queue properties.

PARAMETER	DESCRIPTION
Filter	Filter queues based on the queue properties as described in the Filter parameter on queue cmdlets section. You can use any of the filterable queue properties as described in the Queue properties topic.
GroupBy	Groups the queue results. You can group the results by one of the following properties: • **DeliveryType** • **LastError** • **NextHopCategory** • **NextHopDomain** • **NextHopKey** • **Status** • **ServerName** By default, the results are grouped by **NextHopDomain**. For information about these queue properties, see Queue properties.
ResultSize	Limits the queue results to the value you specify. The queues are sorted in descending order based on the number of messages in the queue, and grouped by the value specified by the *GroupBy* parameter. The default value is 1000. This means that by default, the command displays the top 1000 queues grouped by **NextHopDomain**, and sorted by the queues containing the most messages to the queues containing the least messages.
Timeout	The parameter specifies the number of seconds before the operation times out. The default value is `00:00:10` or 10 seconds.

This example returns all non-empty external queues on the servers named Mailbox01, Mailbox02, and Mailbox03.

```
Get-QueueDigest -Server Mailbox01,Mailbox02,Mailbox03 -Include External -Exclude Empty
```

Message filtering parameters

The following table summarizes the filtering parameters that are available on the message management cmdlets.

CMDLET	FILTERING PARAMETERS	COMMENTS
Get-Message	*Filter* *Identity* *Queue* *Server*	You can't use the *Filter*, *Identity*, or *Queue* parameters in the same command. The *Server* parameter specifies the server where you want to run the command. You can use the *Server* and *Filter* parameters in the same command.

CMDLET	FILTERING PARAMETERS	COMMENTS
Remove-Message **Resume-Message** **Suspend-Message**	*Filter* *Identity* *Server*	You need to use either the *Identity* parameter or the *Filter* parameter, but you can't use them both in the same command. The *Server* parameter specifies the server where you want to run the command. You can use the *Server* and *Filter* parameters in the same command.
Redirect-Message	*Server*	This cmdlet drains active messages from all delivery queues on the specified server, so *Server* is the only filtering parameter that's available. For more information, see Redirect messages in queues.
Export-Message	*Identity*	This parameter isn't really a filter, because it uniquely identifies the message. To identify multiple messages for this cmdlet, use **Get-Message** and pipe the results to **Export-Message**. For more information and examples, see Export messages from queues.

Message identity

The *Identity* parameter on the message management cmdlets uniquely identifies a message in one or more queues, so you can't use any other message filtering parameters. The *Identity* parameter uses the basic syntax `<Server>\<Queue>\<MessageInteger>`.

The following table describes the syntax you can use with *Identity* parameter on the message management cmdlets.

IDENTITY PARAMETER VALUE	DESCRIPTION
`<Server>\<Queue>\<MessageInteger>` or `<Queue>\<MessageInteger>`	A message in a specific queue on the specified or local server. `<Queue>` is the identity of the queue as described in the Queue identity section: • **Persistent queue name** • **Delivery queue name** • **Queue integer** • **Shadow queue identity** `<MessageInteger>` is the unique integer value that's assigned to the message when it first enters the queue database on the server. If the message is sent to multiple recipients that require multiple queues, all copies of the message in all queues in the queue database have the same integer value. However, you need to run the **Get-Message** cmdlet to find this value in the **Identity** or **MessageIdentity** properties.
`<Server>\*\<MessageInteger>` or `*\<MessageInteger>` or `<MessageInteger>`	All copies of the message in all queues in the queue database on the specified or local server.

Filter parameter on message cmdlets

You can use the *Filter* parameter with the **Get-Message**, **Remove-Message**, **Resume-Message**, and

Suspend-Message cmdlets to identify one or more messages based on the properties of the messages. The *Filter* parameter creates an OPath filter with comparison operators to restrict the command to messages that meet the filter criteria. You can use the logical operator `-and` to specify multiple conditions for the match. Here's a generic example of the syntax:

```
Get-Message -Filter "<Property1> -<ComparisonOperator> '<Value1>' -and <Property2> -<ComparisonOperator>
'<Value2>'..."
```

For a complete list of message properties you can use with the *Filter* parameter, see Message properties).

For a list of comparison operators you can use with the *Filter* parameter, see the Comparison operators to use when filtering queues or messages section in this topic.

For examples of procedures that use the *Filter* parameter to view and manage messages, see Procedures for messages in queues.

Queue parameter

The *Queue* parameter is available only on the **Get-Message** cmdlet. You can use this parameter to get all messages in a specific queue, or all messages from multiple queues by using the wildcard character (*). When you use the *Queue* parameter, use the queue identity format `<Server>\<Queue>` as described in the Queue identity section in this topic.

Comparison operators to use when filtering queues or messages

When you create a queue or message filter expression by using the *Filter* parameter, you need to include an comparison operator for the property value to match. The comparison operators that you can use, and how each operator functions are described in the following table. For all operators, the values compared aren't case sensitive.

OPERATOR	FUNCTION	CODE EXAMPLE

`-eq`	Exact match of the specified value.	Show all queues that have a status of Retry:
		`Get-Queue -Filter "Status -eq 'Retry'"`
		Show all messages that have a status of Retry:
		`Get-Message -Filter "Status -eq 'Retry'"`

`-ne`	Does not match the specified value.	Show all queues that don't have a status of Active:
		`Get-Queue -Filter "Status -ne 'Active'"`
		Show all messages that don't have a status of Active:
		`Get-Message -Filter "Status -ne 'Active'"`

`-gt`	Greater than the specified integer or date/time value.	Show queues that currently contain more than 1,000 messages:
		`Get-Queue -Filter "MessageCount -gt 1000"`
		Show messages that currently have a retry count that's more than 3:
		`Get-Message -Filter "RetryCount -gt 3"`

OPERATOR	FUNCTION	CODE EXAMPLE

`-ge`	Greater than or equal to the specified integer or date/time value.	Show queues that currently contain 1,000 or more messages:
		`Get-Queue -Filter "MessageCount -ge 1000"`
		Show messages that currently have a retry count that's 3 or more:
		`Get-Message -Filter "RetryCount -ge 3"`

`-lt`	Less than the specified integer or date/time value.	Show queues that currently contain less than 1,000 messages:
		`Get-Queue -Filter "MessageCount -lt 1000"`
		Show messages that have an SCL that's less than 6:
		`Get-Message -Filter "SCL -lt 6"`

`-le`	Less than or equal to the specified integer or date/time value.	Show queues that currently contain 1,000 or fewer messages:
		`Get-Queue -Filter "MessageCount -le 1000"`
		Show messages that have an SCL that's 6 or less:
		`Get-Message -Filter "SCL -le 6"`

`-like`	Contains the specified text. You need to include the wildcard character (*) in the text string.	Show queues that have a destination to any SMTP domain that ends in Contoso.com:
		`Get-Queue -Filter "Identity -like '*contoso.com'"`
		Show messages that have a subject that contains the text "payday loan":
		`Get-Message -Filter "Subject -like '*payday loan*'"`

You can specify a filter that evaluates multiple expressions by using the logical operator `-and`. The queues or messages must match all of the filter conditions to be included in the results.

This example displays a list of queues that have a destination to any SMTP domain name that ends in Contoso.com and that currently contain more than 500 messages.

```
Get-Queue -Filter "Identity -like '*contoso.com*' -and MessageCount -gt 500"
```

This example displays a list of messages that are sent from any email address in the contoso.com domain that have an SCL value that's greater than 5.

```
Get-Message -Filter "FromAddress -like '*Contoso.com*' -and SCL -gt 5"
```

Advanced paging parameters

When you use the Exchange Management Shell to view queues and messages in queues, your query retrieves one page of information at a time. The advanced paging parameters control the size of the results, and the order that the results are displayed in. All advanced paging parameters are optional and can be used with or without other filtering parameters on the **Get-Queue** and **Get-Message** cmdlets. If you don't specify any advanced paging parameters, the query returns the results in ascending order of identity.

By default, when you specify a sort order, the **Identity** property is always included and sorted in ascending order, because the other available queue or message properties aren't unique.

You can use the *BookmarkIndex* and *BookmarkObject* parameters to mark a position in the sorted results. If the bookmark object no longer exists when you retrieve the next page of results, the results start with the closest item to the bookmark, which depends on the sort order that you specify.

The advanced paging parameters are described in the following table.

PARAMETER	DESCRIPTION
BookmarkIndex	Specifies the position in the results where the displayed results start. The value of this parameter is a 1-based index in the total results. If the value is less than or equal to zero, the first complete page of results is returned. If the value is set to `Int.MaxValue`, the last complete page of results is returned. You can't use this parameter with the *BookmarkObject* parameter.
BookmarkObject	Specifies the object in the results where the displayed results start. If you specify a bookmark object, that object is used as the point to start the search. The rows before or after that object (depending on the value of the *SearchForward* parameter) are retrieved. You can't use this parameter with the *BookmarkIndex* parameter.
IncludeBookmark	Specifies whether to include the bookmark object in the results. Valid values are: `$true`: The bookmark object is included in the results. This is the default value. `$false`: The bookmark object isn't included in the results. Use this value when you run a query for a limited result size, and then specify the last item as the bookmark for the next query. This prevents the bookmark object from being included in both results.
ResultSize	Specifies the number of results to display per page. If you don't specify a value, the default result size of 1,000 objects is used. Exchange limits the results to 250,000.
ReturnPageInfo	This is a hidden parameter. It returns information about the total number of results and the index of the first object of the current page. The default value is `$false`.
SearchForward	Specifies the direction of the search. **Bookmark specified**: Search forward or backward in the results relative to the bookmark index or object. **No bookmark specified**: Search forward or backward in the results from the first or last item in the results. Valid values are: `$true`: Search forward from the first item in the results, or from the specified bookmark. If there are no results beyond the bookmark, the query returns the last full page of results. This is the default value. `$false`: Search backward from the last item in the results, or from the specified bookmark. If there is less than a full page of results beyond the bookmark, the query returns the first full page of results.

PARAMETER	DESCRIPTION
SortOrder	Specifies the message properties that control the sort order of the results. The order that the properties are specified indicates a descending order of precedence (the results are sorted by the first property, then those results are sorted by the second property, and son on). This parameter uses the syntax: `<+|-><Property1>,<+|-><Property2>...`, where `+` sorts the property in ascending order, and `-` sorts the property in descending order. If you don't use this parameter, the results are sorted by the **Identity** property in ascending order.

This example shows how to use the advanced paging parameters in a query. The command returns the first 500 messages on the specified server. The results are sorted first in ascending order by sender address, and then in descending order by message size.

```
Get-Message -Server mailbox01.contoso.com -ResultSize 500 -SortOrder +FromAddress,-Size
```

This example returns the first 500 messages on the specified server in the specified sort order, sets a bookmark object, excludes the bookmark object from the results, and retrieves the next 500 messages in the same sort order.

1. Run the following command to retrieve the first page of results.

    ```
    $Results=Get-Message -Server mailbox01.contoso.com -ResultSize 500 -SortOrder +FromAddress,-Size
    ```

2. To set the bookmark object, run the following command to save the last element of the first page to a variable.

    ```
    $Temp=$Results[$results.length-1]
    ```

3. To retrieve the next 500 objects on the specified server, and to exclude the bookmark object, run the following command.

    ```
    Get-Message -Server mailbox01.contoso.com -BookmarkObject:$Temp -IncludeBookmark $false -ResultSize
    500 -SortOrder +FromAddress,-Size
    ```

Change the location of the queue database

8/3/2020 • 5 minutes to read • Edit Online

Exchange Server uses an Extensible Storage Engine (ESE) database for queue message storage. All the different queues are stored in a single ESE database. Queues exist on Exchange Mailbox servers and Edge Transport servers. For more information about queues, see Queues and messages in queues.

The location of the queue database and the queue database transaction logs is controlled by keys in the `%ExchangeInstallPath%Bin\EdgeTransport.exe.config` XML application configuration file. This file is associated with the Exchange Transport service. The following table explains each key in more detail.

KEY	DESCRIPTION
QueueDatabasePath	Specifies the location of the queue database files. The files are: • Mail.que • Trn.chk The default location is `%ExchangeInstallPath%TransportRoles\data\Queue` .
QueueDatabaseLoggingPath	Specifies the location of the queue database transaction log files. The files are: • Trn.log • Trntmp.log • Trn *nnn*.log • Trnres00001.jrs • Trnres00002.jrs • Temp.edb Note that Temp.edb is used to verify the queue database schema when the Exchange Transport service starts. Although Temp.edb isn't a transaction log file, it's kept in the same location as the transaction log files. The default location is `%ExchangeInstallPath%TransportRoles\data\Queue` .

What do you need to know before you begin?

- Estimated time to complete: 15 minutes.
- Exchange permissions don't apply to the procedures in this topic. These procedures are performed in the operating system of the Exchange server.
- When you stop or restart the Exchange Transport service, mail flow on the server is interrupted.
- When you change the location of the queue database or the transaction logs, the existing queue database and transaction log files aren't moved. A new queue database and new transaction logs are created at the new location. The old files are left at the old location, but they're no longer used. If you want to reuse the old queue database or transaction log files at the new location, you need to move the files to the new location while the Exchange Transport service is stopped.
- The folder for the queue database and transaction logs needs the following permissions:
  - Network Service: Full Control
  - System: Full Control

  - Administrators: Full Control

    If the folder doesn't exist, but the parent folder has these permissions, the new folder is created automatically.

- Any customized Exchange or Internet Information Server (IIS) settings that you made in Exchange XML application configuration files on the Exchange server (for example, web.config files or the EdgeTransport.exe.config file) **will be overwritten** when you install an Exchange CU. Be sure save this information so you can easily re-apply the settings after the install. After you install the Exchange CU, you need to re-configure these settings.

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

- Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Command Prompt to create a new queue database and transaction logs in a new location

1. Create the folder where you want to keep the queue database and transaction logs. Make sure that the correct permissions are applied to the folder.

2. In a Command prompt window, open the EdgeTransport.exe.config file in Notepad by running the following command:

```
Notepad %ExchangeInstallPath%Bin\EdgeTransport.exe.config
```

3. Find and modify the following keys in the `<appSettings>` section.

```
<add key="QueueDatabasePath" value="<LocalPath>" />
<add key="QueueDatabaseLoggingPath" value="<LocalPath>" />
```

   For example, to create a new queue database and transaction logs in D:\Queue\QueueDB, use the following values:

```
<add key="QueueDatabasePath" value="D:\Queue\QueueDB" />
<add key="QueueDatabaseLoggingPath" value="D:\Queue\QueueDB" />
```

   When you're finished, save and close the EdgeTransport.exe.config file.

4. Restart the Exchange Transport service by running the following command:

```
net stop MSExchangeTransport && net start MSExchangeTransport
```

How do you know this worked?

To verify that you've successfully created a new queue database and new transaction logs in the new location, do these steps:

1. Verify the new database files Mail.que and Trn.chk exist at the new location.

2. Verify the new transaction log files Trn.log, Trntmp.log, Trnres00001.jrs, Trnres00002.jrs, and Temp.edb files exist at the new location.

3. If you can delete the old queue database and transaction log files from the old location after the Exchange Transport service has started, the old queue database is no longer being used.

Use the Command Prompt to move the existing queue database and transaction logs to a new location

> **NOTE**
>
> There is also a script to move the queue database and transaction logs, it can be found in the %ExchangeInstallPath%Scripts folder and it's called Move-TransportDatabase.ps1. You have to specify the following parameters: queueDatabasePath, queueDatabaseLoggingPath, iPFilterDatabasePath, iPFilterDatabaseLoggingPath and temporaryStoragePath.

Although you'll need to move the existing queue database to preserve any undelivered messages in it, you typically don't need to move the existing transaction logs because:

- An ordinary shutdown of the Exchange Transport service writes all uncommitted transaction log entries to the queue database.
- Circular logging is used, so transaction logs that contain previously committed database changes aren't preserved.

1. Create the folder where you want to keep the queue database and transaction logs. Make sure that the correct permissions are applied to the folder.

2. In a Command prompt window, open the EdgeTransport.exe.config file in Notepad by running the following command:

   ```
   Notepad %ExchangeInstallPath%Bin\EdgeTransport.exe.config
   ```

3. Find and modify the following keys in the `<appSettings>` section:

   ```
   <add key="QueueDatabasePath" value="<LocalPath>" />
   <add key="QueueDatabaseLoggingPath" value="<LocalPath>" />
   ```

   For example, to change the location of the queue database and transaction logs to D:\Queue\QueueDB, use the following values:

   ```
   <add key="QueueDatabasePath" value="D:\Queue\QueueDB" />
   <add key="QueueDatabaseLoggingPath" value="D:\Queue\QueueDB" />
   ```

   When you're finished, save and close the EdgeTransport.exe.config file.

4. Stop the Exchange Transport service by running the following command:

   ```
   net stop MSExchangeTransport
   ```

5. Move the existing database files Mail.que and Trn.chk from the old location to the new location.

6. Move the existing transaction log files Trn.log, Trntmp.log, Trn *nnnnn*.log, Trnres00001.jrs, Trnres00002.jrs, and Temp.edb from the old location to the new location.

7. Start the Exchange Transport service by running the following command:

```
net start MSExchangeTransport
```

How do you know this worked?

To verify that you've successfully moved the existing queue database and transaction logs to the new location, do these steps:

1. Verify the queue database files Mail.que and Trn.chk exist in the new location.
2. Verify the transaction log files Trn.log, Trntmp.log, Trnres00001.jrs, Trnres00002.jrs, and Temp.edb files exist in the new location.
3. Verify there are no queue database or transaction log files in the old location.

Message retry, resubmit, and expiration intervals

8/3/2020 • 7 minutes to read • Edit Online

In Exchange Server, messages that can't be successfully delivered are subject to various retry, resubmit, and expiration deadlines based on the message's source and destination. *Retry* is a renewed connection attempt with the destination. *Resubmit* is the act of sending messages back to the Submission queue for the categorizer to reprocess. The message *expires* after all delivery efforts have failed over a specified period of time. After a message expires, the sender is notified of the delivery failure, and the message is deleted from the queue.

In all three cases of retry, resubmit, or expire, you can manually intervene before the automatic actions are performed on the messages.

For instructions on how to configure these intervals, see Configure message retry, resubmit, and expiration intervals.

Configuration options for message retry

When a the Transport service on a Mailbox server or an Edge Transport server can't connect to the next hop, the queue is put in a status of Retry. Connection attempts continue until the queue expires or a connection is made.

Configuration options for automatic message retry in the EdgeTransport.exe.config file

The automatic message retry interval settings that are available in the `%ExchangeInstallPath%Bin\EdgeTransport.exe.config` XML application configuration file are described in the following table.

> **NOTE**
>
> Any customized Exchange or Internet Information Server (IIS) settings that you made in Exchange XML application configuration files on the Exchange server (for example, web.config files or the EdgeTransport.exe.config file) **will be overwritten** when you install an Exchange CU. Be sure save this information so you can easily re-apply the settings after the install. After you install the Exchange CU, you need to re-configure these settings.

AUTOMATIC MESSAGE RETRY KEY NAME	DEFAULT VALUE	DESCRIPTION
MailboxDeliveryQueueRetryInterval	`00:05:00` (5 minutes)	How frequently the queues try to connect to the Mailbox Transport Delivery service for a destination mailbox database that can't be successfully reached. To specify a value, enter it as a time span: `dd.hh:mm:ss` where `dd` = days, `hh` = hours, `mm` = minutes, and `ss` = seconds. A valid value is a timespan from `00:00:01` (one second) through `1.00:00:00` (one day).

AUTOMATIC MESSAGE RETRY KEY NAME	DEFAULT VALUE	DESCRIPTION
QueueGlitchRetryCount	4	The number of connection attempts that are immediately tried when a transport server has trouble connecting with the destination server. Such connection problems are typically caused by very brief network outages. A valid value is an integer from 0 through 15. Typically, you don't need to modify this key unless the network is unreliable and continues to experience many accidentally dropped connections.
QueueGlitchRetryInterval	`00:01:00` (1 minute)	The connection interval between each connection attempt that's specified by the *QueueGlitchRetryCount* key. Typically, you don't need to modify this parameter unless the network is unreliable and continues to experience many accidentally dropped connections.

Configuration options for automatic message retry in the Exchange admin center and the Exchange Management Shell

The automatic message retry interval settings that are available in the Exchange admin center (EAC) and the Exchange Management Shell are described in the following table.

AUTOMATIC MESSAGE RETRY SETTING	DEFAULT VALUE	EXCHANGE MANAGEMENT SHELL CONFIGURATION	EXCHANGE ADMIN CENTER CONFIGURATION ON MAILBOX SERVERS
Message retry interval: The retry interval for individual messages that have a status of Retry.	15 minutes (`00:15:00`) We recommend that you don't modify the default value unless you're directed to do so by Microsoft Customer Service and Support, or specific product documentation.	Cmdlet: **Set-TransportService** cmdlet Parameter: *MessageRetryInterval*	n/a
Outbound connection failure retry interval: The retry interval for outbound connection attempts that have previously failed. The previously failed connection attempts are controlled by the transient failure retry count and interval values.	Transport service on Mailbox servers: 10 minutes (`00:10:00`) Edge Transport Servers: 30 minutes (`00:30:00`)	Cmdlet: **Set-TransportService** Parameter: *OutboundConnectionFailureRetryInterval*	**Servers** > select server > **Edit** (✎) > **Transport limits** > **Retry** section > **Outbound connection failure retry interval (seconds)**

AUTOMATIC MESSAGE RETRY SETTING	DEFAULT VALUE	EXCHANGE MANAGEMENT SHELL CONFIGURATION	EXCHANGE ADMIN CENTER CONFIGURATION ON MAILBOX SERVERS
Transient failure retry count: The number of connection attempts that are tried after the queue glitch retry count and interval values have failed. These failures can be caused by server restarts or cached DNS lookup failures. A valid value is an integer from 0 through 15. The value 0 means the next connection attempt is controlled by the outbound connection failure retry interval.	6	Cmdlet: **Set-TransportService** Parameter: *TransientFailureRetryCount*	**Servers** > select server > **Edit** (✎) > **Transport limits** > **Retries** section > **Transient failure retry attempts**
Transient failure retry interval: The connection interval between each connection attempt that's specified by the transient failure retry count value.	Transport service on Mailbox servers: 5 minutes (`00:05:00`) Edge Transport servers: 10 minutes (`00:10:00`)	Cmdlet: **Set-TransportService** Parameter: *TransientFailureRetryInterval*	**Servers** > select server > **Edit** (✎) > **Transport limits** > **Retries** section > **Transient failure retry interval (minutes)**

Configuration options for manual message retry

When a delivery queue is in the status of Retry, you can manually force an immediate connection attempt by using Queue Viewer in the Exchange Toolbox or the **Retry-Queue** cmdlet in the Exchange Management Shell. The manual retry attempt overrides the next scheduled retry time. If the connection isn't successful, the retry interval timer is reset. The delivery queue must be in a status of Retry for this action to have any effect. For more information, see Retry queues.

Configuration options for delay DSN messages

After each message delivery failure, the Transport service on the Edge Transport server or the Mailbox server generates a delay delivery status notification (DSN) message and queues it for delivery to the sender of the undeliverable message. This delay DSN message is sent only after a delay notification interval has passed (the default is 4 hours), and only if the message wasn't successfully delivered during that time. This delay prevents the sending of unnecessary delay DSN messages due to temporary message transmission failures that are ultimately resolved. You can selectively enable or disable the sending of delay DSN notification messages for messages that originate inside or outside the Exchange organization.

The configuration options that are available for delay DSN notification messages are described in the following table.

DELAY DSN SETTING	DEFAULT VALUE	EXCHANGE MANAGEMENT SHELL CONFIGURATION	EXCHANGE ADMIN CENTER CONFIGURATION ON MAILBOX SERVERS
Delay notification timeout: How long the server waits before it sends a delay DSN message to the sender. This value should always be greater than the transient failure retry count multiplied by the transient failure retry interval (the default total is 30 minutes on a Mailbox server, and one hour on an Edge Transport server).	4 hours (`4:00:00`)	Cmdlet: **Set-TransportService** Parameter: *DelayNotificationTimeOut*	**Servers** > select server > **Edit** (✎) > **Transport limits** > **Notifications** section > **Notify sender when message is delayed after (hours)**
External delay DSN enabled: Specifies whether delay DSN messages can be sent to external message senders (senders who are outside the Exchange organization). *ExternalDelayDSNEnabled*	`$true`	Cmdlet: **Set-TransportConfig** Parameter: *ExternalDelayDSNEnabled*	Not available
Internal delay DSN enabled: Specifies whether delay DSN messages can be sent to internal message senders (message senders who are inside the Exchange organization).	`$true`	Cmdlet: **Set-TransportConfig** Parameter: *InternalDelayDSNEnabled*	Not available

Configuration options for message resubmission

Message resubmission sends undelivered messages back to the Submission queue to be reprocessed by the categorizer. For more information about the categorizer and the Submission queue, see Understanding the Transport service on Mailbox servers.

Automatic message resubmission

Undelivered messages in delivery queues are automatically resubmitted if the delivery queue is in the status of Retry and has been unable to successfully deliver any messages for a specified period of time. That period of time is controlled by the *MaxIdleTimeBeforeResubmit* key in the `%ExchangeInstallPath%Bin\EdgeTransport.exe.config` XML application configuration file. The default value is `12:00:00` or 12 hours.

> **NOTE**
>
> Any customized Exchange or Internet Information Server (IIS) settings that you made in Exchange XML application configuration files on the Exchange server (for example, web.config files or the EdgeTransport.exe.config file) **will be overwritten** when you install an Exchange CU. Be sure save this information so you can easily re-apply the settings after the install. After you install the Exchange CU, you need to re-configure these settings.

Manual Message Resubmission

You can manually resubmit messages by using the following methods:

- Resubmit a delivery queue that has the status of Retry, or resubmit the Unreachable queue. For more information, see Resubmit queues.

- Resubmit messages in the poison message queue. For more information, see Resubmit messages in the poison message queue.

- Suspend a queue, suspend the messages in the queue, export the messages to files, and copy the files to the Replay directory on any Mailbox server or Edge Transport server. For more information, see Export messages from queues.

Configuration options for message expiration

The *message expiration timeout interval* specifies the maximum length of time that an Edge Transport server or Mailbox server (the Transport service) tries to deliver a failed message. If the message can't be successfully delivered before the expiration timeout interval has passed, a non-delivery report (also known as an NDR or bounce message) that contains the original message or the message headers is delivered to the sender.

Automatic message expiration

The message expiration timeout interval is described in the following table.

DEFAULT VALUE	EXCHANGE MANAGEMENT SHELL CONFIGURATION	EXCHANGE ADMIN CENTER CONFIGURATION ON MAILBOX SERVERS
2 days (`2.00:00:00`)	Cmdlet: **Set-TransportService** Parameter: *MessageExpirationTimeOut*	**Servers** > select server > **Edit** (✎) > **Transport limits** > **Message expiration** section > **Maximum time since submission (days)**

Manual Message Expiration

Although you can't manually force messages to expire, you can manually remove messages from any queue (except the Submission queue) with or without an NDR. For more information, see Remove messages from queues.

Configure message retry, resubmit, and expiration intervals

8/3/2020 • 9 minutes to read • Edit Online

In Exchange Server, you can configure message retry, resubmit, and expiration intervals in the Transport service on Mailbox servers and Edge Transport servers. For detailed descriptions of these settings, see Message retry, resubmit, and expiration intervals.

What do you need to know before you begin?

- Estimated time to complete each procedure: less than 5 minutes
- You can only use the Exchange admin center (EAC) on Mailbox servers. For more information about the EAC, see Exchange admin center in Exchange Server. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Transport service" and "Edge Transport severs" entries in the Mail flow permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use EdgeTransport.exe.config to configure the queue glitch retry count, the queue glitch retry interval, the mailbox delivery queue retry interval, and the maximum idle time before resubmit interval

- **Queue glitch retry count**: The number of connection attempts that are immediately tried when the Transport service has trouble connecting to the destination server. Typically, you don't need to modify this key unless the network is unreliable and continues to experience many accidentally dropped connections.
- **Queue glitch retry interval**: The interval between each queue glitch retry. Typically, you don't need to modify this key unless the network is unreliable and continues to experience many accidentally dropped connections.
- **Mailbox delivery queue retry interval**: How frequently a queue try to connect to the Mailbox Transport Delivery service for a destination mailbox database that can't be successfully reached.
- **Max idle time before resubmit**: How long undelivered messages in delivery queues the status of Retry wait before they're resubmitted.

To configure these intervals, you modify keys in the %ExchangeInstallPath%Bin\EdgeTransport.exe.config XML application configuration file on Mailbox servers or Edge Transport servers. Changes you save to this file are applied after you restart the Exchange Transport service. When you restart this service, mail flow on the server is temporarily interrupted.

NOTE

Any customized Exchange or Internet Information Server (IIS) settings that you made in Exchange XML application configuration files on the Exchange server (for example, web.config files or the EdgeTransport.exe.config file) **will be overwritten** when you install an Exchange CU. Be sure save this information so you can easily re-apply the settings after the install. After you install the Exchange CU, you need to re-configure these settings.

1. In a Command prompt window on the Mailbox server or Edge Transport server, open the EdgeTransport.exe.config file in Notepad by running this command:

```
Notepad %ExchangeInstallPath%Bin\EdgeTransport.exe.config
```

2. Locate the following keys in the `<appSettings>` section.

```
<add key="QueueGlitchRetryCount" value="<Integer>" />
<add key="QueueGlitchRetryInterval" value="<hh:mm:ss>" />
<add key="MailboxDeliveryQueueRetryInterval" value="<hh:mm:ss>" />
<add key="MaxIdleTimeBeforeResubmit" value="<hh:mm:ss>" />
```

This example changes the queue glitch retry count to 6, the queue glitch retry interval to 30 seconds, the mailbox delivery queue retry interval to 3 minutes, and the maximum idle time before resubmit interval to 6 hours.

```
<add key="QueueGlitchRetryCount" value="6" />
<add key="QueueGlitchRetryInterval" value="00:00:30" />
<add key="MailboxDeliveryQueueRetryInterval" value="00:03:00" />
<add key="MaxIdleTimeBeforeResubmit" value="6:00:00" />
```

3. When you're finished, save and close the EdgeTransport.exe.config file.

4. Restart the Exchange Transport service by running this command:

```
net stop MSExchangeTransport && net start MSExchangeTransport
```

How do you know this worked?

To verify that you've configured these intervals, do these steps:

1. Open the EdgeTransport.exe.config file in Notepad by running this command:

```
Notepad %ExchangeInstallPath%Bin\EdgeTransport.exe.config
```

2. Verify the values of the following keys in the `<appSettings>` section.

```
<add key="QueueGlitchRetryCount" value="<Integer>" />
<add key="QueueGlitchRetryInterval" value="<hh:mm:ss>" />
<add key="MailboxDeliveryQueueRetryInterval" value="<hh:mm:ss>" />
<add key="MaxIdleTimeBeforeResubmit" value="<hh:mm:ss>" />
```

Configure the transient failure retry attempts, the transient failure retry interval, and the outbound connection failure retry interval

- **Transient failure retry attempts**: The number of connection attempts that are tried after the connection attempts controlled by the *QueueGlitchRetryCount* and *QueueGlitchRetryInterval* keys have failed. A valid value is 0 through 15, and the default value is 6. If you set the value to 0, the next connection attempt is controlled by the outbound connection failure retry interval.

- **Transient failure retry interval**: The interval between each transient failure retry attempt. On Mailbox servers, the default value is 5 minutes. On Edge Tranport Servers, the default value is 10 minutes.

- **Outbound connection failure retry interval**: The retry interval for outgoing connection attempts that have previously failed (the transient failure retry attempts and the transient failure retry interval). On Mailbox servers, the default value is 10 minutes. On Edge Tranport Servers, the default value is 30 minutes.

Use the EAC to configure the transient failure retry attempts, the transient failure retry interval, or the outbound connection failure retry interval on Mailbox servers

1. In the EAC, go to **Servers** > **Servers**, select the server, and then click **Edit** ✏.

2. In the server properties window that opens, click **Transport limits**.

3. In the **Retries** section, enter a value for any of these settings:

   - **Outbound connection failure retry interval (seconds)**

   - **Transient failure retry interval (minutes)**

   - **Transient failure retry attempts**

   When you're finished, click **Save**.

Use the Exchange Management Shell to configure the transient failure retry attempts, the transient failure retry interval, and the outbound connection failure retry interval on Mailbox severs or Edge Transport servers

To configure the intervals in the Transport service on Mailbox servers or Edge Transport servers, use this syntax:

```
Set-TransportService -Identity <ServerIdentity> -TransientFailureRetryCount <Integer> -
TransientFailureRetryInterval <hh:mm:ss> -OutboundConnectionFailureRetryInterval <dd.hh:mm:ss>
```

To configure the intervals in the Front End Transport service on Mailbox servers, use this syntax:

```
Set-FrontEndTransportService -Identity <ServerIdentity> -TransientFailureRetryCount <Integer> -
TransientFailureRetryInterval <hh:mm:ss>
```

This example changes the following values on the Mailbox server named Mailbox01:

- The number of transient failure retry attempts is set to 8.

- The transient failure retry interval is set to 1 minute.

- The outbound connection failure retry interval is set to 45 minutes.

```
Set-TransportService -Identity Mailbox01 -TransientFailureRetryCount 8 -TransientFailureRetryInterval 00:01:00
-OutboundConnectionFailureRetryInterval 00:45:00
```

How do you know this worked?

To verify that you've configured these intervals, do any of these steps:

- On a Mailbox server, open the EAC and go to **Servers** > **Servers**, select the server, and then click **Edit** ✏. In the server properties window that opens, click **Transport limits**, and verify the values in the **Retries** section.

- In the Exchange Management Shell on a Mailbox server or Edge Transport server, run this command to verify the property values:

  ```
  Get-TransportService | Format-List Name,TransientFailureRetry*,OutboundConnectionFailureRetryInterval
  ```

- In the Exchange Management Shell on a Mailbox serve, run this command to verify the property values:

  ```
  Get-FrontEndTransportService | Format-List Name,TransientFailureRetry*
  ```

Use the Exchange Management Shell to configure the message retry interval

The message retry interval specifies how long to wait between sending attempts for individual messages in queues that have a status of Retry. The default value is 15 minutes, and we recommend that you don't change the default value unless you're directed to do so by Microsoft Customer Service and Support, or specific product documentation.

To configure the message retry interval, use this syntax:

```
Set-TransportService -Identity <ServerIdentity> -MessageRetryInterval <dd.hh:mm:ss>
```

This example changes the message retry interval to 20 minutes on the Mailbox server named Mailbox01.

```
Set-TransportService -Identity Mailbox01 -MessageRetryInterval 00:20:00
```

How do you know this worked?

To verify that you've configured the message retry interval on a Mailbox server or Edget Transport server, run this command in the Exchange Management Shell to verify the **MessageRetryInterval** property value:

```
Get-TransportService | Format-List Name,MessageRetryInterval
```

Configure the delay DSN timeout settings

- **Delay DSN message notification timeout interval**: How long to wait before sending delay DSN messages to senders. This setting applies to the Transport service on a Mailbox server or an Edge Transport server.

Note: This value should always be greater than the transient failure retry count multiplied by the transient failure retry interval (the default total is 30 minutes on a Mailbox server, and one hour on an Edge Transport server).

- **Internal and external delay DSN settings**: Specifies whether delay DSN messages can be sent to internal or external message senders (senders who are inside or outside the Exchange organization). This setting applies to the Transport service on all Mailbox servers in the organization.

Use the EAC to configure the delay DSN message notification timeout interval on Mailbox servers

1. In the EAC, click **Servers** > **Servers**, select the server, and then click **Edit** ✎.

2. In the server properties window that opens, click **Transport limits**.

3. In the **Notifications** section, enter a value for **Notify sender when message is delayed after (hours)**, and then click **Save**.

Use the Exchange Management Shell to configure the delay DSN message notification timeout interval on

Mailbox servers or Edge Transport servers

To configure the delay DSN message notification timeout interval, use this syntax:

```
Set-TransportService -Identity <ServerIdentity> -DelayNotificationTimeout <dd.hh:mm:ss>
```

This example changes the delay DSN message notification timeout interval to 6 hours on the Mailbox server named Mailbox01.

```
Set-TransportService -Identity Mailbox01 -DelayNotificationTimeout 06:00:00
```

Use the Exchange Management Shell to enable or disable the sending of delay DSN notifications to external or internal message senders

To configure the delay DSN notification settings, use this syntax:

```
Set-TransportConfig -ExternalDelayDSNEnabled <$true | $false> -InternalDelayDSNEnabled <$true |$false>
```

This example prevents the sending of delay DSN notification messages to external senders.

```
Set-TransportConfig -ExternalDelayDSNEnabled $false
```

This example prevents the sending of delay DSN notification messages to internal senders.

```
Set-TransportConfig -InternalDelayDSNEnabled $false
```

How do you know this worked?

To verify that you've configured the delay DSN timeout settings, do any of these steps:

- On a Mailbox server, open the EAC and go to **Servers** > **Servers**, select the server, and then click **Edit** ✏. In the server properties window that opens, click **Transport limits**, and verify the **Notify sender when message is delayed after (hours)** value in the **Notifications** section.
- In the Exchange Management Shell on a Mailbox server or Edge Transport server, run these commands to verify the property values:

  ```
  Get-TransportService | Format-List Name,DelayNotificationTimeout
  ```

  ```
  Get-TransportConfig | Format-List *DelayDSNEnabled
  ```

Configure the message expiration timeout interval

The message expiration timeout interval specifies how long to wait before the message expires and is returned to the sender in a non-delivery report (also known as an NDR or bounce message). This setting applies to the Transport service on a Mailbox server or an Edge Transport server.

Use the EAC to configure the message expiration timeout interval on Mailbox servers

1. In the EAC, click **Servers** > **Servers**, select the server, and then click **Edit** ✏.
2. In the server properties window that opens, click **Transport limits**.
3. In the **Message expiration** section, enter a value for **Maximum time since submission (days)**, and

then click **Save**.

Use the Exchange Management Shell to configure the message expiration timeout interval on Mailbox servers or Edge Transport servers

To configure the message expiration timeout interval, use the following syntax.

```
Set-TransportService -Identity <ServerIdentity> -MessageExpirationTimeout <dd.hh:mm:ss>
```

This example changes the message expiration timeout interval to 4 days on the Exchange server named Mailbox01.

```
Set-TransportService -Identity Mailbox01 -MessageExpirationTimeout 4.00:00:00
```

How do you know this worked?

To verify that you've configured the message expiration timeout interval, do any of these steps:

- On a Mailbox server, open the EAC and go to **Servers** > **Servers**, select the server, and then click **Edit** ✏. In the server properties window that opens, click **Transport limits**, and verify the **Maximum time since submission (days)** value in the **Message expiration** section.

- In the Exchange Management Shell on a Mailbox server or Edge Transport server, run this command to verify the **MessageExpirationTimeout** property value:

  ```
  Get-TransportService | Format-List Name,MessageExpirationTimeout
  ```

DSNs and NDRs in Exchange Server

8/3/2020 • 12 minutes to read • Edit Online

When there's a problem delivering a message, Exchange sends an NDR to the message sender that indicates there was a problem. NDRs include a code that indicates why the message wasn't delivered, and possible solutions to help get the message delivered.

The information that's included in NDRs is designed to be easy to read and helpful for both users and administrators. In some cases, senders can identify and fix their own problems (for example, when there's a typo in the recipient's email address). In other cases, an administrator may need to fix an issue in the Exchange environment, or notify the administrators in the destination domain about problems in their messaging environment.

For procedures related to NDRs in Exchange Server, see Procedures for DSNs and NDRs in Exchange Server.

If you need help with NDRs in Microsoft 365, Office 365, or Exchange Online, see Email non-delivery reports in Exchange Online.

Information in NDRs

This is an example of an NDR:

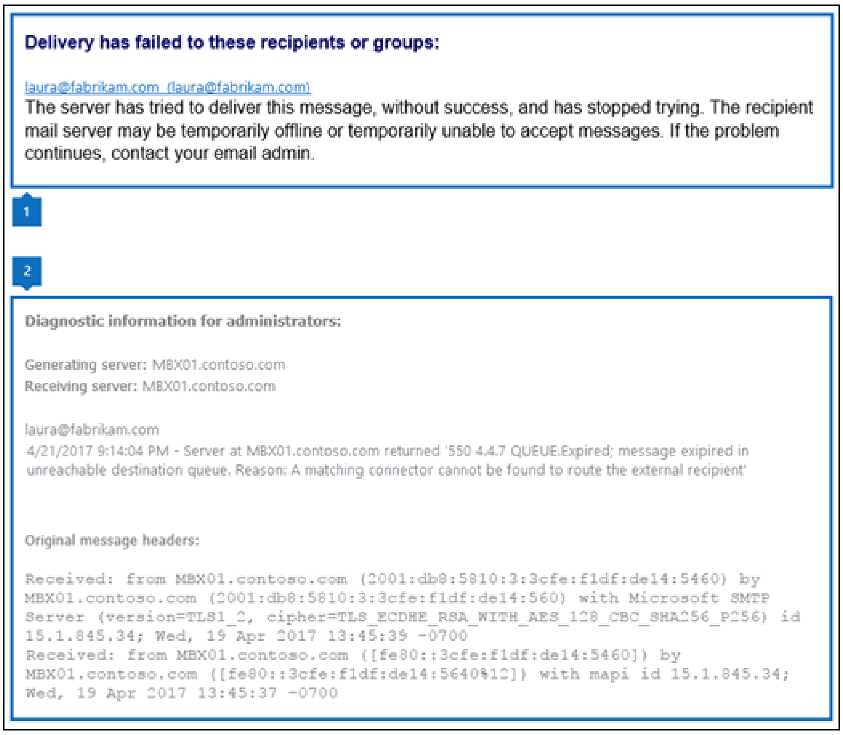

Delivery has failed to these recipients or groups:

laura@fabrikam.com (laura@fabrikam.com)
The server has tried to deliver this message, without success, and has stopped trying. The recipient mail server may be temporarily offline or temporarily unable to accept messages. If the problem continues, contact your email admin.

Diagnostic information for administrators:

Generating server: MBX01.contoso.com
Receiving server: MBX01.contoso.com

laura@fabrikam.com
4/21/2017 9:14:04 PM - Server at MBX01.contoso.com returned '550 4.4.7 QUEUE.Expired; message exipired in unreachable destination queue. Reason: A matching connector cannot be found to route the external recipient'

Original message headers:

```
Received: from MBX01.contoso.com (2001:db8:5810:3:3cfe:f1df:de14:5460) by
MBX01.contoso.com (2001:db8:5810:3:3cfe:f1df:de14:560) with Microsoft SMTP
Server (version=TLS1_2, cipher=TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256_P256) id
15.1.845.34; Wed, 19 Apr 2017 13:45:39 -0700
Received: from MBX01.contoso.com ([fe80::3cfe:f1df:de14:5460]) by
MBX01.contoso.com ([fe80::3cfe:f1df:de14:5640%12]) with mapi id 15.1.845.34;
Wed, 19 Apr 2017 13:45:37 -0700
```

The information in an NDR is separated into two sections:

1. **User information section**: This section appears first and attempts to explain (in non-technical terms) why delivery of the message failed, and possible steps to successfully deliver the message.

   - The text that's displayed in this section is inserted by the Exchange server that generated the NDR.
   - When applicable, the fully qualified domain name (FQDN) of the server that rejected the message is included in the user information section (for example mbx01.contoso.com).
   - If delivery failed for multiple recipients, the email address and reason for failure is listed for each

recipient is listed..

2. **Diagnostic information for administrators section**: This section provides deeper technical information to help administrators troubleshoot the issues that caused the delivery failure.

   A key piece of information in this section is the enhanced status code (for example, 4.4.7).

   - The enhanced status code is returned by the server that generated the NDR (the source server that couldn't deliver the message, or the destination server that rejected the message).

   - The enhanced status code determines the text that's displayed in the user information section (the code value isn't altered by Exchange).

   You can use the **New-SystemMessage** cmdlet in the Exchange Management Shell to modify the text that appears in user information section for a given enhanced status code (including different text in different languages). By creating custom explanations, you can provide specific content for your environment, such as contact information for your help desk, or links to your Intranet for self-service support. For more information, see Procedures for DSNs and NDRs in Exchange Server.

   - The Common enhanced status codes section in this topic explains what the numbers mean, the codes that you're likely to encounter, and suggestions to fix the underlying problem that prevented the message from being delivered.

   The following information is also available in this section:

   - **Generating server**: The messaging server that created the NDR. If a remote server isn't listed below the sender's email address, the generating server is also the server that rejected the original email message. If message delivery fails between senders and recipients in the Exchange organization, the same server typically rejects the original message and generates the NDR.

   - **The rejected recipients**: The recipient's email address in the original message that couldn't be delivered. If delivery fails for multiple recipients, the email address of each recipient is listed. This field also contains the following sub-fields for each email address:

   - **Remote server**: The FQDN of the server that rejected the original message during SMTP transmission (delivery failed after the message body was sent, but before the server acknowledged receiving the message). This field isn't present when:

   - The server that rejected the message also generated the NDR. This is typical for delivery failures between senders and recipients in the same Exchange organization.

   - The remote server acknowledged receiving the original message, but the message was rejected for other reasons (for example, content restrictions).

   - **Enhanced status code**

   - **SMTP response**: The US-ASCII text string that's returned by the messaging server that rejected the original message. This is typically a short explanation of the enhanced status code. This string is not rewritten by Exchange.

   - **Original message headers**: This area contains the message header of the rejected message. These header fields can provide useful diagnostic information (for example, server hops in the message routing path, or whether the **To** field matches the email address of the rejected recipient).

Common enhanced status codes

Enhanced status codes are defined in RFC 3463, and use the syntax *<class>*. *<subject>*. *<detail>*:

- *<class>*: **4** indicates a temporary delivery error. **5** indicates a permanent delivery error.

- *<subject>*: The RFC categorizes the values like this:
  - **1**: Addressing
  - **2**: Mailbox (the recipient)
  - **3**: Mail system (the destination mail system)
  - **4**: Network and routing
  - **5**: Mail delivery protocol
  - **7**: Security or policy
- *<detail>*: A 1 to 3 digit number that further classifies the error.

The following tables contain the enhanced status codes that are returned in NDRs for the most common message delivery failures.

> **NOTE**
>
> For information about enhanced status codes in Microsoft 365 or Office 365 and hybrid environments, see Email non-delivery reports in Exchange Online.

Temporary delivery failures

ENHANCED STATUS CODE	DESCRIPTION	POSSIBLE CAUSES AND SOLUTIONS
4.3.1	`Insufficient system resources`	Free disk space is low (for example, the disk that holds the queue database doesn't have the required amount of free space). For more information, see Understanding back pressure. To move the queue database to different disk, see Change the location of the queue database. Available memory is low (for example, Exchange installed on a virtual machine that's configured to use dynamic memory). Always use static memory on Exchange virtual machines. For more information, see Exchange memory requirements and recommendations.

ENHANCED STATUS CODE	DESCRIPTION	POSSIBLE CAUSES AND SOLUTIONS

4.3.2	`Service not available`
	or
	`Service not active`
		You've configured a custom Receive connector in the Transport (Hub) service on a Mailbox server that listens on port 25. Typically, custom Receive connectors that listen on port 25 belong in the Front End Transport service on the Mailbox server.
		Important Exchange server components are inactive. You can confirm this by running the following command in the Exchange Management Shell: `Get-ServerComponent -Identity <ServerName>`
		.
		To restart all inactive components, run the following command: `Set-ServerComponentState -Identity <ServerName> -Component ServerWideOffline -State Active -Requester Maintenance`
		.
		Incompatible transport agents (in particular, after an Exchange update). After you identify the transport agent, disable it or uninstall it. For more information, see Troubleshoot transport agents.

4.4.1	`Connection timed out`	Transient network issues that might eventually correct themselves. The Exchange server periodically tries to connect to the destination server to deliver the message. After multiple failures, the message is returned to the sender in an NDR with a permanent failure code. For more information about configuring the queue retry and failure intervals, see Configure message retry, resubmit, and expiration intervals. To manually retry a queue, see Retry queues. Firewall or Internet service provider (ISP) restrictions on TCP port 25.
4.4.2	`Connection dropped`	Transient network issues or server problems that might eventually correct themselves. The sending server will retry delivery of the message, and will generate further status reports. The message size limit for the connection has been reached, or the message submission rate for the source IP address has exceeded the configured limit. For more information, see Message rate limits and throttling. Antispam, SMTP proxy, or firewall configuration issues are blocking email from the Exchange server.

ENHANCED STATUS CODE	DESCRIPTION	POSSIBLE CAUSES AND SOLUTIONS

4.4.7	`Message delayed`
	or
	`Queue expired; Message expired`
		Send connector configuration issues. For example:
		• The Send connector is configured to use DNS routing when it should be using smart host routing, or vice-versa. Use nslookup to verify that the destination domain is reachable from the Exchange server.
		• The FQDN that the Send connector provides to HELO or EHLO requests doesn't match the host name in your MX record (for example, mail.contoso.com). Some messaging systems are configured to compare these value in an effort to reduce spam. The default value on a Send connector is blank, which means the FQDN of the Exchange server is used (for example, exchange01.contoso.com).
		The Mailbox Transport Delivery service isn't started on the destination server (which prevents the delivery of the message to the mailbox).
		The destination messaging system has issues with Transport Neutral Encryption Format (TNEF) messages (also known as rich text format or RTF in Outlook). For example, meeting requests or messages with images embedded in the message body.
		If the destination domain uses the Sender Policy Framework (SPF) to check message sources, there may be SPF issues with your domain (for example, your SPF record doesn't include all email sources for your domain).

Permanent delivery failures

ENHANCED STATUS CODE	DESCRIPTION	POSSIBLE CAUSES AND SOLUTIONS.
5.1.0	`Sender denied`	Replying to old messages, or messages that were exported as files (important recipient attributes might have changed). Verify that the recipient's email address is correct. Malformed or missing attributes in contact entries. The sender is blocked by sender filtering (directly, or the sender is on a user's Blocked Senders list, and the Sender Filter agent is configured to use safelist aggregation. For more information, see Sender filtering and Safelist aggregation.

ENHANCED STATUS CODE	DESCRIPTION	POSSIBLE CAUSES AND SOLUTIONS.

5.1.1	`RESOLVER.ADR.ExRecipNotFound; not found`
	or
	`User unknown`
		The recipient's email address is incorrect (the recipient doesn't exist in the destination messaging system). Verify the recipient's email address.
		You recreated a deleted mailbox, and internal users are addressing email messages in Outlook or Outlook on the web using old entries in their autocomplete cache (the X.500 values or **LegacyExchangeDN** values for the recipient are now different). Tell users to delete the entry from their autocomplete cache and select the recipient again.

5.1.3	`STOREDRV.Submit; invalid recipient address`	The recipient's email address is incorrect (for example, it contains unsupported characters or invalid formatting).
5.1.4	`Recipient address reserved by RFC 2606`	Receive connectors reject SMTP connections that contain the top level domains defined in RFC 2606 (.test, .example, .invalid, or .localhost), This behavior is controlled by the *RejectReservedTopLevelRecipientDomains* parameter on the New-ReceiveConnector and Set-ReceiveConnector cmdlets.
5.1.5	`Recipient address reserved by RFC 2606`	Receive connectors reject SMTP connections that contain the second level domains defined in RFC 2606 (example.com, example.net, or example.org). This behavior is controlled by the *RejectReservedSecondLevelRecipientDomains* parameter on the New-ReceiveConnector and Set-ReceiveConnector cmdlets.
5.1.6	`Recipient addresses in single label domains not accepted`	Receive connectors reject SMTP connections that contain single label domains (for example, chris@contoso instead of chris@contoso.com) This behavior is controlled by the *RejectSingleLabelRecipientDomains* parameter on the New-ReceiveConnector and Set-ReceiveConnector cmdlets.

5.1.7	`Invalid address`
	or
	`Unknown sender address`
		There's a problem with the sender's email address. Verify the sender's email address.

ENHANCED STATUS CODE	DESCRIPTION	POSSIBLE CAUSES AND SOLUTIONS.
5.1.8	`Access denied, bad outbound sender`	The sender has exceeded a message rate limit (for example, an application server is configured to relay a large number of messages through Exchange. For more information, see Message rate limits and throttling and Allow anonymous relay on Exchange servers.
5.2.1	`Content Filter agent quarantined this message`	The message was quarantined by content filtering. To configure exceptions to content filtering, see Use the Exchange Management Shell to configure recipient and sender exceptions for content filtering.
5.2.2	`Mailbox full`	The recipient's mailbox has exceeded its storage quota and is no longer able to accept new messages. For more information about configuring mailbox quotas, see Configure storage quotas for a mailbox.
5.2.3	`RESOLVER.RST.RecipSizeLimit; message too large for this recipient`	The message is too large. Send the message again without any attachments, or configure a larger message size limit for the recipient. For more information, see Recipient limits.
5.3.0	`Too many related errors`	The message was determined to be malformed, and was moved to the poison message queue. For more information, see Types of queues.
5.3.2	`STOREDRV.Deliver: Missing or bad StoreDriver MDB properties`	You're using the ABP Routing agent, and the recipient isn't a member of the global address list that's specified in their address book policy (ABP). For more information, see Use the Exchange Management Shell to install and configure the Address Book Policy Routing Agent and Address book policies in Exchange Server.
5.3.3	`Unrecognized command`	Receive connectors that are used for internal mail flow are missing the required Exchange Server authentication mechanism. For more information about authentication on Receive connectors, see Receive connector authentication mechanisms.

ENHANCED STATUS CODE	DESCRIPTION	POSSIBLE CAUSES AND SOLUTIONS.
5.3.4	`Message size exceeds fixed maximum message size`	The message is too large. This error can be generated by the source or destination messaging system. Send the message again without any attachments, or configure a larger message size limit. For more information, see Message size and recipient limits in Exchange Server.
5.3.5	`System incorrectly configured`	A mail loop was detected. Verify that the **FQDN** property on the Receive connector doesn't match the FQDN of another server, service, or device that's used in mail flow in your organization (by default, the Receive connector uses the FQDN of the Exchange server).
5.4.4	`SMTPSEND.DNS.NonExistentDomain; nonexistent domain`	There's a DNS or network adapter configuration issue on the Exchange server. • Verify the internal and external DNS lookup settings for the Exchange by running these commands in the Exchange Management Shell: Get-TransportService | Format-List Name,ExternalDNS*,InternalDNS* Get-FrontEndTransportService | Format-List Name,ExternalDNS*,InternalDNS* You can configure these settings by using the *InternalDNS** and *ExternalDNS** parameters on the **Get-TransportService** and **Get-FrontEndTransportService** cmdlets. By default, these settings are used by Send connectors (the default value of the *UseExternalDNSServersEnabled* parameter value is `$false`). • Check the priority (order) of the network adapters in the operating system of the Exchange server.

ENHANCED STATUS CODE	DESCRIPTION	POSSIBLE CAUSES AND SOLUTIONS.
5.4.6	`Hop count exceeded - possible mail loop`	A configuration error has caused an email loop. By default, after 20 iterations of an email loop, Exchange interrupts the loop and generates an NDR. Verify that Inbox rules for the recipient and sender, or forwarding rules on the recipient 's mailbox aren't causing this (the message generates a message, which generates another message, and the process continues indefinitely). Verify the mailbox doesn't have a **targetAddress** property value in Active Directory (this property corresponds to the *ExternalEmailAddress* parameter for mail users in Exchange). If you remove Exchange servers, or modify settings related to mail routing an mail flow, be sure to restart the Microsoft Exchange Transport and Exchange Frontend Transport services.
5.5.2	`Send hello first`	SMTP commands are sent out of sequence (for example, a server sends an SMTP command like AUTH or MAIL FROM before identifying itself with the EHLO command). After establishing a connection to a messaging server, the first SMTP command must always be EHLO or HELO.
5.5.3	`Too many recipients`	The combined total of recipients on the To, Cc, and Bcc lines of the message exceeds the total number of recipients allowed in a single message for the organization, Receive connector, or sender. For more information, see Message size and recipient limits in Exchange Server.

5.7.1	`Unable to relay`
	or
	`Client was not authenticated`
		You have an application server or device that's trying to relay messages through Exchange. For more information, see Allow anonymous relay on Exchange servers.
		The recipient is configured to only accept messages from authenticated (typically, internal) senders. For more information, see Configure message delivery restrictions for a mailbox.

ENHANCED STATUS CODE	DESCRIPTION	POSSIBLE CAUSES AND SOLUTIONS.

5.7.3	`Cannot achieve Exchange Server authentication`
	or
	`Not Authorized`
		A firewall or other device is blocking the Extended SMTP command that's required for Exchange Server authentication (X-EXPS).
		Internal email traffic is flowing through connectors that aren't configured to use the Exchange Server authentication method . Verify the remote IP address ranges on any custom Receive connectors.

5.7.900 to 5.7.999	`Delivery not authorized, message refused`	The message was rejected by a mail flow rule (also known as a transport rule). This enhanced status code range is available when the rule is configured to reject messages (otherwise, the default code that's used is 5.7.1). For more information, see Mail flow rule actions in Exchange Server.

Procedures for DSNs and NDRs in Exchange Server

8/3/2020 • 15 minutes to read • Edit Online

Like previous versions of Exchange, Exchange Server uses delivery status notifications (also known as DSNs, non-delivery reports, NDRs, or bounce messages) to provide delivery status and failure notification messages to message senders. For more information about NDRs, see DSNs and NDRs in Exchange Server.

You can use the default NDRs that are included in Exchange, or you can use the Exchange Management Shell to create NDRs with custom text to meet the needs of your organization. The custom NDR text replaces the default text for a given enhanced status code or quota event. If you remove the custom NDR, the default NDR text is used (you can't completely remove a default NDR). You can also disable custom NDRs to preserve them, but not use them (the default NDR text is used).

What do you need to know before you begin?

- Estimated time to complete each procedure: less than 10 minutes.
- The main focus of this topic is custom NDR text that replaces the text of default NDRs that are used by Exchange. You can create new NDRs for other enhanced status code values (for example, 5.999.999), but no one will see these NDRs if the enhanced status code isn't used by Exchange. You can use a range of custom enhanced status codes as part of an action for a mail flow rule (also known as a transport rule). For more information, see Mail flow rule actions in Exchange Server.
- The procedures in this topic are available on Mailbox servers and Edge Transport servers.
- You can't use the Exchange admin center (EAC) for most of the procedures in this topic. You need to use the Exchange Management Shell. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "DSNs" entry in the Mail flow permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to view all default NDRs

To output the list of all default NDRs in all languages to an HTML file named C:\My Documents\Default NDRs.html, run this command:

```
Get-SystemMessage -Original | Select-Object -Property Identity,DsnCode,Language,Text | ConvertTo-Html | Set-Content -Path "C:\My Documents\Default NDRs.html"
```

Note: You should output the list to a file, because the list is very long, and you'll receive errors if you don't have the required language packs installed.

For detailed syntax and parameter information, see Get-SystemMessage.

Use the Exchange Management Shell to view custom NDRs

To view a summary list of all custom NDRs in your organization, run this command:

```
Get-SystemMessage
```

Note: By default, there are no custom NDRs, so this command returns no results.

To view detailed information for a custom NDR, use this syntax:

```
Get-SystemMessage -Identity <NDRIdentity>
```

For an explanation of the available *<NDRIdentity>* values, see the Identity values for NDRs section in this topic.

This example returns detailed information for the custom NDR for the enhanced status code 5.1.2 that's sent to internal senders in English. If there's no custom NDR for this combination of language, audience, and enhanced status code, you'll receive an error.

```
Get-SystemMessage En\Internal\5.1.2 | Format-List
```

This example returns detailed information for the custom English NDR for the **ProhibitSendReceive** quota on mailboxes. If there's no custom NDR for this combination of language and quota, you'll receive an error.

```
Get-SystemMessage En\ProhibitSendReceiveMailBox | Format-List
```

For detailed syntax and parameter information, see Get-SystemMessage.

Create custom NDRs

Use the Exchange Management Shell to create custom NDRs for enhanced status codes

To create a custom NDR for an enhanced status code, use this syntax:

```
New-SystemMessage -Internal <$true | $false> -Language <Locale> -DSNCode <x.y.z> -Text "<NDR text>"
```

The values are:

- **Internal**: Controls whether the NDR is sent to internal or external senders. For internal senders, use the value `$true`. For external senders, use the value `$false`. For example, in the custom text for internal senders, you can include help desk contact information that you wouldn't want to include in NDRs for external senders.

- **Language**: For the list of available languages, see the Supported languages for NDRs section in this topic.

- **DSNCode**: The enhanced status code. Valid values are 4. *x*. *y* or 5. *x*. *y* where *x* and *y* are one to three digit numbers.

- **Text**: You can use plain text or HTML formatting. For more information, see the HTML tags and special characters in NDRs section in this topic.

This example creates a custom plain text NDR for the enhanced status code 5.1.2 that's sent to external senders in English.

```
New-SystemMessage -Internal $false -Language En -DSNCode 5.1.2 -Text "You tried to send a message to a
disabled mailbox that's no longer accepting messages. Please contact your System Administrator for more
information."
```

This example creates a custom HTML NDR for the enhanced status code 5.1.2 that's sent to internal senders in English.

```
New-SystemMessage -DSNCode 5.1.2 -Internal $true -Language En -Text 'You tried to send a message to a
<B>disabled</B> mailbox. Please visit <A HREF="https://it.contoso.com">Internal Support</A> or contact
&quot;InfoSec&quot; for more information.'
```

For detailed syntax and parameter information, see New-SystemMessage.

Use the Exchange Management Shell to create custom NDRs for quotas

To create a custom NDR for quotas, use this syntax:

```
New-SystemMessage -Language <Locale> -QuotaMessageType <Quota> -Text "<NDR text>"
```

The values are:

- **Language**: For the list of available languages, see Supported languages for NDRs.
- **QuotaMessageType**: For a list of the available quotas, see Identity values for NDRs.
- **Text**: You can use plain text or HTML formatting. For more information, see HTML tags and special characters in NDRs.

This example creates a custom English plain text NDR for the **ProhibitSendReceive** quota on mailboxes.

```
New-SystemMessage -Language En -QuotaMessageType ProhibitSendReceiveMailBox -Text "Your mailbox is full, and
can't send or receive messages. Delete any unwanted large messages (messages with attachments) and empty your
Deleted Items folder"
```

For detailed syntax and parameter information, see New-SystemMessage.

How do you know this worked?

To verify that you have successfully created a custom NDR, do these steps:

- Run the following command and verify the property values:

  ```
  Get-SystemMessage | Format-List Identity,DsnCode,Language,Text
  ```

- Send a test message that will generate the custom NDR that you configured.

Use the Exchange Management Shell to modify custom NDRs

To modify custom NDRs, use this syntax:

```
Set-SystemMessage -Identity <NDRIdentity> [-Text "<NDR text>"] [-Original]
```

For an explanation of the available *<NDRIdentity>* values, see the Identity values for NDRs section in this topic. For an explanation of the *<NDR text>* values, see the HTML tags and special characters in NDRs section in this topic.

This example changes the text in the custom NDR for the enhanced status code 5.1.2 that's sent to internal senders in English.

```
Set-SystemMessage -Identity En\Internal\5.1.2 -Text "The mailbox you tried to send an email message to is
disabled and is no longer accepting messages. Please contact the Help Desk at extension 123 for assistance."
```

This example changes the text in the custom English NDR for the **ProhibitSendReceive** quota on mailboxes.

```
Set-SystemMessage -Identity En\ProhibitSendReceiveMailBox -Text "Your mailbox is full. Delete large messages
and empty your Deleted Items folder."
```

This example disables the specified custom NDR. The custom NDR is preserved, and appears in the results of **Get-SystemMessage**, but the default NDR is used instead.

```
Set-SystemMessage -Identity En\Internal\5.1.2 -Original
```

Note: If there's no corresponding default NDR, you receive an error when you use the *Original* switch.

For detailed syntax and parameter information, see Set-SystemMessage.

How do you know this worked?

To verify that you have successfully modified a custom NDR, replace *<NDRIdentity>* with the appropriate value, and run this command to verify the property values:

```
Get-SystemMessage -Identity <NDRIdentity> | Format-List
```

Use the Exchange Management Shell to remove custom NDRs

To remove a custom NDR, use this syntax:

```
Remove-SystemMessage -Identity <NDRIdentity>
```

For an explanation of the available *<NDRIdentity>* values, see the Identity values for NDRs section in this topic.

This example removes the custom NDR for the enhanced status code 5.1.2 that's sent to internal senders in English.

```
Remove-SystemMessage -Identity En\Internal\5.1.2
```

This example removes the custom English NDR for the **ProhibitSendReceive** quota on mailboxes.

```
Remove-SystemMessage -Identity En\ProhibitSendReceiveMailBox
```

For detailed syntax and parameter information, see Remove-SystemMessage.

How do you know this worked?

To verify that you have successfully removed a custom NDR, run this command to verify the custom NDR isn't listed:

```
Get-SystemMessage
```

Forward copies of NDRs to the Exchange recipient mailbox

You can configure your Exchange organization to send copies of NDRs to the Exchange recipient. However, by default, no mailbox is assigned to the Exchange recipient, so any messages that are sent to the Exchange recipient are discarded. To send copies of NDRs to the Exchange recipient mailbox, you need to:

1. Assign a mailbox to the Exchange recipient.

2. Specify the enhanced status codes that you want to monitor (not quotas).

Step 1: Use the Exchange Management Shell to assign a mailbox to the Exchange recipient

Note: Due to the high volume of messages, we recommend using a dedicated mailbox for the Exchange recipient. For more information about creating mailboxes, see Create shared mailboxes in the Exchange admin center and Create user mailboxes in Exchange Server.

To assign a mailbox to the Exchange recipient, use this syntax:

```
Set-OrganizationConfig -MicrosoftExchangeRecipientReplyRecipient <MailboxIdentity>
```

This example assigns the existing mailbox named "Contoso System Mailbox" to the Exchange recipient.

```
Set-OrganizationConfig -MicrosoftExchangeRecipientReplyRecipient "Contoso System Mailbox"
```

Step 2: Specify the enhanced status codes that you want to monitor

- You can use the EAC or the Exchange Management Shell.

- By default, even though there are no enhanced status codes specified, NDRs for these codes are automatically sent to the Exchange recipient:

  - `5.1.4`

  - `5.2.0`

  - `5.2.4`

  - `5.4.4`

  - `5.4.6`

  - `5.4.8`

- You can only specify enhanced status codes. You can't specify quotas.

Use the EAC to specify the enhanced status codes to monitor

For more information about the EAC, see Exchange admin center in Exchange Server.

1. In the EAC, go to **Mail flow** > **Receive connectors**.

2. Click **More options** (•••) and select **Organization transport settings**.

3. In the **Organization transport settings** window that opens, click the **Delivery** tab. In the **DSN codes** section, do one or more of these steps:

   - To add entries, type the enhanced status code that you want to monitor (4. *<y.z>* or 5. *<y.z>*), and then click **Add** (+). Repeat this step as many times as you need to.

   - To modify an existing entry, select it click **Edit** (✏), and then modify it inline.

   - To remove an existing entry, select it and then click **Remove** (—).

When you're finished, click **Save**.

Use the Exchange Management Shell to specify the enhanced status codes to monitor

To add enhanced status codes to monitor, which replaces any existing values, use this syntax:

```
Set-TransportConfig -GenerateCopyOfDSNFor <x.y.z>,<x.y.z>...
```

This example configures the Exchange organization to forward all NDRs for the enhanced status code values 5.7.1, 5.7.2, and 5.7.3 to the Exchange recipient.

```
Set-TransportConfig -GenerateCopyOfDSNFor 5.7.1,5.7.2,5.7.3
```

To add or remove entries without modifying any existing values, use this syntax:

```
Set-TransportConfig -GenerateCopyOfDSNFor @{Add="<x.y.z>","<x.y.z>"...; Remove="<x.y.z>","<x.y.z>"...}
```

This example adds the enhanced status code 5.7.5 and removes 5.7.1 from the existing list of NDRs that are forwarded to the Exchange recipient.

```
Set-TransportConfig -GenerateCopyOfDSNFor @{Add="5.7.5"; Remove="5.7.1"}
```

How do you know this worked?

To verify that you've successfully configured copies of NDRs to be sent to the Exchange recipient mailbox,

- Run the following command and verify the property values:

  ```
  Get-TransportConfig | Format-List GenerateCopyOfDSNFor
  ```

- Monitor the Exchange recipient mailbox to see if NDRs that contain the specified enhanced status codes are delivered there.

Identity values for NDRs

The identity of an NDR uses one of these formats:

- **NDRs for enhanced status codes**: *<Language>*\<Internal | External>\ *<DSNcode>*. For example, `En\Internal\5.1.2` or `Ja\External\5.1.2`.
  - **<DSNcode>**: Valid values are 4. *x*. *y* or 5. *x*. *y* where *x* and *y* are one to three digit numbers. To generate a list of the enhanced status codes that are used by Exchange, see the Use the Exchange Management Shell to view all default NDRs section earlier in this topic.
  - **Internal or External**: You can use different text in NDRs for internal or external senders.
  - **<Language>**: For the list of supported languages, see the Supported languages for NDRs section in this topic.
- **NDRs for quotas**: *<Language>*\ *<QuotaMessageType>*. For example, `En\ProhibitSendReceiveMailBox`.
  - **<Language>**: For the list of supported languages, see the Supported languages for NDRs section in this topic.
  - **<QuotaMessageType>**: Valid values are:

Mailbox size quotas:

- **ProhibitSendReceiveMailBox**: A mailbox exceeds its `ProhibitSendReceiveQuota` limit.
- **ProhibitSendMailbox**: A mailbox exceeds its `ProhibitSendQuota` limit.
- **WarningMailbox**: A mailbox exceeds its `IssueWarningQuota` limit when it has a `ProhibitSendQuota` or `ProhibitSendReceiveQuota` limit configured.
- **WarningMailboxUnlimitedSize**: A mailbox exceeds its `IssueWarningQuota` limit when it doesn't have a `ProhibitSendQuota` or `ProhibitSendReceiveQuota` limit configured.

Public folder size quotas:

- **ProhibitPostPublicFolder**: A public folder exceeds its `ProhibitPostQuota` limit.
- **WarningPublicFolder**: A public folder exceeds its `IssueWarningQuota` limit when it has a `ProhibitPostQuota` limit configured.
- **WarningPublicFolderUnlimitedSize**: A public folder exceeds its `IssueWarningQuota` limit when it doesn't have a `ProhibitPostQuota` limit configured.

Maximum number of messages in a mailbox folder:

- **ProhibitReceiveMailboxMessagesPerFolderCount**: A mailbox exceeds its `MailboxMessagesPerFolderCountReceiveQuota` limit.
- **WarningMailboxMessagesPerFolderCount**: A mailbox exceeds its `MailboxMessagesPerFolderCountWarningQuota` limit when it has a `ailboxMessagesPerFolderCountReceiveQuota` limit configured.
- **WarningMailboxMessagesPerFolderUnlimitedCount**: A mailbox exceeds its `MailboxMessagesPerFolderCountWarningQuota` limit when it doesn't have a `MailboxMessagesPerFolderCountReceiveQuota` limit configured.

Maximum number of subfolders in a mailbox folder:

- **ProhibitReceiveFolderHierarchyChildrenCountCount**: A mailbox exceeds its `FolderHierarchyChildrenCountReceiveQuota` limit.
- **WarningFolderHierarchyChildrenCount**: A mailbox exceeds its `FolderHierarchyChildrenCountWarningQuota` limit when it has a `FolderHierarchyChildrenCountReceiveQuota` limit configured.
- **WarningFolderHierarchyChildrenUnlimitedCount**: A mailbox exceeds its `FolderHierarchyChildrenCountWarningQuota` limit when it doesn't have a `FolderHierarchyChildrenCountReceiveQuota` limit configured.
- **ProhibitReceiveFoldersCount**: A mailbox exceeds its `FoldersCountReceiveQuota` limit.
- **WarningFoldersCount**: A mailbox exceeds its `FoldersCountWarningQuota` limit when it has a `FoldersCountReceiveQuota` limit configured.
- **WarningFoldersCountUnlimited** A mailbox exceeds its `FoldersCountWarningQuota` limit when it doesn't have a `FoldersCountReceiveQuota` limit configured.

Maximum number of levels (depth) in a mailbox folder:

- **ProhibitReceiveFolderHierarchyDepth**: A mailbox exceeds its `FolderHierarchyDepthWarningQuota` limit.

- **WarningFolderHierarchyDepth**: A mailbox exceeds its `FolderHierarchyDepthWarningQuota` limit when it has a `FolderHierarchyDepthReceiveQuota` limit configured.
- **WarningFolderHierarchyDepthUnlimited**:: A mailbox exceeds its `FolderHierarchyDepthWarningQuota` limit when it doesn't have a `FolderHierarchyDepthReceiveQuota` limit configured.

Supported languages for NDRs

This table lists the supported language that codes you can use in custom NDRs.

LANGUAGE CODE	LANGUAGE
af	Afrikaans
am-ET	Amharic (Ethiopia)
ar	Arabic
as-IN	Assamese (India)
bg	Bulgarian
bn-BD	Bengali (Bangladesh)
bn-IN	Bengali (India)
bs-Cyrl-BA	Bosnian (Cyrillic, Bosnia and Herzegovina)
bs-Latn-BA	Bosnian (Latin, Bosnia and Herzegovina)
ca	Catalan
cs	Czech
cy-GB	Welsh (Great Britain)
da	Danish
de	German
el	Greek
en	English
es	Spanish
et	Estonian
eu	Basque
fa	Persian

LANGUAGE CODE	LANGUAGE
fi	Finnish
fil-PH	Filipino (Philippines)
fr	French
ga-IE	Irish (Ireland)
gl	Galician
gu	Gujarati
ha-Latn-NG	Hausa (Latin, Nigeria)
he	Hebrew
hi	Hindi
hr	Croatian
hu	Hungarian
hy	Armenian
id	Indonesian
ig-NG	Igbo (Nigeria)
is	Icelandic
it	Italian
iu-Latn-CA	Inuktitut (Latin, Canada)
ja	Japanese
ka	Georgian
kk	Kazakh
km-KH	Khmer (Cambodia)
kn	Kannada
ko	Korean
kok	Konkani
ky	Kyrgyz

LANGUAGE CODE	LANGUAGE
lb-LU	Luxembourgish (Luxembourg)
lo-LA	Lao (Lao People's Democratic Republic)
lt	Lithuanian
lv	Latvian
mi-NZ	Maori (New Zealand)
mk	Macedonian
ml-IN	Malayalam (India)
mr	Marathi
ms	Malay
ms-BN	Malay (Brunei Darussalam)
mt-MT	Maltese (Malta)
ne-NP	Nepali (Nepal)
nl	Dutch
nn-NO	Norwegian (Nynorsk)
no	Norwegian
nso-ZA	Sesotho sa Leboa (South Africa)
or-IN	Oriya (India)
pa	Punjabi
pl	Polish
ps-AF	Pashto (Afghanistan)
pt	Portuguese
pt-PT	Portuguese (Portugal)
qut-GT	K'iche (Guatemala)
quz-PE	Quechua (Peru)
ro	Romanian

LANGUAGE CODE	LANGUAGE
ru	Russian
rw-RW	Kinyarwanda (Rwanda)
si-LK	Sinhala (Sri Lanka)
sk	Slovak
sl	Slovenian
sq	Albanian
sr	Serbian
sr-Cyrl-CS	Serbian (Cyrillic, Serbia)
sv	Swedish
sw	Kiswahili
ta	Tamil
te	Telugu
th	Thai
tn-ZA	Setswana (South Africa)
tr	Turkish
tt	Tatar
uk	Ukrainian
ur	Urdu
uz	Uzbek
vi	Vietnamese
wo-SN	Wolof (Senegal)
xh-ZA	isiXhosa (South Africa)
yo-NG	Yoruba (Nigeria)
zh-Hans	Chinese (Simplified)
zh-Hant	Chinese (Traditional)

LANGUAGE CODE	LANGUAGE
zh-HK	Chinese (Hong Kong)
zu-ZA	isiZulu (South Africa)

To control the languages that are used in NDRs, you use these parameters on the **Set-TransportConfig** cmdlet:

- *ExternalDsnDefaultLanguage*: Specifies the default language to use on external NDRs. The default value is blank (`$null`), which means the default Windows server language is used.
- *InternalDsnDefaultLanguage*: Specifies the default language to use on internal NDRs. The default value is blank (`$null`), which means the default Windows server language is used.
- *ExternalDsnLanguageDetectionEnabled*
  - `$true` : Exchange tries to send an external NDR in the same language as the original message. This is the default value.
  - `$false` : Language detection is disabled for external NDRs, The NDR language is determined by the *ExternalDsnDefaultLanguage* parameter.
- *InternalDsnLanguageDetectionEnabled*
  - `$true` : Exchange tries to send an internal NDR in the same language as the original message. This is the default value.
  - `$false` : Language detection is disabled for internal NDRs, The NDR language is determined by the *InternalDsnDefaultLanguage* parameter.

HTML tags and special characters in NDRs

The custom text that you include in an NDR can contain a maximum of 512 characters, which includes text and HTML tags. For example, you can include a detailed description of the problem, contact information for your help desk, and a link to your support department's web site.

To control whether Exchange uses HTML or plain text in NDRs, you use these parameters on the **Set-TransportConfig** cmdlet:

- *ExternalDsnSendHtml*
  - `$true` : Use HTML tags in NDRs for external senders. This is the default value.
  - `$false` : Use plain text in NDRs for external senders.
- *InternalDsnSendHtml*
  - `$true` : Use HTML tags in NDRs for internal senders. This is the default value
  - `$false` : Use plain text in NDRs for internal senders.

This table describes the HTML tags that you can use in the NDR text.

DESCRIPTION	HTML TAGS
Bold	`<B>` and `</B>`
Italic	`<EM>` and `</EM>`

DESCRIPTION	HTML TAGS
Line break	`<BR>`
Paragraph	`<P>` and `</P>`
Hyperlink	`<A HREF="url">` and `</A>` **Note**: Because this tag contains double quotation marks, you need to use single quotation marks (not double quotation marks) around the complete text string if you use this tag in your custom text. Otherwise, you'll receive an error.

Certain characters in an NDR require escape codes to identify them literally, and not by their function in the NDR. These characters are described in the following table:

CHARACTER	ESCAPE CODE
<	`&lt;`
>	`&gt;`
"	`&quot;`
&	`&amp;`

For example, if you want the NDR to display the text `Please contact the Help Desk at <1234>.`, you need to the value `"Please contact the Help Desk at &lt;1234&gt;."`

This is an example of a custom NDR text value that uses HTML tags and escape codes.

```
'You tried to send a message to a <B>disabled</B> mailbox. Please visit <A
HREF="https://it.contoso.com">Internal Support</A> or contact &quot;InfoSec&quot; for more information.'
```

Content conversion

8/3/2020 • 14 minutes to read • Edit Online

Content conversion is the process of correctly formatting a message for each recipient. The decision to perform content conversion on a message depends on the destination and format of the message. They types of content conversion that occur in Exchange 2016 and Exchange 2019 are unchanged from Exchange 2013:

- **Message conversion for external recipients**: This type of content conversion includes the Transport Neutral Encapsulation Format (TNEF) conversion options and message encoding options for external recipients. Messages sent to recipients inside the Exchange organization don't require this type of content conversion. This type of content conversion is handled by the categorizer in the Transport service on a Mailbox server. Categorization on each message happens after a newly arrived message is put in the Submission queue. In addition to recipient resolution and routing resolution, content conversion is performed on the message before the message is put in a delivery queue. If a single message contains multiple recipients, the categorizer determines the appropriate encoding for each message recipient. Content conversion tracing doesn't capture any content conversion failures that the categorizer encounters as it converts messages sent to external recipients.

- **MAPI conversion for internal recipients**: his type of content conversion is handled by the Mailbox Transport service. The Mailbox Transport service exists on Mailbox servers to transmit messages between mailbox databases on the local server, and the Transport service on Mailbox servers. Specifically, the Mailbox Transport Submission service transmits messages from the sender's Outbox to the Transport service on a Mailbox server. The Mailbox Transport Delivery service transmits messages from the Transport service on a Mailbox server to the recipient's Inbox. The Mailbox Transport Submission service converts all outgoing messages from MAPI and the Mailbox Transport Delivery service converts all incoming messages to MAPI. Content conversion tracing captures these MAPI conversion failures. For more information, see Managing Content Conversion Tracing.

Exchange and Outlook message formats

The following list describes the basic message formats available in Exchange and Outlook:

- **Plain text**: A plain text message uses only US-ASCII text as described in RFC 5322. The message can't contain different fonts or other text formatting. The following two formats can be used for a plain text message:

  - The message headers and the message body are composed of US-ASCII text. Attachments must be encoded by using *Uuencode*. Uuencode represents Unix-to-Unix encoding and defines an encoding algorithm to store binary attachments in the body of an email message by using US-ASCII text characters.

  - The message is MIME-encoded with a **Content-Type** value of `text/plain`, and a **Content-Transfer-Encoding** value of `7bit` for the text parts of a multipart message. Any message attachments are encoded by using Quoted-printable or Base64 encoding. By default, when you compose and send a plain text message in Outlook, the message is MIME-encoded with a **Content-Type** value of `text/plain`.

- **HTML**: An HTML message supports text formatting, background images, tables, bullet points, and other graphical elements. By definition, an HTML-formatted message must be MIME-encoded to preserve these formatting elements.

- **Rich text format (RTF)**: RTF supports text formatting and other graphical elements. RTF is synonymous with TNEF (TNEF and RTF can be used interchangeably). The rich text message format is completely different from the rich text document format that's available in Word.

- **TNEF**: The Transport Neutral Encapsulation Format is a Microsoft-specific format for encapsulating MAPI message properties. A TNEF message contains a plain text version of the message and an attachment that packages the original formatted version of the message. Typically, this attachment is named Winmail.dat. The Winmail.dat attachment includes the following information:

  - Original formatted version of the message (for example, fonts, text sizes, and text colors)

  - OLE objects (for example, embedded pictures or embedded Office documents)

  - Special Outlook features (for example, custom forms, voting buttons, or meeting requests)

  - Regular message attachments that were in the original message

    The resulting plain text message can be represented in the following formats:

  - RFC 5322-compliant message composed of only US-ASCII text with a Winmail.dat attachment encoded in Uuencode

  - Multipart MIME-encoded message that has a Winmail.dat attachment

    Outlook and other email clients that fully understand TNEF process the Winmail.dat attachment and display the original message content without ever displaying the Winmail.dat attachment. Email clients that don't understand TNEF may present TNEF messages in any of the following ways:

  - The plain text version of the message is displayed, and the message contains an attachment named Winmail.dat, Win.dat, or some other generic name such as Att *nnnnn*.dat or Att *nnnnn*.eml where the *nnnnn* placeholder represents a random number.

  - The plain text version of the message is displayed. The TNEF attachment is ignored or removed. The result is a plain text message.

  - Messaging servers that understand TNEF can be configured to remove TNEF attachments from incoming messages. The result is a plain text message. Moreover, some email clients may not understand TNEF, but recognize and ignore TNEF attachments. The result is a plain text message.

    There are third-party utilities that can help convert Winmail.dat attachments.

    TNEF is understood by all versions of Exchange since Exchange Server version 5.5.

- **Summary Transport Neutral Encapsulation Format (STNEF)**: STNEF is equivalent to TNEF. However, STNEF messages are encoded differently than TNEF messages. Specifically, STNEF messages are always MIME-encoded, and always have the **Content-Transfer-Encoding** value `Binary`. Therefore, there's no plain text representation of the message, and there's no distinct Winmail.dat attachment contained in the body of the message. The whole message is represented by using only binary data. Messages that have a **Content-Transfer-Encoding** value of **Binary** can only be transferred between messaging servers that support and advertise the **BINARYMIME** and **CHUNKING** SMTP extensions as defined in RFC 3030. The messages are always transferred between messaging servers by using the **BDAT** command, instead of the standard **DATA** command.

  STNEF is understood by all versions of Exchange since Exchange 2000. STNEF is automatically used for all messages transferred between Exchange servers in the organization since native mode Exchange Server 2003.

  Exchange never sends STNEF messages to external recipients. Only TNEF messages can be sent to recipients outside the Exchange organization.

Content conversion options for external recipients

The content conversion options that you can set in an Exchange organization for external recipients can be described in the following categories:

- **TNEF conversion options**: These conversion options specify whether TNEF should be preserved or removed from messages that leave the Exchange organization.
- **Message encoding options**: These options specify message encoding options, such as MIME and non-MIME character sets, message encoding, and attachment formats.

These conversion and encoding options are independent of one another. For example, whether TNEF messages can leave the Exchange organization isn't related to the MIME encoding settings or plain text encoding settings of those messages.

You can specify the content conversion at various levels of the Exchange organization as described in the following list:

- **Remote domain settings**: Remote domains define the settings for outgoing message transfers between the Exchange organization and external domains.. Even if you don't create remote domain entries for specific domains, there's a predefined remote domain named Default that applies to all remote address spaces (*). For more information about remote domains, see Remote Domains.
- **Mail user and mail contact settings**: Mail users and mail contacts are similar because both have external email addresses and contain information about people outside the Exchange organization. The main difference is mail users have accounts that they can use to log on to Active Directory and access resources in the organization. For more information, see Recipients.
- **Outlook settings**: You can set these message formatting and encoding options in Outlook:
  - **Message format**: You can set the default message format for all messages. You can override the default message format as you compose a specific message.
  - **Internet message format**: You can control whether TNEF messages are sent to remote recipients or whether they are first converted to a more compatible format. You can also specify various message encoding options for messages sent to remote recipients. These settings don't apply to messages sent to recipients in the Exchange organization.
  - **Internet recipient message format (Outlook 2010 or earlier)**: You can control whether TNEF messages are sent to specific contacts in your Contacts folder. These conversion options aren't available for recipients in the Exchange organization.
  - **Internet recipient message encoding options (Outlook 2010 or earlier)**: You can control the MIME or plain text encoding options for specific contacts in your Contacts folder. These conversion options aren't available for recipients in the Exchange organization.
  - **International options**: You can control the character sets used in messages.

  For more information about these settings, see TNEF conversion options and Message encoding options in Exchange Server.

Understanding the structure of email messages

To better understand the content conversion options for external recipients, you need to understand the structure of email messages. An SMTP message is based on plain 7-bit US-ASCII text to compose and send email messages. A standard SMTP message consists of the following elements:

- **Message envelope**: The message envelope is defined in RFC 5321. The message envelope contains information required to transmit and deliver the message. Recipients never see the message envelope, because it's generated by the message transmission process and isn't actually part of the message contents.

- **Message contents**: The message contents are defined in RFC 5322. The message contents consist of the following elements:
  - **Message header**: The message header is a collection of header fields. Header fields consist of a field name, followed by a colon (:) character, followed by a field body, and ended by a carriage return/line feed (CR/LF) character combination.

    A field name must be composed of printable US-ASCII text characters except the colon (:) character. Specifically, ASCII characters that have values from 33 through 57 and 59 through 126 are permitted.

    A field body may be composed of any US-ASCII characters, except for the carriage return (CR) character and the line feed (LF) character. However, a field body may contain the CR/LF character combination when used in *header folding*. Header folding is the separation of a single header field body into multiple lines as described in section 2.2.3 of RFC 5322. Other field body syntax requirements are described in sections 3 and 4 of RFC 5322.

  - **Message body**: The message body is a collection of lines of US-ASCII text characters that appears after the message header. The message header and the message body are separated by a blank line that ends with the CR/LF character combination. The message body is optional. Any line of text in the message body must be less than 998 characters. The CR and LF characters can only appear together to indicate the end of a line.

When SMTP messages contain elements that aren't plain US-ASCII text, the message must be encoded to preserve those elements. The MIME standard defines a method of encoding content in messages that isn't text. MIME allows for text in other character sets, attachments without text, multipart message bodies, and header fields in other character sets. MIME is defined in RFC 2045, RFC 2046, RFC 2047, RFC 4288, RFC 4289, and RFC 2049. MIME defines a collection of header fields that specifies additional message attributes. The following sections describe some important MIME header fields.

MIME-Version header field

Default value: `1.0`

This header field is the first MIME header field that appears in a MIME-formatted message. This header field appears after the other standard RFC 5322 header fields, but before any other MIME header fields. MIME-aware email clients use this header field to identify a MIME-encoded message. When this header field is absent, MIME-aware email clients identify the message as plain text.

Content-Type header field

Default value: `text/plain`

This header field identifies the media type of the message content as described in RFC 2046. A media type consists of:

- **A type**:
  - Types that begin with `x-` aren't standard. The Internet Assigned Numbers Authority (IANA) maintains a list of registered media types. For more information, see MIME Media Types.
  - The *multipart* media type allows for multiple message parts in the same message by using sections defined by different media types. Some **Content-Type** field values include `text/plain`, `text/html`, `multipart/mixed`, and `multipart/alternative`.
- **A subtype**: Subtypes that begin with `vnd.` are vendor-specific.
- **One or more optional parameters**: For example, a `charset=` parameter that defines the MIME character

encoding.

Content-Transfer-Encoding header field

Default value: `7bit`

This header field can describe the following information about a message:

- The encoding algorithm used to transform any non-US-ASCII text or binary data that exists in the message body.
- An indicator that describes the current condition of the message body.

There can be multiple values of the **Content-Transfer-Encoding** header field in a MIME message. When the **Content-Transfer-Encoding** header field appears in the message header, it applies to the whole body of the message. When the **Content-Transfer-Encoding** header field appears in one of the parts of a multipart message, it applies only to that part of the message.

When an encoding algorithm is applied to the message body data, the message body data is transformed into plain US-ASCII text. This transformation allows the message to travel through older messaging servers that only support messages in US-ASCII text. The **Content-Transfer-Encoding** header field values that indicate an encoding algorithm was used on the message body are:

- `Quoted-printable`: Uses printable US-ASCII characters to encode the message body data. If the original message text is mostly US-ASCII text, Quoted-printable encoding gives somewhat readable and compact results. All printable US-ASCII text characters except the equal sign (=) character can be represented without encoding.
- `Base64`: Based primarily on the privacy-enhanced mail (PEM) standard defined in RFC 4648. Base64 encoding uses the 64-character alphabet encoding algorithm and output padding characters defined by PEM to encode the message body data. A Base64 encoded message is typically 33 percent larger than the original message. Base64 encoding creates a predictable increase in message size and is optimal for binary data and non-US-ASCII text.

Typically, you won't see multiple encoding algorithms used in the same message.

When no encoding algorithm has been used on the message body, the **Content-Transfer-Encoding** header field merely identifies the current condition of the message body data. The **Content-Transfer-Encoding** header field values that indicate that no encoding algorithms were used on the message body are:

- `7bit`: Indicates that the message body data is already in the RFC 5322 format. Specifically, this means that the following conditions must be true:
  - All lines of text must be less than 998 characters long.
  - All characters must be US-ASCII text that have character values from 1 through 127.
  - The CR and LF characters can only be used together to indicate the end of a line of text.

    The whole message body may be 7-bit, or part of the message body in a multipart message may be 7-bit. If the multipart message contains other parts that have any binary data or non-US-ASCII text, that part of the message must be encoded using the Quoted-printable or Base64 encoding algorithms.

    Messages that have 7-bit bodies can travel between messaging servers by using the standard DATA command.
- `8bit`: Indicates that the message body data contains non-US-ASCII characters. Specifically, this means that the following conditions must be true:

  - All lines of text must be less than 998 characters long.
  - One or more characters in the message body have values larger than 127.
  - The CR and LF characters can only be used together to indicate the end of a line of text.

    The whole message body may be 8-bit, or part of the message body in a multipart message may be 8-bit. If the multipart message contains other parts that have binary data, that part of the message must be encoded using the Quoted-printable or Base64 encoding algorithms.

    Messages that have 8-bit bodies can only travel between messaging servers that support the **8BITMIME** SMTP extension as defined in RFC 6152, such as Exchange 2000 Server or later. Specifically, this means that the following conditions must be true:

  - The **8BITMIME** keyword must be advertised in the server's EHLO response.
  - Messages are still transferred by using the SMTP standard **DATA** command. However, the `BODY=8BITMIME` parameter must be added to the end of the **MAIL FROM** command.

- `Binary`: Indicates that the message body contains non-US-ASCII text or binary data. Specifically, this means that the following conditions are true:

  - Any sequence of characters is allowed.
  - There is no line length limitation.
  - Binary message elements don't require encoding.

    Messages that have binary bodies can only travel between messaging servers that support the **BINARYMIME** SMTP extension as defined in RFC 3030, such as Exchange 2000 Server or later. Specifically, this means that the following conditions must be true:

  - The **BINARYMIME** keyword must be advertised in the server's EHLO response.
  - The **BINARYMIME** SMTP extension can only be used with the **CHUNKING** SMTP extension. *Chunking* enables large message bodies to be sent in multiple, smaller chunks. Chunking is also defined in RFC 3030. The **CHUNKING** keyword must also be advertised in the server's EHLO response.
  - Messages are transferred using the **BDAT** command instead of the standard **DATA** command.
  - The `BODY=BINARYMIME` parameter must be added to the end of the **MAIL FROM** command when the message has a message body.

The values `7bit`, `8bit`, and `Binary` never exist together in the same multipart message (the values are mutually exclusive). The `Quoted-printable` or `Base64` values may appear in a 7-bit or 8-bit multipart message body, but never in a binary message body. If a multipart message body contains different parts composed of 7-bit and 8-bit content, the whole message is classified as 8-bit. If a multipart message body contains different parts composed of 7-bit, 8-bit, and binary content, the whole message is classified as binary.

Content-Disposition header field

Default value: `Attachment`

This header field instructs a MIME-enabled email client on how it should display an attached file, and is described in RFC 2183. Valid values are:

- `Inline`: The attachment is displayed in the message body.
- `Attachment`: The attached file appears as a regular attachment separate from the message body. Other parameters are also with this values (for example, `Filename`, `Creation-date`, and `Size`).

Message encoding options in Exchange Server

8/3/2020 • 9 minutes to read • Edit Online

The message encoding options in Exchange Server let you specify message characteristics such as MIME and non-MIME character sets, binary encoding, and attachment formats. You can specify message encoding options in the following locations:

- Remote domain settings
- Mail contact and mail user settings
- Outlook settings:
  - Message format
  - Internet message format
  - Internet recipient message format (Outlook 2010 or earlier)
  - Message character set encoding options
- Outlook on the web (formerly known as Outlook Web App) message format settings

Typically, the default settings for these message encoding options will work fine. However, you might need to change the messaging encoding options for recipients that are using older email clients or messaging systems. They'll likely tell you if messages from your Exchange environment appear to have formatting issues.

For more information about content conversion in Exchange, see Content conversion. For TNEF (also known as or Rich Text) settings, see TNEF conversion options.

Remote domain settings

Remote domains specify settings for messages sent to domains that are external to your Exchange organization. For more information, see Remote Domains.

When you configure message encoding options for a remote domain, the settings are applied to all messages that are sent to recipients in that domain. Some settings are available in the Exchange admin center (EAC), but most are only available in the Exchange Management Shell. The message encoding settings are described in this table:

SETTING	EAC CONFIGURATION	EXCHANGE MANAGEMENT SHELL CONFIGURATION
MIME character set: The specified character set is only used for MIME messages that don't contain a character set. This setting won't overwrite character sets that are already specified in outgoing messages. **Non-MIME character set**: This setting is used if either of these conditions are true: • Incoming messages from a remote domain are missing the value of the *charset=* setting in the MIME **Content-Type:** header field. • Outgoing messages to a remote domain are missing the value of the MIME character set.	**Mail flow** > **Remote domains** > **Add** ＋, or select an existing remote domain, and then click **Edit** ✎ > **Supported character set** section.	Cmdlet: **Set-RemoteDomain** Parameters: *CharacterSet* and *NonMimeCharacterSet*
Content type: Valid values are: `MimeHtmlText`: All messages are converted to MIME messages that use HTML formatting, unless the original message is a text message. If the original message is a text message, the outgoing message will be a MIME message that uses text formatting. This is the default value. `MimeText`: All messages are converted to MIME messages that use text formatting. `MimeHtml`: All messages are converted to MIME messages that use HTML formatting.	n/a	Cmdlet: **Set-RemoteDomain** Parameter: *ContentType*
Line wrap size: You can specify the maximum number of characters that can exist on a single line of text in the body of the email message. Older email clients might prefer 78 characters per line.	n/a	Cmdlet: **Set-RemoteDomain** Parameter: *LineWrapSize* The default value is `Unlimited`, which means the email client is responsible for setting the line wrap size in new messages.

Mail contact and mail user settings

Mail contacts and mail users represent users that have external email addresses in your Exchange organization. For more information, see Recipients.

When you configure message encoding options for a mail contact or a mail user, the settings are only applied to messages that are sent to that specific recipient. All settings are only available in the Exchange Management Shell in these cmdlets:

- Enable-MailContact, New-MailContact, or Set-MailContact.
- Enable-MailUser, New-MailUser, or Set-MailUser.

The message encoding settings for mail contacts and mail users are described in this list:

- *UsePreferMessageFormat parameter**: Specifies whether the message format settings for the mail contact or mail user override the corresponding settings for the remote domain. Valid values are:

  - `$true`: Messages sent to the Mail contact or mail user use the message format that's configured for the Mail contact or mail user.

  - `$false`: Messages sent to the Mail contact or mail user use the message format that's configured for the remote domain (the default remote domain or a specific remote domain) or configured by the message sender. This is the default value.

- **MessageFormat parameter**: This parameter specifies the message format for messages sent to the mail contact or mail user. Valid values are `Text` or `Mime`, and the default value is `Mime`.

- **MessageBodyFormat parameter**: This parameter specifies the message body format for messages sent to the mail contact or mail user. Valid values are `Text`, `Html`, or `TextAndHtml`, and the default value is `TextAndHtml`.

  The *MessageFormat* and *MessageBodyFormat* parameters are interdependent:

  - If the *MessageFormat* value is `Mime`, the *MessageBodyFormat* value can be `Text`, `Html`, or `TextAndHtml`.

  - If the *MessageFormat* value is `Text`, the *MessageBodyFormat* value can only be `Text`.

- **MacAttachmentFormat parameter**: Specifies the message attachment format for Apple Macintosh operating system clients. Valid values are `BinHex`, `UuEncode`, `AppleSingle`, or `AppleDouble`, and the default value is `BinHex`.

  The *MessageFormat* and *MacAttachmentFormat* parameters are interdependent:

  - If the *MessageFormat* value is `Text`, the *MacAttachmentFormat* value can be `BinHex` or `UuEncode`.

  - If the *MessageFormat* value is `Mime`, the *MacAttachmentFormat* value can be `BinHex`, `AppleSingle`, or `AppleDouble`.

Outlook settings

As a sender, you can specify the message encoding in Outlook by using any of these methods:

- Configure the default message format to plain text or HTML.

- Configure the message format to plain text or HTML as you're composing the message by using the **Format** area in the **Format Text** tab.

- Configure the message encoding options for messages sent to all external recipients. These options are called *Internet message format* options, and they only apply to remote recipients (not to recipients in the Exchange organization).

- Configure the message encoding options for messages sent to specific external recipients (Outlook 2010 or earlier). These options are called *Internet recipient message format* options, and they only apply to remote recipients in your Contacts folder (not to recipients in the Exchange organization).

For instructions on configuring these settings in Outlook, see Change the message format to HTML, Rich Text Format, or plain text.

By default, Outlook uses automatic character set message encoding by scanning the whole text of the outgoing message to determine the appropriate encoding to use for the message. This setting applies to internal and external recipients. However, you can bypass the automatic selection and specify a preferred encoding for outgoing messages at **File** > **Options** > **Advanced** > **International options**.

Outlook on the web settings

As a sender, you can specify message encoding options in Outlook on the web by using either of these methods:

- Configure the default message format as plain text or HTML in the **Message format** section at **Settings** > **Options** > **Mail** > **Layout**.

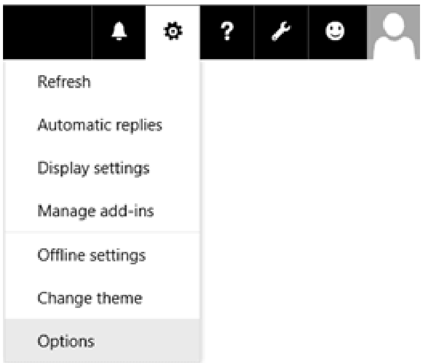

- Configure the message format to plain text or HTML as you're composing the message by clicking **More options** •••, and selecting **Switch to plain text** (if the current format is HTML) or **Switch to HTML** (if the current format is plain text).

Order of precedence for message encoding options

Some message encoding options are available in remote domain settings, Mail contact or mail user settings, and Outlook or Outlook on the web settings. Message encoding options for outgoing messages sent to external recipients are described in the following list from highest priority to lowest priority:

1. Mail contact or mail user settings (if the use preferred message format setting is enabled)

2. Outlook or Outlook on the web settings

3. Remote domain settings

A setting at a higher level overrides the corresponding setting at a lower level. For example, Mail contact or mail user settings override the corresponding setting for a remote domain. Unique settings are unaffected (there's no higher or lower priority setting that conflicts).

The order of precedence for message encoding options are described in the following sections.

Order of precedence for message character sets

The following table describes the order of precedence from highest priority to lowest priority for message character set encoding options.

SOURCE	SETTING	VALUES
Outlook	**Preferred encoding for outgoing messages**	**Automatically select encoding for outgoing messages** enabled or disabled (enabled by default). **Preferred encoding for outgoing messages** set to the specified character set. This is the encoding option that's used if you disable **Automatically select encoding for outgoing messages**
Remote domain	MIME character set and non-MIME character set	The specified MIME and non-MIME character sets (which can be the same).

Notes:

- When you configure the non-MIME character set for a remote domain, the character set is assigned to incoming or outgoing messages to and from the remote domain that don't contain a specified character set.
- The value of the Windows ANSI code page for the Exchange server is used to assign a character set to these types of messages:
  - Internal messages that don't contain a specified character set.
  - Internal messages that contain a specified character set, but don't contain a specified server code page.
- If a message contains a specified but invalid character set, the Exchange server tries to replace the invalid character set with a valid one.

Order of precedence for plain text message encoding options

The following table describes the order of precedence from highest priority to lowest priority for plain text message encoding options.

Note: Only plain text message settings are included here (not plain text settings for MIME encoded messages).

SOURCE	SETTING	VALUES
Mail contact or mail user	Use the preferred message format	If the value `$true`, the plain text message encoding settings for the mail contact or mail user override the corresponding settings in Outlook. If the value is `$false`, the plain text message encoding settings for the mail contact or mail user are ignored (the corresponding settings in Outlook are used).
Mail contact or mail user	Message format	Text
Mail contact or mail user	Message body format	Text
Mail contact or mail user	Mac attachment format	`BinHex` or UUEncode
Outlook 2010 or earlier	Internet recipient message format (settings on a specific contact)	Send plain text only Open a contact in the Contacts folder > double-click the email address > click **View more options for interacting with this person** > select **Outlook properties**, In the **E-mail Properties** dialog that opens, select **Send Plain Text only** in the **Internet format** field.
Outlook	Internet message format	Plain text options for external messages at **File** > **Options** > **Mail** > **Message format**: **Encode attachments in UUENCODE format when sending plain-text messages** (not selected by default) **Automatically wrap text at nn characters** (the default value is 76).

SOURCE	SETTING	VALUES
Remote domain	Line wrap size	132 characters or less, or the value `Unlimited`. The default value is `Unlimited`.

Order of precedence for MIME message encoding options

The following table describes the order of precedence from highest priority to lowest priority for MIME message encoding options.

SOURCE	SETTING	VALUES
Mail contact or mail user	Use the preferred message format	If the value `$true`, the MIME message encoding settings for the mail contact or mail user override the corresponding settings in Outlook. If the value is `$false`, the MIME text message encoding settings for the mail contact or mail user are ignored (the corresponding settings in Outlook, Outlook on the web, or remote domains are used).
Mail contact or mail user	Message format	MIME
Mail contact or mail user	Message body format	Text, HTML, or `TextAndHtml` (the default value is `TextAndHtml`).
Mail contact or mail user	Mac attachment format	`BinHex`, `AppleSingle`, or `AppleDouble` (the default value is `BinHex`).
Outlook or Outlook on the web	Message format	Plain text or HTML
Remote domain	Content type	`MimeHtmlText` (the default value), `MimeText`, or `MimeHtml`

TNEF conversion options

8/3/2020 • 5 minutes to read • Edit Online

TNEF, also known as the Transport Neutral Encapsulation Format, Outlook Rich Text Format, or Exchange Rich Text Format, is a Microsoft-specific format for encapsulating MAPI message properties. All versions of Outlook fully support TNEF. Outlook on the web (formerly known as Outlook Web App) translates TNEF into MAPI and displays the formatted messages. Other email clients that don't support TNEF typically display TNEF formatted messages as plain text messages with Winmail.dat or Win.dat attachments. For more information about TNEF, see Exchange and Outlook message formats.

Administrators can specify whether TNEF should be preserved or removed from messages that leave their Exchange organization. You can specify TNEF conversion options in the following locations:

- Remote domain settings
- Mail contact and mail user settings
- Outlook settings:
  - Message format
  - Internet message format
  - Internet recipient message format (Outlook 2010 or earlier)

Typically, the default TNEF conversion options will work fine (by default, TNEF messages are converted to HTML for external recipients). However, you might need to force plain text conversion for recipients that are using older email clients or messaging systems. They'll likely tell you if TNEF messages from your Exchange environment appear to have formatting issues.

For more information about other content conversion in Exchange, see Content conversion.

TNEF conversion options for remote domains

Remote domains specify settings for messages sent to domains that are external to your Exchange organization. For more information, see Remote Domains.

When you configure TNEF conversion options for a remote domain, the settings are applied to all messages sent to recipients in that domain. You can use the Exchange admin center (EAC) or the Exchange Management Shell to configure these options:

- In the EAC, go to **Mail flow** > **Remote domains** > **Add** +, or select an existing remote domain, and then click **Edit** ✎ > **Use rich-text format** section.
- In the Exchange Management Shell, use the *TnefEnabled* parameter on the **Set-RemoteDomain** cmdlet.

The TNEF conversion options for remote domains are described in this table:

SETTING	VALUE IN THE EAC	VALUE IN EXCHANGE MANAGEMENT SHELL
Use TNEF for all messages sent to the remote domain.	**Always**	`$true`

SETTING	VALUE IN THE EAC	VALUE IN EXCHANGE MANAGEMENT SHELL
Never use TNEF for any messages sent to the remote domain.	**Never**	`$false`
TNEF messages aren't specifically allowed or prevented for recipients in the remote domain. This is the default value. Whether TNEF messages are sent to recipients in the remote domain depends on the specific setting on the mail contact or mail user, or the setting specified by the sender in Outlook.	**Follow user settings**	`$null` (blank)

TNEF conversion options for mail contacts and mail users

Mail contacts and mail users represent users in your Exchange organization that have external email addresses For more information, see Recipients.

When you configure TNEF conversion options for a mail contact or a mail user, those options are applied to all messages sent to that specific recipient. You use the *UseMapiRichTextFormat* parameter on the **Set-MailUser** and **Set-MailContact** cmdlets in the Exchange Management Shell. Valid values are:

- `Always` : TNEF is used for all messages sent to the recipient.
- `Never` : TNEF is never used for any messages sent to the recipient.
- `UseDefaultSettings` : This is the default value. TNEF messages aren't specifically allowed or prevented for the mail user or mail contact. Whether TNEF messages are sent to the recipient depends on the TNEF conversion setting for the remote domain, or the TNEF conversion setting that's configured by the sender in Outlook.

TNEF conversion options in Outlook

Senders can control the default conversion options for TNEF messages sent to all external recipients. These options are called *Internet message format* options. The options only apply to external recipients, and not to recipients in the Exchange organization.

Note: The following options define how Outlook rich text messages are handled when sent to external recipients. If the messages are HTML or plain text, these settings don't apply.

The following TNEF conversion options are available in Outlook:

- **Convert to HTML format**: This is the default option. TNEF messages sent to external recipients are converted to HTML. Any formatting in the message should closely resemble the original message. MIME-encoded HTML messages are supported by most email clients.
- **Convert to Plain Text format**: Any TNEF messages sent to remote recipients are converted to plain text. Any formatting in the message is lost.
- **Send using Outlook Rich Text Format**: Any TNEF messages sent to remote recipients remain TNEF messages.

Senders in Outlook 2010 or earlier can also control the default TNEF message conversion options for TNEF messages sent to specific external recipients. These options are called *Internet recipient message format* options. The options only apply to external recipients stored in your Contacts folder, and not to recipients in the Exchange organization. The following list describes the TNEF conversion options for an external recipient in your Contacts folder:

- **Let Outlook decide the best sending format**: This is the default setting. This setting forces Outlook to use the TNEF conversion option that's specified by the default Internet format as described in the previous list (**Convert to HTML format**, **Convert to Plain Text format**, or **Send using Outlook Rich Text Format**). Therefore, the TNEF message may be left as TNEF, converted to HTML, or converted to plain text (the default result is converted to HTML). If you want to make sure that the TNEF message remains TNEF for the contact, you should change this setting to **Send using Outlook Rich Text format**.
- **Send Plain Text only**: Any TNEF messages sent to the recipient are converted to plain text. Any formatting in the message is lost.
- **Send using Outlook Rich Text format**: Any TNEF messages sent to remote recipients remain TNEF messages.

To configure the TNEF conversion settings in Outlook, see Change the message format to HTML, Rich Text Format, or plain text.

Order of precedence for TNEF conversion options

The TNEF conversion options for messages sent to external recipients are described in the following list from highest priority to lowest priority:

1. Remote domain settings
2. Mail user or mail contact settings
3. Outlook settings

The setting at a higher level overrides the setting at a lower level. The TNEF setting on the remote domain overrides the TNEF setting on the mail contact or mail user, or the setting in Outlook. For example, suppose you send a Rich Text message in Outlook, but the recipient is in a domain where the remote domain setting specifically doesn't allow TNEF messages. The message received by the recipient will be plain text or HTML, but not TNEF.

Note: Exchange never sends Summary Transport Neutral Encoding Format (STNEF) messages to external recipients. Only TNEF messages can be sent to recipients outside the Exchange organization.

Message size and recipient limits in Exchange Server

8/3/2020 • 14 minutes to read • Edit Online

You can apply limits to messages that move through your organization. You can set the maximum size of an entire message as a whole, or the size of individual parts of a message, or both. For example, you could restrict the maximum size of the message header or attachments, or set a maximum number of recipients that can be added to the message. You can apply these limits to your entire Exchange organization, to specific mail transport connectors, specific servers, and to individual mailboxes.

This topic only talks about message and recipient size limits. If you want to know more about how to control how many messages are sent over time, how many connections are allowed over time, and how long Exchange will wait before closing a connection, see Message rate limits and throttling.

As you plan the message size limits for your Exchange organization, consider the following questions:

- What size limits should I impose on all incoming messages?
- What size limits should I impose on all outgoing messages?
- What is the mailbox quota for my organization, and how do the message size limits that I have chosen relate to the mailbox quota size?
- Are there users in my organization who need to send or receive messages that are larger than the maximum allowed size?
- Does my organization include other messaging systems or separate business units that require different message size limits?

This topic provides guidance to help you answer these questions and to apply the appropriate message size limits in the appropriate locations.

Types of message size limits

The following list describes the basic types of message size limits, and the message components that they apply to.

- **Whole message size limits**: Specifies the maximum size of a message, which includes the message header, the message body, and any attachments. Exchange uses the custom **X-MS-Exchange-Organization-OriginalSize:** message header to record the original size of the message as it enters the Exchange organization. Whenever the message size is checked, the lower value of the current message size or the original message size header is used. The size of the message can change because of content conversion, encoding, and transport agent processing.

  For any message size limit, you need to set a value that's larger than the actual size you want enforced. This accounts for the Base64 encoding of attachments and other binary data. Base64 encoding increases the size of the message by approximately 33%, so the value you specify should be approximately 33% larger than the actual message size you want enforced. For example, if you specify a maximum message size value of 64 MB, you can expect a realistic maximum message size of approximately 48 MB.

- **Attachment size limits**: Specifies the maximum size of a single attachment in a message. The message might contain many smaller attachments that greatly increase its overall size. However, the attachment size limit applies only to the size of an individual attachment. While you can't limit the number of attachments on a message, you can use the maximum message size limit to control the maximum total of attachments on the message.

- **Recipient limits**: Specifies the total number of recipients that are allowed in a message. This includes the total number of recipients in the **To:**, **Cc:**, and **Bcc:** fields. A distribution group counts as a single recipient.

- **Message header size limits**: Specifies the maximum size of all message header fields in a message. The size of the message body or attachments isn't considered. Because the header fields are plain text, the size of the header is determined by the number of characters in each header field and by the total number of header fields. Each text character consumes 1 byte.

Scope of limits

The following tables show the message limits at the Organization, Connector, Server, and Mailbox levels, including information about how to configure the limits in the Exchange admin center (EAC) or the Exchange Management Shell. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.

Organizational limits

Organizational limits apply to all Exchange 2019 servers, Exchange 2016 servers, Exchange 2013 Mailbox servers, and Exchange 2010 Hub Transport servers that exist in your organization. On Edge Transport servers, any organizational limits that you configure are applied to the local server.

> **NOTE**
>
> Organizational limits also apply to external senders and external recipients (anonymous or unauthenticated senders or recipients):
>
> - For inbound messages from external senders, Exchange applies the organizational maximum send message size limit (the maximum receive message size limit as described in the Recipient limits section is applied to the internal recipient).
>
> - For outbound messages to external recipients, Exchange applies the organization maximum receive message size limit (the maximum send message size limit as described in the Recipient limits section is applied to the internal sender).
>
> Therefore, a message size must be within the message size limits for both the sender and the recipient. This concept is also explained in the Order of precedence and placement of message size limits section later in this topic.

SIZE LIMIT	DEFAULT VALUE	EAC CONFIGURATION	EXCHANGE MANAGEMENT SHELL CONFIGURATION
Maximum size of a message received	10 MB	**Mail flow** > **Receive connectors** > **More options** ••• > **Organization transport settings** > **Limits** tab > **Maximum receive message size (MB)**	Cmdlet: **Set-TransportConfig** Parameter: *MaxReceiveSize*
Maximum size of a message sent	10 MB	**Mail flow** > **Receive connectors** > **More options** ••• > **Organization transport settings** > **Limits** > **Maximum send message size (MB)**	Cmdlet: **Set-TransportConfig** Parameter: *MaxSendSize*

SIZE LIMIT	DEFAULT VALUE	EAC CONFIGURATION	EXCHANGE MANAGEMENT SHELL CONFIGURATION
Maximum number of recipients in a message	500	**Mail flow** > **Receive connectors** > **More options** ••• > **Organization transport settings** > **Limits** **Maximum number of recipients**	Cmdlet: **Set-TransportConfig** Parameter: *MaxRecipientEnvelopeLimit*
Maximum attachment size for a message that matches the conditions of the mail flow rule (also known as a transport rule)	Not configured	**Mail flow** > **Rules** > **Add** + > **Create a new rule**, or select an existing rule, and then click **Edit** ✏. Click **More options**. Use the condition **Apply this rule if** > **Any attachment** > **size is greater than or equal to**, and enter a value in kilobytes (KB).	Cmdlets: **New-TransportRule**, **Set-TransportRule** Parameter: *AttachmentSizeOver*
Maximum message size for a message that matches the conditions of the mail flow rule	Not configured	**Mail flow** > **Rules** > **Add** + > **Create a new rule**, or select an existing rule, and then click **Edit** ✏. Click **More options**. Use the condition **Apply this rule if** > **The message** > **size is greater than or equal to**, and enter a value in kilobytes (KB).	Cmdlets: **New-TransportRule**, **Set-TransportRule** Parameter: *MessageSizeOver*

To see the values of these organizational limits, run the following commands in the Exchange Management Shell:

```
Get-TransportConfig | Format-List MaxReceiveSize,MaxSendSize,MaxRecipientEnvelopeLimit
```

```
Get-TransportRule | where {($_.MessageSizeOver -ne $null) -or ($_.AttachmentSizeOver -ne $null)} | Format-Table Name,MessageSizeOver,AttachmentSizeOver
```

Connector limits

Connector limits apply to any messages that use the specified Send connector, Receive connector, Delivery Agent connector, or Foreign connector for message delivery.

You can assign specific message size limits to the Active Directory site links in your organization. The Transport service on Mailbox servers uses Active Directory sites, and the costs that are assigned to the Active Directory IP site links as one of the factors to determine the least-cost routing path between Exchange servers in the organization.

You can assign specific message size limits to the Delivery Agent connectors and Foreign connectors that are used to send non-SMTP messages in your organization.

SIZE LIMIT	DEFAULT VALUE	EAC CONFIGURATION	EXCHANGE MANAGEMENT SHELL CONFIGURATION
Maximum size of a message sent through the Receive connector	36 MB	**Mail flow** > **Receive connectors** > **Edit** ✏ > **General** > **Maximum receive message size (MB)**	Cmdlets: **New-ReceiveConnector**, **Set-ReceiveConnector** Parameter: *MaxMessageSize*
Maximum size of all header fields in a message sent through the Receive connector	256 KB	Not available	Cmdlets: **New-ReceiveConnector**, **Set-ReceiveConnector** Parameter: *MaxHeaderSize*
Maximum number of recipients in a message sent through the Receive connector	**Transport service on Mailbox servers** Default *<ServerName>*: 5000 Client Proxy *<ServerName>*: 200 **Front End Transport service on Mailbox servers** Default Frontend *<ServerName>*: 200 Outbound Proxy Frontend *<ServerName>*: 200 Client Frontend *<ServerName>*: 200 If the number of recipients is exceeded in a message from an anonymous sender (for example, an Internet sender), the message is accepted for the first 200 recipients. Most messaging servers will continue to resend the message in groups of 200 recipients until the message is delivered to all recipients.	Not available	Cmdlets: **New-ReceiveConnector**, **Set-ReceiveConnector** Parameter: *MaxRecipientsPerMessage*
Maximum size of a message sent through the Send connector	10 MB	**Mail flow** > **Send connectors** > **Edit** ✏ > **General** tab > **Maximum send message size (MB)**	Cmdlets: **New-SendConnector**, **Set-SendConnector** Parameter: *MaxMessageSize*
Maximum size of a message sent through the Active Directory site link	Unlimited	Not available	Cmdlet: **Set-AdSiteLink** Parameter: *MaxMessageSize*
Maximum size of a message sent through the Delivery Agent connector	Unlimited	Not available	Cmdlets: **New-DeliveryAgentConnector**, **Set-DeliveryAgentConnector** Parameter: *MaxMessageSize*
Maximum size of a message sent through the Foreign connector	Unlimited	Not available	Cmdlet: **Set-ForeignConnector** Parameter: *MaxMessageSize*

To see the values of these connector limits, run the following command in the Exchange Management Shell:

```
Get-ReceiveConnector | Format-Table Name,Max*Size,MaxRecipientsPerMessage; Get-SendConnector | Format-Table
Name,MaxMessageSize; Get-AdSiteLink | Format-Table Name,MaxMessageSize; Get-DeliveryAgentConnector | Format-
Table Name,MaxMessageSize; Get-ForeignConnector | Format-Table Name,MaxMessageSize
```

Server limits

Server limits apply to specific Mailbox servers or Edge Transport servers. You can set these message size limits independently on each Mailbox server or Edge Transport server.

SIZE LIMIT	DEFAULT VALUE	EAC CONFIGURATION	EXCHANGE MANAGEMENT SHELL CONFIGURATION
Maximum size for a message sent by Outlook on the web clients	35 MB	Not available	You configure this value in web.config XML application configuration files on the Mailbox server. For more information, see Configure client-specific message size limits.
Maximum size for a message sent by Exchange ActiveSync clients	10 MB	Not available	You configure this value in web.config XML application configuration files on the Mailbox server. For more information, see Configure client-specific message size limits.
Maximum size for a message sent by Exchange Web Services clients	64 MB	Not available	You configure this value in web.config XML application configuration files on the Mailbox server. For more information, see Configure client-specific message size limits.

The pickup directory that's available on Edge Transport servers and Mailbox servers also has messages size limits that you can configure. Typically, the pickup directory isn't used in everyday mail flow. It's is used by administrators for mail flow testing, or by applications that need to create and submit their own messages files. For more information, see Configure the Pickup Directory and the Replay Directory.

- Maximum size of all header fields in a message file placed in the pickup directory: 64 KB.
- Maximum number of recipients in a message file placed in the pickup directory: 100.

Recipient limits

Recipient limits apply to a specific user object, such as a mailbox, mail contact, mail user, distribution group, or a mail-enabled public folder.

SIZE LIMIT	DEFAULT VALUE	EAC CONFIGURATION	EXCHANGE MANAGEMENT SHELL CONFIGURATION

SIZE LIMIT	DEFAULT VALUE	EAC CONFIGURATION	EXCHANGE MANAGEMENT SHELL CONFIGURATION
Maximum size of a message that can be sent to the specific recipient	Site mailbox provisioning policies: 36 MB All other recipient types: unlimited	For mailboxes: **Recipients** > **Mailboxes** > **Edit** ✏ > **Mailbox features** > **Mail flow** section > **Message size restrictions** section > **View details** > **Received messages** section > **Maximum message size (KB)** For mail users: **Recipients** > **Contacts** > **Edit** ✏ > **Mail flow settings** > **Message size restrictions** > **View details** > **Received messages** section > **Maximum message size (KB)** This setting available in the EAC for other types of recipients.	Cmdlets: **Set-DistributionGroup** **Set-DynamicDistributionGroup** **Set-Mailbox** **Set-MailContact** **Set-MailUser** **Set-MailPublicFolder** **New-SiteMailboxProvisioningPolicy** **Set-SiteMailboxProvisioningPolicy** Parameter: *MaxReceiveSize*
Maximum size of a message that can be sent by the specific sender	Unlimited	For mailboxes: **Recipients** > **Mailboxes** > **Edit** ✏ > **Mailbox features** > **Mail flow** section > **Message size restrictions** section > **View details** > **Sent messages** section > **Maximum message size (KB)** For mail users: **Recipients** > **Contacts** > **Edit** ✏ > **Mail flow settings** > **Message size restrictions** section > **View details** > **Sent messages** section > **Maximum message size (KB)** This setting available in the EAC for other types of senders.	Cmdlets: **Set-DistributionGroup** **Set-DynamicDistributionGroup** **Set-Mailbox** **Set-MailContact** **Set-MailUser** **Set-MailPublicFolder** Parameter: *MaxSendSize*
Maximum number of recipients in a message that's sent by the specific sender	Unlimited	For mailboxes: **Recipients** > **Mailboxes** > **Edit** ✏ > **Mailbox features** > **Mail flow** section > **View details** > **Recipient limit** section > **Maximum recipients** This setting isn't available in the EAC for mail users.	Cmdlets: **Set-Mailbox**, **Set-MailUser** Parameter: *RecipientLimits*

To see the values of these limits, run the corresponding **Get-** cmdlet for the recipient type in the Exchange Management Shell.

For example, to see the limits that are configured on a specific mailbox, run the following command:

```
Get-Mailbox <MailboxIdentity> | Format-List MaxReceiveSize,MaxSendSize,RecipientLimits
```

To see the limits that are configured on all user mailboxes, run the following command:

```
$mb= Get-Mailbox -ResultSize unlimited; $mb | where {$_.RecipientTypeDetails -eq 'UserMailbox'} | Format-
Table Name,MaxReceiveSize,MaxSendSize,RecipientLimits
```

Order of precedence and placement of message size limits

The order of precedence for message size limits is the most restrictive limit is enforced. The only question is where that limit is enforced. The goal is to reject messages that are too large as early in the transport pipeline as possible. For example, it's a waste of system resources for the Internet Receive connector to accept large messages that are eventually rejected because of a lower organizational limit. Make sure that your organization, server, and connector limits are configured in a way that minimizes any unnecessary processing of messages. You do this by keeping the limits the same in all locations, or by configuring more restrictive limits where messages enter your Exchange organization.

An exception to the order is message size limits on mailboxes and messages size limits in mail flow rules. Exchange checks the maximum message size that's allowed on mailboxes before mail flow rules process messages. For example, your organization's message size limit is 50 MB, you configure a 35 MB limit on a mailbox, and you configure a mail flow rule to find and reject messages larger than 40 MB. If an external sender sends a 45 MB message to the mailbox, the message is rejected before the mail flow rule is able to evaluate the message.

Recipient limits between authenticated senders and recipients (typically, internal message senders and recipients) are exempt from the organizational message size restrictions. Therefore, you can configure specific senders and recipients to exceed the default message size limits for your organization. For example, you can allow specific mailboxes to send and receive larger messages than the rest of the organization by configuring custom send and receive limits for those mailboxes.

However, this exemption applies only to messages sent between authenticated senders and recipients (typically, internal senders and recipients). For messages sent between anonymous senders and recipients (typically, Internet senders or Internet recipients), the organizational limits apply. For example, suppose your organizational message size limit is 10 MB, but you configured the users in your marketing department to send and receive messages up to 50 MB. These users will be able to exchange large messages with each other, but not with Internet senders and recipients (unauthenticated senders and recipients).

How recipient limits work together

The recipient limit on a message is enforced in two places:

- At the protocol level during email transfer where the Receive connector *MaxRecipientsPerMessage* is enforced.
- At the Transport level during categorization where *MaxRecipientEnvelopeLimit* is enforced.

There is also the mailbox level *RecipientLimits*, which overrides the Transport level *MaxRecipientEnvelopeLimit* and is also enforced during message categorization. If the mailbox level *RecipientLimits* is set to `unlimited` (the default value), then the maximum number of recipients per message for the mailbox is controlled by the Transport level *MaxRecipientEnvelopeLimit*.

For inbound email, the Receive connector *MaxRecipientsPerMessage* is verified first. However, if the number of recipients exceeds the limit, the message is not rejected; the connection receives the error, `452 4.5.3 Too many recipients`. Most mail servers understand this error and they will continue to resend the message in another connection until the message is delivered to all recipients.

The Receive connector *MaxRecipientsPerMessage* applies to authenticated and anonymous SMTP client submissions. However, when an Exchange server relays email through another Exchange server in the same organization, the Receive connector *MaxRecipientsPerMessage* is bypassed.

When the message is accepted and email is sent to the categorizer, the mailbox level *RecipientLimits* (if it is not set to `unlimited`) or Transport level *MaxRecipientEnvelopeLimit* are checked. If the number of recipients exceeds this limit, the message is rejected and a bounce message is sent with the error `550 5.5.3 RESOLVER.ADR.RecipLimit; too many recipients`.

Here is an example scenario:

The receive connector `MaxRecipientsPerMessage` is set to 100 and the Transport level `MaxRecipientEnvelopeLimit` is set to 500. Now, if someone sends an inbound email to 1000 recipients, the email will typically be accepted because the Receive connector limit will force the sending server to send email in 10 chunks with 100 recipients on each message, which is lower than the transport categorizer setting `MaxRecipientEnvelopeLimit`.

Messages exempt from size limits

The following list shows the types of messages that are generated by Mailbox servers or Edge Transport servers that are exempted from all message size limits except the organizational limit for the maximum number of recipients that are allowed in a message:

- System messages
- Agent-generated message
- Delivery status notification (DSN) messages (also known as non-delivery reports, NDRs, or bounce messages). However, you can use the *ExternalDsnMaxMessageAttachSize* and *InternalDsnMaxMessageAttachSize* parameters on the **Set-TransportConfig** cmdlet to limit the size of original messages that are included in DSN messages (hence, the effective size of the DSN message itself).
- Journal report messages
- Quarantined messages

Message rate limits and throttling

8/3/2020 • 7 minutes to read • Edit Online

Message throttling refers to a group of limits that are set on the number of messages and connections that can be processed by an Exchange server. These limits include message processing rates, SMTP connection rates, and SMTP session timeout values. These limits work together to protect an Exchange server from being overwhelmed by accepting and delivering messages. Although a large backlog of messages and connections may be waiting to be processed, the message throttling limits enable the Exchange server to process the messages and connections in an orderly manner.

> **NOTE**
> *Back pressure* is another feature that helps to avoid overwhelming the system resources of an Exchange server. Key resources, such as available hard disk space and memory utilization are monitored, and when the utilization level exceeds the specified threshold, the server gradually stops accepting new connections and messages. For more information, see Understanding back pressure. There are also static limits that are available on messages, such as the maximum message size, the size of individual attachments, and the number of recipients. For more information about message size limits, see Message size and recipient limits in Exchange Server.

You can set the message rate limits and throttling options in the following locations:

- Mailbox servers and Edge Transport servers. Collectively, we'll refer to these as *transport servers*.
- Send connectors
- Receive connectors
- Users

Message throttling on transport servers

The following table shows the message throttling options that are available on Mailbox servers and Edge Transport servers.

RATE LIMIT	DEFAULT VALUE	EXCHANGE MANAGEMENT SHELL CONFIGURATION	EAC CONFIGURATION
Maximum concurrent mailbox deliveries: The maximum number of delivery threads that the Transport service and the Mailbox Transport Delivery service can have open at the same time to deliver message to mailboxes.	20 We recommend that you don't modify this value unless you're directed to do so by Microsoft Customer Service and Support.	Cmdlet: **Set-TransportService** and **Set-MailboxTransportService** Parameter: *MaxConcurrentMailboxDeliveries*	Not available

RATE LIMIT	DEFAULT VALUE	EXCHANGE MANAGEMENT SHELL CONFIGURATION	EAC CONFIGURATION
Maximum concurrent mailbox submissions: The maximum number of submission threads that the Transport service and the Mailbox Transport Submission service can have open at the same time to send messages from mailboxes.	20 We recommend that you don't modify this value unless you're directed to do so by Microsoft Customer Service and Support.	Cmdlet: **Set-TransportService** and **Set-MailboxTransportService** Parameter: *MaxConcurrentMailboxSubmissions*	Not available
Maximum connection rate per minute: The maximum rate that connections are allowed to be opened with the Transport service.	1200	Cmdlet: **Set-TransportService** Parameter: *MaxConnectionRatePerMinute*	Not available
Maximum concurrent connections: The maximum number of outbound connections that the Transport service can have open at a time.	1000 This value must be greater than or equal to the *MaxPerDomainOutboundConnections* value.	Cmdlet: **Set-TransportService** Parameter: *MaxOutboundConnections*	**Servers** > **Servers** > **Properties** ✏ > **Transport limits** section > **Maximum concurrent connections**. **Note**: In the EAC, you can only set the values 100, 1000, 5000, or unlimited.
Maximum concurrent connections per domain: The maximum number of outbound connections that the Transport service can have open to a single domain at a time.	20 This value must be less than or equal to the *MaxOutboundConnections* value.	Cmdlet: **Set-TransportService** Parameter: *MaxPerDomainOutboundConnections*	**Servers** > **Servers** > **Properties** ✏ > **Transport limits** section > **Maximum concurrent connections per domain**. **Note**: In the EAC, you can only set the values 100, 1000, 5000, or unlimited.

To see the values of these server message throttling settings, run the following command in the Exchange Management Shell:

```
Write-Host "Transport service:" -ForegroundColor yellow; Get-TransportService | Format-List
MaxConcurrent*,MaxConnection*,Max*OutboundConnections; Write-Host "Mailbox Transport service:" -
ForegroundColor yellow; Get-MailboxTransportService | Format-List MaxConcurrent*
```

NOTE

The Pickup directory and the Replay directory that are available on Edge Transport servers and Mailbox servers also have messages rate limits that you can configure. Typically, the Pickup directory and the Replay directory aren't used in everyday mail flow. For more information, see Configure the Pickup Directory and the Replay Directory. The maximum number of message files per minute that can be processed by the Pickup directory and the Replay directory is 100. Each directory can independently process message files at this rate.

Message throttling on Send connectors

The following table shows the message throttling options that are available on Send connectors. Send connectors exist in the Transport service on Mailbox servers and on Edge Transport servers. For more information, see Send connectors.

RATE LIMIT	DEFAULT VALUE	EXCHANGE MANAGEMENT SHELL CONFIGURATION	EAC CONFIGURATION
Connection inactivity time out: The maximum amount of time that an open SMTP connection with a source messaging server can remain idle before the connection is closed.	`00:10:00` (10 minutes)	Cmdlet: **New-SendConnector** and **Set-SendConnector** Parameter: *ConnectionInactivityTimeOut*	Not available
Maximum messages per connection: The maximum number of messages that can be sent over a single connection	20	Cmdlet: **New-SendConnector** and **Set-SendConnector** Parameter: *SmtpMaxMessagesPerConnection*	Not available

To see the values of these Send connector throttling settings, run the following command in the Exchange Management Shell:

```
Get-SendConnector | Format-List Name,ConnectionInactivityTimeout,SmtpMaxMessagesPerConnection
```

Message throttling on Receive connectors

The following table shows the message throttling options that are available on Receive connectors. Receive connectors are available in the Front End Transport service on Mailbox servers, the Transport service on Mailbox servers, and on Edge Transport servers. For more information, see Receive connectors.

RATE LIMIT	DEFAULT VALUE	EXCHANGE MANAGEMENT SHELL CONFIGURATION	EAC CONFIGURATION
Connection time out: The maximum amount of time that an SMTP connection with a source messaging server can remain open, even when the source messaging server is transmitting data.	`00:10:00` (10 minutes) for Receive connectors on Mailbox servers. `00:05:00` (1 minute) for Receive connectors on Edge Transport servers. This value must be greater than the *ConnectionInactivityTimeOut* value.	Cmdlet: **New-ReceiveConnector** and **Set-ReceiveConnector** Parameter: *ConnectionTimeout*	Not available
Connection inactivity time out: The maximum amount of time that an open SMTP connection with a source messaging server can remain idle before the connection is closed.	`00:05:00` (5 minutes) for Receive connectors on Mailbox servers. `00:01:00` (1 minute) for Receive connectors on Edge Transport servers. This value must be less than the *ConnectionTimeout* value.	Cmdlet: **New-ReceiveConnector** and **Set-ReceiveConnector** Parameter: *ConnectionInactivityTimeOut*	Not available

RATE LIMIT	DEFAULT VALUE	EXCHANGE MANAGEMENT SHELL CONFIGURATION	EAC CONFIGURATION
Maximum inbound connections: The maximum number of inbound SMTP connections that are allowed at the same time.	5000	Cmdlet: **New-ReceiveConnector** and **Set-ReceiveConnector** Parameter: *MaxInboundConnection*	Not available
Maximum inbound connections per source: The maximum number of inbound SMTP connections that are allowed from a source messaging server at the same time.	`unlimited` on the default Receive connector named Default *<ServerName>* in the Transport service on Mailbox servers. 20 on other Receive connectors on Mailbox servers and Edge Transport servers.	Cmdlet: **New-ReceiveConnector** and **Set-ReceiveConnector** Parameter: *MaxInboundConnectionPerSource*	Not available
Maximum inbound connection percentage per source: The maximum percentage of inbound SMTP connections that are allowed from a source messaging server at the same time.	100 percent on the default Receive connector named Default *<ServerName>* in the Transport service on Mailbox servers. 2 percent on other Receive connectors on Mailbox servers and Edge Transport servers.	Cmdlet: **New-ReceiveConnector** and **Set-ReceiveConnector** Parameter: *MaxInboundConnectionPercentagePerSource*	Not available
Message rate limit: The maximum number of messages per minute that can be sent by a single source.	`unlimited` on the following default Receive connectors: • Default *<ServerName>* in the Transport service on Mailbox servers. • Default Frontend *<ServerName>* in the Front End Transport service on Mailbox servers. • Outbound Proxy Frontend *<ServerName>* in the Front End Transport service on Mailbox servers. 5 on the following default Receive connectors: • Client Proxy *<ServerName>* in the Transport service on Mailbox servers. • Client Frontend *<ServerName>* in the Front End Transport service on Mailbox servers. 600 on the default Receive connector named Default internal Receive connector *<ServerName>* on Edge Transport servers.	Cmdlet: **New-ReceiveConnector** and **Set-ReceiveConnector** Parameter: *MessageRateLimit*	Not available

RATE LIMIT	DEFAULT VALUE	EXCHANGE MANAGEMENT SHELL CONFIGURATION	EAC CONFIGURATION
Message rate source: This indicates how the message submission rate is calculated. Valid values are: `User` : The rate is calculated for sending user (based on how user authenticates in the SMTP session). • `IPAddress` : The rate is calculated for sending hosts. • `All` : The rate is calculated for both sending users and sending hosts.	`IPAddress` on the following default Receive connectors: • Default *<ServerName>* in the Transport service on Mailbox servers. • Default Frontend *<ServerName>* in the Front End Transport service on Mailbox servers. • Outbound Proxy Frontend *<ServerName>* in the Front End Transport service on Mailbox servers. • Default internal Receive connector *<ServerName>* on Edge Transport servers. `User` on the following default Receive connectors: • Client Proxy *<ServerName>* in the Transport service on Mailbox servers. • Client Frontend *<ServerName>* in the Front End Transport service on Mailbox servers.	Cmdlet: **New-ReceiveConnector** and **Set-ReceiveConnector** Parameter: *MessageRateSource*	Not available
Tarpit interval: The amount of time to artificially delay SMTP responses to unauthenticated remote servers that appear to be abusing the connection. Authenticated connections are never delayed in this manner.	`00:00:05` (5 seconds)	Cmdlet: **New-ReceiveConnector** and **Set-ReceiveConnector** Parameter: *TarpitInterval*	Not available

To see the values of these Receive connector message throttling settings, run the following command in the Exchange Management Shell:

```
Get-ReceiveConnector | Format-List Name,Connection*,MaxInbound*,MessageRate*,TarpitInterval
```

Message throttling on users

The Microsoft Exchange Throttling service tracks resource settings for specific uses and caches the information in memory. Mail flow throttling settings are also known as a *budget*. Restarting the Microsoft Exchange Throttling service resets the mail flow throttling budgets.

Each mailbox has a *ThrottlingPolicy* setting. The default value for this setting is blank (`$null`). You can use the *ThrottlingPolicy* parameter on the **Set-Mailbox** cmdlet to configure a throttling policy for a mailbox.

For more information, see the following topics:

- User workload management in Exchange Server
- Change User Throttling Settings for Specific Users

- Change User Throttling Settings for All Users in Your Organization

Understanding back pressure

8/3/2020 • 15 minutes to read • Edit Online

Back pressure is a system resource monitoring feature of the Microsoft Exchange Transport service that exists on Mailbox servers and Edge Transport servers. Back pressure detects when vital system resources, such as hard drive space and memory, are overused, and takes action to prevent the server from becoming completely overwhelmed and unavailable. For example, when a system resource utilization level on the Exchange server is determined to be too high, the server delays accepting new messages. If the resource utilization gets worse, the server stops accepting new messages to work exclusively on processing all existing messages, and might even stop processing outgoing messages. When the system resource utilization returns to an acceptable level, the Exchange server resumes normal operation by accepting new messages and processing outgoing messages.

Monitored resources

The following system resources are monitored by back pressure:

- **DatabaseUsedSpace[%ExchangeInstallPath%TransportRoles\data\Queue]**: Hard drive utilization for the drive that holds the message queue database.
- **PrivateBytes**: The memory that's used by the EdgeTransport.exe process.
- **QueueLength[SubmissionQueue]**: The number of messages in the Submission queue.
- **SystemMemory**: The memory that's used by all other processes.
- **UsedDiskSpace[%ExchangeInstallPath%TransportRoles\data\Queue]**: Hard drive utilization for the drive that holds the message queue database transaction logs.
- **UsedDiskSpace[%ExchangeInstallPath%TransportRoles\data]**: Hard drive utilization for the drive that's used for content conversion.
- **UsedVersionBuckets[%ExchangeInstallPath%TransportRoles\data\Queue\mail.que]**: The number of uncommitted message queue database transactions that exist in memory.

For each monitored system resource on a Mailbox server or Edge Transport server, the following levels of resource utilization or *pressure* are defined:

- **Low or Normal**: The resource isn't overused. The server accepts new connections and messages.
- **Medium**: The resource is slightly overused. Back pressure is applied to the server in a limited manner. Mail from senders in the organization's authoritative domains can flow. However, depending on the specific resource under pressure, the server uses tarpitting to delay server response or rejects incoming **MAIL FROM** commands from other sources.
- **High**: The resource is severely overused. Full back pressure is applied. All message flow stops, and the server rejects all new incoming **MAIL FROM** commands.

Transition levels define the low, medium and high resource utilization values depending on whether the resource pressure is increasing or decreasing. Typically, a resource utilization level that's lower than the original level is required as the resource utilization decreases. In other words, there really isn't a static value for low, medium and high resource pressure. You need to know if the utilization is increasing or decreasing before you can determine the next change in resource utilization level.

The following sections explain how Exchange handles the situation when a specific resource is under pressure.

Hard drive utilization for the drive that holds the message queue database

Resource: DatabaseUsedSpace[%ExchangeInstallPath%TransportRoles\data\Queue]

Description: Monitors the percentage of total drive space that's consumed by all files on the drive that holds the message queue database. Note that the message queue database file contains unused space, so an accurate description of the total drive space that's consumed by all files is drive size - free disk space - free space in the database.

To change the default location of the message queue database, see Change the location of the queue database.

Pressure transitions (%):

- LowToMedium: 96
- MediumToHigh: 99
- HighToMedium: 97
- MediumToLow: 94

Comments::

The default high level of hard drive utilization is calculated by using the following formula:

100 * (*<hard drive size in MB>* - 500 MB) / *<hard drive size in MB>*

This formula accounts for the fact that there's unused space in the message queue database

1 GB = 1024 MB. The result is rounded down to the nearest integer.

For example, if your message queue database is located on a 1 terabyte (TB) drive (1048576 MB), the high level of utilization is 100*(1048576-500)/1048576) or 99%.

As you can see from the formula and the rounding down behavior, the hard drive needs to be very small before the formula calculates a high utilization value that's less than 99%. For example, a 98% value for high utilization requires a hard drive of approximately 25 GB or less.

Memory used by the EdgeTransport.exe process

Resource: PrivateBytes

Description: Monitors the percentage of memory that's used by the EdgeTransport.exe process that's part of the Microsoft Exchange Transport service. This doesn't include virtual memory in the paging file, or memory that's used by other processes.

Pressure transitions (%):

- LowToMedium: 72
- MediumToHigh: 75
- HighToMedium: 73
- MediumToLow: 71

Comments:

By default, the high level of memory utilization by the EdgeTransport.exe process is 75 percent of the total physical memory or 1 terabyte, whichever is less. The results are always rounded down to the nearest integer.

Exchange keeps a history of the memory utilization of the EdgeTransport.exe process. If the utilization doesn't go down to low level for a specific number of polling intervals, known as the *history depth*, Exchange rejects incoming messages until the resource utilization goes back to the low level. By default, the history depth for

EdgeTransport.exe memory utilization s 30 polling intervals.

Number of messages in the Submission queue

Resource: QueueLength[SubmissionQueue]

Description: Monitors the number of messages in the Submission queue. Typically, message enter the Submission queue from Receive connectors. For more information, see Mail flow and the transport pipeline. A large number of messages in the Submission queue indicates the categorizer is having difficulty processing messages.

Pressure transitions:

- **LowToMedium**: 9999
- **MediumToHigh**: 15000
- **HighToMedium**: 10000
- **MediumToLow**: 2000

Comments:

When the Submission queue is under pressure, the Exchange throttles incoming connections by delaying acknowledgement of incoming messages. Exchange reduces the rate of incoming message flow by *tarpitting*, which delays the acknowledgment of the SMTP **MAIL FROM** command to the sending server. If the pressure condition continues, Exchange gradually increases the tarpitting delay. After the Submission queue utilization returns to the low level, Exchange reduces the acknowledgment delay and eases back into normal operation. By default, Exchange delays message acknowledgments for 10 seconds when under Submission queue pressure. If the resource pressure continues, the delay is increased in 5-second increments up to 55 seconds.

Exchange keeps a history of Submission queue utilization. If the Submission queue utilization doesn't go down to the low level for a specific number of polling intervals, known as the *history depth*, Exchange stops the tarpitting delay and rejects incoming messages until the Submission utilization goes back to the low level. By default, the history depth for the Submission queue is in 300 polling intervals.

Memory used by all processes

Resource: SystemMemory

Description: Monitors the percentage of memory that's used by all processes on the Exchange server. This doesn't include virtual memory in the paging file.

Pressure transitions (%):

- **LowToMedium**: 88
- **MediumToHigh**: 94
- **HighToMedium**: 89
- **MediumToLow**: 84

Comments:

When the server reaches the high level of memory utilization, *message dehydration* occurs. Message dehydration removes unnecessary elements of queued messages that are cached in memory. Typically, complete messages are cached in memory for increased performance. Removal of the MIME content from these cached messages reduces the amount of memory that's used at the expense of higher latency, because the messages are now read directly from the message queue database. By default, message dehydration is enabled.

Hard drive utilization for the drive that holds the message queue database transaction logs

Resource: UsedDiskSpace[%ExchangeInstallPath%TransportRoles\data\Queue]

Description: Monitors the percentage of total drive space that's consumed by all files on the drive that holds the message queue database transaction logs. To change the default location, see Change the location of the queue database.

Pressure transitions (%):

- LowToMedium: 89
- MediumToHigh: 99
- HighToMedium: 90
- MediumToLow: 80

Comments::

The default high level of hard drive utilization is calculated by using the following formula:

100 * (*<hard drive size in MB>* - 1152 MB) / *<hard drive size in MB>*

1 GB = 1024 MB. The result is rounded down to the nearest integer.

For example, if your queue database is located on a 1 terabyte (TB) drive (1048576 MB), the high level of utilization is 100*(1048576-1152)/1048576) or 99%.

As you can see from the formula and the rounding down behavior, the hard drive needs to be fairly small before the formula calculates a high utilization value that's less than 99%. For example, a 98% value for high utilization requires a hard drive of approximately 56 GB or less.

The `%ExchangeInstallPath%Bin\EdgeTransport.exe.config` application configuration file contains the *DatabaseCheckPointDepthMax* key that has the default value `384MB`. This key controls the total allowed size of all uncommitted transaction logs that exist on the hard drive. The value of this key is used in the formula that calculates high utilization. If you customize this value, the formula becomes:

100 * (*<hard drive size in MB>* - Min(5120 MB, 3* *DatabaseCheckPointDepthMax*)) / *<hard drive size in MB>*

> **NOTE**
>
> The value of the *DatabaseCheckPointDepthMax* key applies to all transport-related Extensible Storage Engine (ESE) databases that exist on the Exchange server. On Mailbox servers, this includes the message queue database, and the sender reputation database. On Edge Transport servers, this includes the message queue database, the sender reputation database, and the IP filter database that's used by the Connection Filtering agent.

Hard drive utilization for the drive that's used for content conversion

Resource: UsedDiskSpace[%ExchangeInstallPath%TransportRoles\data]

Description: Monitors the percentage of total drive space that's consumed by all files on the drive that's used for content conversion. The default location of the folder is `%ExchangeInstallPath%TransportRoles\data\Temp` and is controlled by the *TemporaryStoragePath* key in the `%ExchangeInstallPath%Bin\EdgeTransport.exe.config` application configuration file.

Pressure transitions (%):

- LowToMedium: 89
- MediumToHigh: 99
- HighToMedium: 90

- **MediumToLow**: 80

Comments:

The default high level of hard drive utilization is calculated by using the following formula:

100 * (*<hard drive size in MB>* - 500 MB) / *<hard drive size in MB>*

1 GB = 1024 MB. The result is rounded down to the nearest integer.

For example, if your message queue database is located on a 1 terabyte (TB) drive (1048576 MB), the high level of utilization is 100*(1048576-500)/1048576) or 99%.

As you can see from the formula and the rounding down behavior, the hard drive needs to be very small before the formula calculates a high utilization value that's less than 99%. For example, a 98% value for high utilization requires a hard drive of approximately 25 GB or less.

Number of uncommitted message queue database transactions in memory

Resource: **UsedVersionBuckets[%ExchangeInstallPath%TransportRoles\data\Queue\mail.queue]**

Description: Monitors the number of uncommitted transactions for the message queue database that exist in memory.

Pressure transitions:

- **LowToMedium**: 999
- **MediumToHigh**: 1500
- **HighToMedium**: 1000
- **MediumToLow**: 800

Comments::

A list of changes that are made to the message queue database is kept in memory until those changes can be committed to a transaction log. Then the list is committed to the message queue database itself. These outstanding message queue database transactions that are kept in memory are known as *version buckets*. The number of version buckets may increase to unacceptably high levels because of an unexpectedly high volume of incoming messages, spam attacks, problems with the message queue database integrity, or hard drive performance.

When version buckets are under pressure, the Exchange server throttles incoming connections by delaying acknowledgment of incoming messages. Exchange reduces the rate of incoming message flow by *tarpitting*, which delays the acknowledgment of the SMTP **MAIL FROM** command to the sending server. If the resource pressure condition continues, Exchange gradually increases the tarpitting delay. After the resource utilization returns to normal, Exchange gradually reduces the acknowledgement delay and eases back into normal operation. By default, Exchange delays message acknowledgments for 10 seconds when under resource pressure. If the pressure continues, the delay is increased in 5-second increments up to 55 seconds.

When the version buckets are under high pressure, the Exchange server also stops processing outgoing messages.

Exchange keeps a history of version bucket resource utilization. If the resource utilization doesn't go down to the low level for a specific number of polling intervals, known as the *history depth*, Exchange stops the tarpitting delay and rejects incoming messages until the resource utilization goes back to the low level. By default, the history depth for version buckets is in 10 polling intervals.

Actions taken by back pressure when resources are under pressure

The following table summarizes the actions taken by back pressure when a monitored resource is under pressure.

RESOURCE UNDER PRESSURE	UTILIZATION LEVEL	ACTIONS TAKEN
DatabaseUsedSpace	Medium	Reject incoming messages from non-Exchange servers. Reject message submissions from the Pickup directory and the Replay directory. Message resubmission is paused. Shadow Redundancy rejects messages. For more information about Shadow Redundancy, see Shadow redundancy in Exchange Server.
DatabaseUsedSpace	High	All actions taken at the medium utilization level. Reject incoming messages from other Exchange servers. Reject message submissions from mailbox databases by the Microsoft Exchange Mailbox Transport Submission service on Mailbox servers.
PrivateBytes	Medium	Reject incoming messages from non-Exchange servers. Reject message submissions from the Pickup directory and the Replay directory. Message resubmission is paused. Shadow Redundancy rejects messages. Processing messages after a server or Transport service restart (also known as *boot scanning*) is paused. Start message dehydration.
PrivateBytes	High	All actions taken at the medium utilization level. Reject incoming messages from other Exchange servers. Reject message submissions from mailbox databases by the Microsoft Exchange Mailbox Transport Submission service on Mailbox servers.
QueueLength[SubmissionQueue]	Medium	Introduce or increment the tarpitting delay to incoming messages. If normal level isn't reached for the entire Submission queue history depth, take the following actions: • Reject incoming messages from non-Exchange servers. • Reject message submissions from the Pickup directory and the Replay directory. • Message resubmission is paused. • Shadow Redundancy rejects messages. • Boot scanning is paused.

RESOURCE UNDER PRESSURE	UTILIZATION LEVEL	ACTIONS TAKEN
QueueLength[SubmissionQueue]	High	All actions taken at the medium utilization level. Reject incoming messages from other Exchange servers. Reject message submissions from mailbox databases by the Microsoft Exchange Mailbox Transport Submission service on Mailbox servers. Flush enhanced DNS cache from memory. Start message dehydration.
SystemMemory	Medium	Start message dehydration. Flush caches.
SystemMemory	High	All actions taken at the medium utilization level.
UsedDiskSpace (message queue database transaction logs)	Medium	Reject incoming messages from non-Exchange servers. Reject message submissions from the Pickup directory and the Replay directory. Message resubmission is paused. Shadow Redundancy rejects messages.
UsedDiskSpace (message queue database transaction logs)	High	All actions taken at the medium utilization level. Reject incoming messages from other Exchange servers. Reject message submissions from mailbox databases by the Microsoft Exchange Mailbox Transport Submission service on Mailbox servers.
UsedDiskSpace (content conversion)	Medium	Reject incoming messages from non-Exchange servers. Reject message submissions from the Pickup directory and the Replay directory.
UsedDiskSpace (content conversion)	High	All actions taken at the medium utilization level. Reject incoming messages from other Exchange servers. Reject message submissions from mailbox databases by the Microsoft Exchange Mailbox Transport Submission service on Mailbox servers.

RESOURCE UNDER PRESSURE	UTILIZATION LEVEL	ACTIONS TAKEN
UsedVersionBuckets	Medium	Introduce or increment the tarpitting delay to incoming messages. If normal level isn't reached for the entire version bucket history depth, take the following actions: • Reject incoming messages from non-Exchange servers. • Reject message submissions from the Pickup directory and the Replay directory.
UsedVersionBuckets	High	All actions taken at the medium utilization level. Reject incoming messages from other Exchange servers. Reject message submissions from mailbox databases by the Microsoft Exchange Mailbox Transport Submission service on Mailbox servers. Stop processing outgoing messages. Remote delivery is paused.

View back pressure resource thresholds and utilization levels

You can use the **Get-ExchangeDiagnosticInfo** cmdlet in the Exchange Management Shell to view the resources that are being monitored, and the current utilization levels. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.

To view the back pressure settings on an Exchange server, run the following command:

```
[xml]$bp=Get-ExchangeDiagnosticInfo [-Server <ServerIdentity> ] -Process EdgeTransport -Component
ResourceThrottling; $bp.Diagnostics.Components.ResourceThrottling.ResourceTracker.ResourceMeter
```

To see the values on the local server, you can omit the *Server* parameter.

Back pressure configuration settings in the EdgeTransport.exe.config file

All configuration options for back pressure are done in the `%ExchangeInstallPath%Bin\EdgeTransport.exe.config` XML application configuration file. However, few of the settings exist in the file by default.

Caution

These settings are listed only as a reference for the default values. We strongly discourage any modifications to the back pressure settings in the EdgeTransport.exe.config file. Modifications to these settings might result in poor performance or data loss. We recommend that you investigate and correct the root cause of any back pressure events that you may encounter.

General back pressure settings

KEY NAME	DEFAULT VALUE
ResourceMeteringInterval	`00:00:02` (2 seconds)
DehydrateMessagesUnderMemoryPressure	true

DatabaseUsedSpace settings

KEY NAME	DEFAULT VALUE (%)
DatabaseUsedSpace.LowToMedium	96
DatabaseUsedSpace.MediumToHigh	99
DatabaseUsedSpace.HighToMedium	97
DatabaseUsedSpace.MediumToLow	94

PrivateBytes settings

KEY NAME	DEFAULT VALUE (%)
PrivateBytes.LowToMedium	72
PrivateBytes.MediumToHigh	75
PrivateBytes.HighToMedium	73
PrivateBytes.MediumToLow	71
PrivateBytesHistoryDepth	30

QueueLength[SubmissionQueue] settings

KEY NAME	DEFAULT VALUE
QueueLength[SubmissionQueue].LowToMedium	9999
QueueLength[SubmissionQueue].MediumToHigh	15000
QueueLength[SubmissionQueue].HighToMedium	10000
QueueLength[SubmissionQueue].MediumToLow	2000
SubmissionQueueHistoryDepth	300 (after 10 minutes)

SystemMemory settings

KEY NAME	DEFAULT VALUE (%)
SystemMemory.LowToMedium	88
SystemMemory.MediumToHigh	94
SystemMemory.HighToMedium	89
SystemMemory.MediumToLow	84

UsedDiskSpace settings (message queue database transaction logs)

KEY NAME	DEFAULT VALUE (%)
UsedDiskSpace[%ExchangeInstallPath%TransportRoles\data\Queue].LowToMedium	89
UsedDiskSpace[%ExchangeInstallPath%TransportRoles\data\Queue].MediumToHigh	99
UsedDiskSpace[%ExchangeInstallPath%TransportRoles\data\Queue].HighToMedium	90
UsedDiskSpace[%ExchangeInstallPath%TransportRoles\data\Queue].MediumToLow	80

NOTE

Values that contain only `UsedDiskSpace` (for example, `UsedDiskSpace.MediumToHigh`) apply to the message queue database transaction logs and to content conversion.

UsedDiskSpace settings (content conversion)

KEY NAME	DEFAULT VALUE (%)
UsedDiskSpace[%ExchangeInstallPath%TransportRoles\data].LowToMedium	89
UsedDiskSpace[%ExchangeInstallPath%TransportRoles\data].MediumToHigh	99
UsedDiskSpace[%ExchangeInstallPath%TransportRoles\data].HighToMedium	90
UsedDiskSpace[%ExchangeInstallPath%TransportRoles\data].MediumToLow	80
TemporaryStoragePath	`%ExchangeInstallPath%TransportRoles\data\Temp`

UsedVersionBuckets settings

KEY NAME	DEFAULT VALUE
UsedVersionBuckets.LowToMedium	999
UsedVersionBuckets.MediumToHigh	1500
UsedVersionBuckets.HighToMedium	1000
UsedVersionBuckets.MediumToLow	800
VersionBucketsHistoryDepth	10

Back pressure logging information

The following list describes the event log entries that are generated by specific back pressure events in Exchange:

- **Event log entry for an increase in any resource utilization level**

  Event Type: Error

  Event Source: MSExchangeTransport

  Event Category: Resource Manager

  Event ID: 15004

  Description: Resource pressure increased from *<Previous Utilization Level>* to *<Current Utilization Level>*.

- **Event log entry for a decrease in any resource utilization level**

  Event Type: Information

  Event Source: MSExchangeTransport

  Event Category: Resource Manager

  Event ID: 15005

  Description: Resource pressure decreased from *<Previous Utilization Level>* to *<Current Utilization Level>*.

- **Event log entry for critically low available disk space**

  Event Type: Error

  Event Source: MSExchangeTransport

  Event Category: Resource Manager

  Event ID: 15006

  Description: The Microsoft Exchange Transport service is rejecting messages because available disk space is below the configured threshold. Administrative action may be required to free disk space for the service to continue operations.

- **Event log entry for critically low available memory**

  Event Type: Error

  Event Source: MSExchangeTransport

  Event Category: Resource Manager

  Event ID: 15007

  Description: The Microsoft Exchange Transport service is rejecting message submissions because the service continues to consume more memory than the configured threshold. This may require that this service be restarted to continue normal operation.

Use Telnet to test SMTP communication on Exchange servers

8/3/2020 • 10 minutes to read • Edit Online

You can use Telnet to test Simple Mail Transfer Protocol (SMTP) communication between messaging servers. SMTP is the protocol that's used to send email messages from one messaging server to another. Using Telnet can be helpful if you're having trouble sending or receiving messages because you can manually send SMTP commands to a messaging server. In return, the server will reply with responses that would be returned in a typical connection. These results can sometimes help you to figure out why you can't send or receive messages.

You can use Telnet to test SMTP communication to:

- Test mail flow from the Internet into your Exchange organization.
- Test mail flow from your Exchange to another messaging server on the Internet.

> **TIP**
>
> Did you know that, instead of using Telnet to test SMTP connectivity, you can use the Microsoft Remote Connectivity Analyzer at https://testconnectivity.microsoft.com/? With the Remote Connectivity Analyzer, you can choose the connectivity test you want to do, in this case **Inbound SMTP Email**, and follow the instructions shown. It'll step you through the information you need to enter, run the test for you, and then give you the results. Give it a try!

What do you need to know before you begin?

- Estimated time to complete: 15 minutes
- Exchange permissions don't apply to the procedures in this topic. These procedures are performed in the operating system of the Exchange server or a client computer.
- This topic shows you how to use Telnet Client, which is included with Windows. Third-party Telnet clients might require syntax that's different from what's shown in this topic.
- The steps in this topic show you how to connect to an Internet-facing server that allows anonymous connections using TCP port 25. If you're trying to connect to this server from the Internet, you need to make sure your Exchange server is reachable from the Internet on TCP port 25. Similarly, if you're trying to reach a server on the Internet from your Exchange server, you need to make sure your Exchange server can open a connect to the Internet on TCP port 25.
- You might notice some Receive connectors that use TCP port 2525. These are internal Receive connectors and aren't used to accept anonymous SMTP connections.
- If you're testing a connection on a remote messaging server, you should run the steps in this topic on your Exchange server. Remote messaging servers are often set up to make sure the IP address where the SMTP connection is coming from matches the domain in the sender's email address.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Step 1: Install the Telnet Client on your computer

On most versions of Windows, you'll need to install the Telnet client before you can use it. To install it, see Install Telnet Client.

Step 2: Find the FQDN or IP address of the destination SMTP server

To connect to an SMTP server by using Telnet on port 25, you need to use the fully-qualified domain name (FQDN) (for example, mail.contoso.com) or the IP address of the SMTP server. If you don't know the FQDN or IP address, you can use the Nslookup command-line tool to find the MX record for the destination domain.

NOTE

Network policies might prevent you from using the Nslookup tool to query public DNS servers on the Internet. As an alternative, you can use one of the freely-available DNS lookup or MX record lookup web sites on the Internet.

1. At a command prompt, type nslookup, and then press Enter. This command opens the Nslookup session.

2. Type set type=mx, and then press Enter.

3. Type the name of the domain for which you want to find the MX record. For example, to find the MX record for the fabrikam.com domain, type fabrikam.com., and then press Enter.

   NOTE

   When you use a trailing period (.), you prevent any default DNS suffixes from being unintentionally added to the domain name.

   The output of the command looks like this:

```
fabrikam.com mx preference=10, mail exchanger = mail1.fabrikam.com
fabrikam.com mx preference=20, mail exchanger = mail2.fabrikam.com
mail1.fabrikam.com internet address = 192.168.1.10
mail2 fabrikam.com internet address = 192.168.1.20
```

```
You can use any of the host names or IP addresses that are associated with the MX records as the destination
SMTP server. A lower value for preference (preference = 10 vs. 20) indicates a preferred SMTP server.
Multiple MX records and different values of preference are used for load balancing and fault tolerance.
```

4. When you're ready to end the Nslookup session, type exit, and then press Enter.

Step 3: Use Telnet on Port 25 to test SMTP communication

In this example, we're going to use the following values. When you run the commands on your server, replace these values with ones for your organization's SMTP server, domain, etc.

- **Destination SMTP server**: mail1.fabrikam.com

- **Source domain**: contoso.com
- **Sender's e-mail address**: chris@contoso.com
- **Recipient's e-mail address**: kate@fabrikam.com
- **Message subject**: Test from Contoso
- **Message body**: This is a test message

> **TIP**
>
> The commands in the Telnet Client aren't case-sensitive. The SMTP command verbs in this example are capitalized for clarity. You can't use the backspace key in the Telnet session after you connect to the destination SMTP server. If you make a mistake as you type an SMTP command, you need to press Enter, and then type the command again. Unrecognized SMTP commands or syntax errors result in an error message that looks like this: `500 5.3.3 Unrecognized command`

1. Open a Command Prompt window, type `telnet`, and then press Enter.

   This command opens the Telnet session.

2. Type `set localecho`, and then press Enter.

   This **optional** command lets you view the characters as you type them, and it might be required for some SMTP servers.

3. Type `set logfile <filename>`, and then press Enter.

   This **optional** command enables logging and specifies the log file for the Telnet session. If you only specify a file name, the log file is located in the current folder. If you specify a path and file name, the path needs to be on the local computer, and you might need to enter the path and file name in the Windows DOS 8.3 format (short name with no spaces). The path needs to exist, but the log file is created automatically.

4. Type `OPEN mail1.fabrikam.com 25`, and then press Enter.

5. Type `EHLO contoso.com`, and then press Enter.

6. Type `MAIL FROM:<chris@contoso.com>`, and then press Enter.

7. Type `RCPT TO:<kate@fabrikam.com> NOTIFY=success,failure`, and then press Enter.

   The optional NOTIFY command specifies the particular delivery status notification (DSN) messages (also known as bounce messages, nondelivery reports, or NDRs) that the SMTP is required to provide. In this example, you're requesting a DSN message for successful or failed message delivery.

8. Type `DATA`, and then press Enter.

9. Type `Subject: Test from Contoso`, and then press Enter.

10. Press Enter again.

    A blank line is needed between the **Subject:** field and the message body.

11. Type `This is a test message`, and then press Enter.

12. Type a period (.), and then press Enter.

13. To disconnect from the SMTP server, type `QUIT`, and then press Enter.

14. To close the Telnet session, type `quit`, and then press Enter.

Here's what a successful session using the steps above looks like:

```
C:\Windows\System32> telnet
Microsoft Telnet> set localecho
Microsoft Telnet> set logfile c:\TelnetTest.txt
Microsoft Telnet> OPEN mail1.fabrikam.com 25
220 mail1.fabrikam.com Microsoft ESMTP MAIL Service ready at Fri, 5 Aug 2016 16:24:41 -0700
EHLO contoso.com
250-mail1.fabrikam.com Hello [172.16.0.5]
250-SIZE 37748736
250-PIPELINING
250-DSN
250-ENHANCEDSTATUSCODES
250-STARTTLS
250-X-ANONYMOUSTLS
250-AUTH NTLM
250-X-EXPS GSSAPI NTLM
250-8BITMIME
250-BINARYMIME
250-CHUNKING
250 XRDST
MAIL FROM: <chris@contoso.com>
250 2.1.0 Sender OK
RCPT TO: <kate@fabrikam.com> NOTIFY=success,failure
250 2.1.5 Recipient OK
DATA
354 Start mail input; end with <CRLF>.<CRLF>
Subject: test

This is a test message.
.
250 2.6.0 <c89b4fcc-3ad1-4758-a1ab-1e820065d622@mail1.fabrikam.com> [InternalId=5111011082268,
Hostname=mail1.fabrikam.com] Queued mail for delivery
QUIT
221 2.0.0 Service closing transmission channel
```

Step 4: Success and error messages in the Telnet Session

This section provides information about the success and failure responses to the commands that were used in the previous example.

NOTE

The three-digit SMTP response codes that are defined in RFC 5321 are the same for all SMTP messaging servers, but the text descriptions in the responses might be slightly different.

SMTP reply codes

SMTP servers respond to commands with a variety of numerical reply codes in the format of x.y.z where:

- X indicates whether the command was good, bad, or incomplete.
- Y indicates the kind of response that was sent.
- Z provides additional information about the command

When a response is received by the server that opened the connection, it can tell whether the remote server accepted the command and is ready for the next one, or if an error occurred.

The first digit (X) is particularly important to understand because it indicates the success or failure of the command that was sent. Here are its possible values, and their meanings.

REPLY CODE	MEANING
2.y.z	The command that was sent was successfully completed on the remote server. The remote server is ready for the next command.
3.y.z	The command was accepted but the remote server needs more information before the operation can be completed. The sending server needs to send a new command with the needed information.
4.y.z	The command wasn't accepted by the remote server for a reason that might be temporary. The sending server should try to connect again later to see if the remote server can successfully accept the command. The sending server will continue to retry the connection until either a successful connection is completed (indicated by a 2.y.z code) or fails permanently (indicated by a 5.y.z code). An example of a temporary error is low storage space on the remote server. Once more space is made available, the remote server should be able to successfully accept the command.
5.y.z	The command wasn't accepted by the remote server for a reason that is isn't recoverable. The sending server won't retry the connection and will send a non-delivery report back to the user who sent the message. An example of an unrecoverable error is a message that's sent to an email address that doesn't exist.

The table above is based on information provided by RFC 5321 (Simple Mail Transfer Protocol), section 4.2.1. Additional information, including descriptions of the second (Y) and third (Z) digits of SMTP reply codes is included in this section, and in sections 4.2.2 and 4.2.3.

OPEN command

Successful response: `220 mail1.fabrikam.com Microsoft ESMTP MAIL Service ready at <day-date-time>`

Failure response: `Connecting to mail1.fabrikam.com...Could not open connection to the host, on port 25: Connect failed`

Possible reasons for failure

- The destination SMTP service is unavailable.
- Restrictions on the destination firewall.
- Restrictions on the source firewall.
- Incorrect FQDN or IP address for the destination SMTP server.
- Incorrect port number.

EHLO command

Successful response: `250 mail1.fabrikam.com Hello [<sourceIPaddress>]`

Failure response: `501 5.5.4 Invalid domain name`

Possible reasons for failure

- Invalid characters in the domain name.
- Connection restrictions on the destination SMTP server.

NOTE

EHLO is the Extended Simple Message Transfer Protocol (ESMTP) verb that's defined in RFC 5321. ESMTP servers can advertise their capabilities during the initial connection. These capabilities include the maximum accepted message size and supported authentication methods. HELO is the older SMTP verb that is defined in RFC 821. Most SMTP messaging servers support ESMTP and EHLO. If the non-Exchange server that you're trying to connect to doesn't support EHLO, you can use HELO instead.

MAIL FROM command

Successful response: `250 2.1.0 Sender OK`

Failure response: `550 5.1.7 Invalid address`

Possible reasons for failure: A syntax error in the sender's e-mail address.

Failure response: `530 5.7.1 Client was not authenticated`

Possible reasons for failure: The destination server doesn't accept anonymous message submissions. You receive this error if you try to use Telnet to submit a message directly to a Mailbox server that doesn't have a Receive connector that's configured to accept anonymous connections.

RCPT TO command

Successful response: `250 2.1.5 Recipient OK`

Failure response: `550 5.1.1 User unknown`

Possible reasons for failure: The specified recipient doesn't exist.

Collaboration

8/3/2020 • 3 minutes to read • Edit Online

Exchange Server provides the following rich features that can help your end users collaborate in email:

- Site mailboxes (deprecated in SharePoint 2019)
- Public folders
- Shared mailboxes
- Distribution groups

Each of these features has a different user experience and feature set and should be used based on what the user needs to accomplish and what your organization can provide. For example, site mailboxes provide great documentation collaboration features. However site mailboxes rely on SharePoint Server, so if you aren't planning on deploying SharePoint, you should use public folders to share documents.

This topic compares these collaboration features to help you decide which features to offer your users.

Site mailboxes

A site mailbox is functionally comprised of a SharePoint 2013 or later site membership (owners and members), shared storage through an Exchange 2016 or later mailbox for email messages, and a SharePoint site to store and share information. Essentially, site mailboxes bring Exchange email and SharePoint documents together. For users, a site mailbox serves as a central filing cabinet for the project, providing a place to file project email and documents that can be accessed and edited only by site members. In addition, site mailboxes have a specified lifecycle and are optimized to be used for projects that have set start and end dates. To fully implement site mailboxes, end users must use Outlook 2013 or later.

To learn more, see Site mailboxes.

Public folders

Public folders are designed for shared access and provide an easy and effective way to collect, organize, and share information with other people in your workgroup or organization.

Public folders organize content in a deep hierarchy that's easy to browse. Users discover interesting and relevant content by browsing through branches of the hierarchy that are relevant to them. Users always see the full hierarchy in their Outlook folder view. Public folders are a great technology for distribution group archiving. A public folder can be mail-enabled and added as a member of the distribution group. Email sent to the distribution group is automatically added to the public folder for later reference. Public folders also provide simple document sharing and don't require SharePoint Server to be installed in your organization. Finally, end users can use public folders with Outlook 2007 or later.

To learn more, see Public folders.

Shared mailboxes

A shared mailbox is a mailbox that multiple designated users can access to read and send email messages and to share a common calendar. Shared mailboxes can provide a generic email address (such as info@contoso.com or sales@contoso.com) that customers can use to inquire about your company. If the shared mailbox has the Send As permission assigned when a delegated user responds to the email message, it can appear as though the mailbox (for example, sales@contoso.com) is responding, not the actual user.

To learn more, see Shared mailboxes.

Groups

Groups (also called distribution groups) are a collection of two or more recipients that appears in the shared address book. When an email message is sent to a group, it's received by all members of the group. Distribution groups can be organized by a particular discussion subject (such as "Dog Lovers") or by users who share a common work structure that requires them to communicate frequently.

To learn more, see Recipients.

Which one to use?

The following table gives you a quick glance at each of the collaboration features to help you decide which one to use.

	SITE MAILBOXES	PUBLIC FOLDERS	SHARED MAILBOXES	GROUPS
Type of group	Users who work together as a team on a specific project with definitive start and end dates.	With the proper permissions, everyone in your organization can access and search public folders. Public folders are ideal for maintaining history or distribution group conversations.	Delegates working on behalf of a virtual identity, and they can respond to email as that shared mailbox identity. Example: support@tailspintoys.com	Users who need to send email to a group of recipients with a common interest or characteristic.
Ideal group size	Small	Large	Small	Large
Access	Site mailbox owners and members.	Accessible by anyone in your organization.	Users can be granted Full Access and/or Send As permissions. If granted Full Access permissions, users must also add the shared mailbox to their Outlook profile to access the shared mailbox.	For distribution groups, members, must be manually added. For dynamic distribution groups, members are added based on filtering criteria.
Shared calendar?	No	Yes	Yes	No
Email arrives in user's personal Inbox?	No. Email arrives in the site mailbox.	No. Email arrives in the public folder.	No. Email arrives in the Inbox of the shared mailbox.	Yes. Email arrives in the Inbox of a distribution group member.
Supported clients	Outlook 2013 or later SharePoint 2013	Outlook 2007 or later	Outlook 2007 or later Outlook Web App	Outlook 2007 or later Outlook Web App

Site mailboxes

8/3/2020 • 4 minutes to read • Edit Online

Email and documents are traditionally kept in two unique and separate data repositories. Most organizations collaborate using both mediums. The challenge is that both email and documents are accessed using different clients. This usually results in a reduction in user productivity and a degraded user experience.

The *site mailbox*, first introduced in Exchange 2013, is a solution for this problem. Site mailboxes improve collaboration and user productivity by allowing access to both Microsoft SharePoint documents and Exchange email using the same client interface. A site mailbox is functionally comprised of SharePoint site membership (owners and members), shared storage through an Exchange 2016 or Exchange 2019 mailbox for email messages and a SharePoint site for documents, and a management interface that addresses provisioning and lifecycle needs.

Site mailboxes require Exchange 2016 or later and SharePoint Server 2013 or later integration and configuration. For more information about how to configure your Exchange Server organization to work with your SharePoint organization, see the following topics:

- Configure site mailboxes in SharePoint Server.
- Plan Exchange Server integration with SharePoint and Skype for Business

For more information about collaboration features in Exchange Server, see Collaboration.

How do site mailboxes work?

When one project member files mail or documents using the site mailbox, any project member can then access the content. Site mailboxes are surfaced in Outlook 2013 or later and give users easy access to the email and documents for the projects they care about. Additionally, the same set of content can be accessed directly from the SharePoint site itself. With site mailboxes, the content is kept where it belongs. Exchange stores the email, providing users with the same message view for email conversations that they use every day for their own mailboxes. Meanwhile, SharePoint stores the documents, bringing document coauthoring and versioning to the table. Exchange synchronizes just enough metadata from SharePoint to create the document view in Outlook (e.g. document title, last modified date, last modified author, size).

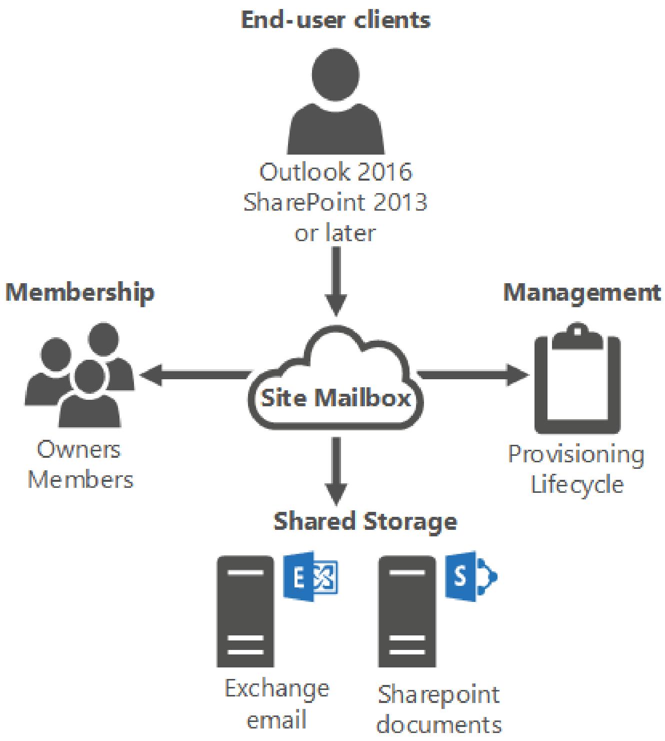

Site mailbox provisioning policies

Site mailbox quotas can be set by using the **SiteMailboxProvisioningPolicy** cmdlets in the Exchange Management Shell. The Site mailbox provisioning policies only apply to the email that is sent to and from the site mailbox and the size of the site mailbox on the Exchange server. The document repository settings are configured in SharePoint. Although you can create multiple site mailbox provisioning policies using the **New-SiteMailboxProvisioningPolicy** cmdlet, only the default provisioning policy will be applied to all site mailboxes. You can't apply multiple policies within your organization. The provisioning policies allow you to set the following quotas:

QUOTA	DESCRIPTION	DEFAULT SETTING
IssueWarningQuota	The *IssueWarningQuota* parameter specifies the site mailbox size that triggers a warning message to the site mailbox	4.5 GB
MaxReceiveSize	The *MaxReceiveSize* parameter specifies the maximum size of email messages that can be received by the site mailbox.	36 MB
ProhibitSendReceiveQuota	The *ProhibitSendReceiveQuota* parameter specifies the size at which the site mailbox can no longer send or receive messages.	5 GB

For more information about how to configure site mailbox provisioning policies, see Manage site mailbox provisioning policies.

Lifecycle policy and retention

The lifecycle of a site mailbox is managed through SharePoint. It is through SharePoint that you should perform all site mailbox tasks such as creating and removing site mailboxes. In addition, you can create a SharePoint Lifecycle policy to manage the lifecycle of a site mailbox. For example, you can create a lifecycle policy in SharePoint that automatically closes all site mailboxes after 6 months. If the user still requires the use of the site mailbox, the user can reactivate the site mailbox through SharePoint. We recommend that you use the Lifecycle application is in the farm. Manually deleting active site mailboxes from Exchange will result in orphaned site mailboxes. .

When the lifecycle application in SharePoint closes a site mailbox, the site mailbox is retained for the period stated in the lifecycle policy in the closed state. The mailbox can then be reactivated by an end-user or by an administrator from SharePoint. After the retention period, the Exchange site mailbox that is housed in the mailbox database will have its name prepended with **MDEL:** to indicate that it has been marked for deletion. You will need to manually remove these site mailboxes from the mailbox database in order to free storage space and the alias. If you don't have the SharePoint Lifecycle Policy enabled, you'll lose the ability to determine which site mailboxes are marked for deletion. Until the site mailbox has been removed by an administrator, the content of the mailbox is still recoverable.

You can use the following command to search for and remove site mailboxes that have been marked for deletion.

```
Get-Mailbox MDEL:* | ?{$_.RecipientTypeDetails -eq "TeamMailbox"} | Remove-Mailbox -Confirm:$false
```

Site mailboxes don't support retention at the item-level. Retention works on a project-level for site mailboxes, so when the entire site mailbox is deleted, the retained items will be deleted.

Compliance

Using the eDiscovery Console in SharePoint, site mailboxes can be part of the In-Place eDiscovery scope as you can do keyword searches against user mailboxes or site mailboxes. In addition, you can put a site mailbox on legal hold. For more info, see In-Place eDiscovery in Exchange Server.

Backup and restore

Backup and Restore for the Exchange site mailboxes housed on the mailbox server will use the same backup and restore method that you use for all Exchange mailboxes. For more information, see Database availability groups.

For SharePoint documents, you should backup and restore into the same place. If you restore your SharePoint content to same URLs, then the site mailbox will continue to work and no additional configuration is needed. If you restore to a different URL, then you'll need to run **Set-SiteMailbox** cmdlet to update the *SharePointURL* property. We recommend that you don't restore SharePoint to a new forest.

Public folders

8/3/2020 • 10 minutes to read • Edit Online

Public folders are designed for shared access and provide an easy and effective way to collect, organize, and share information with other people in your workgroup or organization. Public folders help make content in a deep hierarchy easier to browse. Users will see the full hierarchy in Outlook, which makes it easy for them to find the content they're interested in.

Public folders are available in the following Outlook clients:

- Outlook on the web (formerly known as Outlook Web App) for Exchange 2016 or later
- Supported versions of Outlook for Exchange Server.
- Outlook for Mac 2016 and Outlook for Mac for Office 365.

Public folders can also be used as an archiving method for distribution groups. When you mail-enable a public folder and add it as a member of the distribution group, email sent to the group is automatically added to the public folder for later reference.

Public folders aren't designed to do the following:

- **Data archiving**: Users who have mailbox limits sometimes use public folders instead of mailboxes to archive data. This practice isn't recommended because it affects storage in public folders and undermines the goal of mailbox limits. Instead, we recommend that you use In-Place Archiving in Exchange 2016 as your archiving solution.
- **Document sharing and collaboration**: Public folders don't provide versioning or other document management features, such as controlled check-in and check-out functionality and automatic notifications of content changes. Instead, we recommend that you use SharePoint as your documentation sharing solution.

To learn more about public folders and other collaboration methods in Exchange, see Collaboration.

To browse some frequently asked questions about public folders in Exchange, see FAQ: Public folders.

For more information about the limits and quotas for public folders, see Limits for public folders.

Public folder architecture

Public folders use a mailbox infrastructure to take advantage of the existing high availability and storage technologies of the mailbox database. Public folder architecture uses specially designed mailboxes to store both the public folder hierarchy and the content. This also means that there's no longer a public folder database as there was in earlier version of Exchange. High availability for the public folder mailboxes is provided by a database availability group (DAG). To learn more about DAGs, see Database availability groups.

The main architectural components of public folders are the public folder mailboxes, which can reside in one or more mailbox databases.

Public folder mailboxes

There are two types of public folder mailboxes: the *primary hierarchy mailbox* and *secondary hierarchy mailboxes*. Both types of mailboxes can contain content:

- **Primary hierarchy mailbox**: The primary hierarchy mailbox is the one writable copy of the public folder hierarchy. The public folder hierarchy is copied to all other public folder mailboxes, but these will be read-only copies.

- **Secondary hierarchy mailboxes**: Secondary hierarchy mailboxes contain public folder content as well and a read-only copy of the public folder hierarchy.

> **NOTE**
>
> Retention policies aren't supported for public folder mailboxes.

There are two ways you can manage public folder mailboxes:

- In the Exchange admin center (EAC), navigate to **Public folders** > **Public folder mailboxes**.
- In the Exchange Management Shell, use the ***-Mailbox** set of cmdlets. The following parameters have been added to the New-Mailbox cmdlet to support public folder mailboxes:
  - *PublicFolder*: This parameter is used with the **New-Mailbox** cmdlet to create a public folder mailbox. When you create a public folder mailbox, a new mailbox is created with the mailbox type of `PublicFolder`. For more information, see Create a public folder mailbox.
  - *HoldForMigration*: This parameter is used only if you're migrating public folders from Exchange 2010 to Exchange 2016. For more information, see Migrate public folders later in this topic.
  - *IsHierarchyReady*: This parameter indicates whether the public folder mailbox is ready to serve the public folder hierarchy to users. It's set to `$True` only after the entire hierarchy has been synced to the public folder mailbox. If the parameter is set to $False, users won't use it to access the hierarchy. However, if you set the *DefaultPublicFolderMailbox* property on a user mailbox to a specific public folder mailbox, the user will still access the specified public folder mailbox even if the *IsHierarchyReady* parameter is set to `$False`.
  - *IsExcludedFromServingHierarchy*: This parameter prevents users from accessing the public folder hierarchy on the specified public folder mailbox. For load-balancing purposes, users are equally distributed across public folder mailboxes by default. When this parameter is set on a public folder mailbox, that mailbox isn't included in this automatic load balancing and won't be accessed by users to retrieve the public folder hierarchy. However, if you set the *DefaultPublicFolderMailbox* property on a user mailbox to a specific public folder mailbox, the user will still access the specified public folder mailbox even if the *IsExcludedFromServingHierarchy* parameter is set for that public folder mailbox.

A secondary hierarchy mailbox will serve only public folder hierarchy information to users if it's specified explicitly on the users' mailboxes using the *DefaultPublicFolderMailbox* property, or if the following conditions are met:

- The *IsHierarchyReady* property on the public folder mailbox is set to `$True`.
- The *IsExcludedFromServingHierarchy* property on the public folder mailbox is set to `$False`.

Public folder hierarchy

The public folder hierarchy contains the folders' properties and organizational information, including tree structure. Each public folder mailbox contains a copy of the public folder hierarchy. There's only one writeable copy of the hierarchy, which is in the primary public folder mailbox. For a specific folder, the hierarchy information is used to identify the following:

- Permissions on the folder
- The folder's position in the public folder tree, including its parent and child folders

NOTE

The hierarchy doesn't store information about email addresses for mail-enabled public folders. The email addresses are stored on the directory object in Active Directory.

Hierarchy synchronization

The public folder hierarchy synchronization process uses Incremental Change Synchronization (ICS), which provides a mechanism to monitor and synchronize changes to an Exchange store hierarchy or content. The changes include creating, modifying, and deleting folders and messages. When users are connected to and using content mailboxes, synchronization occurs every 15 minutes. If no users are connected to content mailbox, synchronization will be triggered less often (every 24 hours).If a write operation such as a creating a folder is performed on the primary hierarchy, synchronization is triggered immediately (synchronously) to the content mailbox.

IMPORTANT

Because there's only one writeable copy of the hierarchy, folder creation is proxied to the hierarchy mailbox by the content mailbox users are connected to.

In a large organization, when you create a new public folder mailbox, the hierarchy must synchronize to that public folder before users can connect to it. Otherwise, users may see an incomplete public folder structure when connecting with Outlook. To allow time for this synchronization to occur without users attempting to connect to the new public folder mailbox, set the *IsExcludedFromServingHierarchy* parameter on the **New-Mailbox** cmdlet when creating the public folder mailbox. This parameter prevents users from connecting to the newly created public folder mailbox. When synchronization is complete, run the Set-Mailbox cmdlet with the *IsExcludedFromServingHierarchy* parameter set to `false`, indicating that the public folder mailbox is ready to be connected to. You can use also the Get-PublicFolderMailboxDiagnostics cmdlet to view the sync status by the *SyncInfo* and the *AssistantInfo* properties.

For more information, see Create a public folder.

Public folder content

Public folder content can include email messages, posts, documents, and eForms. The content is stored in the public folder mailbox but isn't replicated across multiple public folders mailboxes. All users access the same public folder mailbox for the same set of content. Although a full text search of public folder content is available, public folder content isn't searchable across public folders and the content isn't indexed by Exchange Search.

NOTE

Outlook on the web is supported, but with limitations. You can add and remove public folders to your Favorites through Outlook, and then perform item-level operations such as creating, editing, deleting posts, and replying to posts through Outlook on the web. However, you can't create or delete public folders from Outlook on the web. Also, only Mail, Post, Calendar, and Contact public folders can be added to the Favorites list in Outlook on the web.

Migrate public folders

- You can migrate public folders in the following scenarios:
- From Exchange 2010 to Exchange 2016 or to Exchange Online.
- From Exchange 2016 or later to Exchange Online.

If you already have Exchange 2010 SP3 public folders in your organization prior to installing Exchange 2016, you must migrate those public folders to Exchange 2016. To do this, use the **PublicFolderMigrationRequst** cmdlets. For more information, see Use batch migration to migrate Exchange 2010 public folders to Exchange 2016. If your organization is moving to Exchange Online, you can migrate your public folders to the cloud and upgrade them at the same time. For details, see Use batch migration to migrate legacy public folders to Microsoft 365, Office 365, and Exchange Online and Use batch migration to migrate Exchange Server public folders to Exchange Online.

Due to the changes in how public folders are stored, Exchange 2010 mailboxes are unable to access the public folder hierarchy on Exchange 2016 or on Exchange Online. However, user mailboxes on Exchange 2016 can connect to Exchange 2010 public folders. Exchange 2016 public folders and legacy public folders can't exist in your Exchange organization simultaneously. This effectively means that there's no coexistence between versions. Migrating public folders to Exchange Server 2016 or Exchange Online is currently a one-time cutover process.

For this reason, we recommend that prior to migrating your Exchange 2010 public folders, you should first migrate your Exchange 2010 mailboxes to Exchange 2016 or Exchange Online. For more information about migrating mailboxes, see Mailbox moves in Exchange Server, Migrate email using the Exchange cutover method, and Perform a staged migration of email to Microsoft 365 or Office 365.

Public folder moves

You can move public folders to a different public folder mailbox, and you can move public folder mailboxes to different mailbox databases. To move public folders to different public folder mailboxes, use the **PublicFolderMoveRequest** set of cmdlets. Subfolders under the public folder that's being moved won't be moved by default. If you want to move a branch of public folders, you can use the `Move-PublicFolderBranch.ps1` script that's installed by default with Exchange. For more information, see Move a Public Folder to a different Public Folder Mailbox.

In addition to moving public folders, you can move public folder mailboxes to different mailbox databases by using the **MoveRequest** set of cmdlets. This is the same set of cmdlets that are used for moving regular mailboxes. For more information, see Move a public folder mailbox to a different mailbox database.

PublicFolderMoveRequest cmdlets and the **MoveRequest** cmdlets use the Mailbox Replication Service to move public folders asynchronously. That means that the cmdlet doesn't do the actual work and, during most of the move, the public folder and public folder mailboxes will still be available to users. Because the Mailbox Replication Service performs mailbox moves, import and export requests, and public folder move requests, it's important to consider throttling and workload management.

Public folder quotas

By default, new public folder mailboxes automatically inherit the size limits of the mailbox database. As a result, to accurately evaluate the current storage quota status for the public folder mailbox using the Get-Mailbox cmdlet, you first need to review the value of the *UseDatabaseQuotaDefaults* property:

- If the value is `True`, the per-mailbox settings are ignored and the mailbox database limits are used.

- If the value is `False`, the per-mailbox settings are used.

If the *UseDatabaseQuotaDefaults* property is `True` and the *ProhibitSendQuota*, *ProhibitSendReceiveQuota*, and *IssueWarningQuota* properties are `unlimited`, the mailbox size isn't really unlimited. Instead, you need to use the Get-MailboxDatabase cmdlet and review the corresponding mailbox database storage limits to find out what the limits for the mailbox are. The default mailbox database quota limits are:

- *IssueWarningQuota*: 1.9 GB

- *ProhibitSendQuota*: 2 GB

- *ProhibitSendReceiveQuota*: 2.3 GB

To find the mailbox database quotas, run the Get-MailboxDatabase cmdlet.

To set the quotas on a public folder mailbox, use the Set-OrganizationConfig cmdlet with the *DefaultPublicFolderIssueWarningQuota* and *DefaultPublicFolderProhibitPostQuota* parameters.

Disaster recovery

Public folders are built on mailbox infrastructure and use the same mechanisms for availability and redundancy. Every public folder mailbox can have multiple redundant copies with automatic failover, just like regular mailboxes. To learn more, see Plan for high availability and site resilience.

In addition to the overall disaster recovery scenario, you can also restore public folders in the following situations:

- **Soft-deleted public folder restore**: The public folder was deleted but is still within the retention period.
- **Soft-deleted public folder mailbox restore**: The public folder mailbox was deleted and is still within the mailbox retention period.
- **Public folder mailbox restore from a recovery database**: You can recover an individual public folder mailbox from backup when the deleted mailbox retention period has elapsed. You then extract data from the restored mailbox and copy it to a target folder or merge it with another mailbox.

In all of these situations, the public folder or public folder mailbox is recoverable by using the **MailboxRestoreRequest** cmdlets.

For more information, see Restore public folders and public folder mailboxes from failed moves.

FAQ about public folder migration

8/3/2020 • 7 minutes to read • Edit Online

This article contains frequently asked questions about public folder migrations.

FAQs and more information

To learn more about public folders, see Public folders.

For more information on public folder migrations, see:

- Use batch migration to migrate Exchange 2010 public folders to Exchange 2016
- Migrate public folders from Exchange 2013 to Exchange 2016 or Exchange 2019
- Use batch migration to migrate legacy public folders to Microsoft 365, Office 365, or Exchange Online
- Use batch migration to migrate Exchange Server public folders to Exchange Online
- Use batch migration to migrate Exchange Server public folders to Microsoft 365 Groups

What are the supported public folder migration scenarios?

The following list details the available options for migrating public folders to Exchange or Exchange Online.

- Exchange 2010 public folders (SP3 RU8 or later) can be migrated to Exchange 2016, Exchange Online, or Microsoft 365 groups.
- Exchange 2013 public folders (CU15 or later) can be migrated to Exchange 2016, Exchange 2019, Exchange Online, or Microsoft 365 groups.
- Exchange 2016 public folders (CU4 or later) can be migrated to Exchange Online or Microsoft 365 groups.
- Exchange 2019 public folders can be migrated to Exchange Online or Microsoft 365 groups.

Currently only migrations to Exchange 2016 or Exchange 2019 in the same Active Directory forest are supported. Cross-forest migrations will be supported in the future.

After migration to Exchange 2016, what happens to the hierarchy on the source Exchange 2010 servers?

During the finalization stage in migration, a lock is placed on the source server to make it inaccessible to users. This lock remains in place to prevent users from accessing the source public folders after migration completes. Although you can release this lock, we don't recommend doing so because the changes can't be synced to Exchange 2016.

When you migrate public folders, what happens to existing public folder rules?

Public folder rules are migrated along with the data and are kept as public folder rules. They aren't converted to mailbox rules.

What happens if hierarchy changes are performed on the source after the initial .csv file was generated? How would these reflect on the destination?

The .csv file is used to determine the mapping between the source hierarchy and the destination mailbox. It contains only the top-level folders. Child folders under the top-level folders are automatically migrated. Therefore, if a new child folder is added, it's migrated during the process. If a new top-level folder is created, it will be created in the mailbox that contains the writable copy of the hierarchy.

For the migration of a geo-distributed hierarchy, how can I make sure that the public folders are created in the location nearest to the target users?

As part of the migration process, a .csv file is generated (using the `publicfoldertomailboxmapgenerator.ps1` script).

This file contains the folder-to-mailbox mapping for the new hierarchy. You can use this .csv file to create public folder mailboxes in the appropriate geographic location and modify the file to place the required folders in the appropriate mailbox so they are near the target users.

The input .csv file can be generated by running the script `AggregatePFData.ps1`, located in the directory < *Exchange Installation Directory*>\V15\Scripts. Run the script as follows:

```
.\AggregatePFData.ps1 | Select-Object -property @{Name="FolderName"; Expression = {$_.Identity}},
@{Name="FolderSize"; Expression = {$_.TotalItemSize.Value.ToBytes()}} | Export-CSV -Path <Path followed by the
name of the CSV>
```

Do existing public folder permissions migrate?

Yes, permissions automatically migrate at the folder level with the data. You don't have to perform this step separately.

Are public folders going away?

No. Public folders are great for Outlook integration, simple sharing scenarios, and for allowing large audiences to access the same data.

Which clients support public folders?

The currently supported Outlook clients for Exchange Server can access public folders. However, users with mailboxes on Exchange 2016 servers can't connect to Exchange 2010 public folders using Exchange Web Services (EWS) clients (for example, Outlook 2016 for Mac). We recommend that you migrate Exchange 2010 public folders to Exchange 2016 to maintain access for those users.

Can public folders be accessed using smart phones or mobile phones?

Public folder access works from Outlook for Windows desktop and Outlook for Mac. However, smart phone client apps including Outlook for Android or Outlook for iOS do not support connecting to public folders.

If you would like to have functionality similar to public folders with content accessible on mobile devices, consult Learn about Microsoft 365 Groups for an alternative.

Are there any limitations in the clients?

Outlook on the web (formerly known as Outlook Web App) is supported, but with some limitations. You can add and remove public folders to your Favorites (if they are Mail, Post, Calendar, or Contact public folders) and perform item level operations, such as creating, editing, deleting posts, and replying to posts. But, you can't do the following in Outlook on the web:

- Create or delete public folders
- Drag-and-drop content
- Access public folders located on servers running previous versions of Exchange

> **NOTE**
>
> You can only create public folder rules that contain the element **reply using a specific template** in mail-enabled public folders. It is possible that pre-existing rules containing **reply using a specific template** will continue to work on non-mail-enabled public folders, but on those folders you cannot create new rules with this template element, or edit existing rules with this element.

In a hybrid scenario, Outlook on the web isn't supported for cross-premises public folders. Users must be in the same location as the public folders to access them with Outlook on the web. Outlook 2016 for Mac users can access public folders in a hybrid scenario if the following conditions are true:

- You've followed the procedures at Hybrid Deployment procedures.
- The April 2016 update for Outlook 2016 for Mac has been installed on all clients.

How can I store a very large hierarchy in a public folder mailbox?

For more information about public folder storage limits, see Limits for public folders.

How can I view the hierarchy public folder mailbox?

Run the following command:

```
Get-OrganizationConfig | Format-List RootPublicFolderMailbox
```

For detailed syntax and parameter information, see Get-OrganizationConfig.

How can I create content mailboxes for public folders using Exchange Management Shell cmdlets?

Run the following command to create the first master hierarchy public folder mailbox and the secondary hierarchy mailboxes.

```
New-Mailbox -PublicFolder -Name <name of public folder>
```

For more detail, see Create a public folder.

In Exchange 2010 there was an option for each mailbox database to specify its public folder database. How does this work now?

There's no longer a database-level setting. Instead, Exchange has a mailbox-level ability to specify the public folder mailbox, but by default Exchange auto-calculates the per-user hierarchy mailbox.

How are public folder metric tools being used in Exchange?

You can use Get-PublicFolderStatistics and Get-PublicFolderItemStatistics cmdlets to get public folder metrics data. This same solution hase been available since Exchange 2010, so nothing has changed here. Public folders don't require additional reporting add-ons.

Can public folders distinguish between internal versus third-party access to public folders?

Starting in Exchange 2013, public folder permissions are managed by using role-based access control (RBAC); access control lists (ACLs) no longer used. You can use Get-PublicFolderStatistics and Get-PublicFolderItemStatistics cmdlets to keep track of accounts that are performing administrative tasks and then audit access accordingly. To learn more about RBAC, see Understanding Role Based Access Control.

Does mailbox audit logging work against public folders?

No. Not at this time.

If you would like to have functionality similar to public folders with audit logging, consult Learn about Microsoft 365 Groups for an alternative.

What are the limits on public folders? What are the recommendations?

For more information about public folder limits, see Limits for public folders.

What are the recommendations for splitting public folder mailboxes? Should they stay on the same database?

In previous versions of Exchange, you could split public folders across public folder databases. You can decide whether to split the content of a public folder mailbox to a mailbox on the same mailbox database or a different database. Typically, a split is recommended to be on a separate database, because you want to balance storage and I\O.

Can you set retention policies on public folders?

Just like in previous versions of Exchange, you can set retention limits on items. For details, see Limits for public folders.

Can you specify which users can use a specific public folder mailbox?

In Exchange 2010, you could specify which users had access to specific public folders. In Exchange 2013 or later, you can set the default public folder mailbox per user. To do so, run the Set-Mailbox cmdlet with the *DefaultPublicFolderMailbox* parameter. For example:

```
Set-Mailbox -Identity kweku@contoso.com -DefaultPublicFolderMailbox "PF_Administration"
```

If the master hierarchy goes down, what's the user impact?

If the master hierarchy public folder mailbox goes down, users can view but not write to public folders. To help prevent the hierarchy from going down, we recommend that you include your public folders in a database availability group (DAG). To learn about DAGs, see Database availability groups.

Can you change which public folder mailbox is the master hierarchy mailbox?

No. If you try to change the master hierarchy mailbox, you'll receive an error.

Do public folders have full text searching capabilities?

Yes, full text search has been available for public folders since Exchange 2013. However, you can't search across multiple public folders.

Limits for public folders

8/3/2020 • 2 minutes to read • Edit Online

In Exchange Server, public folders are based on a mailbox architecture that benefits from the resiliency of a Database Availability Group (DAG) and other mailbox enhancements. However, there are limits and performance considerations that you should take into account.

Limits

The following table lists the limits for public folders in on-premises Exchange Server. Unless the limits are specifically stated as recommended, the values listed in this table are the supported limits for public folders.

> **IMPORTANT**
> Looking for Exchange Online limits for Microsoft 365 or Office 365? See Exchange Online Limits.

ITEM	LIMITS	NOTES
Total number of public folder mailboxes	1,000	1,000 is the limit for Exchange Server 2016 CU2 or later. Although you can create more than 1,000 public folder mailboxes, it is not officially supported. See Create a public folder mailbox.
Total public folders in hierarchy	1,000,000	Although you can create more than 1,000,000 public folders, it is not officially supported. For any deployment of 100,000 or more public folders, we recommend reading Considerations when deploying public folders.
Sub-folders under the parent folder	10,000	Although you can create more than 1,000 sub-folders under a parent folder, it is not recommended. The limit can be enforced with the *FolderHierarchyChildrenCountReceiveQuota* parameter on the Set-Mailbox cmdlet.
Folder depth	300	The folder depth is the number levels of nested folders that can exist in one branch of a public folder tree. The limit can be enforced with the *FolderHierarchyDepthReceiveQuota* parameter on the Set-Mailbox cmdlet.
Maximum messages per public folder	1 million	The limit can be enforced with the *MailboxMessagesPerFolderCountRecieveQuota* parameter on the Set-Mailbox cmdlet.

ITEM	LIMITS	NOTES
Maximum individual public folder size	10 GB	This limit doesn't include subfolders beneath a single folder. See Configure storage quotas for a mailbox.
Public folder mailbox size	100 GB	Although public folder mailbox size can exceed 100 GB, it is not officially supported. See Configure storage quotas for a mailbox.
Number of user logons per public folder mailbox	2,000 concurrent user logons	We recommend that you configure your hierarchy so that you have no more than 2,000 users per public folder mailbox. For example, if you have 20,000 users, you should have 10 public folder mailboxes.
Moved item retention	14 days recommended	Use the *DefaultPublicFolderMovedItemRetention* parameter on the Set-OrganizationConfig cmdlet.
Age limit	We recommend that you set this as the same default that you use for regular mailboxes.	These settings can be set at the following levels: **Organizational level**: Use the *DefaultPublicFolderAgeLimit* parameter on the Set-OrganizationConfig cmdlet. **Folder level**: Use the *AgeLimit* parameter on the Set-PublicFolder cmdlet.
Deleted item retention	We recommend that you set this as the same default that you use for regular mailboxes.	These settings can be set at the following levels: **Organizational level**: Use the *DefaultPublicFolderMovedItemRetention* parameter on the Set-OrganizationConfig cmdlet. **Mailbox level**: Use the *RetainDeletedItemsFor* on the Set-Mailbox cmdlet. **Folder level**: Use the *RetainDeleteItemsFor* parameter on the Set-PublicFolder cmdlet.
Maximum number of public folders that can be migrated from Exchange 2010 to Exchange 2016	500,000	This is the maximum number of public folders you can move to Exchange from Exchange 2010 in a single migration. Although you can attempt to migrate more than 500,000 folders, it is not officially supported. For details on migrating public folders, see Use batch migration to migrate public folders from Exchange 2010 to Exchange 2016.

Considerations when deploying public folders

8/3/2020 • 2 minutes to read • Edit Online

Although there are many advantages to using Exchange public folders, there are some things to consider before implementing them in your organization.

Deployment considerations for public folders

This article contains factors to consider before you deploy public folders in your organization, especially if you plan to have a large number of public folders. Exchange Server supports up to one million public folders.

- Activity in a public folder directly impacts the load that's placed on the public folder mailbox where the folder is located. To avoid client connectivity issues, such as high latency or the inability to access a public folder, we recommend you do the following:

  - Don't let public folder mailboxes exceed 50% of the mailbox size limit. If this happens consider using the `Split-PublicFolderMailbox.ps1` script located in C:\Program Files\Microsoft\Exchange Server\V15\Scripts folder on the Exchange server to move some public folders to a new public folder mailbox.

  - Consider moving heavily-used public folders to a dedicated public folder mailbox.

  - Exclude heavily-used public folders from serving public folder hierarchy. You can do this by setting the *IsExcludedFromServingHierarchy* property on the public folder mailbox using the **Set-Mailbox** cmdlet.

  - For large organizations with many public folders, consider adding additional public folder mailboxes to distribute the load of servicing public folder hierarchy requests.

- Place the primary public folder mailbox in a DAG to improve availability of the mailbox. The primary public folder mailbox is the authoritative copy of the public folder hierarchy.

- Place secondary public folder mailboxes in a DAG or back up the mailboxes frequently.

- Place public folder mailboxes in the geographical location that's nearest the users that will access the public folder content in them.

- Improve public folder hierarchy access times by using the DefaultPublicFolderMailbox property on the users' mailboxes to specify a public folder mailbox close to them. This will prevent those users from retrieving the public folder hierarchy from a public folder mailbox in other geographical locations.

- In deployments with more than 50 secondary public folder mailboxes, we recommend that you don't store public folder content in the primary public folder mailbox. This dedicates the primary public folder mailbox to synchronizing the hierarchy with the secondary public folder mailboxes.

- Exchange 2016 doesn't support public folder databases. As a result, Outlook on the web users will not be able to access Exchange 2010 public folders. Exchange 2016 users can access Exchange 2010 public folders with Outlook or Outlook for Mac.

- Outlook on the web is supported, but with limitations. You can add and remove public folders from your Favorites and perform item-level operations such as creating, editing, deleting posts, and replying to posts. However, you can't create or delete public folders from Outlook on the web. Also, only Mail, Post, Calendar, and Contact public folders can be added to the Favorites list in Outlook on the web.

- Although a full text search of public folder content is available, public folder content isn't searchable across public folders and the content isn't indexed by Exchange Search.
- You must use Outlook 2010 or later to access public folders on Exchange servers.
- Retention policies aren't supported for public folder mailboxes.

Migrate your public folders to Microsoft 365 Groups

8/3/2020 • 8 minutes to read • Edit Online

This article provides a comparison of public folders and Microsoft 365 Groups, and how one or the other might be the best solution for your organization. Public folders have been around as long as Exchange, whereas Groups were introduced more recently. If you want to migrate some or all of your public folders to Groups, this article describes how the process works, and provides links to the articles that walk you through the process, step by step.

What are public folders?

Public folders contain different kinds of data and are organized in a hierarchical structure.

Public folders are not recommended for the following situations:

- **Archiving data**: Users with mailbox limits sometimes use public folders instead of mailboxes to archive data. This practice isn't recommended because it affects storage in public folders and undermines the goal of mailbox limits.
- **Document sharing and collaboration**: Public folders don't provide document management features, such as versioning, controlled check-in and check-out functionality, and automatic notifications of content changes.

What are Microsoft 365 Groups?

Microsoft 365 groups let you choose a set of people who you wish to collaborate with, and then easily set up a collection of resources for those people to share. You don't have to worry about manually assigning permissions to those resources, because adding members to your group automatically gives the members the permissions they need to access the tools and resources your group provides. Groups are also the new and improved experience for those tasks that were previously handled by distribution lists and shared mailboxes.

For the full Groups story, see Learn about Microsoft 365 Groups.

Should you migrate your public folders to Microsoft 365 Groups?

Microsoft 365 Groups is the latest collaboration offering from Microsoft, which means there are many reasons why they would be a preferable solution over public folders, a much older technology. In Outlook, for example, Groups can replace mail-enabled public folders altogether. Compiling a list of every scenario in which Microsoft 365 Groups works better than public folders is impossible, but here are the highlights:

- **Collaboration over email**: Groups in Outlook has a dedicated **Conversations** space that stores all the emails and lets users collaborate over them. The group can even be set up to receive messages from people outside the group or from outside the organization. If you're currently using mail-enabled public folders to store project-related discussions, for example, or purchase orders that need to be viewed by a team of people, using groups would be an improvement. Groups are also better for situations when you simply want to broadcast information to a set of users.
- **Collaboration over documents**: In Outlook, Groups has a dedicated **Files** tab that displays all files from the group's SharePoint team site, as well as from mail attachments. You get one view of all the files, so you don't have to go searching for them like you would in public folders. Co-authoring also becomes easier. If you're using public folders for storing files meant to be consumed by multiple people, consider migrating to Groups.

- **Shared calendar**: Upon creation every group gets a shared calendar (see Calendar sharing in Microsoft 365. Any member of the group can create events on that calendar. When you favorite a group, that group's calendar can be displayed alongside your personal calendar. You can also subscribe to a group's events, in which case events created in that group appear in your personal calendar. If you're using public folders to host calendars for your team, such as a schedule or a timetable, Groups would be an improved experience.

- **Simplified permissions**: When you assign users to a group, they immediately get the permissions they need, whereas with public folders you need to manually assign the proper permissions. Members can be added as "owners" or "members." Owners have full rights in the group, including the ability to perform group management tasks. Members can also create content and edit files like owners, but members cannot delete content that they have not created. If the public folders' permissions model is too overwhelming for you and you want something simple and quick, Microsoft 365 Groups is the way to go.

- **Mobile and Web presence**: Public folders can't be accessed through mobile devices and have a limited set of functionality on the Web. Microsoft 365 Groups, on the other hand, is accessible through Outlook mobile apps and has a richer set of features on the Web. If your team is on the move and requires mobile access, then you should be using Microsoft 365 Groups.

- **Access to a wide range of Microsoft 365 or Office 365 apps**: When you create a group, you unlock access to a wide range of apps from the Microsoft 365 or Office 365 suite. You get a SharePoint team site for storing files and a plan on Planner to track your tasks. Microsoft 365 Groups is the membership service that combines elements of the entire Microsoft 365 or Office 365 suite.

While Microsoft 365 Groups offers many advantages, you should be aware of a few major differences that you'll notice after leaving the public folders experience. These are primarily:

- **Folder hierarchy**: While public folders are often used to organize content in deep rooted hierarchy, Microsoft 365 Groups has a flat structure. All emails in the group reside in the Conversations space and all the documents go into the **Files** tab. Also, you can't create sub-folders in Microsoft 365 groups.

- **Granular permission roles**: While public folders have a variety of permission roles, Microsoft 365 Groups only provides two: owner and member.

Before you move to Groups, it's also a good idea to make note of the various limits that come with creating and maintaining groups. See *How do I manage my groups?* in Learn about Microsoft 365 Groups for more information.

Migrating public folders to Microsoft 365 Groups

If you decide to switch to Microsoft 365 Groups, you can use a process known as *batch migration* to move your email and calendar content from your existing public folders to Groups. The specific steps for running a batch migration depends on which version of Exchange currently hosts your public folder hierarchy. At the end of this article, you will find links to instructions that walk you through the batch migration process.

> **NOTE**
>
> When you finish migrating a mail-enabled public folder to a particular group in Microsoft 365 or Office 365, all the emails addressed to the public folder will at that point be received by the group.

Key benefits of batch migrations are:

- **Mailbox Replication Service (MRS)-based migration**: The migration process uses migration batch cmdlets. Migration to multiple groups can be triggered together in a single migration batch. There are also scripts available to assist in the migration process.

- **Supports mail and calendar public folders**: Copied emails and posts will appear as in Groups as group conversations, and copied calendar items will be visible in group calendars. Other public folder types, such as tasks and contacts, are currently not supported for this migration.

- **On-premises public folders can be migrated directly to Microsoft 365 Groups**: This migration does not require you to first move your public folders to Microsoft 365 or Office 365 and then move to Groups. The MRS data copy cmdlets read the public folder data directly from your on-premises environment and then copy the data to Microsoft 365 Groups. Note that Exchange public folders will require an MRS Proxy-based endpoint.

- **Not an "all or nothing" migration**: You get to choose specific public folders to migrate to Groups, and only those chosen public folders get migrated.

- **One-shot data copy**: Batch migrations are designed to be a simple one-time data copy from source public folders to target groups, without the complexities of incremental synchronization and finalization.

- **Merges public folder data with existing data in a group**: The data copy will merge the public folder content with the existing group's content, if any. If there is a need for incremental data copy, you can simply run the data copy as many times as you need. This will copy incremental data over to the group.

Overview of batch migrations

The following steps outline the overall process of migrating your public folder content to Microsoft 365 Groups in a batch migration. The specific details are contained in the articles listed below.

1. **Select source**: Choose the public folders that you want to migrate. You can choose any folder containing mail or calendar content.

2. **Create target**: Create corresponding groups for your folders, with the desired configurations, such as members, privacy settings, and data classification.

3. **Copy data**: Use the migration batch cmdlets to copy data from public folders to Groups.

4. **Lock source**: Lock the public folders once you have verified the data in Groups.

5. **Cutover**: Copy any new data that has been created between steps 3 and 4.

Note that your public folders and their corresponding groups will remain online for your users during steps 1 through 3 above. After step 3, you can evaluate whether or not to proceed with the rest of the migration, based on the Groups experience and whether or not it suits your users and your organization. You can roll back your migration and resume using public folders at that point. If you do proceed with the migration, after step 5 completes, you can delete the original public folders. Even post-migration it is possible to roll back to public folders, provided you have saved your backup files from the migration process and you have not deleted your original public folders.

Batch migration prerequisites and step-by-step instructions

The following prerequisites are required in your Exchange environment before you can run a batch migration. The specific prerequisites depend on which version of Exchange you're currently running.

1. If your public folders are on-premises, your servers need to be running one of the following versions:

   - Exchange 2010 SP3 RU8 or later

   - Exchange 2013 CU15 or later

   - Exchange 2016 CU4 or later

   - Exchange 2019

2. If your public folders are on-premises, you must have an Exchange Hybrid environment set up. See Exchange Server Hybrid Deployments for more information.

Migration instructions

Click one of the links below for step-by-step instructions on running a batch migration.

- Use batch migration to migrate Exchange Server public folders to Microsoft 365 Groups
- Use batch migration to migrate your Exchange Online public folders to Microsoft 365 Groups
- Use batch migration to migrate your Exchange 2013 public folders to Microsoft 365 Groups
- Use batch migration to migrate your Exchange 2010 public folders to Microsoft 365 Groups

Public folder procedures

8/3/2020 • 2 minutes to read • Edit Online

Use one or more of the procedures listed below to get your public folder infrastructure up and running, and to perform other necessary tasks for managing public folders.

Set up public folders in a new organization

Configure legacy on-premises public folders for a hybrid deployment

Configure modern on-premises public folders for a hybrid deployment

Use batch migration to migrate Exchange 2010 public folders to Exchange 2016

Use batch migration to migrate legacy public folders to Microsoft 365 or Office 365 and Exchange Online

Use batch migration to migrate Exchange Server public folders to Exchange Online

Migrate public folders from Exchange 2013 to Exchange 2016 or Exchange 2019

Configure legacy public folders where user mailboxes are on Exchange 2016 servers

Create a public folder mailbox

Create a public folder

Using favorite public folders in Outlook on the web

Mail-enable or mail-disable a public folder

Update the public folder hierarchy

Remove a public folder

Move a public folder mailbox to a different mailbox database

Move a public folder to a different public folder mailbox

Restore public folders and public folder mailboxes from failed moves

View statistics for public folders and public folder items

Configure legacy on-premises public folders for a hybrid deployment

8/3/2020 • 8 minutes to read • Edit Online

In a hybrid deployment, your users can be in Exchange Online, Exchange on-premises, or both, and your public folders are either in Exchange Online or Exchange on-premises. Public folders can only reside in one place, so you must decide whether your public folders will be in Exchange Online or on-premises. They can't be in both locations. Public folder mailboxes are synchronized to Exchange Online by the Directory Synchronization service. However, mail-enabled public folders aren't synchronized across premises.

This article describes how to synchronize mail-enabled public folders when your users are in Microsoft 365 or Office 365 and your Exchange 2010 SP3 or later version public folders are on-premises. However, a Microsoft 365 or Office 365 user who is not represented by a MailUser object on-premises (local to the target public folder hierarchy) won't be able to access legacy or Exchange on-premises public folders.

> **NOTE**
>
> This topic refers to the Exchange 2010 SP3 or later servers as the *legacy Exchange server*.

You will use the following scripts to sync your mail-enabled public folders. The scripts are initiated by a Windows task that runs in the on-premises environment:

- `Sync-MailPublicFolders.ps1` : This script synchronizes mail-enabled public folder objects from your local Exchange on-premises deployment with Microsoft 365 or Office 365. It uses the local Exchange on-premises deployment as master to determine what changes need to be applied to Microsoft 365 or Office 365. The script will create, update, or delete mail-enabled public folder objects on Microsoft 365 or Office 365 Active Directory based on what exists in the local on-premises Exchange deployment.

- `SyncMailPublicFolders.strings.psd1` : This is a support file used by the preceding synchronization script and should be copied to the same location as the preceding script.

When you complete this procedure your on-premises and Microsoft 365 or Office 365 users will be able to access the same on-premises public folder infrastructure.

What hybrid versions of Exchange will work with public folders?

The following table describes the supported version and location combinations of user mailboxes and public folders. "Hybrid not applicable" is still a supported scenario, but is not considered a hybrid scenario because both the public folders and the users are residing in the same location.

SCENARIO	ON-PREMISES EXCHANGE 2010 USER MAILBOX	ON-PREMISES EXCHANGE 2016/2019 USER MAILBOX	EXCHANGE ONLINE USER MAILBOX
On-Premises Exchange 2010 Public Folders	Hybrid not applicable	Hybrid not applicable	Supported
On-Premises Exchange 2013, Exchange 2016, or Exchange 2019 Public Folders	Hybrid not applicable	Hybrid not applicable	Supported

SCENARIO	ON-PREMISES EXCHANGE 2010 USER MAILBOX	ON-PREMISES EXCHANGE 2016/2019 USER MAILBOX	EXCHANGE ONLINE USER MAILBOX
Exchange Online Public Folders	Not supported	Supported	Hybrid not applicable

A hybrid configuration with Exchange 2003 public folders is not supported. If you're running Exchange 2003 in your organization, you must move all public folder databases and replicas to Exchange 2010 SP3 or later. No public folder replicas can remain on Exchange 2003.

Step 1: What do you need to know before you begin?

- These instructions assume that you have used the Hybrid Configuration Wizard to configure and synchronize your on-premises and Exchange Online environments and that the DNS records used for most users' Autodiscover references an on-premises end-point. For more information, see Hybrid Configuration Wizard.
- These instructions assume that Outlook Anywhere is enabled and functional on the on-premises legacy Exchange servers. For information on how to enable Outlook Anywhere, see Outlook Anywhere.
- Implementing legacy public folder coexistence for a hybrid deployment of Exchange with Microsoft 365 or Office 365 may require you to fix conflicts during the import procedure. Conflicts can happen due to non-routable email address assigned to mail enabled public folders, conflicts with other users and groups in Microsoft 365 or Office 365, and other attributes.
- These instructions assume your Exchange Online organization has been upgraded to a version that supports public folders.
- In Exchange Online, you must be a member of the Organization Management role group. This role group is different from the permissions assigned to you when you subscribe to Exchange Online. For details about how to enable the Organization Management role group, see Manage role groups.
- In Exchange 2010, you must be a member of the Organization Management or Server Management Role Based Access Control (RBAC) role groups. For details, see Add Members to a Role Group.
- In order to access public folders cross-premises, users must upgrade their Outlook clients to the November 2012 or later Outlook public update.
  - To download the November 2012 Outlook update for Outlook 2010, see Update for Microsoft Outlook 2010 (KB2687623) 32-Bit Edition.
  - To download the November 2012 Outlook update for Outlook 2007, see Update for Microsoft Office Outlook 2007 (KB2687404) and download in preferred language.
- Outlook 2016 for Mac and Outlook for Mac for Microsoft 365 or Office 365 are supported for cross-premises public folders if the following conditions are true:
  - The April 2016 update for Outlook 2016 for Mac is installed.
  - Exchange 2016 CU2 or later.
  - Exchange 2013 CU14 or later.
- After you have followed the instructions in this article to configure your on-premises public folders for a hybrid deployment, users who are external to your organization won't be able to send messages to your on-premises public folders unless you take additional steps. You can either set the accepted domain for the public folders to Internal Relay (see Manage accepted domains in Exchange Online for more information) or you can disable Directory Based Edge Blocking (DBEB), as described in Use Directory Based Edge Blocking to

Reject Messages Sent to Invalid Recipients.

- In hybrid mode, Exchange Online users can't access public folders using Outlook on the web (formerly known as Outlook Web App).

Step 2: Make remote public folders discoverable

1. If your public folders are on Exchange 2010 servers, you must install Client Access services on all mailbox servers that have a public folder database. This enables the Exchange RpcClientAccess service to run, which enables all clients to access public folders. For more information, see Install Exchange Server 2010.

   > **NOTE**
   >
   > This server doesn't have to be part of the Client Access load balancing. For more information, see Understanding Load Balancing in Exchange 2010.

2. Create an empty mailbox database on each public folder server.

   For Exchange 2010, run the following command in the Exchange Management Shell. This command excludes the mailbox database from the mailbox provisioning load balancer. This prevents new mailboxes from automatically being added to this database.

   ```
   New-MailboxDatabase -Server <PFServerName_with_CASRole> -Name <NewMDBforPFs> -IsExcludedFromProvisioning $true
   ```

   For Exchange 2007, run the following command in the Exchange Management Shell:

   ```
   New-MailboxDatabase -StorageGroup "<PFServerName>\StorageGroup>" -Name <NewMDBforPFs>
   ```

   > **NOTE**
   >
   > We recommend that the only mailbox that you add to this database is the proxy mailbox that you'll create in the next step. No other mailboxes should be created on this mailbox database.

3. Create a proxy mailbox within the new mailbox database, and hide the mailbox from the address book. The SMTP of this mailbox will be returned by AutoDiscover as the *DefaultPublicFolderMailbox* SMTP, so that by resolving this SMTP the client can reach the legacy exchange server for public folder access.

   ```
   New-Mailbox -Name <PFMailbox1> -Database <NewMDBforPFs>
   ```

   ```
   Set-Mailbox -Identity <PFMailbox1> -HiddenFromAddressListsEnabled $true
   ```

4. For Exchange 2010, enable Autodiscover to return the proxy public folder mailboxes.

   ```
   Set-MailboxDatabase <NewMDBforPFs> -RPCClientAccessServer <PFServerName_with_CASRole>
   ```

5. Repeat the preceding steps for every public folder server in your organization.

Step 3: Download the scripts

1. Download the following files from Mail-enabled Public Folders - directory sync script:

   - `Sync-MailPublicFolders.ps1`
   - `SyncMailPublicFolders.strings.psd1`

2. Save the files to the local computer on which you'll be running PowerShell. For example, C:\PFScripts.

Step 4: Configure directory synchronization

The Directory Synchronization service doesn't synchronize mail-enabled public folders. Running the following script will synchronize the mail-enabled public folders across premises. Special permissions assigned to mail-enabled public folders will need to be recreated in the cloud since cross-premise permission are not supported in Hybrid Deployment scenarios.

> **NOTE**
>
> Synchronized mail-enabled public folders will appear as mail contact objects for mail flow purposes and will not be viewable in the Exchange admin center. See the Get-MailPublicFolder command. To recreate the SendAs permissions in the cloud, use the Add-RecipientPermission command.

1. On the legacy Exchange server, run the following command to synchronize mail-enabled public folders from your local on-premises Active Directory to Microsoft 365 or Office 365.

   ```
   Sync-MailPublicFolders.ps1 -Credential (Get-Credential) -CsvSummaryFile:sync_summary.csv
   ```

   Where `Credential` is your Microsoft 365 or Office 365 name and password, and `CsvSummaryFile` is the path to where you would like to log synchronization operations and errors, in .csv format.

> **NOTE**
>
> Before running the script, we recommend that you first simulate the actions that the script would take in your environment by running it as described above with the `-WhatIf` parameter. We also recommend that you run this script daily to synchronize your mail-enabled public folders.

Step 5: Configure Exchange Online users to access on-premises public folders

The final step in this procedure is to configure the Exchange Online organization and to allow access to the legacy on-premises public folders.

You will point to all of the proxy public folder mailboxes that you created in Step 2: Make remote public folders discoverable to enable theExchange Online organization to access the on-premises public folders.

Run the following command in Exchange Online PowerShell. To learn how to use Windows PowerShell to connect to Exchange Online, see Connect to Exchange Online PowerShell.

```
Set-OrganizationConfig -PublicFoldersEnabled Remote -RemotePublicFolderMailboxes
PFMailbox1,PFMailbox2,PFMailbox3
```

You must wait until Active Directory synchronization has completed to see the changes. This process can take up to 3 hours to complete. If you don't want to wait for the recurring synchronizations that occur every three hours, you can force directory synchronization at any time. For detailed steps to force directory synchronization, see Azure AD

Connect sync: Scheduler. Microsoft 365 or Office 365 randomly selects one of the public folder mailboxes that's supplied in this command.

> **IMPORTANT**
>
> A Microsoft 365 or Office 365 user who is not represented by a MailUser object on-premises (local to the target public folder hierarchy) won't be able to access legacy, Exchange 2016, or Exchange 2019 on-premises public folders. See the Knowledge Base article Exchange Online users can't access legacy on-premises public folders for a solution.

How do I know this worked?

Log on to Outlook for a user who is in Exchange Online and perform the following public folder tests:

- View the hierarchy.
- Check permissions
- Create and delete public folders.
- Post content to and delete content from a public folder.

Use batch migration to migrate Exchange 2010 public folders to Exchange 2016

8/3/2020 • 16 minutes to read • Edit Online

Migrate your public folders from Exchange Server 2010 SP3 RU8 to Exchange Server 2016 within the same forest.

We refer to the Exchange 2010 SP3 RU8 or later server as the *legacy Exchange server*.

> **NOTE**
> The batch migration method described in this article is the only supported method for migrating legacy public folders to Exchange Server. The old serial migration method for migrating public folders is being deprecated and is no longer supported by Microsoft.

You'll perform the migration by using the ***MigrationBatch** cmdlets, and the ***PublicFolderMigrationRequest** cmdlets for troubleshooting. In addition, you'll use the following PowerShell scripts:

- `Export-PublicFolderStatistics.ps1` : This script creates the folder name-to-folder size mapping file.
- `Export-PublicFolderStatistics.psd1` : This support file is used by the Export-PublicFolderStatistics.ps1 script and should be downloaded to the same location.
- `PublicFolderToMailboxMapGenerator.ps1` : This script creates the public folder-to-mailbox mapping file.
- `PublicFolderToMailboxMapGenerator.strings.psd1` : This support file is used by the PublicFolderToMailboxMapGenerator.ps1 script and should be downloaded to the same location.
- `Create-PublicFolderMailboxesForMigration.ps1` : This script creates the target public folder mailboxes for the migration. In addition, this script calculates the number of mailboxes necessary to handle the estimated user load, based on the guidelines for the number of user logons per public folder mailbox recommended in Limits for public folders.
- `Create-PublicFolderMailboxesForMigration.strings.psd1` : This support file is used by the Create-PublicFolderMailboxesForMigration.ps1 script and should be downloaded to the same location.

The Step 1: Download the migration scripts section provides details about where to download these scripts. Be sure to download all scripts to the same location.

For additional management tasks related to public folders, see Public folder procedures.

What migration pathways are supported for Exchange Server versions?

Exchange supports moving your public folders from the following legacy versions of Exchange Server:

- Exchange 2010 SP3 RU8 or later

You can't migrate public folders directly from Exchange 2003. If you're running Exchange 2003 in your organization, you need to move all public folder databases and replicas to Exchange 2010 SP3 RU8 or later. No public folder replicas can remain on Exchange 2003. Additionally, mail destined for an Exchange 2016 public folder can't be routed through an Exchange 2003 server.

What do you need to know before you begin?

- Before you begin, we recommend that you read this topic in its entirety as downtime is required for some steps.

- The Exchange 2010 server needs to be running Exchange 2010 SP3 RU8 or later.

- The maximum number of public folders that can be migrated to Exchange 2016 in a single migration is 500,000.

- In Exchange 2016, you need to be a member of the Organization Management role group. For details about how to enable the Organization Management role group, see Manage role groups.

- In Exchange 2010, you need to be a member of the Organization Management or Server Management RBAC role groups. For details, see Add Members to a Role Group.

- Before you migrate, you should consider the Limits for public folders.

- Before you migrate, move all user mailboxes to Exchange 2016, because users with Exchange 2010 mailboxes will not have access to public folders on Exchange 2016. For details, see Mailbox moves in Exchange Server.

- In a multiple-domain environment, mail-enabled public folders will stop working after migration to Exchange 2016 if Exchange is running in a child domain. This is because in Exchange 2016, mail-enabled public folder objects are required to be under the root domain. To resolve this, you need to mail-disable your mail-enabled public folders and then mail-enable them again, which will allow you to move them to the correct domain location.

- After the migration is complete, if you want external senders to send mail to the migrated mail-enabled public folders, the **Anonymous** user needs to be granted at least the **Create Items** permission. If you don't do this, external senders will receive a delivery failure notification and the messages won't be delivered to the migrated mail-enabled public folder. To read more about how to set permissions on the Anonymous user, see Mail-enable or mail-disable a public folder.

- You must use a single migration batch to migrate all of your public folder data. Exchange allows creating only one migration batch at a time. If you attempt to create more than one migration batch simultaneously, the result will be an error.

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **IMPORTANT**
>
> Before you begin your migration, make sure you migrate your arbitration mailbox to the target Exchange server. Otherwise, your migration batch will hang in the **Starting** state. To identify your migration arbitration mailbox, run the following cmdlet:
>
> `Get-Mailbox -Arbitration -Identity Migration.*`

Step 1: Download the migration scripts

1. Download all scripts and supporting files from Public Folders Migration Scripts.

2. Save the scripts to the local computer on which you'll be running PowerShell. For example, C:\PFScripts. Make sure all scripts are saved in the same location.

Step 2: Prepare for the migration

Perform the following prerequisite steps before you begin the migration.

Prerequisite steps on the Exchange 2010 server

1. For verification purposes at the end of migration, we recommend that you first run the following commands on the Exchange 2010 server to take snapshots of your current public folder deployment:

   - Run the following command to take a snapshot of the original source folder structure:

     ```
     Get-PublicFolder -Recurse | Export-CliXML C:\PFMigration\Legacy_PFStructure.xml
     ```

   - Run the following command to take a snapshot of public folder statistics such as item count, size, and owner:

     ```
     Get-PublicFolderStatistics | Export-CliXML C:\PFMigration\Legacy_PFStatistics.xml
     ```

   - Run the following command to take a snapshot of the permissions:

     ```
     Get-PublicFolder -Recurse | Get-PublicFolderClientPermission | Select-Object Identity,User -
     ExpandProperty AccessRights | Export-CliXML C:\PFMigration\Legacy_PFPerms.xml
     ```

2. If the name of a public folder contains a backslash (\ ), migration will create the migrated public folders in the parent public folder. Before you migrate, we recommend that you rename any public folders that have a backslash in the name.

   To locate public folders in Exchange 2010 that have a backslash in the name, run the following command:

   ```
   Get-PublicFolderStatistics -ResultSize Unlimited | Where {($_.Name -like "*\*") -or ($_.Name -like
   "*/*") } | Format-List Name, Identity
   ```

   If any public folders are returned, you can rename them by running the following command:

   ```
   Set-PublicFolder -Identity <public folder identity> -Name <new public folder name>
   ```

3. Make sure there isn't a record of a previously successful migration by running the following command:

   ```
   Get-OrganizationConfig | Format-List PublicFoldersLockedforMigration, PublicFolderMigrationComplete
   ```

   A previously successful migration will set the *PublicFoldersLockedforMigration* or *PublicFolderMigrationComplete* properties to the value `True`, which will cause your new migration request to fail.

   If the property values are `True`, run the following command to change them to `False`:

   ```
   Set-OrganizationConfig -PublicFoldersLockedforMigration $false -PublicFolderMigrationComplete $false
   ```

   > **NOTE**
   >
   > After resetting these properties, you need to wait for Exchange to detect the new settings. This may take up to two hours to complete.

For detailed syntax and parameter information, see the following topics:

- Get-PublicFolder

- Get-PublicFolderDatabase
- Set-PublicFolder
- Get-PublicFolderStatistics
- Get-PublicFolderClientPermission
- Get-OrganizationConfig
- Set-OrganizationConfig

Prerequisite steps on the Exchange 2016 server

1. Make sure there are no existing public folder migration requests. If there are, clear them or your own migration request will fail. This step isn't required in all cases; it's only required if you think there may be an existing migration request in the pipeline.

   > **IMPORTANT**
   >
   > Before removing a migration request, it is important to understand why there was an existing one. Running the following commands will determine when a previous request was made and help you diagnose any problems that may have occurred. You may need to communicate with other administrators in your organization to determine why the change was made.

   - Run the following command to discover any existing batch migration requests:

     ```
     $batch = Get-MigrationBatch | ?{$_.MigrationType.ToString() -eq "PublicFolder"}
     ```

   - Run the following command to remove any existing public folder batch migration requests.

     ```
     $batch | Remove-MigrationBatch -Confirm:$false
     ```

2. Make sure no public folders or public folder mailboxes exist on the Exchange 2016 servers by running the following command:

   ```
   Get-Mailbox -PublicFolder
   ```

   If the command didn't return any public folder mailboxes, continue to Step 3: Generate the .csv files. If the command returned any public folders, run the following command to see if any public folders exist:

   ```
   Get-PublicFolder
   ```

   If you have any public folders, run the following commands to remove them. Make sure you've saved any information that was in the public folders.

   > **NOTE**
   >
   > All information contained in the public folders will be permanently deleted when you remove them.

   ```
   Get-Mailbox -PublicFolder | Where {$_.IsRootPublicFolderMailbox -eq $false} | Remove-Mailbox -
   PublicFolder -Force -Confirm:$false
   ```

```
Get-Mailbox -PublicFolder | Remove-Mailbox -PublicFolder -Force -Confirm:$false
```

For detailed syntax and parameter information, see the following topics:

- Get-MigrationBatch
- Get-Mailbox
- Get-PublicFolder
- Get-MailPublicFolder
- Disable-MailPublicFolder
- Remove-PublicFolder
- Remove-Mailbox

Step 3: Generate the .csv files

1. On the Exchange 2010 server, run the `Export-PublicFolderStatistics.ps1` script to create the folder name-to-folder size mapping file. This script needs to be run by a local administrator. The file will contain two columns: **FolderName** and **FolderSize**. The values for the **FolderSize** column will be displayed in bytes. For example, **\PublicFolder01,10000**.

   ```
   .\Export-PublicFolderStatistics.ps1 <Folder to size map path> <FQDN of source server>
   ```

   - *FQDN of source server* equals the fully qualified domain name of the Mailbox server where the public folder hierarchy is hosted.

   - *Folder to size map path* equals the file name and path on a network shared folder where you want the .csv file saved. Later in this topic, you'll need to access this file from the Exchange 2016 server. If you specify only the file name, the file will be generated in the current PowerShell directory on the local computer.

2. Run the `PublicFolderToMailboxMapGenerator.ps1` script to create the public folder-to-mailbox mapping file. This file is used to calculate the correct number of public folder mailboxes on the Exchange 2016 server.

   > **NOTE**
   >
   > If the name of a public folder contains a backslash ****, the public folders will be created in the parent public folder. We recommend that you review the .csv file and edit any names that contain a backslash.

   ```
   .\PublicFolderToMailboxMapGenerator.ps1 <Maximum mailbox size in bytes> <Folder to size map path>
   <Folder to mailbox map path>
   ```

   - *Maximum mailbox size in bytes* equals the maximum size you want to set for the new public folder mailboxes. When specifying this setting, be sure to allow for expansion so the public folder mailbox has room to grow.

   - *Folder to size map path* equals the file path of the .csv file you created when running the `Export-PublicFolderStatistics.ps1` script.

   - *Folder to mailbox map path* equals the file name and path of the folder-to-mailbox .csv file that you'll create with this step. If you specify only the file name, the file will be generated in the current

PowerShell directory on the local computer.

Step 4: Create the public folder mailboxes in Exchange 2016

Run the following command to create the target public folder mailboxes. The script will create a target mailbox for each mailbox in the .csv file that you generated previously in Step 3 by running the `PublicFoldertoMailboxMapGenerator.ps1` script.

```
.\Create-PublicFolderMailboxesForMigration.ps1 -FolderMappingCsv Mapping.csv -
EstimatedNumberOfConcurrentUsers:<estimate>
```

Mapping.csv is the file generated by the `PublicFoldertoMailboxMapGenerator.ps1` script in Step 3. The estimated number of simultaneous user connections browsing a public folder hierarchy is usually less than the total number of users in an organization.

Step 5: Start the migration request

After you crate the batch migration request in the Exchange Management Shell, you can view the requests and manage them in the Exchange admin center (EAC).

1. On the Exchange 2016 server, run the following command:

   ```
   New-MigrationBatch -Name PFMigration -SourcePublicFolderDatabase (Get-PublicFolderDatabase -Server
   <Source server name>) -CSVData (Get-Content <Folder to mailbox map path> -Encoding Byte) -
   NotificationEmails <email addresses for migration notifications>
   ```

   The `NotificationEmails` parameter is optional.

2. Start the migration in the EAC or in the Exchange Management Shell.

   - In the Exchange Management Shell, run the following command:

     ```
     Start-MigrationBatch PFMigration
     ```

   - In the EAC:

     a. Log into Exchange Online and open the EAC.

     b. Go to **Recipients** > **Migration**.

     c. Select the migration batch you just created, and then click the start button.

   In the EAC, the **Status** column will show the initial batch status as **Created**. The status changes to **Syncing** during migration. When the migration request is complete, the status will be **Synced**. You can double-click a batch to view the status of individual mailboxes within the batch. Mailbox jobs begin with a status of **Queued**. When the job begins the status is **Syncing**, and once `InitialSync` is complete, the status will show **Synced**.

You can view and manage the progress and completion of the migration in the **Recipients** > **Migration** tab in the EAC.

Because the **New-MigrationBatch** cmdlet initiates a mailbox migration request for each public folder mailbox, you can view the status of these requests using the mailbox migration page in the EAC, and you can create migration reports that can be emailed to you.

1. Log into Exchange Online and open the EAC.

2. Go to **Recipients** > **Migration**.

3. Select the migration request that you just created and then click **View Details** in the **Details** pane.

For detailed syntax and parameter information, see the following topics:

- New-MigrationBatch

- Get-PublicFolderDatabase

- Get-PublicFolderMailboxMigrationRequest

- Get-PublicFolderMailboxMigrationRequestStatistics

Step 6: Lock down the public folders on the Exchange 2010 server for final migration (downtime required)

Until this point in the migration, users have been able to access public folders. The next steps will log users off from the Exchange 2010 public folders and lock the folders while the migration completes its final synchronization. Users won't be able to access public folders during this process. Also, any mail sent to mail-enabled public folders will be queued and won't be delivered until the public folder migration is complete.

Before you run the `PublicFoldersLockedForMigration` command as described below, make sure that all jobs are in the **Synced** state. You can do this by running the `Get-PublicFolderMailboxMigrationRequest` command. Continue with this step only after you've verified that all jobs are in the **Synced** state.

On the Exchange 2010 server, run the following command to lock the public folders for finalization.

```
Set-OrganizationConfig -PublicFoldersLockedForMigration:$true
```

For detailed syntax and parameter information, see Set-OrganizationConfig.

If your organization has multiple public folder databases, you'll need to wait until public folder replication is complete to confirm that all public folder databases have picked up the `PublicFoldersLockedForMigration` property value and any pending changes users recently made to folders have converged across the organization. This may take several hours.

Step 7: Finalize the public folder migration (downtime required)

First, run the following cmdlet to change the Exchange 2016 deployment type to **Remote**:

```
Set-OrganizationConfig -PublicFoldersEnabled Remote
```

Once that is done, you can complete the public folder migration by running the following command:

```
Complete-MigrationBatch PFMigration
```

Or, in EAC, you can complete the migration by clicking **Complete this migration batch**.

When you complete the migration, Exchange will perform a final synchronization between the Exchange 2010 server and Exchange 2016. If the final synchronization is successful, the public folders on the Exchange 2016 server will be unlocked and the status of the migration batch will change to **Completing**, and then **Completed**. It is common for the migration batch to take a few hours before its status changes from **Synced** to **Completing**, at which point the final synchronization will begin.

> **NOTE**
>
> If for any reason the migration batch file does not finalize (the *PublicFolderMigrationComplete* property value is `False`) restart the Information Store (IS) on the Exchange 2010 server.

Step 8: Test and unlock the public folder migration

After you finalize the public folder migration, you should run the following test to make sure that the migration was successful. This allows you to test the migrated public folder hierarchy before you switch to using Exchange 2016 public folders.

1. In PowerShell, run the following command to assign some test mailboxes to use any newly migrated public folder mailbox as the default public folder mailbox.

   ```
   Set-Mailbox -Identity <Test User> -DefaultPublicFolderMailbox <Public Folder Mailbox Identity>
   ```

2. Log on to Outlook 2007 or later with the test user identified in the previous step, and then perform the following public folder tests:

   - View the hierarchy.
   - Check permissions.
   - Create and delete public folders.
   - Post content to and delete content from a public folder.

3. If you run into any issues, see Roll back the migration later in this topic. If the public folder content and hierarchy is acceptable and functions as expected, run the following command to unlock the public folders for all other users.

   ```
   Get-Mailbox -PublicFolder | Set-Mailbox -PublicFolder -IsExcludedFromServingHierarchy $false
   ```

   > **IMPORTANT**
   >
   > Don't use the *IsExcludedFromServingHierarchy* parameter after initial migration validation is complete as this parameter is used by the automated storage management service for Exchange Online.

4. On the Exchange 2010 server, run the following command to indicate that the public folder migration is complete:

   ```
   Set-OrganizationConfig -PublicFolderMigrationComplete:$true
   ```

5. After you've verified that the migration is complete, on the Exchange 2016 server, run the following command:

   ```
   Set-OrganizationConfig -PublicFoldersEnabled Local
   ```

6. Finally, if you want external senders to send mail to the migrated mail-enabled public folders, the **Anonymous** user needs to be granted at least the **Create Items** permission. If you don't do this, external senders will receive a delivery failure notification and the messages won't be delivered to the migrated mail-enabled public folder.

You can use the Exchange Management Shell or Outlook to set the permissions on the Anonymous user. To read more about how to set permissions on the Anonymous user, see Mail-enable or mail-disable a public folder.

How do I know this worked?

In Step 2: Prepare for the migration, you were instructed to take snapshots of the public folder structure, statistics, and permissions before the migration began. The following steps will help verify that your public folder migration was successful by taking the same snapshots after the migration is complete. You can then compare the data in both files to verify success.

1. Run the following command to take a snapshot of the new folder structure.

   ```
   Get-PublicFolder -Recurse | Export-CliXML C:\PFMigration\Cloud_PFStructure.xml
   ```

2. Run the following command to take a snapshot of the public folder statistics such as item count, size, and owner.

   ```
   Get-PublicFolderStatistics -ResultSize Unlimited | Export-CliXML C:\PFMigration\Cloud_PFStatistics.xml
   ```

3. Run the following command to take a snapshot of the permissions.

   ```
   Get-PublicFolder -Recurse | Get-PublicFolderClientPermission | Select-Object Identity,User -
   ExpandProperty AccessRights | Export-CliXML  C:\PFMigration\Cloud_PFPerms.xml
   ```

Remove public folder databases from the Exchange 2010 servers

After the migration is complete, and you have verified that your Exchange 2016 public folders are working as expected, you should remove the public folder databases on the Exchange 2010 servers.

For details about how to remove public folder databases from Exchange 2010 servers, see Remove Public Folder Databases.

Roll back the migration

If you run into issues with the migration and need to reactivate your Exchange 2010 public folders, perform the following steps.

Caution

If you roll your migration back to the Exchange 2010 servers, you will lose any email that was sent to mail-enabled public folders or content that was posted to public folders in Exchange 2016 after the migration. To save this content, you need to export the public folder content to a .pst file and then import it to the Exchange 2010 public folders when the rollback is complete.

1. On the Exchange 2010 server, run the following command to unlock the migrated public folders. This process may take several hours.

   ```
   Set-OrganizationConfig -PublicFoldersLockedForMigration $false
   ```

2. On the Exchange 2016 server, run the following commands to remove the public folder mailboxes.

```
Get-Mailbox -PublicFolder | Where {$_.IsRootPublicFolderMailbox -eq $false} | Remove-Mailbox -
PublicFolder -Force -Permanent $true -Confirm:$false
```

```
Get-Mailbox -PublicFolder | Remove-Mailbox -PublicFolder -Force -Permanent $true -Confirm:$false
```

3. On the Exchange 2010 server, run the following command to set the `PublicFolderMigrationComplete` property value to `False`.

   ```
   Set-OrganizationConfig -PublicFolderMigrationComplete $false
   ```

Use batch migration to migrate Exchange Server public folders to Exchange Online

8/3/2020 • 32 minutes to read • Edit Online

Applies to: Exchange Server 2013, **Exchange Server 2016**, and **Exchange Server 2019**

Migrating your Exchange Server public folders to Exchange Online requires Exchange Server 2013 CU15 or later, or Exchange Server 2016 CU4 or later, to be running in your on-premises environment. All versions of Exchange Server 2019 are supported for batch migrations of public folders.

If you have a mixed environment of both Exchange 2013 and Exchange 2016/2019 public folders in your organization, and you want to move them all to Exchange Online, the instructions in this article will work for you, provided your Exchange 2013 servers have CU15 or later installed.

For instructions on migrating Exchange Server 2010 public folders to Exchange Online, see Use batch migration to migrate legacy public folders to Exchange Online.

What do you need to know before you begin?

- We strongly recommend you review FAQ: Public folders before you attempt a migration.

- When you upgrade to Exchange Server 2013 CU15 or later, or to Exchange Server 2016 CU4 or later, you must also prepare Active Directory or your public folder migration will fail. This Active Directory preparation ensures that all relevant PowerShell cmdlets and parameters are available to you for preparing for and running the migration. See Prepare Active Directory and domains for more information.

- In Exchange Online, you need to be a member of the Organization Management role group. This role group is different from the permissions assigned to you when you subscribe to Microsoft 365, Office 365, or Exchange Online. For details about how to enable the Organization Management role group, see Manage role groups.

- In Exchange Server, you need to be a member of the Organization Management or Server Management RBAC role groups. For details, see Add Members to a Role Group.

- Before you begin the public folder migration, if any single public folder in your organization is larger than 25 GB, we recommend that you delete content from that folder to make it smaller, or divide the public folder's content into multiple, smaller public folders. Note that the 25 GB limit cited here only applies to the public folder and not to any child or sub-folders the folder in question may have. If neither option is feasible, we recommend that you do not move your public folders to Exchange Online. See Exchange Online Limits for more information.

  NOTE

  If your current public folder quotas in Exchange Online are less than 25 GB, you can use the Set-OrganizationConfig cmdlet to increase them with the DefaultPublicFolderIssueWarningQuota and DefaultPublicFolderProhibitPostQuota parameters.

- In Microsoft 365, Office 365, and Exchange Online, you can create a maximum of 1000 public folder mailboxes. However, a maximum of 100 public folders is supported for migration from Exchange Server.

- If you intend to migrate users to Microsoft 365 or Office 365, you should complete your user migration prior to migrating your public folders. For more information, see Ways to migrate multiple email accounts to Microsoft 365 or Office 365.

- MRS Proxy needs to be enabled on at least one Exchange server, a server that is also hosting public folder mailboxes. See Enable the MRS Proxy endpoint for remote moves for details.

- To perform the migration procedures in this article, you can't use the Exchange admin center (EAC). Instead, you need to use the Exchange Management Shell on your Exchange servers. In Exchange Online, you need to use Exchange Online PowerShell. For more information, see Connect to Exchange Online PowerShell.

- To run the migration scripts in this article, you must use an account that has basic authentication enabled. Accounts that use multi-factor authentication (MFA) are currently not supported.

- Skipping the migration of deleted items and deleted folders from Exchange Server to Exchange Online is supported. For more information, see the Exchange Team blog post about modern public folder migrations without dumpster data.

- You must use a single migration batch to migrate all of your public folder data. Exchange allows creating only one migration batch at a time. If you attempt to create more than one migration batch simultaneously, the result will be an error. Also note that once the migration batch has a status of "Completed," no more data can be copied over from the source environment.

- We recommend that you don't use Outlook's PST export feature to migrate public folders to Microsoft 365, Office 365, or Exchange Online. Public folder mailbox growth in Exchange Online is managed using an auto-split feature that splits the public folder mailbox when it exceeds size quotas. Auto-split can't handle the sudden growth of public folder mailboxes when you use PST export to migrate your public folders, and you may have to wait for up to two weeks for auto-split to move the data from the primary mailbox. We recommend that instead you use the cmdlet-based instructions in this article to migrate your public folders. If you still decide to migrate public folders using PST export, see Migrate Public Folders to Office 365 by using Outlook PST export later in this article.

- Before you begin, please read this article in its entirety. For some steps there is downtime required. During this downtime, public folders will not be accessible by anyone. Please also review the list of known issues.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Step 1: Download the migration scripts

1. Download all scripts and supporting files from Exchange 2013/2016/2019 Public Folders Migration Scripts and Exchange 2010/2013/2016/EXO Public Folders to Microsoft 365 or Office 365 Pre-Migration Scripts.

2. Save the scripts to the local computer on which you'll be running PowerShell. For example, C:\PFScripts. Make sure all scripts are saved in the same location.

The scripts and files you're downloading are:

- `SourceSideValidations.ps1` : Source Side Validation script scans the public folders at source and reports issues found along with actions required to fix the issues. You'll run this script on the Exchange server on-premises.

- `Sync-ModernMailPublicFolders.ps1` This script synchronizes mail-enabled public folder objects between your Exchange on-premises environment and Microsoft 365 or Office 365. You'll run this script on an on-premises Exchange server.

- `SyncModernMailPublicFolders.strings.psd1` This support file is used by the Sync-ModernMailPublicFolders.ps1 script and should be downloaded to the same location.

- `Export-ModernPublicFolderStatistics.ps1` This script creates the folder name-to-folder size and deleted item size mapping file. You'll run this script on an on-premises Exchange server.

- `Export-ModernPublicFolderStatistics.strings.psd1` This support file is used by the Export-ModernPublicFolderStatistics.ps1 script and should be downloaded to the same location.

- `ModernPublicFolderToMailboxMapGenerator.ps1` This script creates the public folder-to-mailbox mapping file by using the output from the Export-ModernPublicFolderStatistics.ps1 script. You'll run this script on an on-premises Exchange server.

- `ModernPublicFolderToMailboxMapGenerator.strings.psd1` This support file is used by the ModernPublicFolderToMailboxMapGenerator.ps1 script and should be downloaded to the same location.

- `SetMailPublicFolderExternalAddress.ps1` This script updates the `ExternalEmailAddress` of mail-enabled public folders in your on-premises environment to that of their Exchange Online counterparts, so that emails addressed to your mail-enabled public folders post-migration are properly routed to Exchange Online. You need to run this script on an on-premises Exchange server.

- `SetMailPublicFolderExternalAddress.strings.psd1` This support file is used by the Create-PublicFolderMailboxesForMigration.ps1 script and should be downloaded to the same location.

Step 2: Prepare for the migration

> **NOTE**
>
> We strongly recommend running the Source Side Validation script from an on-premises Exchange Mailbox server. The script will scan and report issues that are known to cause migration to be slow, along with guidance to fix these issues. Use the examples as documented here. The script will perform all the following prerequisites.

Perform all prerequisite steps in the following sections before you begin the public folder migration.

General prerequisite steps

For your migration to be successful, you should:

- Make sure that there are no orphaned public folder mail objects in Active Directory. These are objects in Active Directory without a corresponding Exchange object.

- Confirm that the SMTP email addresses configured for public folders in Active Directory match the SMTP email addresses on the Exchange objects.

- Confirm that there are no duplicate public folder objects in Active Directory. This is necessary to avoid having two or more Active Directory objects that are pointing to the same mail-enabled public folder.

Prerequisite steps in the on-premises Exchange 2013, Exchange 2016, or Exchange 2019 server environment

In Exchange Management Shell (on-premises) perform the following steps:

1. Once your migration is complete, it will take some time for DNS caches across the Internet to direct messages to your mail-enabled public folders in their new location in Exchange Online. You can ensure that your newly migrated mail-enabled public folders receive messages during this DNS transition period by creating an accepted domain with a well-known name. To do this, run the following command in your Exchange on-premises environment. In this example, `target domain` is your Microsoft 365, Office 365, or Exchange Online domain, for which a send connector has already been configured by the Hybrid Configuration Wizard.

```
New-AcceptedDomain -Name PublicFolderDestination_78c0b207_5ad2_4fee_8cb9_f373175b3f99 -DomainName
<target domain> -DomainType InternalRelay
```

Example:

```
New-AcceptedDomain -Name PublicFolderDestination_78c0b207_5ad2_4fee_8cb9_f373175b3f99 -DomainName
"contoso.mail.onmicrosoft.com" -DomainType InternalRelay
```

If the accepted domain already exists in your on-premises environment, rename it to `PublicFolderDestination_78c0b207_5ad2_4fee_8cb9_f373175b3f99` and leave the other attributes intact.

To check if the accepted domain is already present in your on-premises environment, run the following:

```
Get-AcceptedDomain | Where {$_.DomainName -eq "<target domain>"}
```

To rename the accepted domain to `PublicFolderDestination_78c0b207_5ad2_4fee_8cb9_f373175b3f99`, run the following:

```
Get-AcceptedDomain | Where {$_.DomainName -eq "<target domain>"} | Set-AcceptedDomain -Name
PublicFolderDestination_78c0b207_5ad2_4fee_8cb9_f373175b3f99
```

NOTE

If you're expecting your mail-enabled public folders in Exchange Online to receive external emails from the Internet, you have to disable Directory Based Edge Blocking (DBEB) in Exchange Online and Exchange Online Protection (EOP). See Use Directory Based Edge Blocking to Reject Messages Sent to Invalid Recipients for more information.

2. If the name of a public folder contains a backslash **** or a forward slash **/**, it may not get migrated to its designated mailbox during the migration process. Before you migrate, rename any such folders to remove these characters.

   a. To locate public folders that have a backslash in the name, run the following command:

   ```
   Get-PublicFolder -Recurse -ResultSize Unlimited | Where {$_.Name -like "*\*" -or $_.Name -like "*/*"} |
   Format-List Name, Identity, EntryId
   ```

   b. If any public folders are returned, you can rename them by running the following command:

   ```
   Set-PublicFolder -Identity "<public folder EntryId>" -Name "<new public folder name>"
   ```

3. (This step is only required only if you are re-doing a previous migration attempt for some reason. If this is not the case, skip to the next step.) Run the following cmdlets to confirm there isn't a record of a previous, successful migration in your organization. If there is, you need to set that value to `$false`.

   Before changing the values, please confirm that the previous migration attempt can be discarded so that you don't accidentally perform a second migration.

   a. Run the following command to check for any previous migrations, and the status of those migrations:

```
Get-OrganizationConfig | Format-List  PublicFolderMailboxesLockedForNewConnections,
PublicFolderMailboxesMigrationComplete
```

b. If any of the above is returned with a value set to `$true`, make them `$false` by running:

```
Set-OrganizationConfig -PublicFolderMailboxesLockedForNewConnections:$false -
PublicFolderMailboxesMigrationComplete:$false
```

4. For the purpose of verifying the success of the migration upon its completion, we recommend that you run the following commands on all appropriate Exchange 2016 or Exchange 2019 servers. This will take snapshots of your current public folder deployment that you can later use to compare with your newly migrated public folders.

   > **NOTE**
   >
   > Depending on the size of your Exchange organization, it could take some time for these commands to run.

   - Run the following command to take a snapshot of the original source folder structure.

     ```
     Get-PublicFolder -Recurse -ResultSize Unlimited | Export-CliXML OnPrem_PFStructure.xml
     ```

   - Run the following command to take a snapshot of public folder statistics such as item count, size, and owner.

     ```
     Get-PublicFolderStatistics -ResultSize Unlimited | Export-CliXML OnPrem_PFStatistics.xml
     ```

   - Run the following command to take a snapshot of public folder permissions.

     ```
     Get-PublicFolder -Recurse -ResultSize Unlimited | Get-PublicFolderClientPermission | Select-
     Object Identity,User,AccessRights -ExpandProperty AccessRights | Export-CliXML
     OnPrem_PFPerms.xml
     ```

   - Run the following command to take a snapshot of your mail-enabled public folders:

     ```
     Get-MailPublicFolder -ResultSize Unlimited | Export-CliXML OnPrem_MEPF.xml
     ```

   - Save the files generated from the preceding commands in a safe place in order to make a comparison at the end of the migration.

5. If you're using Microsoft Azure Active Directory Connect (Azure AD Connect) to synchronize your on-premises directories with Azure Active Directory, you need to do the following (if you aren't using Azure AD Connect, you can skip this step):

   a. On an on-premises computer, open Microsoft Azure Active Directory Connect, and then select **Configure**.

   b. On the **Additional tasks** screen, select **Customize synchronization options**, and then click **Next**.

   c. On the **Connect to Azure AD** screen, enter the appropriate credentials, and then click **Next**. Once connected, keep clicking **Next** until you're on the **Optional Features** screen.

   d. Make sure that **Exchange Mail Public Folders** is not selected. If it isn't selected, you can continue to the next section, *Prerequisite steps in Exchange Online*. If it is selected, click to clear the check box, and then click **Next**.

> **NOTE**
>
> If you don't see **Exchange Mail Public Folders** as an option on the **Optional Features** screen, you can exit Microsoft Azure Active Directory Connect and proceed to the next section, *Prerequisite steps in Exchange Online*.

e. After you have cleared the **Exchange Mail Public Folders** selection, keep clicking **Next** until you're on the **Ready to configure** screen, and then click **Configure**.

Prerequisite steps in Exchange Online

In Exchange Online PowerShell, do the following steps:

1. Make sure there are no existing public folder migration requests. If there are, clear them or your own migration request will fail. This step is only required if you think there may be an existing migration request in the pipeline (one that has failed or that you wish to abort).

   The following example will discover any existing batch migration requests:

   ```
   Get-MigrationBatch | ?{$_.MigrationType.ToString() -eq "PublicFolder"}
   ```

   The following example removes any existing public folder batch migration requests:

   ```
   Remove-MigrationBatch <name of migration batch> -Confirm:$false
   ```

2. Make sure there aren't any existing public folders or public folder mailboxes in Exchange Online. If you do discover public folders in Exchange Online after following the steps below, it's important to determine why they are there and who in your organization started a public folder hierarchy before you begin removing any public folders and public folder mailboxes.

   a. In Exchange Online PowerShell, run the following command to see if any public folders mailboxes exist:

   ```
   Get-Mailbox -PublicFolder
   ```

   b. If the command doesn't return any public folder mailboxes, continue to Step 3: Generate the .csv files. If the command does return any public folders mailboxes, run the following command to see if any public folders exist:

   ```
   Get-PublicFolder -Recurse
   ```

3. If you do have any public folders in Microsoft 365 or Office 365 or Exchange Online, run the following PowerShell command to remove them (after confirming that they are not needed). Make sure that you've saved any information within these public folders before deleting them, because all information will be permanently deleted when you remove the public folders.

   ```
   Get-MailPublicFolder -ResultSize Unlimited | where {$_.EntryId -ne $null}| Disable-MailPublicFolder -
   Confirm:$false
   Get-PublicFolder -GetChildren \ -ResultSize Unlimited | Remove-PublicFolder -Recurse -Confirm:$false
   ```

4. After the public folders are removed, run the following commands to remove all public folder mailboxes:

```
$hierarchyMailboxGuid = $(Get-OrganizationConfig).RootPublicFolderMailbox.HierarchyMailboxGuid
Get-Mailbox -PublicFolder | Where-Object {$_.ExchangeGuid -ne $hierarchyMailboxGuid} | Remove-Mailbox -
PublicFolder -Confirm:$false -Force
Get-Mailbox -PublicFolder | Where-Object {$_.ExchangeGuid -eq $hierarchyMailboxGuid} | Remove-Mailbox -
PublicFolder -Confirm:$false -Force
Get-Mailbox -PublicFolder -SoftDeletedMailbox | % {Remove-Mailbox -PublicFolder $_.PrimarySmtpAddress -
PermanentlyDelete:$true -force}
```

Step 3: Generate the .csv files

Use the previously downloaded scripts to generate the .csv files that will be used in the migration.

1. From the Exchange Management Shell (on-premises), run the `Export-ModernPublicFolderStatistics.ps1` script to create the folder name-to-folder size mapping file. You must have local administrator permissions to run this script. The resulting file will contain three columns: **FolderName**, **FolderSize**, and **DeletedItemSize**. The values for the **FolderSize** and **DeletedItemSize** columns will be displayed in bytes. For example, **\PublicFolder01,10240, 100** means the public folder in the root of your hierarchy named PublicFolder01 is 10240 bytes (10 KB) in size and there are 100 bytes of recoverable items in it.

   ```
   .\Export-ModernPublicFolderStatistics.ps1 <Folder-to-size map path>
   ```

   Example:

   ```
   .\Export-ModernPublicFolderStatistics.ps1 stats.csv
   ```

2. Run the `ModernPublicFolderToMailboxMapGenerator.ps1` script to create a .csv file that maps source public folders to public folder mailboxes in your Exchange Online destination. This file is used to calculate the correct number of public folder mailboxes in Exchange Online.

Note that the file generated by `ModernPublicFolderToMailboxMapGenerator.ps1` will not contain the name of every public folder in your organization. It will contain references to the parent folders of larger folder trees, or the names of folders which themselves are significantly large. You can think of this file as an "exception" file used to make sure certain folder trees and larger folders get placed into specific public folder mailboxes. It is normal to not see every one of your public folders in this file. Child folders of any folder listed in this mapping file will also be migrated to the same public folder mailbox as their parent folder (unless explicitly mentioned on another line within the mapping file that directs them to a different public folder mailbox).

```
.\ModernPublicFolderToMailboxMapGenerator.ps1 <Maximum mailbox size in bytes><Maximum mailbox recoverable item
size in bytes><Folder-to-size map path><Folder-to-mailbox map path>
```

- `Maximum mailbox size in bytes` is the maximum amount of data you want to migrate into any single public folder mailbox in Exchange Online. The maximum size of this field is currently 50 GB, but we recommend you use a smaller size, such as 50% of maximum size, to allow for future growth.

- `Maximum mailbox recoverable items size in bytes` is the recoverable items quota on your Exchange Online mailboxes. The maximum size of public folder mailboxes In Exchange Online is currently 50 GB. We recommend setting *RecoverableItemsQuota* to 15 GB or less.

- `Folder-to-size map path` is the file path of the .csv file you created when you ran the `Export-ModernPublicFolderStatistics.ps1` script.

- `Folder-to-mailbox map path` is the file path of the folder-to-mailbox .csv file that you're creating in this step. If you only specify a file name, the file will be generated in the current PowerShell directory on the local computer.

Example:

```
.\ModernPublicFolderToMailboxMapGenerator.ps1 -MailboxSize 25GB -MailboxRecoverableItemSize 1GB -ImportFile
.\stats.csv -ExportFile map.csv
```

> **NOTE**
>
> The map.csv generated by the script uses generic names for the target public folder mailboxes that will be created in EXO during the next step (for example, Mailbox1 and Mailbox2). We encourage you to change the public folder mailbox names in the map.csv to suit your organization's naming policies. Also, if your on-premises organization already has mailboxes that match the generic names, you should edit the map.csv and provide unique names for the target public folder mailboxes in Exchange Online. Use Notepad or a similar editor to edit the TargetMailbox names in the map.csv

> **NOTE**
>
> We don't support the migration of public folders to Exchange Online when there are more than 100 unique public folder mailboxes in Exchange Online. During migration, you can have up to 100 public folder mailboxes enabled.

Step 4: Create the public folder mailboxes in Exchange Online

Next, in Exchange Online PowerShell, create the target public folder mailboxes that will contain your migrated public folders.

Run the following script to create the target public folder mailboxes. The script will create a target mailbox for each mailbox in the .csv file that you generated previously in *Step 3: Generate the .csv files*, when you ran the `ModernPublicFoldertoMailboxMapGenerator.ps1` script.

```
$mappings = Import-Csv <Folder-to-mailbox map path>
$primaryMailboxName = ($mappings | Where-Object FolderPath -eq "\" ).TargetMailbox;
New-Mailbox -HoldForMigration:$true -PublicFolder -IsExcludedFromServingHierarchy:$false $primaryMailboxName
($mappings | Where-Object TargetMailbox -ne $primaryMailboxName).TargetMailbox | Sort-Object -unique |
ForEach-Object { New-Mailbox -PublicFolder -IsExcludedFromServingHierarchy:$false $_ }
```

- `Folder-to-mailbox map path` is the file path of the folder-to-mailbox.csv file that was generated by the `ModernPublicFoldertoMailboxMapGenerator.ps1` script in *Step 3: Generate the .csv files*.

Step 5: Start the migration request

A number of commands now need to be run both in your Exchange Server on-premises environment and in Exchange Online.

1. From any of your Exchange 2016 or Exchange 2019 servers hosting public folder mailboxes, execute the following script. This script will synchronize mail-enabled public folders from your local Active Directory to Exchange Online. Make sure that you have downloaded the latest version of this script and that you're running it from Exchange Management Shell.

   ```
   .\Sync-ModernMailPublicFolders.ps1 -Credential (Get-Credential) -CsvSummaryFile:sync_summary.csv
   ```

   - `Credential` is your Exchange Online administrative username and password.
   - `CsvSummaryFile` is the file path to where you want your log file of synchronization operations and

errors located. The log will be in .csv format.

> **NOTE**
>
> Use Sync MEPF Script troubleshooting if you see any errors during the `Sync-ModernMailPublicFolders.ps1` script.

2. In Exchange Online PowerShell, pass the credential of a user who has administrator permissions in the Exchange 2013, Exchange 2016, or Exchange 2019 on-premises environment into the variable `$Source_Credential`. The migration request that you run in Exchange Online will use this credential to gain access to your on-premises Exchange servers to copy the public folder content over to Exchange Online.

   ```
   $Source_Credential = Get-Credential <source_domain>\<PublicFolder_Administrator_Account>
   ```

3. In Exchange Online Powershell, pass the Internet routable fully qualified domain name of your Exchange Mailbox Replication Service (MRS) into the variable `$Source_RemoteServer`. The migration request that you run in Exchange Online will use this remote server to copy the public folder content to Exchange Online.

   ```
   $Source_RemoteServer = "<MRS proxy endpoint server>"
   ```

4. On your on-premises Exchange server, open the Exchange Management Shell and find the GUID of the primary hierarchy mailbox with the following command:

   ```
   (Get-OrganizationConfig).RootPublicFolderMailbox.HierarchyMailboxGuid.GUID
   ```

   Note the output of this command. You will need it in the next step. For example:

   > 91edc6dd-478a-497c-8731-b0b793f5a986

5. In Exchange Online PowerShell, run the following commands to create the public folder migration endpoint and the public folder migration request:

   ```
   [byte[]]$bytes = Get-Content -Encoding Byte <folder_mapping.csv>
   $PfEndpoint = New-MigrationEndpoint -PublicFolder -Name PublicFolderEndpoint -RemoteServer
   $Source_RemoteServer -Credentials $Source_Credential
   New-MigrationBatch -Name PublicFolderMigration -CSVData $bytes -SourceEndpoint $PfEndpoint.Identity -
   SourcePfPrimaryMailboxGuid <guid you noted from previous step> -AutoStart -NotificationEmails <email
   addresses for migration notifications>
   ```

   Where `folder_mapping.csv` is the map file that was generated in *Step 3: Generate the .csv files* and `HierarchyMailboxGUID` is the output you noted in the previous step. Be sure to provide the full file path to `folder_mapping.csv`. If the map file was moved for any reason, be sure to use the new location.

   Separate multiple email addresses with commas.

NOTE

If you receive an error that states, `A parameter cannot be found that matches parameter name 'SourcePfPrimaryMailboxGuid'`, then instead use the following command to create the migration batch:

```
New-MigrationBatch -Name PublicFolderMigration -CSVData $bytes -SourceEndpoint $PfEndpoint.Identity
-AutoStart -NotificationEmails <email addresses for migration notifications>
```

6. Finally, start the migration using the following command in Exchange Online PowerShell:

```
Start-MigrationBatch PublicFolderMigration
```

While batch migrations need to be created using the `New-MigrationBatch` cmdlet in Exchange Online PowerShell, the progress and completion of the migration can be viewed and managed in the EAC or by running the Get-MigrationBatch cmdlet. The `New-MigrationBatch` cmdlet initiates a mailbox migration request for each public folder mailbox, and you can view the status of these requests using the mailbox migration page.

To go to the mailbox migration page:

1. Log on to Exchange Online and open the EAC.
2. Navigate to **Recipients**, and then select **Migration**.
3. Select the migration request that was just created and then, on the **Details** pane, select **View Details**.

Before moving on to *Step 6: Lock down the public folders on the Exchange on-premises server*, verify that all data has been copied and that there are no errors in the migration. Once you have confirmed that the batch has moved to the state of **Synced**, run the commands mentioned in *Step 2: Prepare for the migration*, in the final step under **Prerequisite steps in the Exchange Server on-premises environment**, to take a snapshot of the public folders on-premises.

Once these commands have run, you can proceed to the next step. Note that these commands could take a while to complete depending on the number of folders you have. The migration process will synchronize the data from the source (on-premises) environment once every 24 hours.

You can use the following cmdlets to monitor your migration:

- Get-PublicFolderMailboxMigrationRequest
- Get-PublicFolderMailboxMigrationRequestStatistics
- Get-MigrationBatch

Step 6: Lock down the public folders on the Exchange on-premises server (public folder downtime required)

Until this point in the migration process, users have been able to access your on-premises public folders. The following steps will now log off users off from Exchange Server public folders and then lock the folders as the migration process completes its final synchronization. Users won't be able to access public folders during this time, and any messages sent to these mail-enabled public folders will be queued and remain undelivered until the public folder migration is complete.

NOTE

The final sync might take a substantial amount of time, depending on the changes made to the source environment, the size of the public folder deployment, server capacity, and so on. If the folder hierarchy had many corrupt ACLs that were not cleaned up before the migration, there might be a significant delay in completion. It is recommended that you plan for a minimum of 48 hours of downtime for the final sync to complete.

Ensure the migration batch and individual migration requests have successfully synced.

Run the following command in EXO PowerShell for more information:

```
Get-MigrationBatch |?{$_.MigrationType -like "*PublicFolder*"} | ft *last*sync*
```

```
Get-PublicFolderMailboxMigrationRequest | Get-PublicFolderMailboxMigrationRequestStatistics |ft
targetmailbox,*last*sync*
```

The LastSyncedDate (on migration batch) and LastSuccessfulSyncTimestamp (on individual jobs) should be within the last 7 days. If the date is too far in the past, such as more than a month ago, you might want to review public folder migration requests and ensure that all the requests were synced recently.

After you have confirmed that the batch and all migration requests have successfully synced, in your on-premises environment, run the following command to lock the Exchange Server public folders for finalization.

```
Set-OrganizationConfig -PublicFolderMailboxesLockedForNewConnections $true
```

NOTE

If you aren't able to access the `-PublicFolderMailboxesLockedForNewConnections` parameter, it could be because your Active Directory was not prepared during the CU upgrade, as we advised above in *What do you need to know before you begin?* See Prepare Active Directory and domains for more information. Also note that any users who need access to public folders should be migrated first, **before** you migrate the public folders themselves.

If your organization has public folder mailboxes on multiple Exchange servers, you'll need to wait until Active Directory replication is complete. Once complete, you can confirm that all public folder mailboxes have picked up the `PublicFolderMailboxesLockedForNewConnections` flag, and that any pending changes users recently made to their public folders have converged across the organization. All of this could take several hours.

Run the following command in your on-premises environment to ensure that public folders are locked:

```
Get-PublicFolder \
```

The expected result if public folders are locked is:

```
Couldn't find the public folder mailbox. + CategoryInfo : NotSpecified: (:) [Get-PublicFolder],
ObjectNotFoundException
```

Step 7: Finalize the public folder migration (public folder downtime required)

You need to check the following items before you can complete your public folder migration:

1. Confirm that there are no other public folder mailbox moves or public folder moves going on in your on-premises Exchange environment. To do this, use the **Get-MoveRequest** and **Get-PublicFolderMoveRequest** cmdlets to list any existing public folder moves. If there are any moves in progress, or in the **Completed** state, remove them.

2. At this point, we recommend re-running the following script to ensure that any new mail-enabled public folders are synchronized with Exchange Online:

   ```
   .\Sync-ModernMailPublicFolders.ps1 -Credential (Get-Credential) -CsvSummaryFile:sync_summary.csv
   ```

3. To complete the public folder migration, run the following command in Exchange Online PowerShell:

   ```
   Complete-MigrationBatch PublicFolderMigration
   ```

> **IMPORTANT**
>
> After a migration batch is completed, no additional data can be synchornized from the on-premises Exchange servers and Exchange Online.

When you run `Complete-MigrationBatch PublicFolderMigration`, Exchange will perform a final synchronization between your Exchange on-premises organization and Exchange Online. During this period, the status of the migration batch will change from **Synced** to **Completing**, and then finally to **Completed**. If the final synchronization is successful, the public folders in Exchange Online will be unlocked. However, it is strongly recommended that you complete Step 8 and Step 9 of this article before you open up public folders to your users.

It's common for the status of migration batch to remain on **Synced** for a few hours before it switches to **Completing**. For migrations involving a large number of target mailboxes, it's normal to see the status remain in the **Synced** state for more than 24 hours, provided none of the underlying public folder migration requests have failed or were qurantined.

Step 8: Test and unlock public folders in Exchange Online

Once the public folder migration is complete, take the following steps to test the success of the migration, and to officially verify its completion. These final tasks allow you to test the migrated public folder hierarchy before you permanently switch your organization to Exchange Online public folders.

1. In Exchange Online PowerShell, configure some test user mailboxes to use one of your newly migrated public folder mailboxes as their default public folder mailbox:

   ```
   Set-Mailbox -Identity <test user> -DefaultPublicFolderMailbox <public folder mailbox identity>
   ```

   Make sure that your test users have necessary permissions to create public folders.

2. Log on to Outlook with the test user you designated in the previous step, and then perform the following public folder tests. Note that it may take 15 to 30 minutes for changes to take effect. Once Outlook is aware of the changes, it might prompt you to restart a couple of times.

   a. View the hierarchy.

   b. Check permissions.

   c. Create some public folders and then delete them.

   d. Post content to, and delete content from, a public folder.

   If you run into any issues and determine you aren't ready to switch your organization's public folders entirely to Exchange Online, see Roll back a public folder migration from Exchange Server to Exchange Online.

3. Run the following command in Exchange Online PowerShell to unlock your public folders in Exchange Online. After you run the command, it may take approximately 15 to 30 minutes for the changes to take effect. Once Outlook is aware of the changes, it might prompt your users to restart Outlook a couple of times.

   ```
   Set-OrganizationConfig -RemotePublicFolderMailboxes $Null -PublicFoldersEnabled Local
   ```

Step 9: Finalize the migration on-premises

To enable emails to mail-enabled public folders on-premises, perform the following steps:

1. Run the following command in your on-premises environment, to take a backup of the emails in the queue that were sent to your mail-enabled public folders. This backup can be used in scenarios where email delivery to mail-enabled public folders failed for any reason:

   ```
   $Server=Get-TransportService;ForEach ($t in $server) {Get-Message -Server $t -ResultSize Unlimited| ?
   {$_.Recipients -like "*PF.InTransit*"} | ForEach-Object {Suspend-Message $_.Identity -Confirm:$False;
   $Temp="C:\ExportFolder\"+$_.InternetMessageID+".eml"; $Temp=$Temp.Replace("<","_");
   $Temp=$Temp.Replace(">","_"); Export-Message $_.Identity | AssembleMessage -Path $Temp;Resume-message
   $_.Identity -Confirm:$false}}
   ```

2. In your on-premises environment, run the following script to make sure all emails to mail-enabled public folders are correctly routed to Exchange Online. The script will stamp mail-enabled public folders with an `ExternalEmailAddress` that points them to their Exchange Online counterparts:

   ```
   .\SetMailPublicFolderExternalAddress.ps1 -ExecutionSummaryFile:mepf_summary.csv
   ```

3. If your testing is successful, in your on-premises environment, run the following command to indicate that the public folder migration is complete:

   ```
   Set-OrganizationConfig -PublicFolderMailboxesMigrationComplete:$true -PublicFoldersEnabled Remote
   ```

How do I know this worked?

In Step 2: Prepare for the migration, you took snapshots of your on-premises public folder structure, statistics, and permissions. The following steps will help you verify your public folder migration was successful by taking the same snapshots in Exchange Online post-migration. Compare the data in both files to verify success.

1. In Exchange Online PowerShell, run the following command to take a snapshot of the new folder structure:

   ```
   Get-PublicFolder -Recurse -ResultSize Unlimited | Export-CliXML Cloud_PFStructure.xml
   ```

2. In Exchange Online PowerShell, run the following command to take a snapshot of the public folder statistics, including item count, size, and owner:

   ```
   Get-PublicFolder -Recurse -ResultSize Unlimited | Get-PublicFolderStatistics | Export-CliXML
   Cloud_PFStatistics.xml
   ```

3. In Exchange Online PowerShell, run the following command to take a snapshot of the permissions:

```
Get-PublicFolder -Recurse -ResultSize Unlimited | Get-PublicFolderClientPermission | Select-Object
Identity,User,AccessRights | Export-CliXML Cloud_PFPerms.xml
```

4. Exchange Online PowerShell, run the following command to take a snapshot of the mail-enabled public folders:

   ```
   Get-MailPublicFolder -ResultSize Unlimited | Export-CliXML Cloud_MEPF.xml
   ```

Known issues

The following are common public folder migration issues that you may encounter in your organization.

- We don't support the migration of public folders to Exchange Online when there are more than 100 unique public folder mailboxes in Exchange Online.

- Permissions for the root public folder and the EFORMS REGISTRY folder will not be migrated to Exchange Online, and you will have to manually apply them in Exchange Online. To do this, run the following command in your Exchange Online PowerShell. Run the command once for each permission entry that is present on-premises but missing in Exchange Online:

  ```
  Add-PublicFolderClientPermission "\" -User <user> -AccessRights <access rights>
  Add-PublicFolderClientPermission "\NON_IPM_SUBTREE\EFORMS REGISTRY" -User <user> -AccessRights <access
  rights>
  ```

- There is a known issue where some public folder migrations will fail if some public folder mailboxes are not serving the public folder hierarchy. This means the `IsExcludedFromServingHierarchy` parameter on one or more mailboxes is set to `$true`. To avoid this, set all mailboxes in Exchange Online to serve the hierarchy.

- **Send As** and **Send on Behalf** permissions don't get migrated to Exchange Online. If this happens with your migration, use the following commands in your on-premises environment to note who has these permissions.

  To see which public folders have Send As permissions on-premises:

  ```
  Get-MailPublicFolder | Get-ADPermission | ?{$_.ExtendedRights -like "*Send-As*"}
  ```

  To see which public folders have Send on Behalf permissions on-premises:

  ```
  Get-MailPublicFolder | ?{$_.GrantSendOnBehalfTo -ne "$null"} | Format-Table name,GrantSendOnBehalfTo
  ```

  To add Send As permission to a mail-enabled public folder in Exchange Online, in Exchange Online PowerShell type:

  ```
  Add-RecipientPermission -Identity <mail-enabled public folder primary SMTP address> -Trustee <name of
  user to be assigned permission> -AccessRights SendAs
  ```

  Example:

  ```
  Add-RecipientPermission -Identity send1 -Trustee Exo1 -AccessRights SendAs
  ```

  To add Send on Behalf permission to a mail-enabled public folder in Exchange Online, in Exchange Online

  PowerShell type:

  ```
  Set-MailPublicFolder -Identity <name of public folder> -GrantSendOnBehalfTo <user or comma-separated
  list of users>
  ```

  Example:

  ```
  Set-MailPublicFolder send2 -GrantSendOnBehalfTo exo1,exo2
  ```

- Having more than 10,000 folders under the "\NON_IPM_SUBTREE\DUMPSTER_ROOT" folder can cause the migration to fail. Therefore, check the "\NON_IPM_SUBTREE\DUMPSTER_ROOT" folder to see if there are more than 10,000 folders directly under it (immediate children). You can use the following command to find the number of public folders in this location:

  ```
  (Get-PublicFolder -GetChildren "\NON_IPM_SUBTREE\DUMPSTER_ROOT").Count
  ```

  Exchange Online does not support more than 10,000 subfolders, which is why migrations of more than 10,000 folders will fail. We are currently developing a script to unblock such configurations. In the meantime, we suggest waiting to migrate your public folders.

- Migration jobs are not making progress or are stalled. This can happen if there are too many jobs running in parallel, causing jobs to fail with intermittent errors. You can reduce the number of concurrent jobs by modifying `MaxConcurrentMigrations` and `MaxConcurrentIncrementalSyncs` to a smaller number. Use the following example to set these values:

  ```
  Set-MigrationEndpoint <PublicFolderEndpoint> -MaxConcurrentMigrations 30 -MaxConcurrentIncrementalSyncs
  20 -SkipVerification
  ```

- Migration jobs fail with the error "Error: Dumpster of the Dumpster folder." If you see this error, it should be resolved if you stop the batch and then restart it.

- Migration jobs fail with the error "Request was quarantined because of the following error: The given key was not present in the dictionary." This happens when a corrupt item is present in a folder which migration jobs cannot copy. To work around this:

  1. Stop the migration batch.

  2. Identify the folder containing the bad item. The migration report should include references to the folder that was being copied when the error occurred.

  3. In your on-premises environment, move the affected folder to the primary public folder mailbox. You can use the `New-PublicFolderMoveRequest` cmdlet to move folders.

  4. Wait for the folder move to complete. After it is complete, remove the move request. Finally, re-start the migration batch.

Remove public folder mailboxes from your Exchange on-premises environment

After the migration is complete and you have verified that your public folders in Exchange Online are working as expected and contain all expected data, you can remove your on-premises public folder mailboxes.

Be aware that this step is irreversible, because once public folder mailboxes are deleted, they cannot be recovered. Therefore we strongly recommend that, in addition to validating the success of your migration, that you also monitor your Exchange Online public folders for a few weeks before removing the on-premises public folder mailboxes.

Migrate Public Folders to Microsoft 365 or Office 365 by using Outlook PST export

We recommend that you don't use Outlook's PST export feature to migrate public folders to Microsoft 365 or Office 365 or Exchange Online if your on-premises public folder hierarchy is greater than 30 GB. Microsoft 365 and Office 365 online public folder mailbox growth is managed using an auto-split feature that splits the public folder mailbox when it exceeds size quotas. Auto-split can't handle the sudden growth of public folder mailboxes when you use PST export to migrate your public folders and you may have to wait for up to two weeks for auto-split to move the data from the primary mailbox. In addition, consider the following before using Outlook PST to export public folders to Microsoft 365 or Office 365 or Exchange Online:

- Public folder permissions will be lost during this process. Capture the current permissions before migration and manually add them back once the migration is completed.
- If you use complex permissions or have many folders to migrate, we recommend that you use the cmdlet method for migration.
- Any item and folder changes made to the source public folders during the PST export migration will be lost. Therefore, we recommend that you use the cmdlet method if this export and import process will take a long time to complete.

If you still want to migrate your public folders by using PST files, follow these steps to ensure a successful migration.

1. Use the instructions in Step 1: Download the migration scripts to download the migration scripts. You only need to download the `PublicFolderToMailboxMapGenerator.ps1` file.
2. Follow step number 2 of Step 3: Generate the .csv files to create the public folder-to-mailbox mapping file. This file is used to calculate the correct number of public folder mailboxes in Exchange Online.
3. Create the public folder mailboxes that you'll need based on the mapping file. For more information, see Use the EAC to create a public folder mailbox.
4. Use the **New-PublicFolder** cmdlet to create the top-most public folder in each of the public folder mailboxes by using the *Mailbox* parameter.
5. Export and import the PST files using Outlook.
6. Set the permissions on the public folders using the EAC. For more information, follow Step 3: Assign permissions to the public folder in the Set up public folders in a new organization article.

Caution

If you've already started a PST migration and have run into an issue where the primary mailbox is full, you have two options for recovering the PST migration. The first option is to wait for the auto-split to move the data from the primary mailbox. This may take up to two weeks. However, all the public folders in a completely filled public folder mailbox won't be able to receive new content until the auto-split completes. Option two is to create a public folder mailbox in Exchange Server and then use the **[New-PublicFolder]** cmdlet with the *Mailbox* parameter to create the remaining public folders in the secondary public folder mailbox.

Roll back a public folder migration from Exchange Server to Exchange Online

8/3/2020 • 2 minutes to read • Edit Online

If you run into issues with your public folder migration to Exchange Online, or for any other reason need to reactivate your Exchange Server public folders, follow the steps below.

Roll back the migration

Note that if you roll back your migration, you will lose any content that was added to public folders in Exchange Online post-migration, either through clients or via email for mail-enabled public folders. To save this content, you can export the post-migration public folder content to a .pst file, which can then be imported into the on-premises public folders when the rollback is complete.

1. In your Exchange on-premises environment, run the following command to unlock your Exchange Server public folders (note that the unlocking may take several hours):

   ```
   Set-OrganizationConfig -PublicFolderMailboxesLockedForNewConnections:$false -PublicFolderMailboxesMigrationComplete:$false -PublicFoldersEnabled Local
   ```

2. In your Exchange on-premises environment, revert the `ExternalEmailAddress` of any mail-enabled public folder that was updated by SetMailPublicFolderExternalAddress.ps1 (the script used in *Step 8: Test and unlock public folders in Exchange Online* of Use batch migration to migrate Exchange Server public folders to Exchange Online. You can refer to the summary file created by the script to identify the ones that were modified, or use the file OnPrem_MEPF.xml file generated earlier in the same batch migration process to get the original properties for all mail-enabled public folders.

3. In Exchange Online PowerShell, run the following commands to remove all Exchange Online public folders and mailboxes:

   ```
   Get-MailPublicFolder -ResultSize Unlimited | where {$_.EntryId -ne $null}| Disable-MailPublicFolder -Confirm:$false
   Get-PublicFolder -GetChildren \ -ResultSize Unlimited | Remove-PublicFolder -Recurse -Confirm:$false
   $hierarchyMailboxGuid = $(Get-OrganizationConfig).RootPublicFolderMailbox.HierarchyMailboxGuid
   Get-Mailbox -PublicFolder | Where-Object {$_.ExchangeGuid -ne $hierarchyMailboxGuid} | Remove-Mailbox -PublicFolder -Confirm:$false -Force
   Get-Mailbox -PublicFolder | Where-Object {$_.ExchangeGuid -eq $hierarchyMailboxGuid} | Remove-Mailbox -PublicFolder -Confirm:$false -Force
   Get-Mailbox -PublicFolder -SoftDeletedMailbox | Remove-Mailbox -PublicFolder -PermanentlyDelete:$true
   ```

4. Run the following command in your Exchange Online environment to redirect public folder traffic back to on-premises (Exchange Server):

   ```
   Set-OrganizationConfig -PublicFoldersEnabled Remote
   ```

5. See Configure Exchange 2013 public folders for a hybrid deployment for instructions on reconfiguring access to your on-premises public folders, so that your Exchange Online users can access them.

Use batch migration to migrate Exchange Server public folders to Microsoft 365 Groups

8/3/2020 • 19 minutes to read • Edit Online

Through a process known as *batch migration*, you can move some or all of your Exchange Server public folders to Microsoft 365 Groups. Groups is a new collaboration offering from Microsoft that offers certain advantages over public folders. See Migrate your public folders to Microsoft 365 Groups for an overview of the differences between public folders and Groups, and reasons why your organization may or may not benefit from switching to Groups.

This article contains the step-by-step procedures for performing the actual batch migration of your Exchange Server public folders.

What do you need to know before you begin?

Ensure that all of the following conditions are met before you begin preparing your migration.

- The Exchange server needs to be running **Exchange 2016 CU4** or later.
- In Exchange Online, you need to be a member of the Organization Management role group. This role group is different from the permissions assigned to you when you subscribe to Microsoft 365, Office 365, or Exchange Online. For details about how to enable the Organization Management role group, see Manage role groups.
- In Exchange Server, you need to be a member of the Organization Management or Server Management RBAC role groups. For details, see Add Members to a Role Group.
- Before you migrate your public folders to Microsoft 365 Groups, we recommend that you first move user mailboxes to Microsoft 365 or Office 365 for those users who need access to Microsoft 365 Groups after migration. For more information, see Ways to migrate multiple email accounts to Microsoft 365 or Office 365.
- MRS Proxy needs to be enabled on at least one Exchange server, and that server must also be hosting public folder mailboxes. See Enable the MRS Proxy endpoint for remote moves for details.
- You can't use the Exchange admin center (EAC) or the Exchange Management Console (EMC) to perform this procedure. On the Exchange 2016 or Exchange 2019 servers, you need to use the Exchange Management Shell. For Exchange Online, you need to use Exchange Online PowerShell. For more information, see Connect to Exchange Online using remote PowerShell.
- Only public folders of type calendar and mail can be migrated to Microsoft 365 Groups at this time; migration of other types of public folders is not supported. Also, the target groups in Microsoft 365 or Office 365 are expected to be created prior to the migration.
- The batch migration process only copies messages and calendar items from public folders for migration to Microsoft 365 Groups. It doesn't copy other entities of public folders like policies, rules, and permissions.
- Microsoft 365 Groups comes with a 50GB mailbox. Ensure that the sum of public folder data that you're migrating totals less than 50GB. In addition, leave storage space for additional content to be added by your users in the future, post-migration. We recommend migrating public folders no bigger than 25GB in total size.
- This is not an "all or nothing" migration. You can pick and choose specific public folders to migrate, and only those public folders will be migrated. If the public folder being migrated has sub-folders, those sub-folders

will not be automatically included in the migration. If you need to migrate them, you need to explicitly include them. The migration batch allows for a mapping of a maximum two sub-folders to a single Microsoft 365 or Office 365 Group mailbox.

- The public folders will not be affected in any manner by this migration. However, once you use our lock-down script to make the migrated public folders read-only, your users will be forced to use Microsoft 365 groups instead of public folders.
- You must use a single migration batch to migrate all of your public folder data. Exchange allows creating only one migration batch at a time. If you attempt to create more than one migration batch simultaneously, the result will be an error.
- Before you begin, we recommend that you read this article in its entirety, as downtime is required for some steps.

Step 1: Get the scripts

The batch migration to Microsoft 365 groups requires running a number of scripts at different points in the migration, as described below in this article. Download the scripts and their supporting files from this location. After all the scripts and files are downloaded, save them to the same location, such as `c:\PFtoGroups\Scripts`.

Before proceeding, verify you have downloaded and saved all of the following scripts and files:

> **NOTE**
>
> Make sure to save all scripts and files to the same location.

- **AddMembersToGroups.ps1**. This script adds members and owners to Microsoft 365 groups based on permission entries in the source public folders.
- **AddMembersToGroups.strings.psd1**. This support file is used by the script `AddMembersToGroups.ps1`.
- **LockAndSavePublicFolderProperties.ps1**. This script makes public folders read-only to prevent any modifications, and it transfers the mail-related public folder properties (provided the public folders are mail-enabled) to the target groups, which will re-route emails from the public folders to the target groups. This script also backs up the permission entries and the mail properties before modifying them.
- **LockAndSavePublicFolderProperties.strings.psd1**: This support file is used by the script `LockAndSavePublicFolderProperties.ps1`.
- **UnlockAndRestorePublicFolderProperties.ps1**. This script restores access rights and mail properties of the public folders using backup files created by `LockandSavePublicFolderProperties.ps1`.
- **UnlockAndRestorePublicFolderProperties.strings.psd1**. This support file is used by the script `UnlockAndRestorePublicFolderProperties.ps1`.
- **WriteLog.ps1**. This script enables the preceding three scripts to write logs.
- **RetryScriptBlock.ps1**. This script enables the `AddMembersToGroups`, `LockAndSavePublicFolderProperties`, and `UnlockAndRestorePublicFolderProperties` scripts to retry certain actions in the event of transient errors.

For details about `AddMembersToGroups.ps1`, `LockAndSavePublicFolderProperties.ps1`, and `UnlockAndRestorePublicFolderProperties.ps1`, and the tasks they execute in your environment, see Migration scripts later in this article.

Step 2: Prepare for the migration

The following steps are necessary to prepare your organization for the migration:

1. Compile a list of public folders (mail and calendar types) that you want to migrate to Microsoft 365 Groups.

2. Have a list of corresponding target groups for each public folder being migrated. You can either create a new group in Microsoft 365 or Office 365 for each public folder or use an existing group. If you're creating a new group, see Learn about Microsoft 365 Groups to understand the settings a groublic folder that you're migrating has the default permission set to **Author** or above, you should create the corresponding group in Microsoft 365 or Office 365 with the **Public** privacy setting. However, for users to see the public group under the **Groups** node in Outlook, they will still have to join the group.

3. Rename any public folders that contain a backslash (****) in their name. Otherwise, those public folders may not get migrated correctly.

4. You need to have the migration feature **PAW** enabled for your Microsoft 365 or Office 365 organization. To verify this, run the following command in Exchange Online PowerShell:

   ```
   Get-MigrationConfig
   ```

   If the output under **Features** lists **PAW**, then the feature is enabled and you can continue to *Step 3: Create the .csv file*.

   If PAW is not yet enabled for your tenant, it could be because you have some existing migration batches, either public folder batches or user batches. These batches could be in any state, including Completed. If this is the case, please complete and remove any existing migration batches until no records are returned when you run `Get-MigrationBatch`. Once all existing batches are removed, PAW should get enabled automatically. Note that the change may not reflect in `Get-MigrationConfig` immediately, which is okay. Once this step is completed, you can continue creating new batches of user migrations.

Step 3: Create the .csv file

Create a .csv file, which will provide input for one of the migration scripts.

The .csv file needs to contain the following columns:

- **FolderPath**. Path of the public folder to be migrated.

- **TargetGroupMailbox**. SMTP address of the target group in Microsoft 365 or Office 365. You can run the following command to see the primary SMTP address.

  ```
  Get-UnifiedGroup <alias of the group> | Format-Table PrimarySmtpAddress
  ```

An example .csv:

```
"FolderPath","TargetGroupMailbox"
"\Sales","sales@contoso.onmicrosoft.com"
"\Sales\APAC","apacsales@contoso.onmicrosoft.com"
"\Sales\EMEA","emeasales@contoso.onmicrosoft.com"
```

Note that a mail folder and a calendar folder can be merged into a single group in Microsoft 365 or Office 365. However, any other scenario of multiple public folders merging into one group isn't supported within a single migration batch. If you do need to map multiple public folders to the same Microsoft 365 or Office 365 group, you can accomplish this by running different migration batches, which should be executed consecutively, one after another. You can have up to 500 entries in each migration batch.

One public folder should be migrated to only one group in one migration batch.

Step 4: Start the migration request

In this step, you gather information from your Exchange environment, and then you use that information in Exchange Online PowerShell to create a migration batch. After that, you start the migration.

1. On the Exchange 2016 or Exchange 2019 server, find the MRS proxy endpoint server and make note of it. You will need this information later when you run the migration request.

2. In Exchange Online PowerShell, use the information that was returned above in step 1 to run the following commands. The variables in these commands will be the values from step 1.

    a. Pass the credential of a user with administrator permissions in the Exchange Server environment into the variable `$Source_Credential`. When you eventually run the migration request in Exchange Online, you will use this credential to gain access to your Exchange 2016 or Exchange 2019 servers in order to copy the content over to Exchange Online.

    ```
    $Source_Credential = Get-Credential
    <source_domain>\<PublicFolder_Administrator_Account>
    ```

    b. Use the MRS proxy server information from your Exchange Server environment that you noted in Step 1 above and pass that value into the variable `$Source_RemoteServer`.

    ```
    $Source_RemoteServer = "<MRS proxy endpoint>"
    ```

3. In Exchange Online PowerShell, run the following command to create a migration endpoint:

    ```
    $PfEndpoint = New-MigrationEndpoint -PublicFolderToUnifiedGroup -Name PFToGroupEndpoint -RemoteServer
    $Source_RemoteServer -Credentials $Source_Credential
    ```

4. Run the following command to create a new public folder to Microsoft 365 or Office 365 group migration batch. In this command:

    - **CSVData** is the .csv file created above in Step 3: Create the .csv file. Be sure to provide the full path to this file. If the file was moved for any reason, be sure to verify and use the new location.

    - **NotificationEmails** is an optional parameter that can be used to set email addresses that will receive notifications about the status and progress of the migration.

    - **AutoStart** is an optional parameter which, when used, starts the migration batch as soon as it is created.

    - **PublicFolderToUnifiedGroup** is the parameter to indicate that it is a public folder to Microsoft 365 Groups migration batch.

    ```
    New-MigrationBatch -Name PublicFolderToGroupMigration -CSVData (Get-Content <path to .csv file> -
    Encoding Byte) -PublicFolderToUnifiedGroup -SourceEndpoint  $PfEndpoint.Identity [-NotificationEmails
    <email addresses for migration notifications>] [-AutoStart]
    ```

5. Start the migration by running the following command in Exchange Online PowerShell. Note that this step is necessary only if the `-AutoStart` parameter was not used while creating the batch above in step 4.

```
Start-MigrationBatch PublicFolderToGroupMigration
```

While batch migrations need to be created using the **New-MigrationBatch** cmdlet in Exchange Online PowerShell, the progress of the migration can be viewed and managed in Exchange admin center. You can also view the progress of the migration by running the Get-MigrationBatch and Get-MigrationUser cmdlets. The **New-MigrationBatch** cmdlet initiates a migration user for each Microsoft 365 or Office 365 group mailbox, and you can view the status of these requests using the mailbox migration page.

To view the mailbox migration page:

1. In Exchange Online, open Exchange admin center.

2. Navigate to **Recipients**, and then select **Migration**.

3. Select the migration request that was just created and then, on the **Details** pane, select **View Details**.

When the batch status is **Completed**, you can move on to *Step 5: Add members to Microsoft 365 groups from public folders*.

Step 5: Add members to Microsoft 365 groups from public folders

You can add members to the target group in Microsoft 365 or Office 365 manually as required. However, if you want to add members to the group based on the permission entries in public folders, you need to do that by running the script `AddMembersToGroups.ps1` on the Exchange 2016 or Exchange 2019 server as shown in the following command. User mailboxes must be synced to Exchange Online in order to be added as members of a Microsoft 365 or Office 365 group. To know which public folder permissions are eligible to be added as members of a group in Microsoft 365 or Office 365, see Migration scripts later in this article.

In the following command:

- **MappingCsv** is the .csv file created above in *Step 3: Create the .csv file*. Be sure to provide the full path to this file. If the file was moved for any reason, be sure to verify and use the new location.

- **BackupDir** is the directory where the migration log files will be stored.

- **ArePublicFoldersOnPremises** is a parameter to indicate whether public folders are located on-premises or in Exchange Online.

- **Credential** is the Exchange Online username and password.

```
.\AddMembersToGroups.ps1 -MappingCsv <path to .csv file> -BackupDir <path to backup directory> -
ArePublicFoldersOnPremises $true -Credential (Get-Credential)
```

Once users have been added to a group in Microsoft 365 or Office 365, they can begin using it.

Step 6: Lock down the public folders (public folder downtime required)

When the majority of the data in your public folders has migrated to Microsoft 365 Groups, you can run the script `LockAndSavePublicFolderProperties.ps1` on the Exchange 2016 or Exchange 2019 server to make the public folders read-only. This step ensures that no new data is added to public folders before the migration completes.

NOTE

If there are mail-enabled public folders (MEPFs) among the public folders being migrated, this step will copy some properties of MEPFs, such as SMTP addresses, to the corresponding group in Microsoft 365 or Office 365 and then mail-disable the public folder. Because the migrating MEPFs will be mail-disabled after the execution of this script, you will start seeing emails sent to MEPFs instead being received in the corresponding groups. For more details, see Migration scripts later in this article.

In the following command:

- **MappingCsv** is the .csv file created above in *Step 3: Create the .csv file*. Be sure to provide the full path to this file. If the file was moved for any reason, be sure to verify and use the new location.
- **BackupDir** is the directory where the backup files for permission entries, MEPF properties, and migration log files will be stored. This backup will be useful in case you need to roll back to public folders.
- **ArePublicFoldersOnPremises** is a parameter to indicate whether public folders are located on-premises or in Exchange Online.
- **Credential** is the Exchange Online username and password.

```
.\LockAndSavePublicFolderProperties.ps1 -MappingCsv <path to .csv file> -BackupDir <path to backup directory>
-ArePublicFoldersOnPremises $true -Credential (Get-Credential)
```

Step 7: Finalize the public folder to Microsoft 365 Groups migration

After you've made your public folders read-only, you'll need to perform the migration again. This is necessary for a final incremental copy of your data. Before you can run the migration again, you'll have to remove the existing batch, which you can do by running the following command:

```
Remove-MigrationBatch <name of migration batch>
```

Next, create a new batch with the same .csv file by running the following command. In this command:

- **CSVData** is the .csv file created above in *Step 3: Create the .csv file*. Be sure to provide the full path to this file. If the file was moved for any reason, be sure to verify and use the new location.
- **NotificationEmails** is an optional parameter that can be used to set email addresses that will receive notifications about the status and progress of the migration.
- **AutoStart** is an optional parameter which, when used, starts the migration batch as soon as it is created.

```
New-MigrationBatch -Name PublicFolderToGroupMigration -CSVData (Get-Content <path to .csv file> -Encoding
Byte) -PublicFolderToUnifiedGroup -SourceEndpoint $PfEndpoint.Identity [-NotificationEmails <email addresses
for migration notifications>] [-AutoStart]
```

After the new batch is created, start the migration by running the following command in Exchange Online PowerShell. Note that this step is only necessary if the `-AutoStart` parameter was not used in the preceding command.

```
Start-MigrationBatch PublicFolderToGroupMigration
```

After you have finished this step (the batch status is **Completed**), verify that all data has been copied to Microsoft 365 Groups. At that point, provided you're satisfied with the Groups experience, you can begin deleting the migrated public folders from your Exchange Server environment.

> **IMPORTANT**
>
> While there are supported procedures for rolling back your migration and returning to public folders, this isn't possible after the source public folders have been deleted. See How do I roll back to public folders from Microsoft 365 Groups? for more information.

Known issues

The following known issues can occur during a typical public folders to Microsoft 365 Groups migration.

- The script that transfers SMTP address from mail-enabled public folders to Microsoft 365 or Office 365 groups only adds the addresses as secondary email addresses in Exchange Online. Because of this, if you have Exchange Online Protection (EOP) or Centralized Mail Flow setup in your environment, will have issues sending email to the groups (to the secondary email addresses) post-migration.

- If the .csv mapping file has an entry with invalid public folder path, the migration batch displays as **Completed** without throwing an error, and no further data is copied.

Migration scripts

For your reference, this section provides in-depth descriptions for three of the migration scripts and the tasks they execute in your Exchange environment. You can download all scripts and supporting files from this location.

AddMembersToGroups.ps1

This script will read the permissions of the public folders being migrated and then add members and owners to Microsoft 365 groups as follows:

- Users with the following permission roles will be added as members to a group in Microsoft 365 or Office 365. **Permission roles**: Owner, PublishingEditor, Editor, PublishingAuthor, Author

- In addition to the above, users with the following minimum access rights will also be added as members to a group in Microsoft 365 or Office 365. **Access rights**: ReadItems, CreateItems, FolderVisible, EditOwnedItems, DeleteOwnedItems

- Users with access right "Owner" will be added as owners to a group and users with other eligible access rights will be added as members.

- Security groups cannot be added as members of Microsoft 365 groups. Therefore they will be expanded, and then the individual users will be added as members or owners to the groups based on the access rights of the security group.

- When users in security groups that have access rights over a public folder have themselves explicit permissions over the same public folder, explicit permissions will be given preference. For example, consider a case in which a security group called "SG1" has members User1 and User2. Permission entries for the public folder "PF1" are as follows:

  SG1: Author in PF1

  User1: Owner in PF1

  In this case, User1 will be added as an owner to the Microsoft 365 group.

- When the default permission of a public folder being migrated is 'Author' or above, the script will suggest setting the corresponding group's privacy setting as 'Public'.

This script can be run even after the lock-down of public folders, with parameter the `ArePublicFoldersLocked` set to `$true`. In this scenario, the script will read permissions from the back up file created during lock-down.

LockAndSavePublicFolderProperties.ps1

This script makes the public folders being migrated read-only. When mail-enabled public folders are migrated, they will first be mail-disabled and their SMTP addresses will be added to the respective Microsoft 365 groups. Then the permission entries will be modified to make them read-only. A back up of the mail properties of mail-enabled public folders, as well as the permission entries of all the public folders, will be copied, before performing any modification on them.

If there are multiple migration batches, a separate backup directory should be used with each mapping .csv file.

The following mail properties will be stored, along with respective mail-enabled public folders and Microsoft 365 groups:

- PrimarySMTPAddress
- EmailAddresses
- ExternalEmailAddress
- EmailAddressPolicyEnabled
- GrantSendOnBehalfTo
- SendAs Trustee list

The above mail properties will be stored in a .csv file, which can be used in the roll back process (if you want to return to using public folders, see How do I roll back to public folders from Microsoft 365 Groups? for more information). A snapshot of the mail-enabled public folders' properties will also be stored in a file called PfMailProperties.csv. This file is not necessary for the roll back process, but can still be used for your reference.

The following mail properties will be migrated to target groups as part of lock down:

- PrimarySMTPAddress
- EmailAddresses
- SendAs Trustee list
- GrantSendOnBehalfTo

The script ensures that the PrimarySMTPAddress and EmailAddresses of migrating mail-enabled public folders will be added as secondary SMTP addresses of the corresponding Microsoft 365 groups. Also, SendAs and SendOnBehalfTo permissions of users on mail-enabled public folders will be given equivalent permission in the corresponding target groups.

Access rights allowed

Only the following access rights will be allowed for users to ensure that the public folders are made read-only for all users. These are stored in **ListOfAccessRightsAllowed**:

- ReadItems
- CreateSubfolders
- FolderContact
- FolderVisible

The permission entries will be modified as follows:

BEFORE LOCK DOWN	AFTER LOCK DOWN
None	None
AvailabilityOnly	AvailabilityOnly
LimitedDetails	LimitedDetails
Contributor	FolderVisible
Reviewer	ReadItems, FolderVisible
NonEditingAuthor	ReadItems, FolderVisible
Aughor	ReadItems, FolderVisible
Editor	ReadItems, FolderVisible
PublishingAuthor	ReadItems, CreateSubfolders, FolderVisible
PublishingEditor	ReadItems, CreateSubfolders, FolderVisible
Owner	ReadItems, CreateSubfolders, FolderContact, FolderVisible

- Access rights for users without read permissions will be left untouched, and they will continue to be blocked from read rights.
- For users with custom roles, all the access rights that are not in **ListOfAccessRightsAllowed** will be removed. In the event that the users don't have any access rights from the allowed list after filtering, these users' access right will be set to 'None'.

There might be an interruption in sending emails to mail-enabled public folders during the time between when the folders are mail-disabled and their SMTP addresses are added to Microsoft 365 Groups.

UnlockAndRestorePublicFolderProperties.ps1

This script will re-assign permissions back to public folders, based on the back up file taken during public folder lock-down. This script will also mail-enable public folders that had been mail-disabled, after it removes the folders' SMTP addresses from their respective Microsoft 365 groups. There might be slight downtime during this process.

How do I roll back to public folders from Microsoft 365 Groups?

In the event that you change your mind and want to return to using public folders after using Microsoft 365 Groups, the command listed below will restore your environment to the state it was pre-migration. A roll back can be performed as long as the backup files exist and as long as you didn't delete the public folders post-migration.

On your Exchange 2016 or Exchange 2019 server, run the following command. In this command:

- **BackupDir** is the directory where the backup files for permission entries, MEPF properties, and migration log files will be stored. Make sure you use the same location you specified in *Step 6: Lock down the public folders to cut-over (public folder downtime required)*.
- **ArePublicFoldersOnPremises** is a parameter to indicate whether public folders are located on-premises or in Exchange Online.
- **Credential** is the Exchange Online username and password.

```
.\UnlockAndRestorePublicFolderProperties.ps1 -BackupDir <path to backup directory> -ArePublicFoldersOnPremises
$true -Credential (Get-Credential)
```

Be aware that any items added to the Microsoft 365 group, or any edit operations performed in the groups, are not copied back to your public folders. Therefore there will be data loss, assuming new data was added while the public folder was a group.

Note also that it's not possible to restore a subset of public folders, which means all of the public folders there were migrated should be restored.

The corresponding Microsoft 365 groups won't be deleted as part of the roll back process. You must clean or delete those groups manually.

Migrate public folders from Exchange 2013 to Exchange 2016 or Exchange 2019

8/3/2020 • 5 minutes to read • Edit Online

To migrate your Exchange 2013 public folders to Exchange 2016 or Exchange 2019, you need to move all of your Exchange 2013 public folder mailboxes to an Exchange 2016 server or Exchange 2019 server.

Before you move your public folder mailboxes, here are some things you should consider:

- **Capacity**: The size of your public folder mailboxes might vary significantly depending on how many public folders and public folder mailboxes you have. Make sure the target Exchange servers where you'll move your public folder mailboxes have enough storage capacity.

- **Time**: It might take a while to move your public folder mailboxes. The following items could impact how long it takes:

- Public folder mailbox size

- The number of public folder mailboxes

- Network bandwith

The good news is that your public folders will remain available during the public folder mailbox move. There's only a brief time window where the public folders might no be available (as the move completes).

What do you need to know before you begin?

- To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to move public folder mailboxes from Exchange 2013 to Exchange 2016 or Exchange 2019

1. Run the following command to get a list of all Exchange 2013 public folder mailboxes:

```
Get-ExchangeServer | Where {($_.AdminDisplayVersion -Like 'Version 15.0*') -And ($_.ServerRole -Like
'*Mailbox*')} | Get-Mailbox -PublicFolder | Get-MailboxStatistics | Format-Table -Auto
ServerName,DisplayName,TotalItemSize
```

2. Use the following syntax to list all mailbox databases on all Exchange 2016 or Exchange 2019 Mailbox servers:

```
Get-ExchangeServer | Where {($_.AdminDisplayVersion -like '<Version>') -and ($_.ServerRole -Like
"*Mailbox*")} | Get-MailboxDatabase | Format-List Server,Name,EdbFilePath
```

You can use the location information that's returned by this command to check the available free disk space for each mailbox database.

This example returns the locations of all mailbox databases on all Exchange 2016 Mailbox servers.

```
Get-ExchangeServer | where {($_.AdminDisplayVersion -like 'Version 15.1*') -and ($_.ServerRole -Like
'*Mailbox*')} | Get-MailboxDatabase | Format-List Server,Name,EdbFilePath
```

This example returns the locations of all mailbox databases on all Exchange 2019 Mailbox servers.

```
Get-ExchangeServer | where {($_.AdminDisplayVersion -like 'Version 15.2*') -and ($_.ServerRole -Like
'*Mailbox*')} | Get-MailboxDatabase | Format-List Server,Name,EdbFilePath
```

This example returns the locations of all mailbox databases on all Exchange 2016 and Exchange 2019 Mailbox servers.

```
Get-ExchangeServer | where {(($_.AdminDisplayVersion -like 'Version 15.1*') -or ($_.AdminDisplayVersion
-like 'Version 15.2*')) -and ($_.ServerRole -Like '*Mailbox*')} | Get-MailboxDatabase | Format-List
Server,Name,EdbFilePath
```

3. Use the information from the previous steps to decide the target mailbox database and/or Mailbox server (if you have more than one) to move some or all of your public folder mailboxes to. For example, you might not want to move three large public folder mailboxes to a server with low available drive space.

   You can also decide whether you want to move all public folder mailboxes at once, all public folder mailboxes on a specific server, or a specific public folder mailbox.

   Choose the command that fits the kind of move you want to do. Be sure to replace the Exchange server names, database names, and public folder mailbox names with your own.

   - Move all Exchange 2013 public folder mailboxes at once.

     ```
     Get-ExchangeServer | Where {($_.AdminDisplayVersion -Like "Version 15.0*") -And ($_.ServerRole -
     Like "*Mailbox*")} | Get-Mailbox -PublicFolder | New-MoveRequest -TargetDatabase
     Ex2016MbxDatabase
     ```

   - Move all public folder mailboxes on a specific Exchange 2013 server at once.

     ```
     Get-Mailbox -PublicFolder -Server Ex2013Mbx | New-MoveRequest -TargetDatabase Ex2016MbxDatabase
     ```

   - Move a specific Exchange 2013 public folder mailbox.

     ```
     New-MoveRequest "Sales Public Folder Mailbox" -TargetDatabase Ex2016MbxDatabase
     ```

4. To see the status of the move requests you created, run the following command:

   ```
   Get-MoveRequest
   ```

   Depending on the size of the public folder mailboxes you're moving and your available network capacity, it could take several hours or days for the moves to complete.

For a list of possible status values that can be returned, see the next section.

How do you know this worked?

To verify that you've successfully migrated all of your Exchange 2013 public folders to Exchange 2016 or Exchange 2019, do the following steps:

- Check the status of the move requests you created by running the following command in the Exchange Management Shell on an Exchange 2016 or Exchange 2019 Mailbox server:

  ```
  Get-MoveRequest
  ```

  The command will return each move request you created along with one of the following status values:

  - **Completed**: The public folder mailbox was successfully moved to the target mailbox database.
  - **CompletedWithWarning**: The public folder mailbox was moved to the target mailbox database, but one or more issues were encountered during the move. You can find more information by viewing the move report that was delivered to the Administrator mailbox.
  - **CompletionInProgress**: The public folder mailbox move to the target mailbox database is in its final stages. Public folders hosted in this mailbox may be unavailable for a brief period of time while the move is finalized.
  - **InProgress**: The public folder mailbox move to the target mailbox database is underway. Public folders hosted in this mailbox are available during this portion of the move.
  - **Failed**: The public folder mailbox move failed for one or more reasons. You can find more information by viewing the move report that was delivered to the Administrator mailbox.
  - **Queued**: The public folder mailbox move has been submitted but the move hasn't started yet.
  - **Retry**: The migration service is currently having trouble proceeding with the job, but it has not given up, and will continue trying.
  - **AutoSuspended**: The public folder mailbox move is ready to enter its final stages but won't proceed further until you manually resume the move.

    This option can be helpful if you want to choose the time a move will complete. You can automatically suspend a move when you create it by using the *SuspendWhenReadyToComplete* switch on the **New-MoveRequest** cmdlet. To resume the move when you're ready, use the **Resume-MoveRequest** cmdlet.
  - **Suspended**: The public folder mailbox move has been manually suspended by **Suspend-MoveRequest** cmdlet and won't proceed until you manually resume the move. To resume the move when you're ready, use the **Resume-MoveRequest** cmdlet.

- View the location of your public folder mailboxes after their move request has completed by running the following command on an Exchange 2016 or Exchange 2019 server:

  ```
  Get-Mailbox -PublicFolder | Get-MailboxStatistics | Format-Table ServerName,DisplayName,TotalItemSize
  ```

  In the list public folder mailboxes that are returned, verify that they've each been moved to an Exchange 2016 Mailbox server.

Set up public folders in a new organization

8/3/2020 • 3 minutes to read • Edit Online

Public folders in Exchange are based on a mailbox architecture that allows public folders to benefit from things such as the resiliency of a Database Availability Group (DAG) and other mailbox features.

For limits in on-premises Exchange Server, see Limits for public folders.

For additional management tasks related to public folders in Exchange Server, see Public folder procedures.

What do you need to know before you begin?

- Estimated time to complete this task: 30 minutes.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Public folders" entry in the Sharing and collaboration permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Step 1: Create the primary public folder mailbox

The primary public folder mailbox contains a writeable copy of the public folder hierarchy plus content and is the first public folder mailbox that you create for your organization. Subsequent public folder mailboxes will be secondary public folder mailboxes, which will contain a read-only copy of the hierarchy plus content.

For detailed steps, see Create a public folder mailbox.

Step 2: Create your first public folder

For detailed steps, see Create a public folder.

Step 3: Assign permissions to the public folder

After you create the public folder, you'll need to assign the **Owner** permissions level so that at least one user can access the public folder from the client and create subfolders. Any public folders created after this one will inherit the permissions of the parent public folder.

1. In the Exchange admin center (EAC), navigate to **Public folders** > **Public folders**.
2. In the list view, select the public folder.
3. In the details pane, under **Folder permissions**, click **Manage**.
4. In **Public Folder Permissions**, click **Add** +.
5. Click **Browse** to select a user.

6. In the **Permission level** list, select a level. At least one user should be an **Owner**.

7. Click **Save**.

8. You can add multiple users by clicking **Add** + and assigning the appropriate permissions using the steps above. You can also customize the permission level by selecting or clearing the check boxes. When you edit a predefined permission level such as **Owner**, the permission level will change to **Custom**.

For information about how to use the Exchange Management Shell to assign permissions to a public folder, see Add-PublicFolderClientPermission.

Step 4 (Optional): Mail-enable the public folder

If you want users to send mail to the public folder, you can mail-enable the public folder. This step is optional. If you don't mail-enable the public folder, users can post messages to the public folder by dragging into it items from Outlook.

1. In the EAC, navigate to **Public folders** > **Public folders**.

2. In the list view, select the public folder you want to mail-enable.

3. In the details pane, under **Mail settings - Disabled**, click **Enable**.

   A warning displays asking if you're sure you want to enable mail for the public folder. Click **Yes**.

The public folder will be mail-enabled and the name of the public folder will become the alias of the public folder. If you have multiple recipients with that name, the public folder's alias will be appended with a number. For example, if you have a distribution group named SalesTeam and you create a public folder named SalesTeam and then mail-enable it, the alias of that public folder will be SalesTeam1.

For information about how to use the Exchange Management Shell to mail-enable a public folder, see Enable-MailPublicFolder.

Create a public folder mailbox in Exchange Server

8/3/2020 • 3 minutes to read • Edit Online

Before you can create a public folder in Exchange server, you must first create a public folder mailbox. Public folder mailboxes contain the hierarchy information as well as the content for public folders.

The first public folder mailbox that you create in the organization is the primary hierarchy mailbox, which contains the only writable copy of the public folder hierarchy. Any additional public folder mailboxes that you create are secondary hierarchy mailboxes, which contain a read-only copy of the public folder hierarchy. You can create multiple public folder mailboxes for load balancing.

NOTE

For more information about the storage quotas and limits for public folders in on-premises Exchange, see Limits for public folders.

For additional management tasks related to public folders in Exchange Server, see Public folder procedures.

What do you need to know before you begin?

- Estimated time to complete: less than 5 minutes.
- Public folders on Exchange 2010 servers can't exist in the same organization with Exchange 2016 or later public folders. If you try to create a public folder mailbox when you still have legacy public folders, you'll receive the error **An existing Public Folder deployment has been detected. To migrate existing Public Folder data, create new Public Folder mailbox using -HoldForMigration switch.**

  Before you can create public folders in Exchange Server 2016 or later, you need to migrate your Exchange 2010 public folders by following the steps in Use batch migration to migrate public folders from Exchange 2010 to Exchange 2016.
- To move your public folder mailboxes from Exchange 2013 to Exchange 2016 or Exchange 2019, see Migrate public folders from Exchange 2013 to Exchange 2016 or Exchange 2019.
- For more information about the Exchange admin center, see Exchange admin center in Exchange Server. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Public folders" entry in the Sharing and collaboration permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to create a public folder mailbox

1. In the EAC, go to **Public folders** > **Public folder mailboxes**, and then click **Add** +.

2. In the **New public folder mailbox** page that opens, enter the following information:

   - **Name**: Enter the name for the public folder mailbox.

   - **Organizational unit**: Click **Browse** to select the location in Active Directory where the mailbox object is created.

   - **Mailbox database**: Click **Browse** to select the mailbox database where the mailbox is created.

   When you're finished, click **Save**.

Use the Exchange Management Shell to create a public folder mailbox

To create a public folder mailbox, use the following syntax:

```
New-Mailbox -PublicFolder -Name <Name>
```

This example creates the primary hierarchy public folder mailbox named Master Hierarchy, because this is the first public folder mailbox in the organization (the value of the *Name* parameter doesn't determine whether the mailbox is the primary hierarchy public folder mailbox).

```
New-Mailbox -PublicFolder -Name "Master Hierarchy"
```

This example creates a secondary hierarchy public folder mailbox named Istanbul, because this isn't the first public folder mailbox in the organization (the value of the *Name* parameter doesn't determine whether the mailbox is a secondary hierarchy public folder mailbox).

```
New-Mailbox -PublicFolder -Name Istanbul
```

For detailed syntax and parameter information, see New-Mailbox.

How do you know this worked?

To verify that you've successfully created the a public folder mailbox, do any of these steps:

- In the EAC, go to **Public folders** > **Public folder mailboxes** and verify the public folder mailbox is listed. The primary hierarcy public folder mailbox has the value **Primary Hierarchy** for the **Contains** property. All other public folder mailboxes have the value **Secondary Hierarchy** for the **Contains** property.

- In the Exchange Management Shell, run the following command to verify the mailbox is listed, and check the value of the **IsRootPublicFolderMailbox** property to see if the mailbox is the primary hierarchy public folder mailbox (`True`) or a secondary hierarchy public folder mailbox (`False`):

  ```
  Get-Mailbox -PublicFolder | Format-Table -Auto Name,ServerName,Database,IsRootPublicFolderMailbox
  ```

- In the Exchange Management Shell, run the following commands to verify the primary hierarchy public folder mailbox:

  1. Run the following command:

~~~
```PowerShell
Get-OrganizationConfig | Format-List RootPublicFolderMailbox
```
~~~

2. Use the GUID value returned by the first command with **Get-Mailbox** to confirm the mailbox name. You can copy the GUID value by right-clicking in the Exchange Management Shell window, selecting **Mark**, highlighting the GUID value, and then pressing ENTER.

~~~
```PowerShell
Get-Mailbox -PublicFolder -Identity <GUID>
```
~~~

Create a public folder

8/3/2020 • 2 minutes to read • Edit Online

Public folders are designed for shared access and provide an easy and effective way to collect, organize, and share information with other people in your workgroup or organization.

By default, a public folder inherits the settings of its parent folder, including the permissions settings.

For additional management tasks related to public folders, see Public folder procedures.

What do you need to know before you begin?

- Estimated time to complete: 5 minutes.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Public folders" entry in the Sharing and collaboration permissionstopic.
- You can't create a public folder unless you've first created a public folder mailbox. For more information about how to create a public folder mailbox, see Create a public folder mailbox.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

Use the EAC to create a public folder

When using the EAC to create a public folder, you'll only be able to set the name and the path of the public folder. To configure additional settings, you'll need to edit the public folder after it's created.

1. Navigate to **Public folders** > **Public folders**.
2. If you want to create this public folder as a child of an existing public folder, click the existing public folder in the list view. If you want to create a top-level public folder, skip this step.
3. Click **Add** **+**.
4. In **Public Folder**, type the name of the public folder.

> **IMPORTANT**
>
> Don't use a backslash () in the name when creating a public folder.

5. In the **Path** box, verify the path to the public folder. If this isn't the desired path, click **Cancel** and follow Step 2 of this procedure.
6. Click **Save**.

Use the Exchange Management Shell to create a public folder

This example creates a public folder named Reports in the path Marketing\2016.

```
New-PublicFolder -Name Reports -Path \Marketing\2016
```

IMPORTANT

Don't use a backslash () in the name when creating a public folder.

For detailed syntax and parameter information, see New-PublicFolder.

How do you know this worked?

To verify that you've successfully created a public folder, do the following:

- In the EAC, click **Refresh** to refresh the list of public folders. Your new public folder should be displayed in the list.

- In the Exchange Management Shell, run any of the following commands:

```
Get-PublicFolder -Identity \Marketing\2016\Reports | Format-List
```

```
Get-PublicFolder -Identity \Marketing\2016 -GetChildren
```

```
Get-PublicFolder -Recurse
```

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Mail-enable or mail-disable a public folder

8/3/2020 • 3 minutes to read • Edit Online

Public folders are designed for shared access and provide an easy and effective way to collect, organize, and share information with other people in your workgroup or organization. Mail-enabling a public folder allows users to post to the public folder by sending an email message to it. When a public folder is mail-enabled additional settings become available for the public folder in the Exchange admin center (EAC), such as email addresses and mail quotas. In the Exchange Management Shell, before a public folder is mail-enabled, you use the **Set-PublicFolder** cmdlet to manage all of its settings. After the public folder is mail-enabled, you use the **Set-PublicFolder** and the **Set-MailPublicFolder** cmdlets to manage the settings.

If you want users on the Internet to send mail to a mail-enabled public folder, you need to set addition permissions using the **Add-PublicFolderClientPermission** cmdlet.

For additional management tasks related to managing public folders, see Public folder procedures.

What do you need to know before you begin?

- Estimated time to complete: 5 minutes
- To ensure that users on the Internet can send e-mail messages to a mail-enabled public folder, the public folder needs to have at least the *CreateItems* access right granted to the Anonymous account. If you want to learn how to do this, see Allow anonymous users to send email to a mail-enabled public folder later in this article.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Public folders" entry in the Sharing and collaboration permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to mail-enable or mail-disable a public folder

1. Navigate to **Public folders** > **Public folders**.
2. In the list view, select the public folder that you want to mail-enable or mail-disable.
3. In the details pane, under **Mail settings**, click **Enable** or **Disable**.
4. A warning box displays asking if you're sure you want to enable or disable email for the public folder. Click **Yes** to continue.

If you want external users to send mail to this public folder, make sure you follow the steps in Allow anonymous users to send email to a mail-enabled public folder.

Use the Exchange Management Shell to mail-enable a public folder

This example mail-enables the public folder Help Desk.

```
Enable-MailPublicFolder -Identity "\Help Desk"
```

This example mail-enables the public folder Reports under the Marketing public folder, but hides the folder from address lists.

```
Enable-MailPublicFolder -Identity "\Marketing\Reports" -HiddenFromAddressListsEnabled $True
```

If you want external users to send mail to this public folder, make sure you follow the steps in Allow anonymous users to send email to a mail-enabled public folder.

For detailed syntax and parameter information, see Enable-MailPublicFolder.

Use the Exchange Management Shell to mail-disable a public folder

This example mail-disables the public folder Marketing\Reports.

```
Disable-MailPublicFolder -Identity "\Marketing\Reports"
```

For detailed syntax and parameter information, see Disable-MailPublicFolder.

Allow anonymous users to send email to a mail-enabled public folder

You can use either Outlook or the Exchange Management Shell to set permissions on a public folder's Anonymous account. You can't use the EAC to set permissions on the Anonymous account.

Use Outlook to set permissions for the Anonymous account

1. Open Outlook using an account that's been granted Owner permissions on the email-enabled public folder you want anonymous users to send mail to.

2. Navigate to **Public folders - <user's name>**.

3. Navigate to the public folder you want to change.

4. Right-click on the public folder, click **Properties** and then select the **Permissions** tab.

5. Select the **Anonymous** account, select **Create items** under **Write**, and then click **OK**.

Use the Exchange Management Shell to set permissions for the Anonymous account

This example sets the `CreateItems` permission for the Anonymous account on the "Customer Feedback" mail-enabled public folder.

```
Add-PublicFolderClientPermission "\Customer Feedback" -AccessRights CreateItems -User Anonymous
```

For detailed syntax and parameter information, see Add-PublicFolderClientPermission.

View statistics for public folders and public folder items

8/3/2020 • 2 minutes to read • Edit Online

You can use the Exchange Management Shell to retrieve statistics about a public folder, such as the display name, creation time, last user modified time, last user access, and item size. You can use this information to make decisions about deleting or retaining public folders.

> **NOTE**
>
> While you can view some of the quota and usage information in the Exchange admin center (EAC), this information is incomplete, and we recommend that you use the Exchange Management Shell to view public folder statistics. To view quota and usage information for public folders by navigating to **Public Folders** > **Edit** ✎ > **Mailbox usage**.

For additional management tasks related to public folders, see Public Folder Procedures in Exchange Online.

What do you need to know before you begin?

- Estimated time to complete: 1 minute.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Public folders" entry in the Sharing and collaboration permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to retrieve public folder statistics

This example returns the statistics for the public folder Marketing with a piped command to format the list.

```
Get-PublicFolderStatistics -Identity \Marketing | Format-List
```

> **NOTE**
>
> The value for the *Identity* parameter must include the path to the public folder. For example, if the public folder Marketing existed under the parent folder Business, you would provide the following value: `\Business\Marketing`

For detailed syntax and parameter information, see Get-PublicFolderStatistics.

Note that some parameters and settings might be available only in Exchange Online or only in Exchange Server.

Use the Exchange Management Shell to view statistics for public folder

items

You can view the following information about items within a public folder:

- Type of item
- Subject
- Last user modification time
- Last user access time
- Creation time
- Attachments
- Message size

You can use this information to make decisions about what actions to take for your public folders, such as which public folders to delete. For example, you may want to delete a public folder if the items haven't been accessed for over two years, or you may want to convert a public folder that's being used as a document repository to another client access application.

This example returns default statistics for all items in the public folder Pamphlets under the path \Marketing\2013. Default information includes item identity, creation time, and subject.

```
Get-PublicFolderItemStatistics -Identity "\Marketing\2013\Pamphlets"
```

This example returns additional information about the items within the public folder Pamphlets, such as subject, last modification time, creation time, attachments, message size, and the type of item. It also includes a piped command to format the list.

```
Get-PublicFolderItemStatistics -Identity "\Marketing\2010\Pamphlets" | Format-List
```

For detailed syntax and parameter information, see Get-PublicFolderItemStatistics.

Note that some parameters and settings might be available only in Exchange Online or only in Exchange Server.

Use the Exchange Management Shell to export the output of the Get-PublicFolderItemStatistics cmdlet to a .csv file

This example exports the output of the cmdlet to the PFItemStats.csv file that includes the following information for all items within the public folder \Marketing\Reports:

- Subject of the message (`Subject`)
- Date and time that the item was last modified (`LastModificationTime`)
- Whether the item has attachments (`HasAttachments`)
- Type of item (`ItemType)`
- Size of the item (`MessageSize`)

```
Get-PublicFolderItemStatistics -Identity "\Marketing\Reports" | Select
Subject,LastModificationTime,HasAttachments,ItemType,MessageSize | Export-CSV C:\PFItemStats.csv
```

For detailed syntax and parameter information, see Get-PublicFolderItemStatistics.

Shared mailboxes

8/3/2020 • 2 minutes to read • Edit Online

A shared mailbox is a mailbox that multiple users can use to read and send email messages. Shared mailboxes can also be used to provide a common calendar, allowing multiple users to schedule and view vacation time or work shifts.

Why set up a shared mailbox?

- Provides a generic email address (for example, info@contoso.com or sales@contoso.com), that customers can use to inquire about your company.
- Allows departments that provide centralized services to employees (for example, help desk, human resources, or printing services), to respond to employee questions.
- Allows multiple users to monitor and reply to email sent to an email address (for example, an address used specifically by the help desk).

What are shared mailboxes?

A shared mailbox is a type of user mailbox that doesn't have its own username and password. As a result, users can't log into them directly. To access a shared mailbox, users must first be granted Send As or Full Access permissions to the mailbox. Once that's done, users sign into their own mailboxes and then access the shared mailbox by adding it to their Outlook profile. In Exchange 2003 and earlier, shared mailboxes were just a regular mailbox to which an administrator could grant delegate access. Beginning in Exchange 2007, shared mailboxes became their own recipient type:

- **RecipientType**: UserMailbox
- **RecipientTypeDetails**: SharedMailbox

In previous version of Exchange, creating a shared mailbox was a multi-step process in which you had to use the Exchange Management Shell to complete some of the tasks. In Exchange 2013 and later, you can use the Exchange admin center (EAC) to create a shared mailbox in one step. For details, see Create shared mailboxes in the Exchange admin center. In fact, the EAC has a feature area devoted entirely to shared mailboxes. Just navigate to **Recipients** > **Shared mailboxes** to view all the management tasks for shared mailboxes.

You can use the following permissions with a shared mailbox.

- **Full Access**: The Full Access permission lets a user log into the shared mailbox and act as the owner of that mailbox. While logged in, the user can create calendar items; read, view, delete, and change email messages; create tasks and calendar contacts. However, a user with Full Access permission can't send email from the shared mailbox unless they also have Send As or Send on Behalf permission.
- **Send As**: The Send As permission lets a user impersonate the shared mailbox when sending mail. For example, if Kweku logs into the shared mailbox Marketing Department and sends an email, it will look like the Marketing Department sent the email.
- **Send on Behalf**: The Send on Behalf permission lets a user send email on behalf of the shared mailbox. For example, if John logs into the shared mailbox Reception Building 32 and sends an email, it look like the mail was sent by "John on behalf of Reception Building 32". You can't use the EAC to grant Send on Behalf permissions, you must use **Set-Mailbox** cmdlet with the *GrantSendonBehalf* parameter.

NOTE

A shared mailbox is not designed for direct logon. The user account for the shared mailbox itself should stay in a Disabled (or "disconnected") state.

Converting shared mailboxes

In previous versions of Exchange, you could use a regular mailbox as a delegated mailbox. If you have delegated mailboxes, you can use the Exchange Management Shell to convert those delegate mailboxes to shared mailboxes. For details, see Convert a mailbox in Exchange Server.

Create shared mailboxes in the Exchange admin center

8/3/2020 • 3 minutes to read • Edit Online

If your organization uses a hybrid Exchange environment, you should use the on-premises Exchange admin center (EAC) to create and manage shared mailboxes. The Exchange admin center (EAC) is the single unified management console that allows for managing both your on-premises and Exchange Online organizations and allows you to connect and configure features for both organizations. For more information, see Hybrid management in Exchange hybrid deployments.

Use the EAC to create a shared mailbox

For information on limitations, automapping, and getting your users set up, see Create a shared mailbox.

You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "User mailboxes" entry in the Recipients Permissions topic.

1. Go to **Recipients** > **Shared** > **Add** +.

2. Fill-in the required fields:

   - **Display name**

   - **Email address**

3. To grant Full Access or Send As permissions, click **Add** +, and then select the users you want to grant permissions to. You can use the CTRL key to select multiple users. Confused about which permission to use? See Which permissions should you use? later in this topic.

   > **NOTE**
   >
   > The Full Access permission allows a user to open the mailbox as well as create and modify items in it. The Send As permission allows anyone other than the mailbox owner to send email from this shared mailbox. Both permissions are required for successful shared mailbox operation.

4. Click **Save** to save your changes and create the shared mailbox.

Use the EAC to edit shared mailbox delegation

1. Go to **Recipients** > **Shared** > **Edit** ✎.

2. Click **Mailbox delegation**

3. To grant or remove Full Access and Send As permissions, click **Add** + or **Remove** − and then select the users you want to grant permissions to.

   > **NOTE**
   >
   > The Full Access permission allows a user to open the mailbox as well as create and modify items in it. The Send As permission allows anyone other than the mailbox owner to send email from this shared mailbox. Both permissions are required for successful shared mailbox operation.

4. Click **Save** to save your changes.

Use the Exchange Management Shell to create a shared mailbox

This example creates the shared mailbox Sales Department and grants Full Access and Send on Behalf permissions for the security group MarketingSG. Users who are members of the security group will be granted the permissions to the mailbox.

> **NOTE**
>
> This example assumes that you've already created the security group MarketingSG and that security group is mail-enabled. See Manage mail-enabled security groups in Exchange Server.

```
New-Mailbox -Shared -Name "Sales Department" -DisplayName "Sales Department" -Alias Sales | Set-Mailbox -
GrantSendOnBehalfTo MarketingSG | Add-MailboxPermission -User MarketingSG -AccessRights FullAccess -
InheritanceType All
```

For detailed syntax and parameter information, see New-Mailbox.

Which permissions should you use?

You can use the following permissions with a shared mailbox.

- **Full Access**: The Full Access permission lets a user log into the shared mailbox and act as the owner of that mailbox. While logged in, the user can create calendar items; read, view, delete, and change email messages; create tasks and calendar contacts. However, a user with Full Access permission can't send email from the shared mailbox unless they also have Send As or Send on Behalf permission.

- **Send As**: The Send As permission lets a user impersonate the shared mailbox when sending mail. For example, if Kweku logs into the shared mailbox Marketing Department and sends an email, it will look like the Marketing Department sent the email.

- **Send on Behalf**: The Send on Behalf permission lets a user send email on behalf of the shared mailbox. For example, if John logs into the shared mailbox Reception Building 32 and sends an email, it look like the mail was sent by "John on behalf of Reception Building 32". You can't use the EAC to grant Send on Behalf permissions, you must use **Set-Mailbox** cmdlet with the *GrantSendonBehalf* parameter.

More information

For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Email addresses and address books in Exchange Server

8/3/2020 • 2 minutes to read • Edit Online

Exchange uses address books to organize and store email address information for recipients in the organization. The topics that will help you learn about and configure email addresses and address books in Exchange Server are described in the following table.

KEY TERMINOLOGY	DESCRIPTION	TOPIC
Address book policies	The global address list (GAL) is the master list of all recipients in your Exchange organization. Address book policies (ABPs) provide a simpler mechanism for GAL segmentation in organizations that require multiple GALs. An ABP defines a GAL, an offline address book (OAB), a room list, and one or more address lists. You can then assign the ABP to users.	Address book policies in Exchange Server
Address lists	An address list is a subset of a GAL. Each address list is a dynamic collection of one or more types recipients. You can use address lists to help users find the recipients and resources that they need.	Address lists in Exchange Server
Details templates	Details templates control the appearance of recipient properties that are displayed in address lists in Outlook.	Details Templates
Email address policies	Email address policies are the rules that create email addresses for Exchange recipients.	Email address policies in Exchange Server
Hierarchical address books	The hierarchical address book (HAB) presents recipients in the GAL by using your organization's unique business structure (for example, seniority or management hierarchy), which provides an efficient method for locating internal recipients.	Hierarchical Address Books
Offline address books	An offline address book (OAB) is a collection of address lists that can be downloaded and used in Outlook by users that are disconnected from the Exchange organization.	Offline address books in Exchange Server

Email address policies in Exchange Server

8/3/2020 • 7 minutes to read • Edit Online

Email address policies define the rules that create email addresses for recipients in your Exchange organization. Email address policies in Exchange Server 2016 and Exchange Server 2019 are basically unchanged from Exchange Server 2010.

The SMTP domains that are available to use in email address policies are defined by the *accepted domains* that are configured in the Exchange organization (specifically, authoritative domains and internal relay domains). For more information about accepted domains, see Accepted domains in Exchange Server.

The basic components of an email address policy are:

- **Email address templates**: Define the email address format for the recipients (for example `<firstname>@contoso.com` or `<lastname>.<firstname>@contoso.com`).
- **Recipient filter**: Specifies the recipients whose email addresses are configured by the policy.
- **Priority**: Specifies the order to apply the email address policies (important if a recipient is identified by more than one policy).

To configure email address policies, see Procedures for email address policies in Exchange Server.

Email address templates

An email address template contains the **address type** and the **address format**. An email address policy can contain multiple email address templates. One template must define the primary (reply) SMTP email address, and there can be only one primary SMTP email address defined in the policy (it's the **Reply-To:** email address for recipients). Other email address templates in the policy define the additional or *proxy* addresses for recipients.

Address types

Although you'll primarily use SMTP email addresses in email address policies, other email address types are available. The valid address type values are:

- **SMTP**
- **GWISE**: Novell GroupWise. By default, looks for the missing `%ExchangeInstallPath%Mailbox\address\gwise\amd64\gwxpxgen.dll` file to validate the email address format.
- **NOTES**: Lotus Notes. By default, uses the included `%ExchangeInstallPath%Mailbox\address\notes\amd64\ntspxgen.dll` file to validate the email address format.
- **X400**: By default, uses the included `%ExchangeInstallPath%Mailbox\address\notes\amd64\x400prox.dll` file to validate the email address format.

Notes:

- In the Exchange Management Shell, the value `SMTP` specifies the primary email address, and the value `smtp` specifies additional (proxy) addresses.

  In the EAC, only the **Make this format the reply email address** check box controls whether the email address is the primary address or a proxy address. It doesn't matter whether you type SMTP or smtp in the **Enter a custom address type** field. However, in the list of email address templates in the policy, the EAC shows the value **SMTP** (bold and uppercase) for the primary address, and smtp (not bold and lowercase) for proxy addresses.

- The types of email addresses that you can configure in a email address policy are limited compared to those you can configure on individual recipients.
- All non-SMTP email addresses are considered custom address types. Exchange doesn't provide unique dialog boxes or property pages for X.400, Novell GroupWise, or Lotus Notes email address types. Non-SMTP email addresses require the appropriate .dll files.

Address formats

An SMTP email address uses the syntax `chris@contoso.com`, where the value `chris` is the *local part* of the email address, and the value `contoso.com` is the SMTP domain (also known as the *address space* or *name space*). The available SMTP domain values are determined by the accepted domains that are configured for your organization.

You can use email address policies to assign multiple SMTP email addresses to recipients by using different combinations of the local part and domain values. However, only one SMTP email address in a policy can be configured as the primary address.

All SMTP email address formats in the Exchange Management Shell, or custom SMTP email address formats in the EAC require you to use variables to define the local part of the email address. These variables are described in the following table:

VARIABLE	VALUE
%d	Display name
%g	Given name (first name)
%i	Middle initial
%m	Exchange alias
%r*xy*	Replace all occurrences of *x* with *y*
%r*xx*	Remove all occurrences of *x*
%s	Surname (last name)
%*n*g	The first *n* letters of the first name. For example, `%2g` uses the first two letters of the first name.
%*n*s	The first *n* letters of the last name. For example, `%2s` uses the first two letters of the last name.

In addition to variables, you can also use US ASCII text characters that are allowed in Exchange email addresses (for example, periods (`.`) or underscores (`_`). Note that each period needs to be surrounded by other valid characters (for example `%g.%s`).

In the EAC, you can selected from a short list of precanned SMTP email address formats. These address formats are described in the following table, where the example user is named Elizabeth Brunner, and the domain is contoso.com:

EXAMPLE	EXCHANGE MANAGEMENT SHELL EQUIVALENT
`<alias>@contoso.com`	`%m@contoso.com`

EXAMPLE	EXCHANGE MANAGEMENT SHELL EQUIVALENT
`elizabeth.brunner@contoso.com`	`%g.%s@contoso.com`
`ebrunner@contoso.com`	`%1g%s@contoso.com`
`elizabethb@contoso.com`	`%g%1s@contoso.com`
`brunner.elizabeth@contoso.com`	`%s.%g@contoso.com`
`belizabeth@contoso.com`	`%1s%g@contoso.com`
`brunnere@contoso.com`	`%s%1g@contoso.com`

Recipient filters for email address policies

Recipient filters identify the recipients that the email address policy applies to. There are two basic options: **precanned recipient filters** and **custom recipient filters**. These are basically the same recipient filtering options that are used by dynamic distribution groups and address books. The following table summarizes the differences between the two filtering methods.

RECIPIENT FILTERING METHOD	USER INTERFACE	FILTERABLE RECIPIENT PROPERTIES	FILTER OPERATORS
Precanned recipient filters	Exchange admin center (EAC) and the Exchange Management Shell	Limited to: • Recipient type (All recipient types or any combination of user mailboxes, resource mailboxes, mail contacts, mail users, and groups) • Company • Custom Attribute 1 to 15 • State or Province • Department	Property values require an exact match. Wildcards and partial matches aren't supported. For example, "Sales" doesn't match the value "Sales and Marketing". Multiple values of the same property always use the **or** operator. For example, "Department equals Sales or Department equals Marketing". Multiple properties always use the **and** operator. For example, "Department equals Sales and Company equals Contoso".
Custom recipient filters	Exchange Management Shell only	You can use virtually any available recipient attributes.	You use OPATH filter syntax to specify any available Windows PowerShell filter operators. Wildcards and partial matches are supported.

Notes:

- You can't used precanned filters and customized filters at the same time.
- The recipient's location in Active Directory (the organizational unit or container) is available in both precanned and custom recipient filters.

- If you create an email address policy in the Exchange Management Shell that uses custom recipient filters, you can't edit the recipient filters in the EAC.

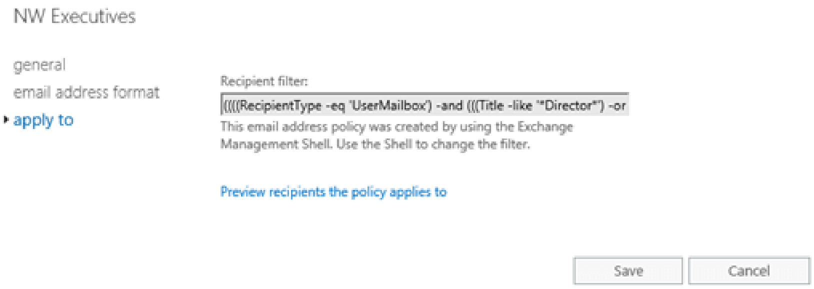

NW Executives

Recipient filter: ((((RecipientType -eq 'UserMailbox') -and (((Title -like '*Director*') -or

This email address policy was created by using the Exchange Management Shell. Use the Shell to change the filter.

- You can prevent individual recipients from being affected by email address policies. For example:
  - In the EAC, in the properties of the recipient, on the **Email address** tab, clear the check box: **Automatically update email addresses based on the email address policy applied to this recipient**.
  - In the Exchange Management Shell, set the *EmailAddressPolicyEnabled* parameter to the value `$false` on the recipient management cmdlet (for example, **Set-Mailbox** or **Set-DistributionGroup**).

Priority of email address policies

If a recipient is identified by multiple email address policies, the recipient's email addresses are only configured by the first email address policy that's evaluated. You configure the order that the policies are evaluated by using the priority of the policy. A lower priority number indicates a higher priority, higher priority policies are evaluated first, and the default email address policy is always evaluated last. You assign a higher priority (lower number) to policies that use the most specific or restrictive recipient filters.

Here are some other issues to consider:

- A recipient can only be affected by one email address policy. After the recipient is matched by the filtering properties of the policy, all other policies are ignored.
- **All** email address policies, including policies that have never been applied, are evaluated based on priority. For example, if you have a priority 1 policy and a priority 2 policy that both identity a recipient, the match in the first policy prevents the second policy from updating the recipient's email addresses, *even if the first policy has never been applied to the recipient*.

Default email address policy

Exchange setup creates a default email address policy that applies email addresses to all recipients in your organization. The properties of the default email address policy are described in the following list:

- **Name**: Default Policy
- **Priority**: Lowest (all other email address policies are evaluated before the default policy).
- **Email address format**
  - **Type**: `SMTP` (primary email address)
  - **Domain**: `<alias>@<ADForestRootFQDN>`. This domain value is used because it's the first accepted domain in the Exchange organization.
- **Apply to**: All recipient types.

You can't delete the default email address policy, and you can't designate another policy as the default. You can modify some properties of the default policy, but the modification options are limited:

- You can't filter recipients by type or properties (applies to all recipient types).
- You can't change the name or priority of the policy.
- You can fully customize the email address templates in the policy (modify, add, or remove templates). For more information, see Modify email address policies.

Apply email address policies

After you create or modify an email address policy in the EAC or the Exchange Management Shell, the policy needs to be applied to the affected recipients.

If the updates affect a large number of recipients (our recommendation is more than 3000), you should use the Exchange Management Shell to apply the updates to the affected recipients. For more information, see Apply email address policies to recipients.

Procedures for email address policies in Exchange Server

8/3/2020 • 17 minutes to read • Edit Online

Email address policies assign email addresses to recipients in your Exchange organization. You use the Exchange admin center (EAC) or the Exchange Management Shell to configure email address policies in Exchange Server.

For more information about email address policies, see Email address policies in Exchange Server.

What do you need to know before begin?

- Estimated time to complete each procedure: 5 minutes.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Email address policies" entry in the Email addresses and address books in Exchange Server topic.
- The procedures in this topic primarily focus on SMTP email addresses in email address policies, but other address types are available. For more information, see Address types.
- Before you can use an SMTP domain in an email address policy, you need to configure the domain as an accepted domain (specifically, an authoritative domain or internal relay domain). For more information, see Accepted domains in Exchange Server.
- To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forum at: Exchange Server.

Create email address policies

After you create an email address policy, you need to apply the policy to recipients. For more information, see the Apply email address policies to recipients section in this topic.

Use the EAC to create email address policies

1. In the EAC, go to **Mail flow** > **Email address policies**, and then click **Add** ✚.
2. In **New Email address policy** windows that opens, configure the following settings:
   - **Policy name**: Enter a unique, descriptive name for the policy.
   - **Email address format**: Click **Add** (✚) to configure an email address template. After you add the first template to define the primary SMTP email address, you can add additional templates for proxy email addresses (SMTP or otherwise), or you can click **Edit** (✎) to modify an existing template. For details about the settings that are available, see the Email address format window in the EAC section in this topic.

You can also click **Remove** (🗑) to delete existing templates.

Notes:

- The first SMTP email address template that you create here defines the primary (**Reply-To:**) SMTP email address. This template has the **Type** value **SMTP** (bold and uppercase), while other SMTP templates for proxy addresses have the **Type** value smtp (not bold and lowercase).

- You can't delete the email address template that defines the primary SMTP email address in the policy. Instead, you can add or modify another template, configure it to as the primary email address, and then delete the original template.

- **Run this policy in this sequence with other policies**: The value that you can select here depends on how many other email address policies you've manually created. For example, for the first email address policy that you create, the only available value is 1. If you create another policy, you can select 1 or 2. Remember, the first email address policy that identifies a recipient configures the recipient's email addresses. All other policies are ignored, even if the first policy is unapplied and can't configure the recipient's email addresses.

- For details about the recipient filters that are available here, see the Recipient filters in the EAC section in this topic.

3. When you're finished, click **Save**. You'll receive a warning message that tells you to click **Apply** in the details pane to apply the policy to recipients. For more information, see the Apply email address policies to recipients section in this topic.

Use the Exchange Management Shell to create email address policies

An email address policy in the Exchange Management Shell requires a **recipient filter**, and one or more **email address templates**. For details about recipient filters, see the Recipient filters in the Exchange Management Shell section in this topic.

Email address templates use the syntax `<Type>:<AddressFormat>`:

- `<Type>`: A valid email address type as described in Address types. For example, `SMTP` for the primary email address, and `smtp` for proxy addresses.

- `<AddressFormat>`: For SMTP email addresses, a domain or subdomain that's configured as accepted domain (authoritative or internal relay), and valid variables and ASCII text characters as described in Address formats. For example: `<alias>@contoso.com` requires the value `%m@contoso.com`, and `<firstname>.<lastname>@contoso.com` requires the value `%g.%s@contoso.com`.

To create an email address policy, use the following syntax:

```
New-EmailAddressPolicy -Name "<Policy Name>" <Precanned recipient filter | Custom recipient filter> [-RecipientContainer <OrganizationalUnit>] [-Priority <AllowedInteger>] -EnabledEmailAddressTemplates "SMTP:<PrimaryEmailAddressFormat>","smtp:<ProxyEmailAddress1>","smtp:<ProxyEmailAddress2>"...
```

This example creates an email address policy with a precanned recipient filter:

- **Name**: Southeast Offices

- **Precanned recipient filter**: All users with mailboxes where the **State or province** value is GA, AL, or LA (Georgia, Alabama, or Louisiana).

- **Primary SMTP email address**: `<last name>.<first two letters of the first name>@contoso.com`

- **Additional proxy email addresses**: `<last name>.<first two letters of the first name>@contoso.net`

- **Priority**: *n*+1, where *n* is the number of manually created email address policies that already exist (we didn't use the *Priority* parameter, and the default value is *n*+1). Remember, the first email address policy that identifies a recipient configures the recipient's email addresses. All other policies are ignored, even if the first policy is unapplied and can't configure the recipient's email addresses.

```
New-EmailAddressPolicy -Name "Southeast Offices" -IncludedRecipients MailboxUsers -ConditionalStateorProvince
"GA","AL","LA" -EnabledEmailAddressTemplates
"SMTP:%s%2g@southeast.contoso.com","smtp:%s%2g@southeast.contoso.net"
```

This example creates an email address policy with a custom recipient filter:

- **Name**: Northwest Executives

- **Custom recipient filter**: All users with mailboxes where the **Title** value contains Director or Manager, and the **State or province** value is WA, OR, or ID (Washington, Oregon, or Idaho).

- **Primary SMTP email address**: `<first two letters of the first name><last name>@contoso.com`

- **Additional proxy email addresses**: None

- **Priority**: 2

```
New-EmailAddressPolicy -Name "Northwest Executives" -RecipientFilter "(RecipientType -eq 'UserMailbox') -and
(Title -like '*Director*' -or Title -like '*Manager*') -and (StateOrProvince -eq 'WA' -or StateOrProvince -eq
'OR' -or StateOrProvince -eq 'ID')" -EnabledEmailAddressTemplates "SMTP:%2g%s@contoso.com" -Priority 2
```

Notes:

- Typically, you use the *EnabledEmailAddressTemplates* parameter to define the primary SMTP email address and one or more proxy addresses (SMTP or otherwise). However, if you're only going to define the primary SMTP email address and no additional proxy addresses, you can use the *EnabledPrimarySMTPAddressTemplate* parameter instead. This parameter doesn't require the `SMTP:` prefix, and you can't use this parameter with the *EnabledEmailAddressTemplates* parameter.

- The *EnabledEmailAddressTemplates* parameter requires at least one template with the `<Type>` value `SMTP` (to define the primary SMTP email address). After that, if you don't include a `<Type>` prefix for a template, the value `smtp` (an SMTP proxy address) is assumed.

For detailed syntax and parameter information, see New-EmailAddressPolicy.

How do you know this worked?

To verify that you've successfully created an email address policy, use either of the following procedures:

- In the EAC, go to **Mail flow** > **Email address policies**, verify that the policy is listed, and the details are correct. Select the policy and click **Edit** (✎) to view details that aren't displayed in the list view.

- In the Exchange Management Shell, run the following command to verify the property values:

  ```
  Get-EmailAddressPolicy | Format-List
  Name,Priority,Enabled*,RecipientFilterType,RecipientContainer,RecipientFilter,IncludedRecipients,Conditi
  onal*
  ```

Modify email address policies

- For the default email address policy, you can't modify the name, priority, or recipient filter settings. You can only modify the email address templates.

- After you modify an email address policy, you need to apply the policy to recipients. For more information, see the Apply email address policies to recipients section in this topic.

- If you created an email address policy in the Exchange Management Shell that uses a custom recipient filter, you can't modify the recipient filter in the EAC. You need to use the Exchange Management Shell.

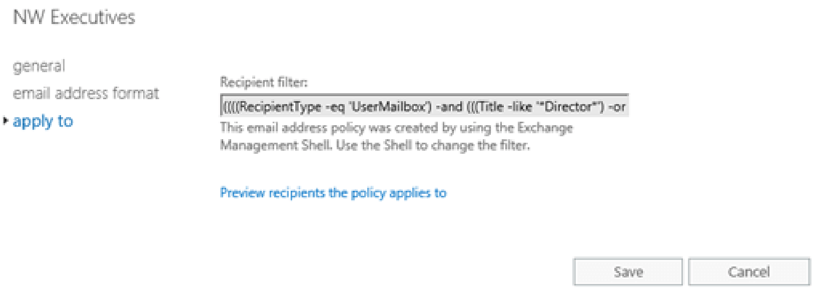

Recipient filter: ((((RecipientType -eq 'UserMailbox') -and (((Title -like '*Director*') -or

This email address policy was created by using the Exchange Management Shell. Use the Shell to change the filter.

- You can't use the EAC or the Exchange Management Shell to replace a custom recipient filter with a precanned recipient filter or vice-versa in an existing email address policy.

Modify email address policies in the EAC

The same settings are available as when you created the policy, although the settings are now located on separate tabs.

1. In the EAC, go to **Mail flow** > **Email address policies**, select the policy from the list, and then click **Edit** (✎).

2. Configure the settings on the following tabs:

   - **General**

   - **Policy name**: A unique, descriptive name for the policy.

   - **Run this policy in this sequence with other policies**: Remember, the first email address policy that identifies a recipient configures the recipient's email addresses. All other policies are ignored, even if the first policy is unapplied and can't configure the recipient's email addresses.

   - **Email address format**: For details about the settings that are available when you click **Add** (+) or **Edit** (✎), see the Email address format window in the EAC section in this topic.

     You can also click **Remove** (🗑) to delete existing email address templates.

   Notes:

   - The **Type** value **SMTP** (bold and uppercase) indicates the primary SMTP email address, and the value smtp (not bold and lowercase) indicates a proxy address.

   - You can't delete the email address template that defines the primary SMTP email address in the policy. Instead, you can add or modify another template, configure it to define the primary email address, and then delete the original template.

   - **Apply to**: For details about the recipient filters that are available here, see the Recipient filters in the EAC section in this topic.

     Note: Even if you configured a custom recipient filter in the Exchange Management Shell, you can still select **Preview recipients the policy applies to** here.

3. When you're finished, click **Save**. You'll receive a warning message that tells you to click **Apply** in the details pane to apply the policy to recipients. For more information, see the Apply email address policies to recipients section in this topic.

Modify email address policies in the Exchange Management Shell

The same basic settings are available as when you created the policy. For more information, see the Use the Exchange Management Shell to create email address policies section in this topic.

To modify an existing email address template, use the following syntax:

```
Set-EmailAddressPolicy -Identity <EmailAdressPolicyIdentity> [-Name <Name>] [<Precanned recipient filter |
Custom recipient filter>] [-RecipientContainer <OrganizationalUnit>] [-Priority <AllowedInteger>] [-
EnabledEmailAddressTemplates <"Type1:AddressFormat1","Type2:AddressFormat2"...] [-
DisabledEmailAddressTemplates <"Type1:AddressFormat1","Type2:AddressFormat2"... | $null>]
```

When you modify the *Conditional* parameter values, you can use the following syntax to add or remove values without affecting other existing values: `@{Add="<Value1>","<Value2>"...; Remove="<Value1>","<Value2>"...}`.

This example modifies the existing email address policy named Southeast Executives by adding the **State or province** value TX (Texas) to the precanned recipient filter.

```
Set-EmailAddressPolicy -Identity "Southeast Executives" -ConditionalStateOrProvince @{Add="TX"}
```

The *DisabledEmailAddressTemplates* parameter specifies inactive email address templates that are no longer used in the policy, and uses the same syntax as the *EnabledEmailAddressTemplates* parameter (except that *DisabledEmailAddressTemplates* can't contain a primary SMTP email address). Typically, this property is only populated if you've migrated from a previous version of Exchange. However, if a domain is specified in this property, you can't remove the corresponding accepted domain.

This example clears the disabled email address templates from the email address policy named Contoso Executives.

```
Set-EmailAddressPolicy -Identity "Contoso Executives" -DisabledEmailAddressTemplates $null
```

For detailed syntax and parameter information, see Set-EmailAddressPolicy.

How do you know this worked?

To verify that you've successfully modified an email address policy, use either of the following procedures:

- In the EAC, go to **Mail flow** > **Email address policies**, and verify the properties are correct. Select the policy and click **Edit** (✎) to view properties that aren't displayed in the list view.

- In the Exchange Management Shell, run the following command to verify the property values:

  ```
  Get-EmailAddressPolicy | Format-List
  Name,Priority,*Template*,RecipientFilterType,RecipientContainer,RecipientFilter,IncludedRecipientsCondit
  ional*
  ```

Apply email address policies to recipients

After you create or modify an email address policy in the EAC or the Exchange Management Shell, you need to apply the policy to the affected recipients.

- If the policy affects more than 3000 recipients, we recommend that you use the Exchange Management Shell. The recipient updates will take a long time, and will prevent you from using the EAC session until the updates are finished.

- If the policy affects less than 3000 recipients, it's OK to use the EAC.

Use the EAC to apply email address policies to recipients

1. In the EAC, go to **Mail flow** > **Email address policies**.

2. Select the email address policy that you want to apply (a policy that has the **Status** value **Unapplied**).

3. In the details pane, click **Apply**.

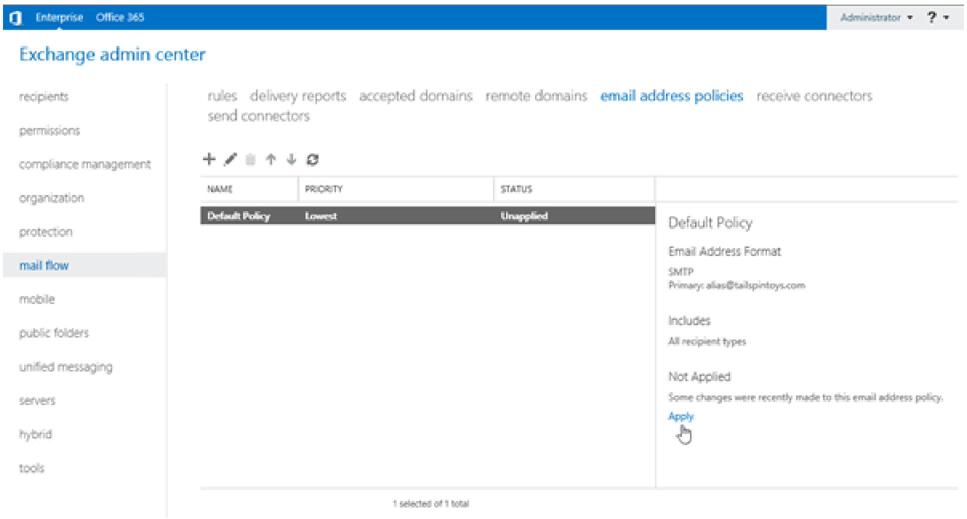

NAME	PRIORITY	STATUS
Default Policy	Lowest	Unapplied

Default Policy

Email Address Format

SMTP
Primary: alias@tailspintoys.com

Includes

All recipient types

Not Applied

Some changes were recently made to this email address policy.

Apply

1 selected of 1 total

4. After you click **Apply**, a warning message that appears. Click **Yes** to apply the policy by using the EAC. A progress bar allows you to monitor the recipient update process. When updates are complete, click **Close**.

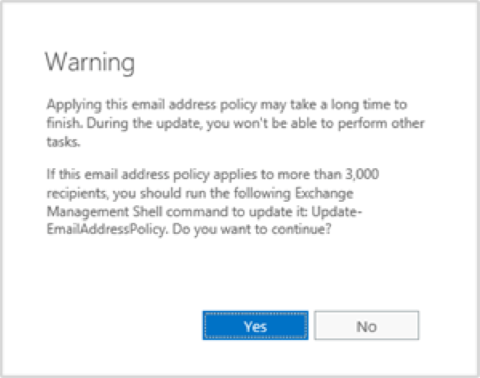

Warning

Applying this email address policy may take a long time to finish. During the update, you won't be able to perform other tasks.

If this email address policy applies to more than 3,000 recipients, you should run the following Exchange Management Shell command to update it: Update-EmailAddressPolicy. Do you want to continue?

Use the Exchange Management Shell to apply email address policies to recipients

To apply an email address policy to recipients, use the following syntax:

```
Update-EmailAddressPolicy -Identity <EmailAddressPolicyIdentity> [-FixMissingAlias] -
[UpdateSecondaryAddressesOnly]
```

This example applies the email address policy named Northwest Executives.

```
Update-EmailAddressPolicy -Identity "Northwest Executives"
```

For detailed syntax and parameter information, see Update-EmailAddressPolicy.

How do you know this worked?

To verify that you've successfully applied an email address policy, use either of the following procedures:

- In the EAC, go to **Mail flow** > **Email address policies**, and verify that the **Status** value of the policy is **Applied**.

- In the Exchange Management Shell, run the following command to verify the **RecipientFilterApplied** property has the value `True`:

```
Get-EmailAddressPolicy | Format-Table -Auto Name,RecipientFilterApplied
```

Remove email address policies

- You can't delete the default email address policy.
- If the policy affects more than 3000 recipients, we recommend that you use the Exchange Management Shell to remove the policy. The recipient updates will take a long time, and will prevent you from using the EAC session until the updates are finished. If removing the policy affects less than 3000 recipients, it's OK to use the EAC.

Use the EAC to remove email address policies

1. In the EAC, go to **Mail flow** > **Email address policies**.
2. Select the email address policy that you want to delete, and then click **Remove** 🗑.
3. Click **Yes** in the warning message that appears. A progress bar allows you to monitor the recipient update process. When updates are complete, click **Close**.

Use the Exchange Management Shell to remove email address policies

To remove an email address policy, use the following syntax:

```
Remove-EmailAddressPolicy -Identity <EmailAddressPolicyIdentity>
```

This example removes the email address policy named Southeast Offices.

```
Remove-EmailAddressPolicy -Identity "Southeast Offices"
```

For detailed syntax and parameter information, see Remove-EmailAddressPolicy.

How do you know this worked?

To verify that you've successfully removed an email address policy, use either of the following procedures:

- In the EAC, go to **Mail flow** > **Email address policies**, and verify that the policy is no longer listed.
- In the Exchange Management Shell, run the following command to verify that the email address policy isn't listed:

  ```
  Get-EmailAddressPolicy
  ```

Reference

Email address format window in the EAC

As you create or modify an email address policy in the EAC, in the **Email address format** section, an **Email address format** window appears when you click **Add** (+) or **Edit** (✎). The following settings are available in this window:

- Precanned SMTP email addresses:
  - **Select an accepted domain**: Select an accepted domain (authoritative domain or internal relay domain) from the drop down list. Note that if you've configured an accepted domain for a domain and all subdomains (for example, `*.contoso.com`), only the root domain (`contoso.com`) is available in the drop down list.

  Or

  - **Specify a custom domain name for the email address**: Select this option when you need to enter a subdomain of a `*.<domain>` accepted domain. For example, if `*.contoso.com` is configured as an authoritative domain, you can type eu.contoso.com in this field.

  And then:

  - **Email address format**: Select one of the available email address templates from the list.

- Custom SMTP or non-SMTP email addresses:

  - Click **More options** and then select **Enter a custom address type**.

  - **Enter a custom address type**: If this is the first email address template that you're configuring in the policy, type SMTP, and then continue to the **Email address parameters** field to define the primary SMTP email address format.

    After you've configured a template in the policy to define the primary SMTP email address, you can type SMTP or another address type value to configure email address templates for additional proxy addresses. For more information about the type values that you can use, see Address types.

  - **Email address parameters**: For SMTP email addresses, this value contains:

  - Valid variables and ASCII text characters as described in Address formats.

  - A domain or subdomain that's configured as an accepted domain (authoritative or internal relay).

    An example value is `%3g.%s@contoso.com` for `<first three letters of the first name>.<last name>` @contoso.com.

- **Make this format the reply email address**: The first email address template in a policy is automatically configured as the primary (reply) email address (you can't uncheck the check box). When you add additional templates to the policy, you can select this check box to define the primary email address.

Recipient filters in the EAC

When you create or modify email address policies in the EAC, the following recipient filter settings are available:

- **Specify the types of recipients this email address policy will apply to**:

  - **All recipient types**

  Or

  - **Only the following recipient types**: Select one or more of the following values:

    - **Users with Exchange mailboxes**

    - **Mail users with external email addresses**

    - **Resource mailboxes**

    - **Mail contacts with external email addresses**

    - **Mail-enabled groups**

- **Create rules to further define the recipients that this email address policy applies to**:

  1. Click **Add rule** and select one of the recipient properties from the drop down list:

     - **Recipient container** (container or organization unit)

     - **State or province**

    - **Company**
    - **Department**
    - Custom attribute 1 to 15

2. Enter a value for the property you selected:

    - If you selected **Recipient container**, a **Select an organizational unit** dialog box appears that allows you to select the container or OU in Active Directory.
    - For other recipient properties, a **Specify words or phrases** dialog appears that allows you to add, edit and remove text values.
    - Property values require an exact match. Wildcards and partial matches aren't supported. For example, the value "Sales" doesn't match "Sales and Marketing".
    - Multiple values of the same property use the **or** operator. For example, "Department equals Sales or Department equals Marketing"

3. After you've selected a property and value, click **Add rule**.

4. Repeat the previous steps to configure more filters. Note that multiple properties use the **and** operator. For example, "Department equals Sales and Company equals Contoso".

    Preview recipients the policy applies to: When you click this setting, a **Preview** dialog appears that shows you the recipients that are identified by the filters you configured.

Notes:

- You can't configure any recipient filter settings in the default email address policy (**All recipient types** is selected).
- If you configure too many recipient filter rules, you can restrict the policy to the point where it doesn't contain any recipients.

Recipient filters in the Exchange Management Shell

In the Exchange Management Shell, you can specify **precanned recipient filters**, or **custom recipient filters**, but not both at the same time.

- **Precanned recipient filters**:

    - Uses the required *IncludedRecipient* parameter with the `AllRecipients` value *or* one or more of the following values: `MailboxUsers`, `MailContacts`, `MailGroups`, `MailUsers`, or `Resources`. You can specify multiple values separated by commas.
    - You can also use any of the optional *Conditional* filter parameters: *ConditionalCompany*, *ConditionalCustomAttribute[1to15]*, *ConditionalDepartment*, and *ConditionalStateOrProvince*.

        You specify multiple values for a *Conditional* parameter by using the syntax `"<Value1>","<Value2>"...`. Multiple values of the same property implies the **or** operator. For example, "Department equals Sales or Marketing or Finance".

- **Custom recipient filters**: Uses the required *RecipientFilter* parameter with an OPATH filter.

    - The basic OPATH filter syntax is `"<Property1> -<Operator> '<Value1>' <Property2> -<Operator> '<Value2>'..."`.
    - Double quotation marks `" "` are required around the whole OPATH filter. Although the filter is a string (not a system block), you can also use braces `{ }`, but only if the filter doesn't contain variables that require expansion..

  - Hyphens (`-`) are required before all operators. Here are some of the most frequently used operators:
  - `and` , `or` , and `not` .
  - `eq` and `ne` (equals and does not equal; not case-sensitive).
  - `lt` and `gt` (less than and greater than).
  - `like` and `notlike` (string contains and does not contain; requires at least one wildcard in the string. For example, `"Department -like 'Sales*'"` .
  - Use parentheses to group `<Property> -<Operator> '<Value>'` statements together in complex filters. For example, `"(Department -like 'Sales*' -or Department -like 'Marketing*') -and (Company -eq 'Contoso' -or Company -eq 'Fabrikam')"` . Exchange stores the filter in the **RecipientFilter** property with each individual statement enclosed in parentheses, but you don't need to enter them that way.
  - For more information, see Additional OPATH syntax information.
  - After you use the **New-EmailAddressPolicy** cmdlet to create a policy that uses custom recipient filters, you can't modify the recipient filters in the EAC. You need to use the **Set-EmailAddressPolicy** cmdlet with the *RecipientFilter* parameter in the Exchange Management Shell.

Note: The *RecipientContainer* (organizational unit) recipient filter parameter is available to both precanned recipient filters and custom recipient filters.

Offline address books in Exchange Server

8/3/2020 • 11 minutes to read • Edit Online

An offline address book (OAB) is a local copy of an address list collection. OABs are used for address book queries by Outlook clients that are configured in cached Exchange mode. OABs are the only option for Outlook clients that are disconnected from the Exchange server, but they're also queried first by connected Outlook clients as a way to help reduce the workload on Exchange servers. You can configure which address lists are included in an OAB, access to specific OABs, how frequently the OABs are generated, and where the OABs are distributed from.

By default, a new installation of Exchange creates an OAB named Default Offline Address Book on the server. This OAB is also the default OAB, which means it's the OAB that's used by mailboxes and mailbox databases that don't have an OAB assigned to them.

OABs in Exchange 2013 and later are improved over OABs in Exchange 2010. These changes were introduced in Exchange 2013:

- Only web-based distribution is supported (public folder distribution is no longer available). Web-based distribution allows:
  - Support for more concurrent downloads by client computers.
  - Reduced bandwidth usage.
  - More control over the OAB distribution points.
- Only OAB version 4 is supported. This version of the OAB is Unicode, and allows clients to receive differential updates, instead of always using full downloads. All versions of Outlook that are supported by Exchange fully support OAB version 4.
- A mailbox assistant (not the Microsoft Exchange System Attendant service) is the process that's responsible for generating OABs. This allows OAB generation to run or pause based on the workload of the server (workload management).
- OAB generation occurs in a designated arbitration mailbox (not on a designated OAB generation server). These mailboxes can use database availability groups (DAGs) to help prevent a single point of failure for OAB generation and downloads.

For OAB procedures, see Procedures for offline address books in Exchange Server.

To learn more about address lists, see Address lists in Exchange Server.

OAB generation

OAB generation is controlled by the mailbox assistant named **OABGeneratorAssistant** that runs under the Microsoft Exchange Mailbox Assistants service. OAB generation occurs in a designated arbitration mailbox that has the `OrganizationCapabilityOABGen` value for the **PersistedCapability** property. An arbitration mailbox with this capability is also known as an *organization mailbox*.

By default, OABs are generated every 8 hours. To change the OAB generation schedule, see Change the offline address book generation schedule in Exchange Server. To manually update an OAB, see Use the Exchange Management Shell to update offline address books.

The arbitration mailbox named `SystemMailbox{bb558c35-97f1-4cb9-8ff7-d53741dc928c}` is the first organization mailbox in your organization. By default, this organization mailbox is responsible for generating all OABs (the first

OAB named Default Offline Address Book, and any new OABs that you create).

You can create additional organization mailboxes to generate OABs. Exchange Server contains the improvements to OAB generation that were introduced in Exchange 2013 Cumulative Update 7 (CU7):

- You can configure multiple OABs to be generated by the same organization mailbox, but you can't configure an OAB to be generated by more than one organization mailbox. If you configured an OAB with multiple organization mailboxes, each copy of the OAB had a different unique identifier. So, a full OAB download was required whenever a client was proxied to a different organization mailbox location.

- You can configure an OAB to allow a read-only copy (also known as a *shadow copy*) to be distributed to all organization mailboxes in the organization (also known as *shadow distribution*). All copies of the OAB have the same unique identifier, so full a OAB download isn't required when a client is proxied to a different organization mailbox location.

  Typically, shadow copies are only required in multi-site Exchange organizations. You configure an organization mailbox in each site, and you configure shadow distribution for an OAB to help prevent cross-site OAB download requests by clients (likely over slow WAN links). To create additional organization mailboxes, see Use the Exchange Management Shell to create organization mailboxes.

  Shadow distribution is described in detail in the next section.

To find all organization mailboxes, and the organization mailbox that's defined for an OAB, see Use the Exchange Management Shell to find organization mailboxes.

The OAB files are generated and stored in the designated organization mailbox, so the destination for OAB download requests is the Mailbox server that holds the active copy of the organization mailbox. The OAB files are copied from the organization mailbox to `%ExchangeInstallPath%ClientAccess\OAB\<OAB GUID>` for retrieval by clients. Clients never connect directly to this backend location. Client requests for the OAB are proxied by the Client Access (frontend) services on a Mailbox server to this backend location.

OAB distribution

By default, Outlook clients are configured to download the OAB every 24 hours, or users can initiate a manual download from Outlook at any time.

OAB distribution to clients depends on Internet Information Services (IIS) virtual directories and the Autodiscover service. The IIS virtual directory that's used for client access to OABs is located in the default web site in the Client Access (frontend) services on the Mailbox server, and is named OAB (Default Web Site). This virtual directory is automatically created when you install Exchange, and is configured to service internal clients at the URL `https://<ServerName>/oab` (for example, `https://mailbox01.contoso.com/oab`). You'll need to manually configure the external URL that's used to distribute OABs to external clients. For more information, see Step 4: Configure external URLs in Configure mail flow and client access on Exchange servers.

In the properties of the OAB, you can configure the OAB virtual directories that are available to distribute the OAB to clients. The default setting restricts OAB distribution to the OAB virtual directories on the server that holds the OAB's organization mailbox. However, the Client Access services on *any* Mailbox server can proxy incoming OAB download requests to the correct location. Therefore, we recommend that you configure all OAB virtual directories to accept requests to download the OAB. For instructions, see Use the Exchange Management Shell to configure any virtual directory in the organization to accept download requests for the OAB.

The Autodiscover service advertises the OAB URLs that you've configured. Autodiscover is supported by all versions of Outlook and virtually all mobile devices that are currently by Exchange. Here's a summary of the OAB distribution process:

1. Outlook receives the OAB URL from Autodiscover, and connects to the Client Access (frontend) services on a Mailbox server.

2. The Client Access services on the Mailbox server that accepted the connection performs these steps:

   a. Queries Active Directory to find the organization mailbox that's responsible for generating the user's OAB (the default OAB, the OAB that's specified for the mailbox database, or the OAB that's specified for the mailbox).

   b. Queries Active Directory again to find the mailbox database that hosts the organization mailbox for the OAB, and the Mailbox server that currently holds the active copy of the database.

   c. Proxies the OAB download request to the identified Mailbox server.

   d. Retrieves the OAB files from the backend location `%ExchangeInstallPath%ClientAccess\OAB\<GUID>` and proxies them back to the client.

If a shadow copy of the OAB exists in an organization mailbox in the local Active Directory site (the site where the user is connecting from), then a local Mailbox server is used to download the OAB. However, synchronization of the shadow copy between organization mailboxes is performed on-demand. Here's how it works:

1. Let's say the organization mailbox doesn't have a suitable shadow copy of the OAB. This can be caused by the following conditions:

   - A client has never requested a download of the shadow copy.

   - The shadow copy is out of date. Shadow copies are aware when an updated copy of the parent OAB has been generated and published (manually, or by the default 8 hour OAB generation schedule). The affected Mailbox servers will stop distributing the outdated shadow copy to clients.

2. The first client tries to download the shadow copy will receive error `0x80190194 (BG_E_HTTP_ERROR_404)` in Outlook. This will trigger a full copy of the OAB from the parent to the shadow copy. The following events are reported:

   - `Event ID: 102`

     `Source: MSExchange OABRequestHandler`

     `Description: The OABRequestHandler has begun downloading the OAB <GUID> from the server <Server>.`

   - `Event ID: 103`

     `Source: MSExchange OABRequestHandler`

     `Description: The OABRequestHandler has finished downloading the OAB <GUID>.`

3. The `OABRequestHandler` will make up to three immediate attempts to copy the OAB files from the Mailbox server that holds the parent OAB generation mailbox. If all three attempts fail, the `OABRequestHandler` will retry the copy after one hour. The following events are reported:

   - `Event ID: 104`

     `Source: MSExchange OABRequestHandler`

     ```
     Description: Download of the OAB <GUID> failed. The job will be re-submitted. The error was:
     BG_ERROR_CONTEXT=BE_ERROR_CONTEXT_REMOTE_FILE; error code=0x80190194
     ```

   - `Event ID: 105`

     `Source: MSExchange OABRequestHandler`

     ```
     Description: Download of the OAB <GUID> has failed too many times. The job will not be
     resubmitted for the next hour.
     ```

4. If the OAB is configured for shadow distribution, but there's no organization mailbox in the local Active

Directory site (the site where the user is connecting from), the Client Access services will proxy the OAB download request back to the Mailbox server that holds the organization mailbox for the parent OAB.

Conditions that cause a full OAB download

The improvements to OABs typically require clients to download OAB updates, not the full and complete OAB. However, full OAB downloads are sometimes required. For example:

- The `Changes.oab` files are greater than or equal to half the size of the full OAB files. Outlook compares the total size of the compressed `Changes.oab` files that are required to update the OAB to the total size of the compressed full OAB files on the server.
- There's no OAB on your computer (for example, during the initial setup of Outlook).
- A differential file is missing on the server. Missing differential files can be caused by the following conditions:
  - You haven't used Outlook to connect to your Exchange mailbox in more than 30 days (by default, the differential files are stored on the server for 30 days).
  - The server couldn't generate the differential file for a day that's required to update your local copy of the OAB.
- A more recent version of the OAB is available on the server (for example, your mailbox was upgraded from Exchange 2010, and your local copy of the OAB is version 3).
- Applying changes to the OAB failed. For example, differential files are corrupted on the server (the server crashed during differential file generation).
- The OAB is not present on your computer (for example, you manually deleted one or more local OAB files).
- A previous full download failed, so Outlook has to start over.
- You initiated a manual download of the full OAB.

OAB planning and deployment

Whether you use a single OAB or multiple OABs, consider the following factors as you plan and implement your OAB strategy:

- Th size of each OAB in your organization. OAB sizes can vary from a few megabytes to hundreds of megabytes. The following factors can affect the size of the OAB:
  - The usage of certificates in your organization. The more public key infrastructure (PKI) certificates, the larger the OAB. PKI certificates range from 1 kilobyte (KB) to 3 KB. They're the single largest contributor to the OAB size.
  - The number of mail recipients in your organization.
  - The number of groups in your organization.
  - User information that your organization adds to each recipient object. For example, some organizations configure full address and contact details for each user.
- The number of OAB downloads.
- The number and frequency of parent distinguished name changes for recipient objects in Active Directory.
- SMTP address mismatches.
- The overall number of changes that you make to Active Directory.

- Recipients that you've hidden in Active Directory by using methods outside of Exchange will be visible in OABs (for example, by using the Windows security descriptor). To effectively hide recipients in OABs, configure the **Hide from address lists** property for the recipient in the Exchange admin center (EAC) or the *HiddenFromAddressListsEnabled* parameter in the corresponding recipient management cmdlet in the Exchange Management Shell. For more information, see Hide recipients from address lists. Or, you can create an address list that doesn't include the hidden recipients, assign the address list to the OAB, and assign the OAB to users (directly or by making the OAB the default). For more information about creating address lists, see Create address lists.

Move OAB generation to another server

In Exchange 2010, moving OAB generation to another server required you to specify a different generation server in the properties of the OAB. But in Exchange 2013, Exchange 2016 and Exchange 2019, OAB generation occurs in a designed organization mailbox, not on a designated server. To move OAB generation to another server, you need to move the organization mailbox. For example:

- Move the existing organization mailbox to a different Exchange 2013, Exchange 2016, or Exchange 2019 server (you can't move the organization mailbox to an Exchange 2010 server).

- Configure the OAB to use an existing organization mailbox on a different server. For more information, see Use the Exchange Management Shell to change the organization mailbox that's responsible for generating an offline address book.

- Create a new organization mailbox on a different server, and configure the OAB to use that organization mailbox. For more information, see Use the Exchange Management Shell to create organization mailboxes.

Remember, you can configure multiple OABs to use the same organization mailbox, but you can't configure an OAB to use more than one organization mailbox. If you need multiple copies of the OAB in different locations (typically, in different Active Directory sites), verify that an organization mailbox is exists in the site, and enable shadow distribution for the OAB. For more information, see Use the Exchange Management Shell to enable shadow distribution for offline address books.

Procedures for offline address books in Exchange Server

8/3/2020 • 12 minutes to read • Edit Online

An offline address book (OAB) in Exchange Server allows Outlook users in cached Exchange mode to access address list and global address list information while they're disconnected from the server. For more information, see Offline address books in Exchange Server.

Here's a list of OAB procedures that are covered in this topic:

- Use the Exchange Management Shell to view offline address books
- Use the Exchange Management Shell to create offline address books
- Use the Exchange Management Shell to modify offline address books:
  - Use the Exchange Management Shell to configure the default offline address book
  - Use the Exchange Management Shell to add and remove address lists from offline address books
  - Use the Exchange Management Shell to change the organization mailbox that's responsible for generating an offline address book
  - Use the Exchange Management Shell to configure any virtual directory in the organization to accept download requests for the OAB
  - Use the Exchange Management Shell to enable shadow distribution for offline address books
- Use the Exchange Management Shell to update offline address books
- Use the Exchange Management Shell to remove offline address books
- Use the Exchange Management Shell to find organization mailboxes
- Use the Exchange Management Shell to create organization mailboxes
- Assign offline address books to mailbox databases
- Use the Exchange Management Shell to assign offline address books to mailboxes

To change the OAB generation schedule, see Change the offline address book generation schedule in Exchange Server.

What do you need to know before you begin?

- Estimated time to complete each procedure: 5 minutes
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Offline address books" entry in the Email address and address book permissions topic.
- You can't do most of these procedures in the Exchange admin center (EAC). You need to use the Exchange Management Shell. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell. For more information about the EAC, see Exchange admin center in Exchange Server.

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to view offline address books

To view OABs, use the following syntax:

```
Get-OfflineAddressBook [-Identity <OABIdentity>]
```

This example returns a summary list of all OABs in your organization.

```
Get-OfflineAddressBook
```

This example returns detailed information about the OAB named Default Offline Address Book.

```
Get-OfflineAddressBook -Identity "Default Offline Address Book" | Format-List
```

This example returns values for the specified properties on all OABs in your organization.

```
Get-OfflineAddressBook | Format-List
Name,GUID,AddressLists,GeneratingMailbox,IsDefault,VirtualDirectories,GlobalWebDistributionEnabled,ShadowMailb
oxDistributionEnabled
```

For detailed syntax and parameter information, see Get-OfflineAddressBook.

Use the Exchange Management Shell to create offline address books

If you've created multiple address lists, you can use OABs to make the address lists available to users when they're offline.

To create new offline address books, use the following syntax:

```
New-OfflineAddressBook -Name "<Name>" -AddressLists "<GlobalAddressListOrAddressList1>","
<GlobalAddressListOrAddressList2>,..." [-GlobalWebDistributionEnabled $true] [-GeneratingMailbox
<OrganizationMailboxIdentity>] [-IsDefault $true] [-ShadowMailboxDistributionEnabled $true]
```

This example creates a new OAB named Contoso Executives OAB with the following properties:

- The Default Global Address List and Contoso Executives Address List are included in the OAB.

- All OAB virtual directories in the organization can accept requests to download the OAB.

- The organization mailbox that's responsible for generating the OAB is `SystemMailbox{bb558c35-97f1-4cb9-8ff7-d53741dc928c}` (we didn't use the *GeneratingMailbox* parameter to specify a different organization mailbox).

- The OAB isn't used by mailboxes and mailbox databases that don't have an OAB specified (we didn't use the *IsDefault* parameter with the value `$true`).

- Shadow distribution for the OAB is disabled (read-only copies of the OAB aren't copied to all other organization mailboxes, because we didn't use the *ShadowMailboxDistributionEnabled* parameter with the value `$true`).

```
New-OfflineAddressBook -Name "Contoso Executives OAB" -AddressLists "Default Global Address List","Contoso
Executives Address List" -GlobalWebDistributionEnabled $true
```

For detailed syntax and parameter information, see New-OfflineAddressBook.

How do you know this worked?

To verify that you've successfully created the OAB, run the following command to verify the property values:

```
Get-OfflineAddressBook | Format-List
Name,AddressLists,GeneratingMailbox,IsDefault,VirtualDirectories,GlobalWebDistributionEnabled
```

Use the Exchange Management Shell to modify offline address books

To modify OABs, use the following syntax:

```
Set-OfflineAddressBook -Identity "<OABIdentity>" [-Name <Name>] [-AddressLists "
<GlobalAddressListOrAddressList1>","<GlobalAddressListOrAddressList2>,..."] [-VirtualDirectories $null] [-
GlobalWebDistributionEnabled $true] [-GeneratingMailbox <OrganizationMailboxIdentity>] [-IsDefault $true] [-
ShadowMailboxDistributionEnabled <$true | $false>]
```

For detailed syntax and parameter information, see Set-OfflineAddressBook.

Use the Exchange Management Shell to configure the default offline address book

By default, the automatically-created OAB named Default Offline Address Book is the default OAB. The default OAB is used by:

- Mailboxes in mailbox databases where the database has no OAB assigned (by default, all databases)
- Mailboxes without an address book policy (ABP) assigned, or where the assigned ABP policy has no OAB defined (by default, there are no ABPs).
- Mailboxes without an OAB assigned (by default, all mailboxes)

This example configures the OAB named Contoso Executives OAB to be the default OAB.

```
Set-OfflineAddressBook -Identity "Contoso Executives OAB" -IsDefault $true
```

Use the Exchange Management Shell to add and remove address lists from offline address books

When you modify the address lists that are configured in an OAB, the values that you specify will *replace* any address lists in the OAB. To add address lists to the OAB, specify the current address lists plus the ones you want to add. To remove address lists from the OAB, specify the current address lists minus the ones you want to remove.

In this example, the OAB named Marketing OAB is already configured with Address List 1 and Address List 2. To keeps those address lists and add Address List 3, run the following command:

```
Set-OfflineAddressBook -Identity "Marketing OAB" -Address Lists "Address List1","Address List 2","Address List
3"
```

Similarly, to keep the OAB configured with Address List 1 and Address 2, but remove Address List 3, run the following command:

```
Set-OfflineAddressBook -Identity "Marketing OAB" -AddressLists "Address List 1","Address List 2"
```

Use the Exchange Management Shell to change the organization mailbox that's responsible for generating an offline address book

Typically, you only need to configure multiple organization mailboxes if you have Exchange servers in different Active Directory sites. You can configure multiple OABs to use the same organization mailbox, but you can't configure an OAB to use more than one organization mailbox. If you need multiple copies of the OAB in different locations, enable shadow distribution for the OAB. For more information, see the Use the Exchange Management Shell to enable shadow distribution for offline address books section in this topic.

This example changes the organization mailbox that's responsible for generating the OAB named Default Offline Address Book.

```
Set-OfflineAddressBook -Identity "Default Offline Address Book" -GeneratingMailbox OABGen2
```

Note: To configure an arbitration mailbox that you can use as an organization mailbox, see the Use the Exchange Management Shell to create organization mailboxes section in this topic.

Use the Exchange Management Shell to configure any virtual directory in the organization to accept download requests for the OAB

The Client Access (frontend) services on any Mailbox server can proxy the OAB download request to the correct location. The OAB files are downloaded from the backend location `%ExchangeInstallPath%ClientAccess\OAB\<OAB GUID>` on the Mailbox server that holds the active copy of the OAB's designated organization mailbox (or from the server that holds a shadow copy of the OAB).

This example modifies the OAB named Default Offline Address Book to allow any virtual directory in the organization to accept requests to download the OAB.

1. Run the following command:

   ```
   Set-OfflineAddressBook -Identity "Default Offline Address Book" -VirtualDirectories $null
   ```

2. Run the following command:

   ```
   Set-OfflineAddressBook -Identity "Default Offline Address Book" -GlobalWebDistributionEnabled $true
   ```

Use the Exchange Management Shell to enable shadow distribution for offline address books

Before you enable shadow distribution to distribute a read-only copy of the OAB to organization mailboxes in different Active Directory sites, verify that an organization mailbox exists in each site. To create organization mailboxes, see the Use the Exchange Management Shell to create organization mailboxes section in this topic.

This example enables shadow distribution for the OAB named Contoso Executives OAB.

```
Set-OfflineAddressBook -Identity "Contoso Executives OAB" -ShadowMailboxDistributionEnabled $true
```

How do you know this worked?

To verify that you've successfully modified the OAB, run the following command to verify the property values:

```
Get-OfflineAddressBook | Format-List
Name,AddressLists,GeneratingMailbox,IsDefault,VirtualDirectories,GlobalWebDistributionEnabled,
```

Use the Exchange Management Shell to update offline address books

Changes in an OAB aren't available to users until the scheduled OAB generation (by default, every 8 hours). If you don't want to wait, you can use the procedures in this topic to immediately update an OAB.

To change the OAB generation schedule, see Change the offline address book generation schedule in Exchange Server.

To update an OAB, use the following syntax:

```
Update-OfflineAddressBook -Identity <OABIdentity>
```

This example updates the OAB named Default Offline Address Book.

```
Update-OfflineAddressBook -Identity "Default Offline Address Book"
```

This example updates all OABs.

```
Get-OfflineAddressBook | Update-OfflineAddressBook
```

For detailed syntax and parameter information, see Update-OfflineAddressBook.

Use the Exchange Management Shell to remove offline address books

To remove OABs, use the following syntax:

```
Remove-OfflineAddressBook -Identity <OABIdentity>
```

This example removes the OAB named Contoso Executives OAB.

```
Remove-OfflineAddressBook -Identity "Contoso Executives OAB"
```

Note: If the removed OAB is the default OAB, you need to create or configure another OAB as the default (the *IsDefault* parameter value is `$true`).

How do you know this worked?

To verify that you've successfully removed the OAB, run the following command to verify that the OAB is gone.

```
Get-OfflineAddressBook
```

Use the Exchange Management Shell to find organization mailboxes

Only organization mailboxes can generate OABs. An organization mailbox is an arbitration mailbox that has the `OrganizationCapabilityOABGen` value in the **PersistedCapability** property. To find the organization mailboxes in your organization, run the following command:

```
Get-Mailbox -Arbitration | where {$_.PersistedCapabilities -like "*OAB*"} | Format-List
Name,ServerName,PersistedCapabilities
```

To find the organization mailbox that's used to generate an OAB, run the following command:

```
Get-OfflineAddressBook | Format-List Name,AddressLists,GeneratingMailbox,IsDefault
```

Use the Exchange Management Shell to create organization mailboxes

Typically, you only need to create multiple arbitration mailboxes in multi-site Exchange organizations. You can have an organization mailbox in each site, and you can configure shadow distribution for an OAB (so a read only copy of the OAB is stored in all organization mailboxes). For more information, see Use the Exchange Management Shell to enable shadow distribution for offline address books.

You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Recipient Provisioning Permissions" section in the Recipients Permissions topic.

1. Create an arbitration mailbox by using the following syntax:

   ```
   New-Mailbox -Arbitration -Name <UniqueName> -UserPrincipalName <UPN> [-Database <DBIdentity>] [-Alias
   <Alias>] [-DisplayName "<DisplayName>"]
   ```

   This example creates a new arbitration mailbox named OAB Gen 2, with the UPN (account name) oabgen2@contoso.com, in the default database.

   ```
   New-Mailbox -Arbitration -Name "OAB Gen 2" -UserPrincipalName oabgen2@contoso.com
   ```

2. Turn the arbitration mailbox into an organization mailbox by using the following syntax:

   ```
   Set-Mailbox -Identity <MailboxIdentity> -Arbitration -OABGen $true -MaxSendSize 1GB
   ```

   This example turns the OAB Gen 2 arbitration mailbox into an organization mailbox.

   ```
   Set-Mailbox -Identity "OAB Gen 2" -Arbitration -OABGen $true -MaxSendSize 1GB
   ```

3. To activate the OAB generation capabilities of the new organization mailbox, run **Update-OfflineAddressBook** for *any* OAB in the organization. For example:

   ```
   Update-OfflineAddressBook -Identity "Default Offline Address Book"
   ```

How do you know this worked?

To verify that you've successfully created an organization mailbox, run the following command and verify the mailbox is returned:

```
Get-Mailbox -Arbitration | where {$_.PersistedCapabilities -like "*OAB*"} | Format-List
Name,ServerName,PersistedCapabilities
```

Assign offline address books to mailbox databases

When you assign an OAB to a mailbox database, all mailboxes in the databases will use that OAB instead of the default OAB, unless the mailbox has an OAB assigned. By default, no OAB is assigned to a mailbox database.

You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Mailbox databases" entry in the Recipients Permissions topic.

Use the EAC to assign an offline address book to a mailbox database

1. Open the EAC, and go to **Servers** > **Databases**. Select the database from the list, and then click **Edit** (✎).

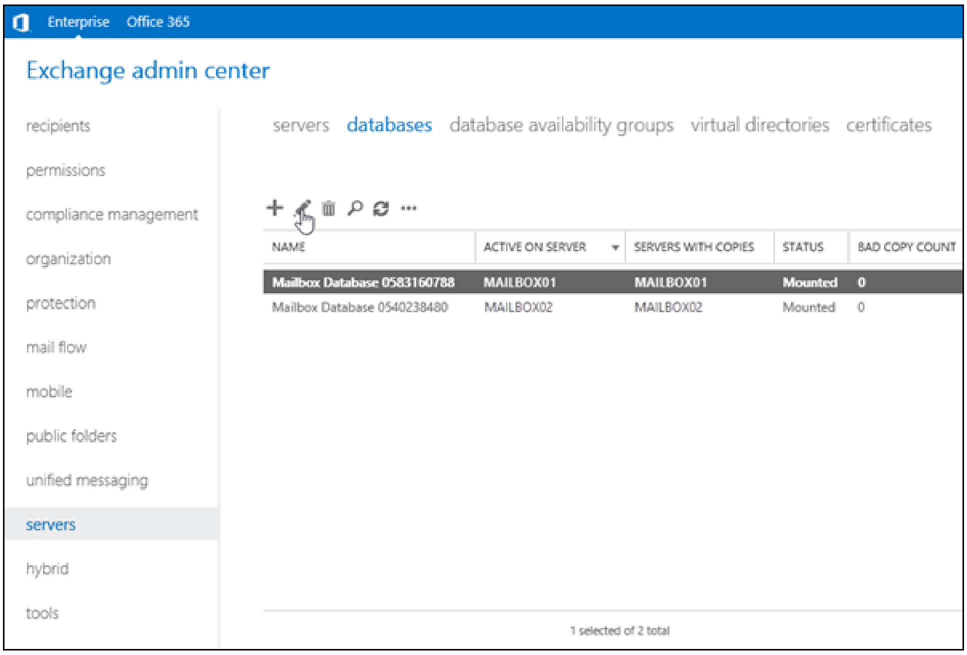

NAME	ACTIVE ON SERVER	SERVERS WITH COPIES	STATUS	BAD COPY COUNT
Mailbox Database 0583160788	MAILBOX01	MAILBOX01	Mounted	0
Mailbox Database 0540238480	MAILBOX02	MAILBOX02	Mounted	0

1 selected of 2 total

2. The **Mailbox Database** window opens. Click the **Client settings** tab, and then click **Browse** next to **Offline address book**.

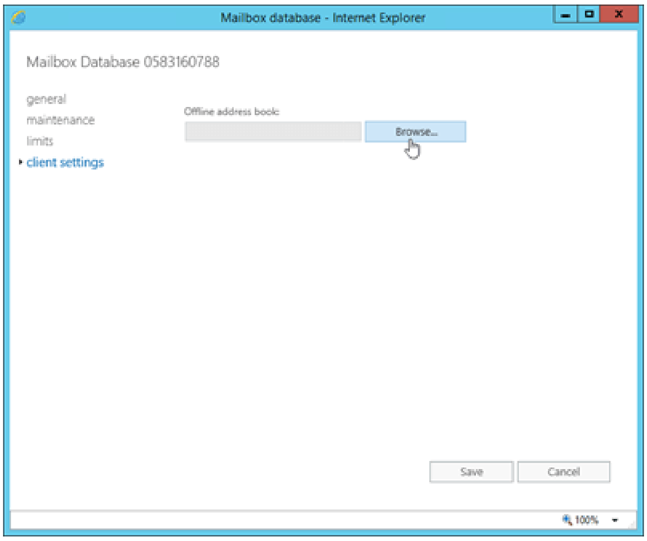

Mailbox Database 0583160788

Offline address book:

3. In the **Select Offline Address Book** window that opens, select the OAB from the list, and click **OK**.

4. Back in the **Mailbox Database** window, click **Save**.

Use the Exchange Management Shell to assign an offline address book to a mailbox database

Use the following syntax:

```
Set-MailboxDatabase -Identity <DatabaseIdentity> -OfflineAddressBook <OABIdentity>
```

This example assigns the OAB named Contoso Executives OAB to the mailbox database named MBX DB02.

```
Set-MailboxDatabase -Identity "MBX DB02" -OfflineAddressBook "Contoso Executives OAB"
```

How do you know this worked?

To verify that you've successfully assigned an OAB to a mailbox database, use either of the following procedures:

- In the EAC, go to **Servers** > **Databases**. Select the database from the list, and then click **Edit** (✎). In the **Mailbox database** window opens, click the **Client settings** tab, and verify that the OAB is listed in **Offline address book**.
- In the Exchange Management Shell, run the following command:

  ```
  Get-MailboxDatabase | Format-Table -Auto Name,OfflineAddressBook
  ```

Use the Exchange Management Shell to assign offline address books to mailboxes

When you assign an OAB to a mailbox, the default OAB and the OAB that's assigned to the mailbox database (if any) aren't used by the mailbox. By default, no OAB is assigned to a mailbox.

Note: If the mailbox has an address book policy (ABP) assigned, and the ABP has an OAB defined, the OAB that's directly assigned to the mailbox will take precedence over the ABP. For more information ABPs, see Address book policies in Exchange Server.

You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Recipient Provisioning Permissions" section in the Recipients Permissions topic.

To assign an OAB to a mailbox, use the following syntax:

```
Set-Mailbox -Identity <MailboxIdentity> -OfflineAddressBook <OABIdentity>
```

This example assigns the OAB named Contoso Executives to the mailbox laura@contoso.com.

```
Set-Mailbox -Identity laura@contoso.com -OfflineAddressBook "Contoso Executives OAB"
```

This example assigns the OAB named Contoso US to a filtered list of mailboxes.

```
$USContoso = Get-User -ResultSize Unlimited -Filter "RecipientType -eq 'UserMailbox' -and Company -eq
'Contoso' -and CountryOrRegion -eq 'US'"; $USContoso | foreach {Set-Mailbox $_.Identity -OfflineAddressBook
"Contoso United States"}
```

How do you know this worked?

To verify that you've successfully assigned an OAB to a mailbox, replace *<MailboxIdentity>* with the identity of the mailbox, and run the following command:

```
Get-Mailbox -Identity "<MailboxIdentity>" | Format-Table -Auto Name,OfflineAddressBook
```

Address lists in Exchange Server

8/3/2020 • 7 minutes to read • Edit Online

An *address list* is a collection of mail-enabled recipient objects from Active Directory. Address lists are based on recipient filters, and are basically unchanged from Exchange 2010. You can filter by recipient type (for example, mailboxes and mail contacts), recipient properties (for example, Company or State or Province), or both. Address lists aren't static; they're updated dynamically. When you create or modify recipients in your organization, they're automatically added to the appropriate address lists. These are the different types of address lists that are available:

- **Global address lists (GALs)**: The built-in GAL that's automatically created by Exchange includes every mail-enabled object in the Active Directory forest. You can create additional GALs to separate users by organization or location, but a user can only see and use one GAL.

- **Address lists**: Address lists are subsets of recipients that are grouped together in one list, which makes them easier to find by users. Exchange comes with several built-in address lists, and you can create more based on you organization's needs.

- **Offline address books (OABs)**: OABs contain address lists and GALs. OABs are used by Outlook clients in cached Exchange mode to provide local access to address lists and GALs for recipient look-ups. For more information, see Offline address books in Exchange Server.

Users in your organization use address lists and the GAL to find recipients for email messages. Here's an example of what address lists look like in Outlook 2016:

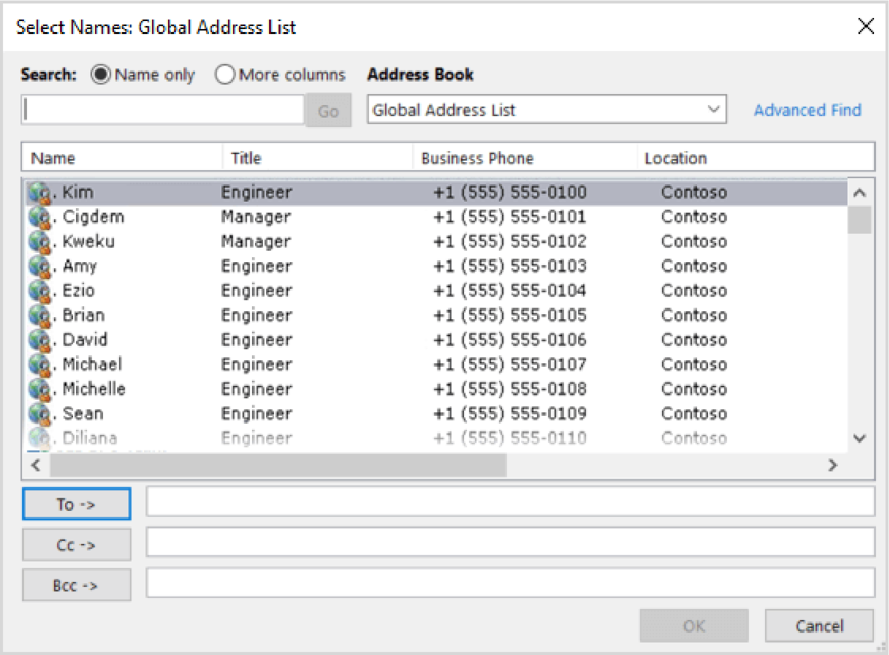

Select Names: Global Address List

Search: ● Name only ○ More columns **Address Book** Global Address List

Name	Title	Business Phone	Location
. Kim	Engineer	+1 (555) 555-0100	Contoso
. Cigdem	Manager	+1 (555) 555-0101	Contoso
. Kweku	Manager	+1 (555) 555-0102	Contoso
. Amy	Engineer	+1 (555) 555-0103	Contoso
. Ezio	Engineer	+1 (555) 555-0104	Contoso
. Brian	Engineer	+1 (555) 555-0105	Contoso
. David	Engineer	+1 (555) 555-0106	Contoso
. Michael	Engineer	+1 (555) 555-0107	Contoso
. Michelle	Engineer	+1 (555) 555-0108	Contoso
. Sean	Engineer	+1 (555) 555-0109	Contoso
. Diliana	Engineer	+1 (555) 555-0110	Contoso

For procedures related to address lists, see Procedures for address lists in Exchange Server.

Recipient filters for address lists

Recipient filters identify the recipients that are included in address lists and GALs. There are two basic options: **precanned recipient filters** and **custom recipient filters**. These are basically the same recipient filtering options that are used by dynamic distribution groups and email address policies. The following table summarizes the differences between the two filtering methods.

RECIPIENT FILTERING METHOD	USER INTERFACE	FILTERABLE RECIPIENT PROPERTIES	FILTER OPERATORS
Precanned recipient filters	**Address lists**: Exchange admin center (EAC) and the Exchange Management Shell **GALs**: Exchange Management Shell only	Limited to: • Recipient type (All recipient types or any combination of user mailboxes, resource mailboxes, mail contacts, mail users, and groups) • Company • Custom Attribute 1 to 15 • Department • State or Province	Property values require an exact match. Wildcards and partial matches aren't supported. For example, "Sales" doesn't match the value "Sales and Marketing". Multiple values of the same property always use the **or** operator. For example, "Department equals Sales or Department equals Marketing". Multiple properties always use the **and** operator. For example, "Department equals Sales and Company equals Contoso".
Custom recipient filters	Exchange Management Shell only	You can use virtually any available recipient attributes. For more information, see Filterable Properties for the -RecipientFilter Parameter.	You use OPATH filter syntax to specify any available Windows PowerShell filter operators. Wildcards and partial matches are supported.

Notes:

- You can't used precanned filters and customized filters at the same time.
- The recipient's location in Active Directory (the organizational unit or container) is available in both precanned and custom recipient filters.
- If an address list uses custom recipient filters instead of precanned filters, you can see the address list in the EAC, but you can't modify or remove it by using the EAC.
- You can hide recipients from all address lists and GALs. For more information, see Hide recipients from address lists.

Global address lists

By default, a new installation of Exchange Server creates an GAL named Default Global Address List that's the primary repository of all recipients in the Exchange organization. Typically, most organizations have only one GAL, because users can only see and use one GAL in Outlook and Outlook on the web (formerly known as Outlook Web App). You might need to create multiple GALs if you want to prevent groups of recipients from seeing each other (for example, you single Exchange organization contains two separate companies). If you plan on creating additional GALs, consider the following issues:

- You can only use the Exchange Management Shell to create, modify, remove, and update GALs.
- The GAL that users see in Outlook and Outlook on the web is named Global Address List, even though the default GAL is named Default Global Address List, and any new GALs that you create will require a unique name (users can't tell which GAL that they're using by name).
- Users can only see a GAL that they belong to (the recipient filter of the GAL includes them). If a user belongs to multiple GALs, they'll still see only one GAL based on the following conditions:

- The user needs permissions to view the GAL. You assign user permissions to GALs by using address book policies (ABPs). For more information, see Address book policies in Exchange Server.
- If a user is still eligible to see multiple GALs, only the largest GAL is used (the GAL that contains the most recipients).
- Each GAL needs a corresponding offline address book (OAB) that includes the GAL. To create OABs, see Use the Exchange Management Shell to create offline address books.

Default address lists

By default, Exchange comes with five built-in address lists and one GAL. These address lists are described in the following table. Note that by default, system-related mailboxes like arbitration mailboxes and public folder mailboxes are hidden from address lists.

NAME	TYPE	DESCRIPTION	RECIPIENT FILTER USED
All Contacts	Address list	Includes all mail contacts in the organization. To learn more about mail contacts, see Recipients.	`"Alias -ne $null -and (ObjectCategory -like 'person' -and ObjectClass -eq 'contact')"`
All Distribution Lists	Address list	Includes all distribution groups and mail-enabled security groups in the organization. To learn more about mail-enabled groups, see Recipients.	`"Alias -ne $null -and ObjectCategory -like 'group'"`
All Rooms	Address list	Includes all room mailboxes. Equipment mailboxes aren't included. To learn more about room and equipment (resource) mailboxes, see Recipients.	`"Alias -ne $null -and (RecipientDisplayType -eq 'ConferenceRoomMailbox' -or RecipientDisplayType -eq 'SyncedConferenceRoomMailbox')"`
All Users	Address list	Includes all user mailboxes, linked mailboxes, remote mailboxes (Microsoft 365 or Office 365 mailboxes), shared mailboxes, room mailboxes, equipment mailboxes, and mail users in the organization. To learn more about these recipient types, see Recipients.	`"((Alias -ne $null) -and (((((((ObjectCategory -like 'person') -and (ObjectClass -eq 'user') -and (-not(Database -ne $null)) -and (-not(ServerLegacyDN -ne $null)))) -or (((ObjectCategory -like 'person') -and (ObjectClass -eq 'user') -and (((Database -ne $null) -or (ServerLegacyDN -ne $null))))))) -and (-not(RecipientTypeDetailsValue -eq 'GroupMailbox')))))"`
Default Global Address List	GAL	Includes all mail-enabled recipient objects in the organization (users, contacts, groups, dynamic distribution groups, and public folders.	`"((Alias -ne $null) -and (((ObjectClass -eq 'user') -or (ObjectClass -eq 'contact') -or (ObjectClass -eq 'msExchSystemMailbox') -or (ObjectClass -eq 'msExchDynamicDistributionList') -or (ObjectClass -eq 'group') -or (ObjectClass -eq 'publicFolder'))))"`

NAME	TYPE	DESCRIPTION	RECIPIENT FILTER USED
Public Folders	Address list	Includes all mail-enabled public folders in your organization. Access permissions determine who can view and use public folders. For more information about public folders, see Public folders.	`"Alias -ne $null -and ObjectCategory -like 'publicFolder'"`

Custom address lists

An Exchange organization might contain thousands of recipients, so the built-in address lists could become quite large. To prevent this, you can create custom address lists to help users find what they're looking for.

For example, consider a company that has two large divisions in one Exchange organization:

- Fourth Coffee, which imports and sells coffee beans.
- Contoso, Ltd, which underwrites insurance policies.

For most day-to-day activities, employees at Fourth Coffee don't communicate with employees at Contoso, Ltd. Therefore, to make it easier for employees to find recipients who exist only in their division, you can create two new custom address lists: one for Fourth Coffee and one for Contoso, Ltd. However, if an employee is unsure about where recipient exists, they can search in the GAL, which contains all recipients from both divisions.

You can also create address lists under other address lists. For example, you can create an address list that contains all recipients in Manchester, and you can create another address list under Manchester named Sales that contains only sales people in the Manchester office. You can also move address lists back to the root, or under other address lists after you've created them. For more information, see Use the Exchange Management Shell to move address lists.

Best practices for creating additional address lists

Although address lists are useful tools for users, poorly planned address lists can cause frustration. To make sure that your address lists are practical for users, consider the following best practices:

- Address lists should make it easier for users to find recipients.
- Avoid creating so many address lists that users can't tell which list to use.
- Use a naming convention and location hierarchy for your address lists so users can immediately tell what the list is for (which recipients are included in the list). If you have difficulty naming your address lists, create fewer lists and remind users that they can find anyone in your organization by using the GAL.

For detailed instructions about creating address lists in Exchange Server, see Create address lists.

Update address lists

After you create or modify an address list, you need to update the membership.

If the address list contains a large number of recipients (our recommendation is more than 3000), you should use the Exchange Management Shell to update the address list (not the EAC). For more information, see Update address lists.

To update a GAL, you always need to use the Exchange Management Shell. For more information, see Use the Exchange Management Shell to update global address lists.

Procedures for address lists in Exchange Server

8/3/2020 • 22 minutes to read • Edit Online

Address lists and global address lists (GALs) are collections of mail-enabled recipient objects from Active Directory. You can create or modify GALs, and update using the tools available in the Exchange admin center (EAC) and the Exchange Management Shell. For more information, see Address lists in Exchange Server.

These are the address list and GAL procedures that you'll find in this topic:

- **Global address list procedures**
  - Use the Exchange Management Shell to update global address lists
  - Use the Exchange Management Shell to view members of global address lists
  - Use the Exchange Management Shell to create global address lists
  - Use the Exchange Management Shell to modify global address lists
  - Use the Exchange Management Shell to remove global address lists
- **Address list procedures**
  - Update address lists
  - View the members of address lists
  - Create address lists
  - Modify address lists
  - Use the Exchange Management Shell to move address lists
  - Remove address lists
  - Hide recipients from address lists

Recipient filters in the EAC

Recipient filters in the Exchange Management Shell

What do you need to know before you begin?

- Estimated time to complete each procedure: 5 minutes.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Address lists" entry in the Email address and address book permissions topic.
- You can do some of the procedures in this topic by using the EAC. For more information about the EAC, see Exchange admin center in Exchange Server. Some procedures require the Exchange Management Shell. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Global address list procedures

All procedures for modifying or updating a GAL require the Exchange Management Shell.

Use the Exchange Management Shell to update global address lists

After you create or modify a GAL, you need to update its membership. Updating a GAL only starts the update process. It may take several hours for the GAL update to be completed.

To update a GAL, use the following syntax:

```
Update-GlobalAddressList -Identity <GALIdentity>
```

This example updates the GAL named Contoso GAL.

```
Update-AddressList -Identity "Contoso GAL"
```

This example updates all GALs in the organization that require updates.

```
Get-GlobalAddressList | where {$_.RecipientFilterApplied -eq $false} | Update-GlobalAddressList
```

For detailed syntax and parameter information, see Update-GlobalAddressList.

How do you know this worked?

To verify that you've successfully updated the GAL, replace *<GALIdentity>* with the name of the address list, and run the following command to verify that the **RecipientFilterApplied** property value is present:

```
Get-AddressList -Identity <GALIdentity> | Format-Table -Auto Name,RecipientFilterApplied
```

Use the Exchange Management Shell to view members of global address lists

- Technically, this procedure returns *all* recipients (including hidden recipients) that match the recipient filters for the GAL. The recipients that are actually visible in the GAL have the **HiddenFromAddressListsEnabled** property value `False`.

- If the GAL isn't up to date (the **RecipientFilterApplied** property has the value `False`), you should update the GAL before you view the members. For more information, see the previous section.

To view the members of a GAL, use the following syntax:

```
$GAL = Get-GlobalAddressList -Identity <GALIdentity>; Get-Recipient -ResultSize unlimited -
RecipientPreviewFilter $GAL.RecipientFilter | select Name,PrimarySmtpAddress,HiddenFromAddressListsEnabled
```

This example returns the members of the GAL named Humongous Insurance.

```
$GAL = Get-GlobalAddressList -Identity "Humongous Insurance"; Get-Recipient -ResultSize unlimited -
RecipientPreviewFilter $GAL.RecipientFilter | select Name,PrimarySmtpAddress,HiddenFromAddressListsEnabled
```

This example exports the results to the file C:\My Documents\Humongous Insurance Export.csv.

```
$GAL = Get-GlobalAddressList -Identity "Humongous Insurance"; Get-Recipient -ResultSize unlimited -
RecipientPreviewFilter $GAL.RecipientFilter | select Name,PrimarySmtpAddress,HiddenFromAddressListsEnabled |
Export-Csv -NoTypeInformation -Path "C:\My Documents\Humongous Insurance Export.csv"
```

Use the Exchange Management Shell to create global address lists

For more information about the requirements and implications of having multiple GALs in your organization, see Global address lists.

For details about recipient filters in the Exchange Management Shell, see the Recipient filters in the Exchange Management Shell section in this topic.

To create a GAL, use the following syntax:

```
New-GlobalAddressList -Name "<GAL Name>" [<Precanned recipient filter | Custom recipient filter>]
```

This example creates a GAL with a precanned recipient filter:

- **Name**: Contoso GAL
- **Precanned recipient filter**: All recipient types where the **Company** value is Contoso.

```
New-GlobalAddressList -Name "Contoso GAL" -IncludedRecipients AllRecipients -ConditionalCompany Contoso
```

This example creates a GAL with a custom recipient filter:

- **Name**: Agency A GAL
- **Custom recipient filter**: All recipient types where the CustomAttribute15 property contains the value AgencyA.

```
New-GlobalAddressList -Name "Agency A GAL" -RecipientFilter "CustomAttribute15 -like '*AgencyA*'"
```

For detailed syntax and parameter information, see New-GlobalAddressList.

How do you know this worked?

To verify that you've successfully created a GAL, use either of the following procedures:

- In the EAC, go to **Organization** > **Address lists**, select the address list, and click **Edit** (✏) to view the details.
- In the Exchange Management Shell, replace *<GAL Name>* with the name of the GAL, and run the following command to verify the property values:

  ```
  Get-GlobalAddressList -Identity "<GAL Name>" | Format-List
  Name,RecipientFilterType,RecipientContainer,RecipientFilter,IncludedRecipients,Conditional*
  ```

Use the Exchange Management Shell to modify global address lists

- The same settings are available as when you created the GAL. For more information, see the previous section.
- After you modify the GAL, you need to update its membership. For more information, see the Use the Exchange Management Shell to update global address lists section in this topic.

- You can't replace a custom recipient filter with a precanned recipient filter or vice-versa in an existing GAL.

To modify a GAL, use the following syntax:

```
Set-GlobalAddressList -Identity <GALIdentity>] [-Name <Name>] [<Precanned recipient filter | Custom recipient
filter>] [-RecipientContainer <OrganizationalUnit>]
```

When you modify the *Conditional* parameter values, you can use the following syntax to add or remove values without affecting other existing values: `@{Add="<Value1>","<Value2>"...; Remove="<Value1>","<Value2>"...}`.

This example modifies the existing GAL named Contoso GAL by adding the **Company** value Fabrikam to the precanned recipient filter.

```
Set-GlobalAddressList -Identity "Contoso GAL" -ConditionalCompany @{Add="Fabrikam"}
```

For detailed syntax and parameter information, see Set-GlobalAddressList.

How do you know this worked?

To verify that you've successfully modified a GAL, use either of the following procedures:

- In the EAC, go to **Organization** > **Address lists**, select the address list, and click **Edit** (✎) to view the details.

- In the Exchange Management Shell, replace *<GAL Name>* with the name of the GAL, and run the following command to verify the property values:

  ```
  Get-GlobalAddressList -Identity "<GAL Name>" | Format-List
  Name,RecipientFilterType,RecipientContainer,RecipientFilter,IncludedRecipients,Conditional*
  ```

Use the Exchange Management Shell to remove global address lists

- You can't remove the GAL named Default Offline Address Book, which is the GAL that's automatically created by Exchange, and the only GAL that has the **IsDefaultGlobalAddressList** property value `True`.

- You can't remove a GAL that's defined in an offline address book (OAB). To modify the address lists that are defined in an OAB, see Use the Exchange Management Shell to add and remove address lists from offline address books.

To remove a GAL, use the following syntax:

```
Remove-GlobalAddressList -Identity <GALIdentity>
```

This example removes the address list named Agency A GAL.

```
Remove-GlobalAddressList -Identity "Agency A GAL"
```

For detailed syntax and parameter information, see Remove-GlobalAddressList.

How do you know this worked?

To verify that you've successfully removed a GAL, use either of the following procedures:

- In the EAC, go to **Organization** > **Address lists**, and verify that the GAL is no longer listed.

- In the Exchange Management Shell, run the following command to verify that the GAL isn't listed:

```
Get-GlobalAddressList
```

Address list procedures

Update address lists

After you create or modify an address list in the EAC or the Exchange Management Shell, you need to update the membership of the address list.

- If the address list contains more than 3000 recipients, we recommend that you use the Exchange Management Shell to update the address list. Updating the membership of the address list will take a long time, and will prevent you from using the EAC session until the address list is fully updated.
- If the address list contains fewer than 3000 recipients, it's OK to use the EAC.

Use the EAC to update address lists

1. In the EAC, go to **Organization** > **Address lists**, and select the address list that you want to update.
   - If the address list needs to be updated, you'll see a **Not up to date** section with an **Update** link in the details pane. Click **Update**.
   - If the address list is already up to date, you'll see **This address list is up to date** in the details pane.
2. After you click **Update**, a warning message that appears. Click **Yes** to update the address list by using the EAC. A progress bar allows you to monitor the update process. When the update is complete, click **Close**.

Use the Exchange Management Shell to update address lists

To update an address list, use the following syntax:

```
Update-AddressList -Identity [<AddressListIdentity>]
```

This example updates the address list named Northwest Executives.

```
Update-AddressList -Identity "Northwest Executives"
```

This example updates the address list named Sales that's located under the address list named North America.

```
Update-AddressList "North America\Sales"
```

This example updates all address lists in the organization that require updates.

```
Get-AddressList | where {$_.RecipientFilterApplied -eq $false} | Update-AddressList
```

For detailed syntax and parameter information, see Update-AddressList.

How do you know this worked?

To verify that you've successfully updated an address list, use either of the following procedures:

- In the EAC, go to **Organization** > **Address lists**, select the address list, and verify that you see **This address list is up to date** (instead of **Not up to date** with an **Update** link) in the details pane.
- In the Exchange Management Shell, replace *<AddressListIdentity>* with the name of the address list, and run the following command to verify the **RecipientFilterApplied** property value:

```
Get-AddressList -Identity <AddressListIdentity> | Format-Table -Auto Name,RecipientFilterApplied
```

View the members of address lists

If the address list isn't up to date, you should update the address list before you view the members. For more information, see the previous section.

Use the EAC to view the members of address lists

1. In the EAC, go to **Organization** > **Address lists**, and select the address list, and then click **Edit** (✎).

2. Click **Preview recipients the address list includes**.

Use the Exchange Management Shell to view members of address lists

- Technically, this procedure returns *all* recipients (including hidden recipients) that match the recipient filters for the address list. The recipients that are actually visible in the address list have the **HiddenFromAddressListsEnabled** property value `False`.

To view the members of an address list, use the following syntax:

```
$AL = Get-AddressList -Identity <AddressListIdentity>; Get-Recipient -ResultSize unlimited -
RecipientPreviewFilter $AL.RecipientFilter | select Name,PrimarySmtpAddress,HiddenFromAddressListsEnabled
```

This example returns the members of the address list named Southeast Offices.

```
$AL = Get-AddressList -Identity "Southeast Offices"; Get-Recipient -ResultSize unlimited -
RecipientPreviewFilter $AL.RecipientFilter | select Name,PrimarySmtpAddress,HiddenFromAddressListsEnabled
```

This example exports the results to the file C:\My Documents\Southeast Offices Export.csv.

```
$AL = Get-AddressList -Identity "Southeast Offices"; Get-Recipient -ResultSize unlimited -
RecipientPreviewFilter $AL.RecipientFilter | select Name,PrimarySmtpAddress,HiddenFromAddressListsEnabled |
Export-Csv -NoTypeInformation -Path "C:\My Documents\Southeast Offices Export.csv"
```

Create address lists

You can create address lists by using the EAC or the Exchange Management Shell. In the EAC, when you create an address list, you're required to include a recipient filter that's based on the recipient type (specific types or all recipients). In the Exchange Management Shell, you aren't required to include a recipient filter that's based on recipient type.

Use the EAC to create address lists

1. In the EAC, go to **Organization** > **Address lists**, and then click **New** (+).

2. In the **Address list** windows that opens, configure the following settings:

   - **Name**: Enter a unique, descriptive name for the address list.

   - **Address list path**: You can create the address list in the root (" ****", also known as All Address Lists), or you can create the address list under an existing address list. To create the address list under an existing address list, click **Browse**, select the address list in the picker window, and then click **OK**.

   - For details about the recipient filters and preview options that are available here, see the Recipient filters in the EAC section in this topic.

3. When you're finished, click **Save**. You'll receive a warning message that tells you to click **Update** in the details pane to update the membership of the address list. For more information, see the Update address lists section in this topic.

Use the Exchange Management Shell to create address lists

You can create address lists with or without recipient filters. For details about recipient filters in the Exchange Management Shell, see the Recipient filters in the Exchange Management Shell section in this topic.

To create an address list, use the following syntax:

```
New-AddressList -Name "<Address List Name>" [-Container <ExistingAddressListPath>] [<Precanned recipient filter | Custom recipient filter>] [-RecipientContainer <OrganizationalUnit>]
```

This example creates an address list with a precanned recipient filter:

- **Name**: Southeast Offices
- **Location**: Under the root ("`\`", also known as All Address Lists) because we didn't use the *Container* parameter, and the default value is "`\`".
- **Precanned recipient filter**: All users with mailboxes where the **State or province** value is GA, AL, or LA (Georgia, Alabama, or Louisiana).

```
New-AddressList -Name "Southeast Offices" -IncludedRecipients MailboxUsers -ConditionalStateorProvince "GA","AL","LA"
```

This example creates an address list with a custom recipient filter:

- **Name**: Northwest Executives
- **Location**: Under the existing address list named North America.
- **Custom recipient filter**: All users with mailboxes where the **Title** value contains Director or Manager, and the **State or province** value is WA, OR, or ID (Washington, Oregon, or Idaho).

```
New-AddressList -Name "Northwest Executives" -Container "\North America"-RecipientFilter "(RecipientType -eq 'UserMailbox') -and (Title -like '*Director*' -or Title -like '*Manager*') -and (StateOrProvince -eq 'WA' -or StateOrProvince -eq 'OR' -or StateOrProvince -eq 'ID')"
```

For detailed syntax and parameter information, see New-AddressList.

How do you know this worked?

To verify that you've successfully created an address list, use either of the following procedures:

- In the EAC, go to **Organization** > **Address lists**, select the address list, and click **Edit** (✎) to view the details.
- In the Exchange Management Shell, replace *[<AddressListPath>*] *<AddressListName>* with the name and (optionally) location of the address list, and run the following command to verify the property values:

  ```
  Get-AddressList -Identity "[<AddressListPath>\]<AddressListName>" | Format-List Name,RecipientFilterType,RecipientContainer,RecipientFilter,IncludedRecipients,Conditional*
  ```

Modify address lists

- If you created an address list with no recipient filters or a custom recipient filter in the Exchange Management Shell, you can't modify the address list in the EAC. You need to use the Exchange Management Shell.
- After you modify an address list, you need to update its membership. For more information, see the Update address lists section in this topic.

- You can't replace a custom recipient filter with a precanned recipient filter or vice-versa in an existing address list.

- You can change the location of an address list by using the **Move-AddressList** cmdlet in the Exchange Management Shell. For more information, see the Use the Exchange Management Shell to move address lists section in this topic.

Modify address lists in the EAC

1. In the EAC, go to **Organization** > **Address lists**, select the address list, and then click **Edit** (✎).

2. In **Address list** windows that opens, configure the following settings:

   - **Display name**: Enter a unique, descriptive name for the address list.

   - For details about the recipient filters and preview options that are available here, see the Recipient filters in the EAC section in this topic.

3. When you're finished, click **Save**. You'll receive a warning message that tells you to click **Update** in the details pane to update the membership of the address list. For more information, see the Update address lists section in this topic.

Modify address lists in the Exchange Management Shell

- The same basic settings are available as when you created the address list. For more information, see the Use the Exchange Management Shell to create address lists section in this topic.

- You can't use this procedure to move an address list. For more information, see the Use the Exchange Management Shell to move address lists section in this topic.

To modify an existing address list, use the following syntax:

```
Set-AddressList -Identity <AddressListIdentity> [-Name <Name>] [<Precanned recipient filter | Custom recipient
filter>] [-RecipientContainer <OrganizationalUnit>]
```

When you modify the *Conditional* parameter values, you can use the following syntax to add or remove values without affecting other existing values: `@{Add="<Value1>","<Value2>"...; Remove="<Value1>","<Value2>"...}`.

This example modifies the existing address list named Southeast Offices by adding the **State or province** value TX (Texas) to the precanned recipient filter.

```
Set-AddressList -Identity "Southeast Offices" -ConditionalStateOrProvince @{Add="TX"}
```

For detailed syntax and parameter information, see Set-AddressList.

How do you know this worked?

To verify that you've successfully modified an address list, use either of the following procedures:

- In the EAC, go to **Organization** > **Address lists**, select the address list, and click **Edit** (✎) to view the details.

- In the Exchange Management Shell, replace *<AddressListIdentity>* with the path\name of the address list, and run the following command to verify the property values:

  ```
  Get-AddressList -Identity "<AddressListIdentity>" | Format-List
  Name,RecipientFilterType,RecipientContainer,RecipientFilter,IncludedRecipients,Conditional*
  ```

Use the Exchange Management Shell to move address lists

You can select the location of an address list when you create an address list in the EAC or the Exchange

Management Shell. But, you can only move an existing address list by using the **Move-AddressList** cmdlet in the Exchange Management Shell. If the source address list contains child address lists under it, the address list hierarchy is moved to the target location that you specify.

To move an address list, use the following syntax:

```
Move-AddressList -Identity "<AddressListIdentity>" -Target "<AddressListIdentity or \>"
```

This example moves the address list named Southeast Offices from the root (" `\` ", also known as All Address Lists) to the address list named North America.

```
Move-AddressList -Identity "Southeast Offices" -Target "North America"
```

For detailed syntax and parameter information, see Move-AddressList.

How do you know this worked?

To verify that you've successfully modified an address list, use either of the following procedures:

- In the EAC, go to **Organization** > **Address lists**, select the address list, and click **Edit** (✎) to view the details.

- In the Exchange Management Shell, replace *<AddressListIdentity>* with the path\name of the address list, and run the following command to verify the property values:

  ```
  Get-AddressList -Identity "<AddressListIdentity>" | Format-List
  Name,RecipientFilterType,RecipientContainer,RecipientFilter,IncludedRecipients,Conditional*
  ```

Remove address lists

If the address list contains more than 3000 recipients, we recommend that you use the Exchange Management Shell to remove the address list. Removing the address list will take a long time, and will prevent you from using the EAC session until the address list is fully removed. If the address list contains less than 3000 recipients, it's OK to use the EAC to remove the address list.

- You can't remove an address list that's defined in an offline address book (OAB). To modify the address lists that are defined in an OAB, see Use the Exchange Management Shell to add and remove address lists from offline address books.

- You can't remove an address list that contains child address lists (you'll receive an error). You first need to do one of the following steps:

  - Use the EAC to remove the parent and all child address lists at the same time.

  - Use the Exchange Management Shell to move all child address lists to another location by using the **Move-AddressList** cmdlet.

Use the EAC to remove address lists

1. In the EAC, go to **Organization** > **Address lists**.

2. Select the address list or lists that you want to remove, and then click **Remove** (🗑). You can select multiple address lists by pressing the CTRL key while selecting each list.

3. Click **Yes** in the warning message that appears. A progress bar allows you to monitor the removal process. When the removal is complete, click **Close**.

Use the Exchange Management Shell to remove address lists

To remove an address list, use the following syntax:

```
Remove-AddressList -Identity "[<AddressListPath>\]<AddressListName>" [-Recursive]
```

This example removes the address list named Southeast Offices and all its children from under the North America address list.

```
Remove-AddressList -Identity "North America\Southeast Offices" -Recursive
```

For detailed syntax and parameter information, see Remove-AddressList.

How do you know this worked?

To verify that you've successfully removed an address list, use either of the following procedures:

- In the EAC, go to **Organization** > **Address lists**, and verify that the address list is no longer listed.

- In the Exchange Management Shell, run the following command to verify that the address list isn't listed:

  ```
  Get-AddressList
  ```

Hide recipients from address lists

Hiding a recipient from address lists doesn't prevent the recipient from receiving email messages; it prevents users from finding the recipient in address lists. The recipient is hidden from **all** address lists and GALs (effectively, they're exceptions to the recipient filters in all address lists). If you want to selectively include the recipient in some address lists but not others, you need to adjust the recipient filters in the address lists to include or exclude the recipient.

Hiding a mailbox from address lists also prevents Outlook from finding the mailbox in GAL when you create a new profile, or add an additional mailbox to an existing profile. To add the hidden mailbox in Outlook, you can temporarily make the mailbox visible in address lists, configure Outlook, and then hide the mailbox from address lists again.

Use the EAC to hide recipients from address lists

1. In the EAC, go to one of the following locations based on the recipient type:

   - **Recipients** > **Mailboxes**: User mailboxes, linked mailboxes, and remote mailboxes.

   - **Recipients** > **Groups**: Distribution groups, mail-enabled security groups, and dynamic distribution groups.

   - **Recipients** > **Resources**: Room and equipment mailboxes.

   - **Recipients** > **Contacts**: Mail users and mail contacts.

   - **Recipients** > **Shared**: Shared mailboxes.

   - **Public folders** > **Public folders**: Mail-enabled public folders.

2. Select the recipient that you want to hide from address lists, and then click **Edit** (✏).

3. The recipient properties window opens. What you do next depends on the recipient type:

   - **Mailboxes, Contacts and Shared**: On the **General** tab, select **Hide from address lists**.

   - **Groups**: On the **General** tab, select **Hide this group from address lists**.

   - **Resources**: On the **General** tab, click **More options**, and then select **Hide from address lists**.

- **Public folders**: On the **General mail properties** tab, select **Hide from Exchange address list**.

  When you're finished, click **Save**.

Use the Exchange Management Shell to hide recipients from address lists

To hide a recipient from address lists, use the following syntax:

```
Set-<RecipientType> -Identity <RecipientIdentity> -HiddenFromAddressListsEnabled $true
```

<RecipientType> is one of these values:

- `DistributionGroup`
- `DynamicDistributionGroup`
- `Mailbox`
- `MailContact`
- `MailPublicFolder`
- `MailUser`
- `RemoteMailbox`

This example hides the distribution group named Internal Affairs from address lists.

```
Set-DistributionGroup -Identity "Internal Affairs" -HiddenFromAddressListsEnabled $true
```

This example hides the mailbox michelle@contoso.com from address lists.

```
Set-Mailbox -Identity michelle@contoso.com -HiddenFromAddressListsEnabled $true
```

Notes:

- To make the recipient visible in address lists again, use the value `$false` for the *HiddenFromAddressListsEnabled* parameter.
- By default, arbitration mailboxes and public folder mailboxes are hidden from address lists. If you use the **Set-Mailbox** cmdlet to change this or any other setting for arbitration or public folder mailboxes, you need to include the *Arbitration* or *PublicFolder* switches.

How do you know this worked?

You can verify that you've successfully hidden a recipient from address lists by using any of the following procedures:

- In the EAC, select the recipient, click **Edit** (✏) and verify the hide from address lists setting is selected.
- In the Exchange Management Shell, run the following command and verify the recipient is listed:

  ```
  Get-Recipient -ResultSize unlimited -Filter "HiddenFromAddressListsEnabled -eq `$true"
  ```

- Open the GAL in Outlook or Outlook on the web (formerly known as Outlook Web App), and verify the recipient isn't visible.

Recipient filters in the EAC

When you create or modify address lists in the EAC, the following recipient filter settings are available:

- **Types of recipients to include**
  - **All recipients**

    Or

  - **Only the following recipient types**: Select one or more of the following values:
  - **Users with Exchange mailboxes**
  - **Mail users with external email addresses**
  - **Resource mailboxes**
  - **Mail contacts with external email addresses**
  - **Mail-enabled groups**
- **Create rules to further define the recipients**

1. Click **Add rule** and select one of the recipient properties from the drop down list:
   - **Recipient container** (container or organization unit)
   - **State or province**
   - **Company**
   - **Department**
   - Custom attribute 1 to 15
2. Enter a value for the property you selected:
   - If you selected **Recipient container**, a **Select an organizational unit** dialog box appears that allows you to select the container or OU in Active Directory.
   - For other recipient properties, a **Specify words or phrases** dialog appears that allows you to add, edit and remove text values.
   - Property values require an exact match. Wildcards and partial matches aren't supported. For example, the value "Sales" doesn't match "Sales and Marketing".
   - Multiple values of the same property use the **or** operator. For example, "Department equals Sales or Department equals Marketing"
3. After you've selected a property and value, click **Add rule**.
4. Repeat the previous steps to configure more filters. Note that multiple properties use the **and** operator. For example, "Department equals Sales and Company equals Contoso".

   Preview recipients the address list includes: When you click this setting, a **Preview** dialog appears that shows you the recipients that are identified by the filters you configured.

Recipient filters in the Exchange Management Shell

In the Exchange Management Shell, you can specify **precanned recipient filters**, or **custom recipient filters**, but not both at the same time.

- **Precanned recipient filters**

  - Uses the required *IncludedRecipient* parameter with the `AllRecipients` value *or* one or more of the following values: `MailboxUsers`, `MailContacts`, `MailGroups`, `MailUsers`, or `Resources`. You can specify multiple values separated by commas.

  - You can also use any of the optional *Conditional* filter parameters: *ConditionalCompany*, *ConditionalCustomAttribute[1to15]*, *ConditionalDepartment*, and *ConditionalStateOrProvince*.

    You specify multiple values for a *Conditional* parameter by using the syntax `"<Value1>","<Value2>"...`. Multiple values of the same property implies the **or** operator. For example, "Department equals Sales or Marketing or Finance".

- **Custom recipient filters**: Uses the required *RecipientFilter* parameter with an OPATH filter.

  - The basic OPATH filter syntax is `"<Property1> -<Operator> '<Value1>' <Property2> -<Operator> '<Value2>'..."`.

  - Double quotation marks `" "` are required around the whole OPATH filter. Although the filter is a string (not a system block), you can also use braces `{ }`, but only if the filter doesn't contain variables that require expansion..

  - Hyphens (`-`) are required before all operators. Here are some of the most frequently used operators:

  - `and`, `or`, and `not`.

  - `eq` and `ne` (equals and does not equal; not case-sensitive).

  - `lt` and `gt` (less than and greater than).

  - `like` and `notlike` (string contains and does not contain; requires at least one wildcard in the string. For example, `"Department -like 'Sales*'"`.

  - Use parentheses to group `<Property> -<Operator> '<Value>'` statements together in complex filters. For example,

    ```
    "(Department -like 'Sales*' -or Department -like 'Marketing*') -and (Company -eq 'Contoso' -or
    Company -eq 'Fabrikam')"
    ```

    . Exchange stores the filter in the **RecipientFilter** property with each individual statement enclosed in parentheses, but you don't need to enter them that way.

  - For more information, see Additional OPATH syntax information.

  - After you use the **New-AddressList** cmdlet to create an address list that uses custom recipient filters, you can't modify the address list in the EAC. You need to use the **Set-AddressList** cmdlet with the *RecipientFilter* parameter in the Exchange Management Shell.

Note: The *RecipientContainer* (organizational unit) recipient filter parameter is available to both precanned recipient filters and custom recipient filters.

Address book policies in Exchange Server

8/3/2020 • 3 minutes to read • Edit Online

Address book policies (ABPs) lets administrators segment users into specific groups to provide customized views of the organization's global address list (GAL). The goal of an ABP is to provide a simpler mechanism for GAL segmentation (also known as *GAL segregation*) in on-premises organizations that require multiple GALs.

An ABP contains these elements:

- One GAL. For more information about GALs, see Global address lists.
- One offline address book (OAB). For more information about OABs, see Offline address books in Exchange Server.
- One room list. Note that this room list is a custom address list that specifies rooms (contains the filter `RecipientDisplayType -eq 'ConferenceRoomMailbox'`). It's not a room finder that you create with the *RoomList* switch on the **New-DistributionGroup** or **Set-DistributionGroup** cmdlet. For more information, see Create and manage room mailboxes.
- One or more address lists. For more information about address lists, see Custom address lists.

For procedures involving ABPs, see Procedures for address book policies in Exchange Server.

Notes:

- ABPs create only a virtual separation of users from a directory perspective, not a legal separation.
- Implementing an ABP is a multi-step process that requires planning. For more information, see Scenario: Deploying address book policies in Exchange Server.

How ABPs work

The following diagram shows how ABPs work. The user is assigned Address Book Policy A that contains a subset of address lists that are available in the organization. When the ABP is created and assigned to the user, the ABP becomes the scope of the address lists that the user is able to view.

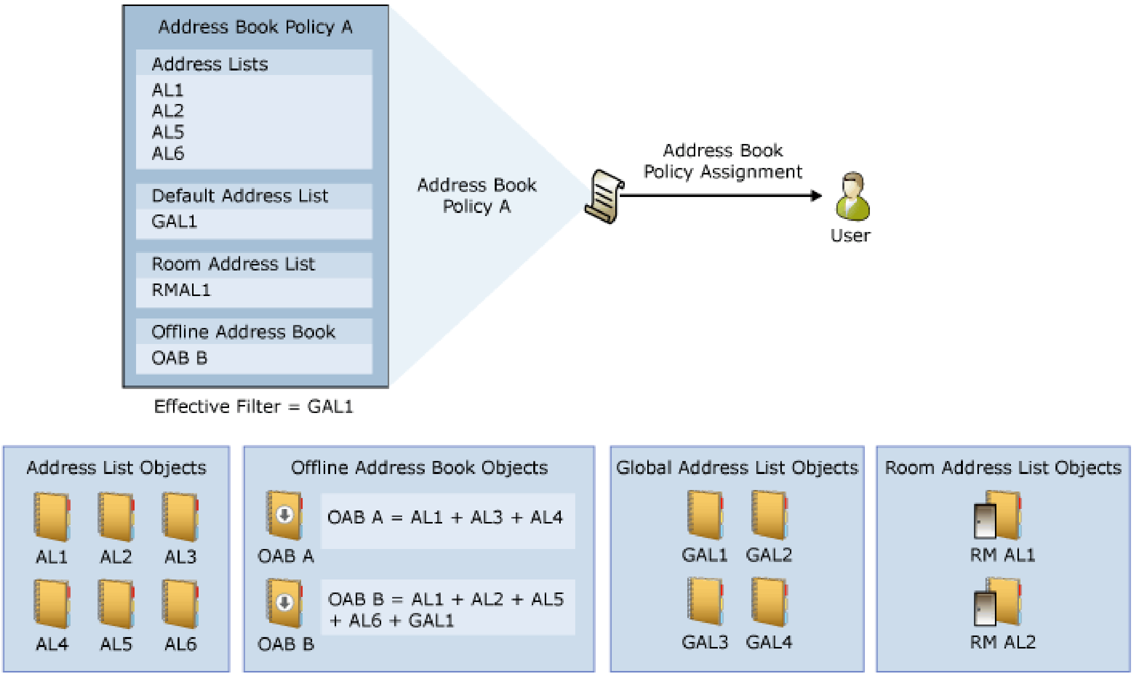

APBs take effect when a user connects to the Client Access (frontend) services on a Mailbox server. If you change an ABP, the updated APB takes effect when a user restarts or reconnects their client app, or you restart the Mailbox server (specifically, the Microsoft Exchange RPC Client Access service in the backend services).

Address Book Policy Routing agent

In an Exchange organization that doesn't use ABPs, the following things occur when a user creates an email message in Outlook or Outlook on the web and sends the message to another recipient in the organization:

1. The email address resolves to the user's display name. For example, if you type sardor@contoso.com in the **To** field, the SMTP email address resolves to **Sarah Dorsey**.

2. After the name resolves, you can view the recipient's contact card by double-clicking on the user's name. The contact card shows the recipient's contact information, such as office and phone number.

If you're using ABPs, and you don't want the users in the ABPs to view each other's potentially private information, you can turn on the Address Book Policy Routing agent. The ABP Routing agent is a Transport agent that controls how recipients are resolved in your organization. When the ABP Routing agent is installed and configured, users that are assigned to different GALs by different ABPs can't view each other's contact cards (they appear as external recipients to each other).

For details about how to turn on the ABP Routing agent, see Use the Exchange Management Shell to install and configure the Address Book Policy Routing Agent.

ABP example

In the following diagram, Fabrikam and Tailspin Toys share the same Exchange organization and the same CEO. The CEO is the only employee common to both companies.

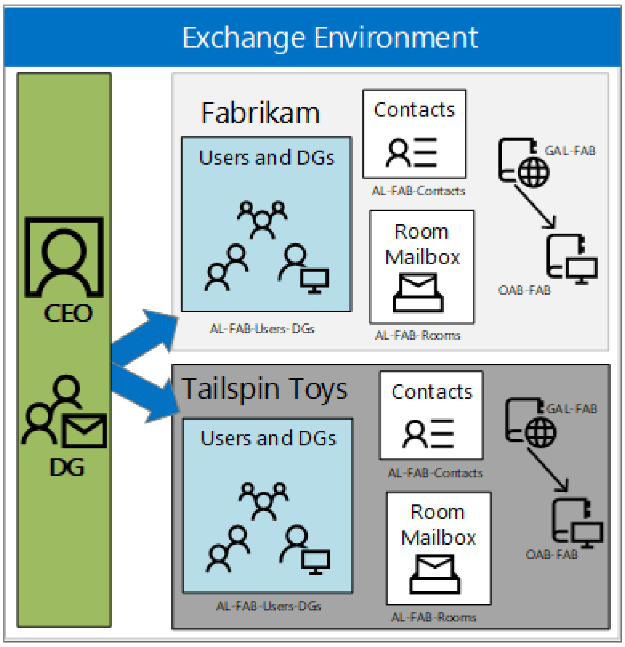

The suggested configuration includes three ABPs:

- One ABP is assigned to Fabrikam employees. The GAL and address lists in the ABP include Fabrikam employees and the CEO.

- One ABP is assigned to Tailspin Toys employees. The GAL and address lists in the ABP include Tailspin Toys employees and the CEO.

- One ABP is assigned to only the CEO. The (default) GAL and address lists in the ABP include all employees (Fabrikam, Tailspin Toys, and the CEO).

Based on this configuration, the ABPs help to enforce these requirements:

- The users in Tailspin Toys can only see Tailspin Toys employees and the CEO when they browse the GAL.
- The users in Fabrikam can only see Fabrikam employees and the CEO when they browse the GAL.
- The CEO can see all Fabrikam and Tailspin Toys employees when she browses the GAL.
- Users who view the CEO's group membership can see only groups that belong to their company. They can't see groups that belong to the other company.

Scenario: Deploying address book policies in Exchange Server

8/3/2020 • 13 minutes to read • Edit Online

The scenarios in this topic describe the deployment solutions for address book policies (ABPs) in three of the most common organization types where multiple entities (companies, government agencies, school classrooms, etc.) share a common Exchange environment. In all scenarios, a recipient filter divides recipients into separate virtual organizations, which then defines the ABPs that are applied to users in those virtual organizations. For more information recipient filters and virtual organizations, see the Considerations and best practices for address book policies section later in this topic.

For more information about ABPs, see Address book policies in Exchange Server. For ABP procedures, see Procedures for address book policies in Exchange Server.

Scenario 1: Two separate companies in one Exchange organization

This scenario applies to companies or divisions that share the same Exchange environment, but have no common employees or management. In addition, the divisions have no special security or privacy concerns.

In this scenario, Contoso and Humongous Insurance are two separate companies that share the same Exchange environment. An ABP for each company lets employees in one company see only members of the same company in the global address list (GAL) in Outlook and Outlook on the web (formerly known as Outlook Web App). All distribution groups belong to one company or the other, and no distribution group contains members from both companies.

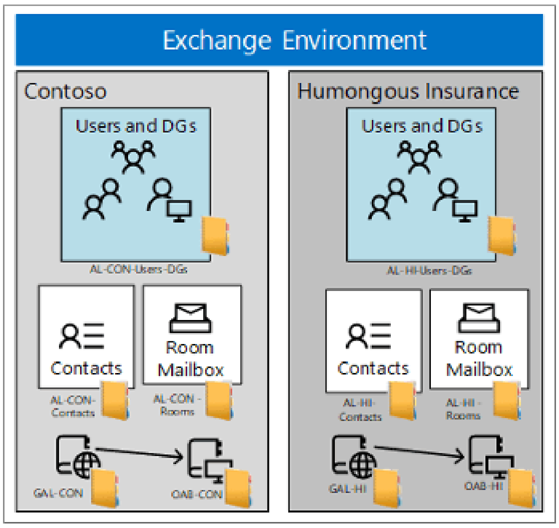

The GAL, offline address book (OAB), room list, and address lists that are required inn the ABPs for this scenario are described in the this table:

ABP ELEMENT	CONTOSO	HUMONGOUS INSURANCE
Global address list	GAL_CON	GAL_HI
Offline address book	OAB_CON	OAB_HI
Room list	AL_CON_Rooms	AL_HI_Rooms

ABP ELEMENT	CONTOSO	HUMONGOUS INSURANCE
Address Lists	AL_CON_Groups AL_CON_Users AL_CON_Contacts	AL_HI_Groups AL_HI_Users AL_HI_Contacts

Scenario 2: Two companies sharing a CEO in one Exchange organization

This scenario applies to companies or divisions that share Exchange environment, and the only employees in common are in upper management.

In this scenario, Fabrikam and Tailspin Toys are separate companies in the same Exchange environment that share the same CEO, who is the only person in common between the two companies. This scenario uses three ABPs that have the following requirements:

- Employees in one company can only see recipients in their company when they browse the GAL, and employees in both companies can see the CEO in the GAL and in distribution groups.
- The CEO can see all recipients in both companies, is able to create distribution groups that span both companies, and the groups are visible in each company's GAL. However, group members only see other members from their respective company (group members from the other company are hidden).
- Employees who look at the CEO's group membership will only see groups in their company. They won't see groups in the other company.
- Each company has a distribution group named Senior Leadership that includes the management of that company and the CEO.
- The names of the three ABPs are: ABP_FAB, ABP_TAIL, and ABP_CEO.

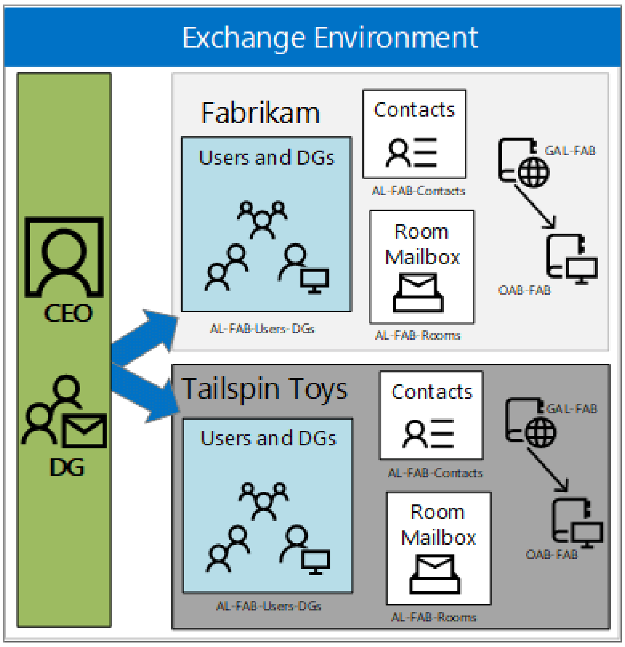

The GAL, OAB, room list, and address lists that are required in the ABPs for this scenario are described in the this table:

ABP ELEMENT	FABRIKAM	TAILSPIN TOYS	CEO
Name	ABP_FAB	AB_TAIL	ABP_CEO

ABP ELEMENT	FABRIKAM	TAILSPIN TOYS	CEO
Global address list	GAL_FAB	GAL_TAIL	Default Global Address Book
Offline address book	OAB_FAB	OAB_TAIL	Default Offline Address Book
Room address list	AL_FAB_Rooms	AL_TAIL_Rooms	All Rooms
Address lists	AL_FAB_Users_DGs AL_FAB_Contacts	AL_TAIL_Users_DGs AL_TAIL_Contacts	AL_FAB_Users_DGs AL_FAB_Contacts AL_TAIL_Users_DGs AL_TAIL_Contacts

For a complete walkthrough of creating the required elements for this scenario, see the Detailed deployment steps for Scenario 2: Two companies sharing a CEO in one Exchange organization section at the end of this topic.

Scenario 3: Education

This scenario is applicable to schools or universities where a division of class rooms is necessary to ensure the privacy of the students, and has the following requirements:

- Students in each class can only see other students in their class, their teacher, and the principal.
- Teachers can only see students in their own classes.
- Teachers can see the principal and all other teachers.
- Distribution groups are created for the parents and faculty that are associated with each class.

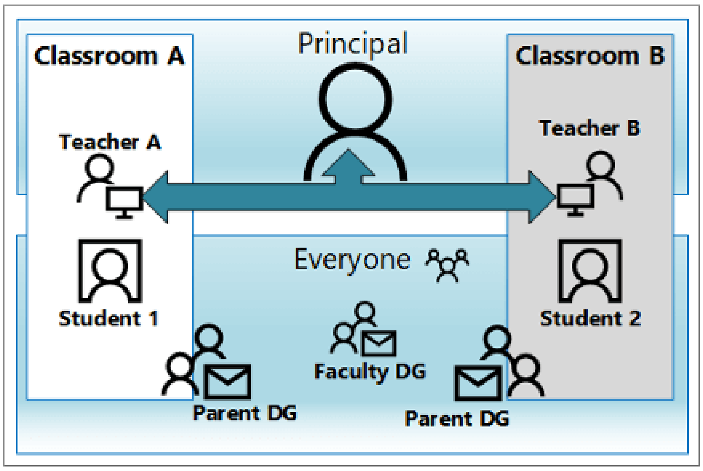

The GAL, OAB, room list, and address lists that are required in the ABPs for this scenario are described in the this table:

ABP ELEMENT	STUDENTS_CLASSA	TEACHERS_CLASSA	PRINCIPAL
Global address list	GAL_StudentsClassA	GAL_TeachersClassA	GAL_Everyone
Offline address book	OAB_StudentsClassA	OAB_TeachersClassA	Default Offline Address Book
Room address list	AL_BlankRoom	AL_BlankRoom	All Rooms

ABP ELEMENT	STUDENTS_CLASSA	TEACHERS_CLASSA	PRINCIPAL
Address Lists	AL_ClassAAL_Principal	AL_ClassAAL_AllTeachersAL_AllGroupsAL_Principal	AL_ClassA AL_ClassB AL_AllTeachers AL_AllStudents AL_AllGroups

Considerations and best practices for address book policies

These are the important issues to consider when you use ABPs in your organization:

- You can't use hierarchical address books (HABs) and ABPs simultaneously. To learn more, see Understanding Hierarchical Address Books.

- A user that's assigned an ABP needs to exist in the GAL that's specified for the ABP.

- If you create ABPs in your organization and don't assign an ABP to some users, those recipients can see *all* address lists.

- To divide recipients into virtual organizations, we recommend using the **CustomAttribute1** to **CustomAttribute15** attributes on recipients. These attributes work better than the other pre-canned conditional attributes such as **Company**, **Department**, or **StateOrProvince** because:

  - Not all recipient types support the **Company**, **Department** or **StateOrProvince** attributes (for example, distribution groups, dynamic distribution groups, and mail-enabled public folders).

  - The **CustomAttribute1** to **CustomAttribute15** attributes aren't configurable by users on their own mailboxes, and are entirely under the control of administrators.

  - Even recipient types that support the **Company**, **Department** or **StateOrProvince** attributes require different cmdlets to configure them.

    For example, to configure values for **Company**, **Department** or **StateOrProvince** on mailboxes, mail users, or mail contacts, you can't use the **Set-Mailbox**, **Set-MailUser**, or **Set-MailContact** cmdlets. Instead, you need to use the **Set-User** and **Set-Contact** cmdlets. In contrast, the *CustomAttribute1* to *CustomAttribute15* parameters are available on the corresponding **Set-*** cmdlets for all recipient types.

    For more information about recipient filtering, see Recipient filtering on Edge Transport servers.

- Client applications that access Active Directory directly through LDAP will bypass the logic that's built into ABPs.

- At a minimum, the GAL that's specified in an ABP must contain all address lists (including the room address list) that are specified in the ABP (it's OK if the ABP contains additional address lists). Don't create a GAL that contains fewer recipients than the address lists in the same ABP.

- We recommend against creating distribution groups that cross virtual organization boundaries. Groups that contain members of multiple virtual organizations lead to these issues:

  - A group member will see the email addresses of all group members if they request a delivery receipt or a read receipt when they send a message to the distribution group.

  - Encrypted messages that are sent to the distribution group can cause issues when some group members don't have valid digital IDs. For example, suppose a distribution group contains three members from Agency A, and two members from Agency B. Furthermore, one of the members from Agency A and two of the members in Agency B have invalid digital IDs. If a member from Agency A sends an encrypted messages to the distribution group, they'll receive a warning that there are three recipients without valid digital IDs. However, only the email address for the member in Agency A will appear in the warning message.

  - ABPs don't apply to all users or processes that use the **Get-Group** cmdlet, so these users will see all members of any group that they have access to.

    Because if this issue, we recommend that you prevent users from managing their own groups in Outlook or Outlook on the web. To do this, remove the MyDistributionGroupMembership RBAC role assignment from the users. For more information, see Manage role assignment policies.

    If you allow users to use Outlook or Outlook on the web to manage groups, visibility to the full group membership list must be OK for the group owners.

- All ABPs must contain a room address list. However, if your organization doesn't use room address lists, you can create an empty room address list.

  Note: The room list that's required for an ABP is an address list that specifies rooms (contains the filter `RecipientDisplayType -eq 'ConferenceRoomMailbox'`). It's not a room finder distribution group that you create with the *RoomList* switch on the **New-DistributionGroup** or **Set-DistributionGroup** cmdlets. For more information, see Create and manage room mailboxes.

- Deploying ABPs doesn't prevent users in one virtual organization from sending email to users in another virtual organization. If you want to prevent users from sending email across virtual organizations, we recommend that you create a mail flow rule (also known as a transport rule) that looks for messages sent between the recipients. For example, to prevent Contoso users from receiving messages from Fabrikam users and vice-versa, but still allow Fabrikam's senior leadership team to send messages to Contoso users, you can create the following mail flow rule in the Exchange Management Shell:

  ```
  New-TransportRule -Name "Ethical Wall: Contoso-Fabrikam" -BetweenMemberOf1 "AllFabrikamEmployees" -
  BetweenMemberOf2 "AllContosoEmployees" -DeleteMessage -ExceptIfFrom seniorleadership@fabrikam.com
  ```

  For more information about mail flow rules, see Mail flow rules in Exchange Server.

- To configure a feature that's similar to address book policies in the Skype for Business or Lync client, you can set the **msRTCSIP-GroupingID** attribute for specific users. For details, see PartitionByOU Replaced with msRTCSIP-GroupingID.

Detailed deployment steps for Scenario 2: Two companies sharing a CEO in one Exchange organization

This section walks you through the deployment steps for Scenario 2: Two companies sharing a CEO in one Exchange organization. If you recall, Fabrikam and Tailspin Toys are separate companies that share the same CEO.

To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.

Step 1: Install and configure the Address Book Policy Routing Agent

The ABP Routing Agent makes users that are assigned different GALs appear as external recipients to each other. For detailed instructions, see Use the Exchange Management Shell to install and configure the Address Book Policy Routing Agent.

Step 2: Define your virtual organizations

In this scenario, the **CustomAttribute15** attribute defines the virtual organizations: the value `FAB` for Fabrikam recipients, the value `TAIL` for Tailspin Toys recipients, and the value `CEO` for the CEO, which is required so Fabrikam and Tailspin users can see the CEO. If you don't include the CEO in the Fabrikam and Tailspin Toys virtual organizations, the CEO can see everyone, but no one can see the CEO. For more information about recipient filtering, see Recipient filtering on Edge Transport servers.

To set the **CustomAttribute15** attribute value for the Fabrikam and Tailspin Toys mailboxes, distribution groups, dynamic distribution groups, mail contacts, and mail users, use the following syntax:

```
$<VariableName> = Get-<RecipientType> -ResultSize Unlimited | where PrimarySMTPAddress -match <fabrikam.com |
tailspintoys.com>
$<VariableName> | foreach {Set-<RecipientType> -Identity ($_.GUID).ToString() -CustomAttribute15 <FAB | TAIL>
```

Notes:

- Using the recipient's GUID value for the *Identity* parameter can help avoid collisions if there are similar usernames in both organizations (for example, julia@fabrikam.com and julia@contoso.com).
- The valid <RecipientType> values for the cmdlet names are Mailbox, DistributionGroup, DynamicDistributionGroup, MailContact, and MailUser. You need to configure the **CustomAttribute15** attribute value for each recipient type separately.

This example sets the value `FAB` for the **CustomAttribute15** attribute on all Fabrikam mailboxes.

```
$FAB_MBX = Get-Mailbox -ResultSize Unlimited | where PrimarySMTPAddress -match fabrikam.com
$FAB_MBX | foreach {Set-Mailbox -Identity ($_.GUID).ToString() -CustomAttribute15 FAB}
```

Step 3: Create the required elements for the address book policies

Create address lists

This organization requires four custom address lists:

- AL_FAB_Users_DGs
- AL_FAB_Contacts
- AL_TAIL_Users_DGs
- AL_TAIL_Contacts

This example creates the address list named AL_FAB_Users_DGs that contains all Fabrikam users, distribution groups, and dynamic distribution groups *and* the CEO.

```
New-AddressList -Name "AL_FAB_Users_DGs" -RecipientFilter "((RecipientType -eq 'UserMailbox') -or
(RecipientType -eq 'MailUniversalDistributionGroup') -or (RecipientType -eq 'DynamicDistributionGroup')) -and
(CustomAttribute15 -eq 'FAB') -or (CustomAttribute15 -eq 'CEO')"
```

This example creates the address list named AL_FAB_Contacts that contains all Fabrikam mail contacts.

```
New-AddressList -Name "AL_FAB_Contacts" -RecipientFilter "(RecipientType -eq 'MailContact') -and
(CustomAttribute15 -eq 'FAB')"
```

This example creates the address list named AL_TAIL_Users_DGs that contains all Tailspin Toys users, distribution groups, and dynamic distribution groups *and* the CEO.

```
New-AddressList -Name "AL_TAIL_Users_DGs" -RecipientFilter "((RecipientType -eq 'UserMailbox') -or
(RecipientType -eq 'MailUniversalDistributionGroup') -or (RecipientType -eq 'DynamicDistributionGroup')) -and
(CustomAttribute15 -eq 'TAIL') -or (CustomAttribute15 -eq 'CEO')"
```

This example creates the address list named AL_TAIL_Contacts that contains all Tailspin Toys mail contacts.

```
New-AddressList -Name "AL_TAIL_Contacts" -RecipientFilter "(RecipientType -eq 'MailContact') -and
(CustomAttribute15 -eq 'TAIL')"
```

For more information, see Create address lists.

Create room lists

This organization requires two custom room lists:

- AL_FAB_Rooms
- AL_TAIL_Rooms

This example creates the room list named AL_FAB_Rooms for Fabrikam room mailboxes.

```
New-AddressList -Name AL_FAB_Rooms -RecipientFilter "(Alias -ne $null) -and (CustomAttribute15 -eq 'FAB') -and
(RecipientDisplayType -eq 'ConferenceRoomMailbox') -or (RecipientDisplayType -eq
'SyncedConferenceRoomMailbox')"
```

This example creates a room list named AL_TAIL_Rooms for Tailspin Toys room mailboxes.

```
New-AddressList -Name AL_TAIL_Rooms -RecipientFilter "(Alias -ne $null) -and (CustomAttribute15 -eq 'TAIL') -
and (RecipientDisplayType -eq 'ConferenceRoomMailbox') -or (RecipientDisplayType -eq
'SyncedConferenceRoomMailbox')"
```

Note: This example creates a blank room list named AL_BlankRoom if the organization doesn't have any room mailboxes (an ABP requires a room list, even if it's empty):

```
New-AddressList -Name AL_BlankRoom -RecipientFilter "(Alias -ne $null) -and ((RecipientDisplayType -eq
'ConferenceRoomMailbox') -or (RecipientDisplayType -eq 'SyncedConferenceRoomMailbox'))"
```

For more information about creating address lists, see Create address lists.

Create GALs

This organization requires two custom GALs:

- GAL_FAB
- GAL_TAIL

This example creates the GAL named GAL_FAB for Fabrikam that includes all Fabrikam recipients *and* allows the Fabrikam users to see the CEO.

```
New-GlobalAddressList -Name "GAL_FAB" -RecipientFilter "(CustomAttribute15 -eq 'FAB') -or (CustomAttribute15 -
eq 'CEO')"
```

This example creates the GAL named GAL_TAIL for Tailspin Toys that includes all Tailspin Toys recipients *and* allows the Tailspin Toys users to see the CEO.

```
New-GlobalAddressList -Name "GAL_TAIL" -RecipientFilter "(CustomAttribute15 -eq 'TAIL') -or (CustomAttribute15
-eq 'CEO')"
```

Note: Don't use a GAL in an ABP that contains recipients that are missing from address lists in the ABP. The combination of all address lists must match the recipients in the GAL.

For more information, see Use the Exchange Management Shell to create global address lists.

Create OABs

This organization requires two custom GALs:

- OAB_FAB
- OAB_TAIL

This example creates the OAB named OAB_FAB for Fabrikam that includes the Fabrikam GAL.

```
New-OfflineAddressBook -Name "OAB_FAB" -AddressLists "GAL_FAB"
```

This example creates the OAB named OAB_TAIL for Tailspin Toys that includes the Tailspin Toys GAL.

```
New-OfflineAddressBook -Name "OAB_TAIL" -AddressLists "GAL_TAIL"
```

Note: If you want users to see all recipients in the virtual organization, make sure that you include the GAL in OAB. Otherwise, you can reduce the download size of the OAB by specifying a reduced list of address lists that are included in the OAB.

For more information, see Use the Exchange Management Shell to create offline address books.

Step 4: Create the address book policies

This organization requires three ABPs:

ABP ELEMENT	FABRIKAM	TAILSPIN TOYS	CEO
Name	ABP_FAB	ABP_TAIL	ABP_CEO
Global address list	GAL_FAB	GAL_TAIL	Default Global Address Book
Offline address book	OAB_FAB	OAB_TAIL	Default Offline Address Book
Room address list	AL_FAB_Rooms	AL_TAIL_Rooms	All Rooms
Address lists	AL_FAB_Users_DGs AL_FAB_Contacts	AL_TAIL_Users_DGs AL_TAIL_Contacts	AL_FAB_Users_DGs AL_FAB_Contacts AL_TAIL_Users_DGs AL_TAIL_Contacts

This example creates the ABP named ABP_FAB that contains the GAL, OAB, room list and address lists for Fabrikam.

```
New-AddressBookPolicy -Name "ABP_FAB" -AddressLists "AL_FAB_Users_DGs","AL_FAB_Contacts" -OfflineAddressBook
"\OAB_FAB" -GlobalAddressList "\GAL_FAB" -RoomList "\AL_FAB_Rooms"
```

This example creates the ABP named ABP_TAIL that contains the GAL, OAB, room list and address lists for Tailspin Toys.

```
New-AddressBookPolicy -Name "ABP_TAIL" -AddressLists "AL_TAIL_Users_DGs","AL_TAIL_Contacts" -
OfflineAddressBook "\OAB_TAIL" -GlobalAddressList "\GAL_TAIL" -RoomList "\AL_TAIL_Rooms"
```

This example creates the ABP named ABP_CEO that contains the GAL, OAB, room list and address lists for the CEO.

```
New-AddressBookPolicy -Name "ABP_CEO" -AddressLists
"AL_FAB_Users_DGs","AL_FAB_Contacts","AL_TAIL_Users_DGs","AL_TAIL_Contacts" -OfflineAddressBook "\Default
Offline Address Book" -GlobalAddressList "\Default Global Address List" -RoomList "\All Rooms"
```

For more information, see Procedures for address book policies in Exchange Server.

Step 5: Assign the address book policies to mailboxes

This example assigns the ABP named ABP_FAB to all Fabrikam mailboxes.

```
$Fab = Get-Mailbox -ResultSize unlimited -Filter "CustomAttribute15 -eq 'FAB'"; $Fab | foreach {Set-Mailbox -
Identity $_.Identity -AddressBookPolicy 'ABP_FAB'}
```

This example assigns the ABP named ABP_TAIL to all Tailspin Toys mailboxes.

```
$Tail = Get-Mailbox -ResultSize unlimited -Filter "CustomAttribute15 -eq 'TAIL'"; $Tail | foreach {Set-Mailbox
-Identity $_.Identity -AddressBookPolicy 'ABP_TAIL'}
```

This example assigns the ABP named ABP_CEO to the CEO named Gabriela Laureano.

```
Set-Mailbox -Identity "Gabriela Laureano" -AddressBookPolicy "ABP_CEO"
```

Note: If the user is already connected to Outlook or Outlook on the web when the ABP is applied to their mailbox, they'll need to close and restart their client application before they can see the new address lists and GAL.

For more information, see Assign address book policies to mailboxes.

Other considerations

After you create or modify an address list or GAL, you need to update the membership.

If the address list contains a large number of recipients (our recommendation is more than 3000), you should use the Exchange Management Shell to update the address list (not the Exchange admin center). For more information, see Update address lists.

To update a GAL, you always need to use the Exchange Management Shell. For more information, see Use the Exchange Management Shell to update global address lists.

Procedures for address book policies in Exchange Server

8/3/2020 • 12 minutes to read • Edit Online

Address book policies (ABPs) allow you to segment users into specific groups to give them customized global address lists (GALs) in Outlook and Outlook on the web (formerly known as Outlook Web App). For more information about ABPs, see Address book policies in Exchange Server.

Note: Implementing an ABP is a multi-step process that requires planning. For more information, see Scenario: Deploying address book policies in Exchange Server.

What do you need to know before you begin?

- Estimated time to complete each procedure: Less than 5 minutes.
- You can assign ABPs to mailboxes in the Exchange admin center (EAC), but all other ABP procedures require the Exchange Management Shell. For more information about accessing and using the EAC, see Exchange admin center in Exchange Server. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Address book policies" entry in the Email address and address book permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.
- Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to view address book policies

To view ABPs, use this syntax:

```
Get-AddressBookPolicy [-Identity <ABPIdentity>]
```

This example returns a summary list of all ABPs in the organization:

```
Get-AddressBookPolicy
```

This example returns detailed information for the ABP named All Fabrikam ABP.

```
Get-AddressBookPolicy -Identity "All Fabrikam ABP" | Format-List
```

For detailed syntax and parameter information, see Get-AddressBookPolicy.

Use the Exchange Management Shell to create address book policies

An ABP requires one global address list (GAL), one offline address book (OAB), one room list, and one or more address lists. To view the available objects, use the **Get-GlobalAddressList**, **Get-OfflineAddressBook**, and **Get-AddressList** cmdlets.

Note: The room list that's required for an ABP is an address list that specifies rooms (contains the filter `RecipientDisplayType -eq 'ConferenceRoomMailbox'`). It's not a room finder distribution group that you create with the *RoomList* switch on the **New-DistributionGroup** or **Set-DistributionGroup** cmdlets.

To create an ABP, use this syntax:

```
New-AddressBookPolicy -Name "<Unique Name>" -GlobalAddressList "<GAL>" -OfflineAddressBook "<OAB>" -RoomList "<RoomList>" -AddressLists "<AddressList1>","<AddressList2>"...
```

This example creates an ABP named All Fabrikam ABP with the these settings:

- **GAL**: All Fabrikam
- **OAB**: Fabrikam-All-OAB
- **Room list**: All Fabrikam Rooms
- **Address lists**: All Fabrikam Mailboxes, All Fabrikam DLs, and All Fabrikam Contacts

```
New-AddressBookPolicy -Name "All Fabrikam ABP" -GlobalAddressList "\All Fabrikam" -OfflineAddressBook \Fabrikam-All-OAB -RoomList "\All Fabrikam Rooms" -AddressLists "\All Fabrikam Mailboxes","\All Fabrikam DLs","\All Fabrikam Contacts"
```

For detailed syntax and parameter information, see New-AddressBookPolicy.

How do you know this worked?

To verify that you've successfully created an ABP, use either of these procedures:

- Run this command in the Exchange Management Shell to verify that the ABP is listed:

  ```
  Get-AddressBookPolicy
  ```

- Replace *<ABPIdentity>* with the name of the ABP, and run this command in the Exchange Management Shell to verify the property values:

  ```
  Get-AddressBookPolicy -Identity "<ABPIdentity>" | Format-List
  ```

Use the Exchange Management Shell to modify address book policies

You use the **Set-AddressBookPolicy** cmdlet to modify an existing ABP. The settings are identical to the settings that are available when you create an ABP.

- The *Name*, *GlobalAddressList*, *OfflineAddressBook*, and *RoomList* parameters all take single values, so the value you specify replaces the existing value.

  This example modifies the ABP named "All Fabrikam ABP" by replacing the OAB with the specified OAB.

  ```
  Set-AddressBookPolicy -Identity "All Fabrikam ABP" -OfflineAddressBook \Fabrikam-OAB-2
  ```

- The *AddressLists* parameter takes multiple values, so you need to decide whether you want to *replace* the existing address lists in the ABP, or *add and remove* address lists without affecting the other address lists in the ABP.

This example replaces the existing address lists in the ABP named Government Agency A with the specified address lists.

```
Set-AddressBookPolicy -Identity "Government Agency A" -AddressLists "GovernmentAgencyA-
Atlanta","GovernmentAgencyA-Moscow"
```

To add address lists to an ABP, you need to specify the new address lists *and* any existing address lists that you want to keep.

This example adds the address list named Contoso-Chicago to the ABP named ABP Contoso, which is already configured to use the address list named Contoso-Seattle.

```
Set-AddressBookPolicy -Identity "ABP Contoso" -AddressLists "Contoso-Chicago","Contoso-Seattle"
```

To remove address lists from an ABP, you need to specify the existing address lists that you want to keep, and omit the address lists that you want to remove.

For example, the ABP named ABP Fabrikam uses the address lists named Fabrikam-HR and Fabrikam-Finance. To remove the Fabrikam-HR address list, specify only the Fabrikam-Finance address list.

```
Set-AddressBookPolicy -Identity "ABP Fabrikam" -AddressLists Fabrikam-Finance
```

For detailed syntax and parameter information, see Set-AddressBookPolicy.

How do you know this worked?

To verify that you've successfully modify an ABP, replace *<ABPIdentity>* with the name of the ABP, and run this command in the Exchange Management Shell to verify the property values:

```
Get-AddressBookPolicy -Identity "<ABPIdentity>" | Format-List
```

Use the Exchange Management Shell to remove address book policies

- You can't remove an ABP if it's assigned to a mailbox. To see if an ABP is assigned to a mailbox, replace *<ABPIdentity>* with the name of the ABP, and run this command in the Exchange Management Shell to get the **DistinguishedName** value:

  `Get-AddressBookPolicy -Identity <ABPIdentity> | Format-List DistinguishedName`

  Then, use the **DistinguishedName** value of the ABP in this command to show all mailboxes where the ABP is assigned:

  `Get-Mailbox -ResultSize unlimited -Filter "AddressBookPolicy -eq '<DistinguishedName>'"`

- To remove ABP assignments from mailboxes, see the Assign address book policies to mailboxes section in this topic.

To remove an ABP, use this syntax:

```
Remove-AddressBookPolicy -Identity <ABPIdentity>
```

This example removes the ABP named ABP_TailspinToys.

```
Remove-AddressBookPolicy -Identity "ABP_TailspinToys"
```

For detailed syntax and parameter information, see Remove-AddressBookPolicy.

How do you know this worked?

To verify that you've successfully removed an ABP, use either of these procedures:

- Run this command in the Exchange Management Shell to verify that the ABP isn't listed:

  ```
  Get-AddressBookPolicy
  ```

- Replace *<ABPIdentity>* with the name of the ABP, and run this command to confirm that an error is returned:

  ```
  Get-AddressBookPolicy -Identity "<ABPIdentity>"
  ```

Assign address book policies to mailboxes

- Users aren't automatically assigned an ABP when you create mailboxes. If you don't assign an ABP to a mailbox, the GAL for your entire organization is visible to the user in Outlook and Outlook on the web.

- To identify your virtual organizations for ABPs, we recommend that you use the CustomAttribute1-15 attributes on mailboxes, contacts, and groups, because these attributes are the most widely available and manageable for all recipient types. For more information, see Scenario: Deploying address book policies in Exchange Server.

- The procedures to assign ABPs to mailboxes or remove the ABP assignments from mailboxes are the same:

  - To assign ABPs to mailboxes, you select the ABP in EAC, or specify the ABP in the Exchange Management Shell.

  - To remove the ABP assignments from mailboxes, you select the value **[No Policy]** in the EAC, or use the value `$null` in the Exchange Management Shell.

Use the Exchange admin center (EAC) to assign an ABP to a single mailbox

1. In the EAC, go to **Recipients** > **Mailboxes**.

2. In the list of mailboxes, find the mailbox that you want to modify. You can:

   - Scroll through the list of mailboxes.

   - Click **Search** 🔍 and enter part of the user's name, email address, or alias.

   - Click **More options** ••• > **Advanced search** to find the mailbox.

   Once you've found the mailbox that you want to modify, select it, and then click **Edit** ✎.

3. On the mailbox properties page that opens, click **Mailbox features**.

4. Click the drop-down arrow in **Address book policy**, and select the ADP that you want to apply.

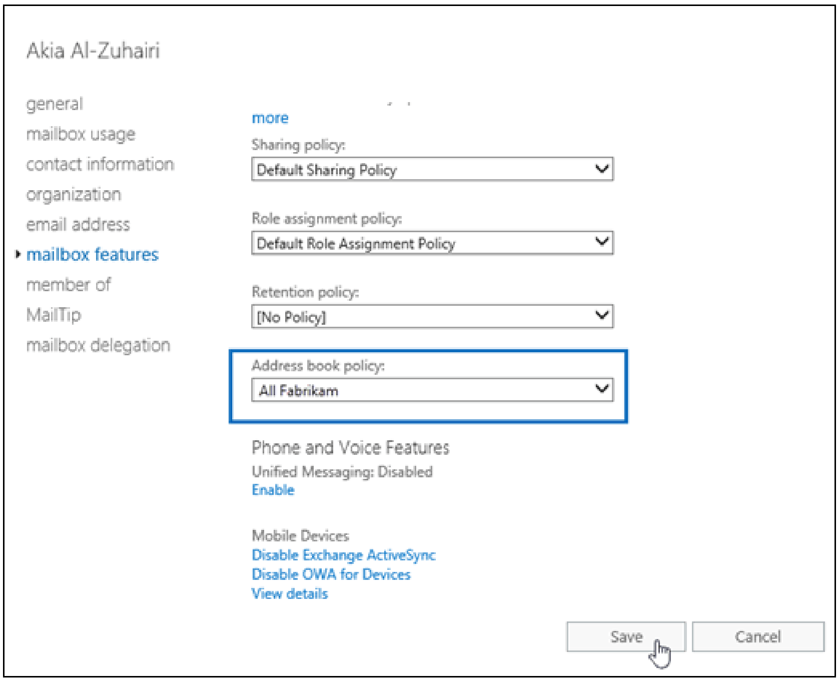

Sharing policy: Default Sharing Policy

Role assignment policy: Default Role Assignment Policy

Retention policy: [No Policy]

Address book policy: All Fabrikam

Phone and Voice Features

Unified Messaging: Disabled

Mobile Devices

When you're finished, click **Save**.

Note: You can also assign an ABP when you create a user mailbox in the EAC by clicking **More options**, and clicking the drop-down arrow in **Address book policy**.

Use the Exchange Management Shell to assign an address book policy to a single mailbox

To assign an ABP to a mailbox, use this syntax:

```
Set-Mailbox -Identity <MailboxIdentity> -AddressBookPolicy <ABPIdentity> or $null
```

This example assigns the ABP named All Fabrikam to mailbox joe@fabrikam.com.

```
Set-Mailbox -Identity joe@fabrikam.com -AddressBookPolicy "All Fabrikam"
```

Note: You can also assign an ABP when you create a user mailbox with the **New-Mailbox** cmdlet by using the *AddressBookPolicy* parameter. If you don't specify an ABP when you create the mailbox, no ABP is assigned (the default value is blank or `$null`).

For detailed syntax and parameter information, see Set-Mailbox.

Use the EAC to assign an address book policy to multiple mailboxes

1. In the EAC, go to **Recipients** > **Mailboxes**.

2. In the list of mailboxes, find the mailboxes that you want to modify. For example:

   a. Click **More options** ••• > **Advanced search**.

   b. In the **Advanced search** window that opens, select **Recipient types** and verify the default value **User mailbox**.

   c. Click **More options**, and then click **Add a condition**.

   d. In the **Select one** drop-down box that appears, select the appropriate **Custom attribute 1** to **Custom attribute 15** values that defines your virtual organizations.

   e. In the **Specify words or phrases** dialog that appears, enter the value that you want to search for, and then click **OK**.

   f. Back on the **Advanced search** window, click **OK**. In the EAC at **Recipients** > **Mailboxes**, click **More options** ••• > **Advanced search** to find user mailboxes.

3. In the list of mailboxes, select multiple mailboxes of the same type (for example, **User**) from the list. For example:

   - Select a mailbox, hold down the Shift key, and select another mailbox that's farther down in the list.

   - Hold down the CTRL key as you select each mailbox.

   After you select multiple mailboxes of the same type, the title of the details pane changes to **Bulk Edit**.

4. In the details pane, scroll down and click **More options**, scroll down to **Address Book Policy**, and then click **Update**.

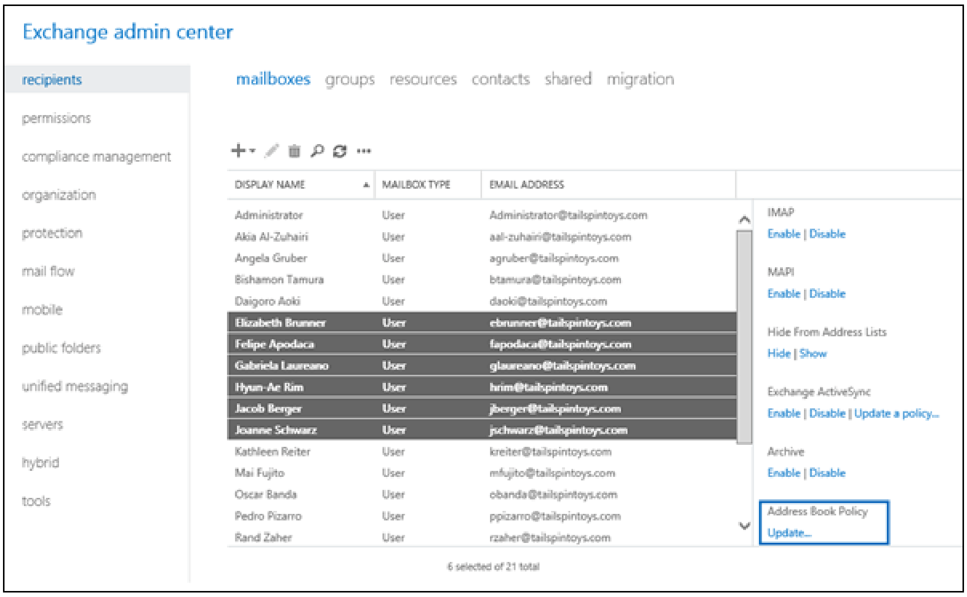

DISPLAY NAME	MAILBOX TYPE	EMAIL ADDRESS
Administrator	User	Administrator@tailspintoys.com
Akia Al-Zuhairi	User	aal-zuhairi@tailspintoys.com
Angela Gruber	User	agruber@tailspintoys.com
Bishamon Tamura	User	btamura@tailspintoys.com
Daigoro Aoki	User	daoki@tailspintoys.com
Elizabeth Brunner	User	ebrunner@tailspintoys.com
Felipe Apodaca	User	fapodaca@tailspintoys.com
Gabriela Laureano	User	glaureano@tailspintoys.com
Hyun-Ae Rim	User	hrim@tailspintoys.com
Jacob Berger	User	jberger@tailspintoys.com
Joanne Schwarz	User	jschwarz@tailspintoys.com
Kathleen Reiter	User	kreiter@tailspintoys.com
Mai Fujito	User	mfujito@tailspintoys.com
Oscar Banda	User	obanda@tailspintoys.com
Pedro Pizarro	User	ppizarro@tailspintoys.com
Rand Zaher	User	rzaher@tailspintoys.com

6 selected of 21 total

5. In the **Bulk assign address book policy** window that opens, select the ABP by clicking the drop-down arrow in **Select Address Book Policy**, and then click **Save**.

Use the Exchange Management Shell to assign an address book policy to multiple mailboxes

You can use the **Get-Mailbox** or **Get-Content** cmdlets to identify the user mailboxes that you want to assign the ABP to. For example:

- Use the *Filter* parameter to create OPATH filters that identify the mailboxes. For more information, see Filterable Properties for the -Filter Parameter.

- Use a text file to specify the mailboxes. The text file contains one mailbox (email address, name, or other unique identifier) on each line like this:

  ```
  ebrunner@tailspintoys.com
  fapodaca@tailspintoys.com
  glaureano@tailspintoys.com
  hrim@tailspintoys.com
  ```

This example assigns the ABP named ABP_EngineeringDepartment to all user mailboxes where the `CustomAttribute11` attribute contains the value Engineering Department.

```
Get-Mailbox -Filter "RecipientType -eq 'UserMailbox' -and CustomAttribute11 -like '*Engineering Department'" |
Set-Mailbox -AddressBookPolicy "ABP_EngineeringDepartment"
```

This example uses the text file C:\My Documents\Accounts.txt to assign the same ABP to the specified user mailboxes.

```
Get-Content "C:\My Documents\Accounts.txt" | foreach {Set-Mailbox $_ -AddressBookPolicy
"ABP_EngineeringDepartment"}
```

For detailed syntax and parameter information, see Get-Mailbox.

How do you know this worked?

To verify that you've successfully assigned an ABP to a mailbox, do any of these steps:

- In the EAC, go to **Recipients** > **Mailboxes** > select the mailbox > click **Edit** ✎ > **Mailbox features** and verify the **Address Book Policy** value.

  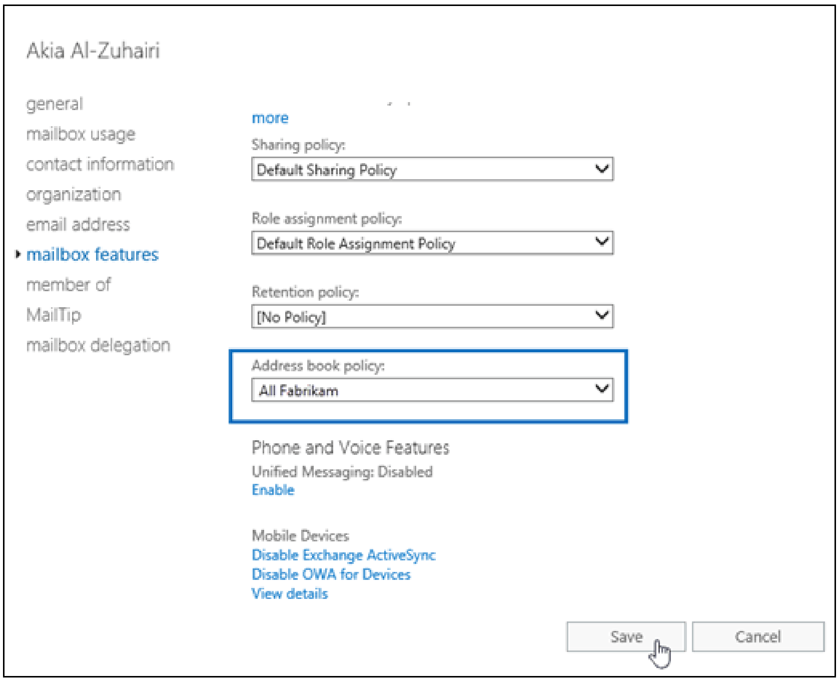

  Akia Al-Zuhairi

  Sharing policy: Default Sharing Policy

  Role assignment policy: Default Role Assignment Policy

  Retention policy: [No Policy]

  Address book policy: All Fabrikam

  Phone and Voice Features

  Unified Messaging: Disabled

  Mobile Devices

- In the Exchange Management Shell, replace *<MailboxIdentity>* with the identity of the mailbox (for example, name, alias, or email address), and run this command:

  ```
  Get-Mailbox -Identity "<MailboxIdentity>" | Format-List AddressBookPolicy
  ```

- In the Exchange Management Shell, use the same filter that you used to identify the mailboxes. For example:

  ```
  Get-Mailbox -Filter "RecipientType -eq 'UserMailbox' -and CustomAttribute11 -like '*Engineering
  Department'" | Format-Table -Auto Name,EmailAddress,AddressBookPolicy
  ```

- In the Exchange Management Shell, replace *<ABPIdentity>* with the name of the ABP, and run this command to get the **DistinguishedName** value:

  ```
  Get-AddressBookPolicy -Identity <ABPIdentity> | Format-List DistinguishedName
  ```

  Then, use the **DistinguishedName** value of the ABP in this command to show all mailboxes where the ABP is assigned:

  ```
  Get-Mailbox -ResultSize unlimited -Filter "AddressBookPolicy -eq '<DistinguishedName>'"
  ```

Use the Exchange Management Shell to install and configure the Address Book Policy Routing Agent

Address Book Policy routing (ABP routing) controls how recipients are resolved in organizations that use ABPs. When ABP routing is enabled, users that are assigned different GALs appear as external recipients to each other.

ABP routing requires that you install and enable the Address Book Policy Routing Agent (ABP Routing Agent) on all Mailbox servers in your organization, and enable ABP routing globally in your organization. After you do this, it might take up to 30 minutes for messages to be processed by the ABP Routing Agent.

You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Transport Agents" entry in the Mail flow permissions topic.

Step 1: Install the ABP Routing agent

To install the ABP Routing Agent on the local Mailbox server, run this command on every Mailbox server in the organization.

```
Install-TransportAgent -Name "ABP Routing Agent" -TransportAgentFactory
"Microsoft.Exchange.Transport.Agent.AddressBookPolicyRoutingAgent.AddressBookPolicyRoutingAgentFactory" -
AssemblyPath
$env:ExchangeInstallPath\TransportRoles\agents\AddressBookPolicyRoutingAgent\Microsoft.Exchange.Transport.Agen
t.AddressBookPolicyRoutingAgent.dll
```

Note: You'll get a warning that the Transport service needs to be restarted for the changes to take effect. But, don't restart the Transport service until you finish Step 2 (so you only have to restart the Transport service once).

For detailed syntax and parameter information, see Install-TransportAgent.

Step 2: Enable the ABP Routing agent

To enable the ABP Routing Agent on the local Mailbox server, run this command on every Mailbox server in the organization.

```
Enable-TransportAgent "ABP Routing Agent"
```

For detailed syntax and parameter information, see Enable-TransportAgent.

Step 3: Restart the Transport service

To restart the Transport service, run this command on every Mailbox server in the organization.

```
Restart-Service MSExchangeTransport
```

For detailed syntax and parameter information, see Get-TransportAgent.

Step 4: Enable ABP routing globally in the Exchange organization

To enable ABP routing globally in the Exchange organization, run this command once on any Mailbox server:

```
Set-TransportConfig -AddressBookPolicyRoutingEnabled $true
```

For detailed syntax and parameter information, see Set-TransportConfig.

Note: To disable ABP routing after you've enabled it, do these steps:

1. Run this command once on any Mailbox server to globally disable ABP routing:

```
Set-TransportConfig -AddressBookPolicyRoutingEnabled $false
```

2. Disable the ABP Routing Agent by running this command on every Mailbox server where the agent is installed:

   ```
   Disable-TransportAgent "ABP Routing Agent"
   ```

3. Run this command on every Mailbox server where the agent is installed:

   ```
   Restart-Service MSExchangeTransport
   ```

How do you know this worked?

To verify that you've successfully installed and configured the ABP Routing Agent, use any of these steps:

- Run this command on a Mailbox server to verify that ABP routing is enabled for the organization:

  ```
  Get-TransportConfig | Format-List AddressBookPolicyRoutingEnabled
  ```

- Run this command on every Mailbox server to verify that the ABP Routing Agent is enabled:

  ```
  Get-TransportAgent "ABP Routing Agent"
  ```

- Have a user that's assigned an ABP send an email message to an user that's assigned a different ABP, and verify that the sender's email address doesn't resolve to their display name.

Messaging policy and compliance in Exchange Server

8/3/2020 • 2 minutes to read • Edit Online

Email has become a reliable and ubiquitous communication medium for information workers in organizations of all sizes. Messaging stores and mailboxes have become repositories of valuable data. It's important for organizations to formulate messaging policies that dictate the fair use of their messaging systems, provide user guidelines for how to act on the policies, and where required, provide details about the types of communication that may not be allowed.

Organizations must also create policies to manage email lifecycle, retain messages for the length of time based on business, legal, and regulatory requirements, preserve email records for litigation and investigation purposes, and be prepared to search and provide the required email records to fulfill eDiscovery requests.

Messaging policy and compliance in Exchange Server

The following table provides an overview of the messaging policy and compliance features in Exchange 2016 and Exchange 2019 and includes links to topics that will help you learn about and use these features.

FEATURE	DESCRIPTION	RESOURCES
In-Place Archiving	*In-Place Archiving* helps you regain control of your organization's messaging data by eliminating the need for personal store (.pst) files and allowing users to store messages in an *archive mailbox* accessible in Outlook 2010 and later and Outlook on the web.	In-Place Archiving in Exchange Server
In-Place Hold and Litigation Hold	When a reasonable expectation of litigation exists, organizations are required to preserve electronically stored information, including email that's relevant to the case. In-Place Hold allows you to search and preserve messages matching query parameters. Litigation Hold only allows you to place all items in a mailbox on hold. For both types of holds, messages are protected from permanent deletion, modification, and tampering and can be preserved indefinitely or for a specified period.	In-Place Hold and Litigation Hold in Exchange Server
In-Place eDiscovery	In-Place eDiscovery allows you to search mailbox data across your Exchange organization, preview search results, copy search results to a Discovery mailbox, or export the results to a PST file	In-Place eDiscovery in Exchange Server

FEATURE	DESCRIPTION	RESOURCES
Administrator audit logging	Administrator audit logs enable you to keep a log of changes made by administrators to Exchange server and organization configuration and to Exchange recipients. You might use administrator audit logging as part of your change control process or to track changes and access to configuration and recipients for compliance purposes.	Administrator audit logging in Exchange Server
Mailbox audit logging	Because mailboxes can potentially contain sensitive, high business impact information and personally identifiable information, it's important that you track who logs on to the mailboxes in your organization and what actions are taken. It's especially important to track access to mailboxes by users other than the mailbox owner (known as delegate users). Using mailbox audit logging, you can log mailbox access by administrators, delegates (including administrators with full access permissions), and mailbox owners.	Mailbox audit logging in Exchange Server
Data loss prevention	Data loss prevention (DLP) in Exchange Server includes 80 sensitive information types that are ready for you to use in your DLP policies.	Sensitive information types in Exchange Server
Mail flow rules (also known as transport rules)	Use mail flow rules to look for specific conditions in messages that pass through your organization and take action on them. You can use conditions and exceptions to define when a mail flow rule is applied, and then apply an action on messages when the conditions are met.	Mail flow rule conditions and exceptions (predicates) in Exchange Server Mail flow rule actions in Exchange Server

In-Place Archiving in Exchange Server

8/3/2020 • 9 minutes to read • Edit Online

In-Place Archiving in Exchange Server helps you regain control of your organization's messaging data by eliminating the need for personal store (.pst) files and allowing users to store messages in an *archive mailbox*. The archive mailbox is an additional mailbox that's enabled for a user's primary mailbox. The archive mailbox is accessible in Outlook and Outlook on the web (formerly known as Outlook Web App). Users can view an archive mailbox and move or copy messages between their primary mailbox and their archive mailbox.

You can provision a user's archive mailbox on the same mailbox database as the user's primary mailbox, a different mailbox database on the same Mailbox server, or on a mailbox database on a different Mailbox server in the same Active Directory site. In Exchange hybrid deployments, you can also provision a cloud-based archive mailbox for primary mailboxes located in your on-premises organization.

Client access to archive mailboxes

The following table lists the client applications that can be used to access archive mailboxes.

CLIENT	ACCESS TO ARCHIVE MAILBOX?
Outlook for Mac for Office 365 Outlook 2016 for Mac or later Microsoft 365 Apps for enterprise Outlook 2013 or later Outlook on the web	Yes. Users can copy or move items from their primary mailbox to their archive mailbox, and can also use retention policies to move items to the archive. Outlook doesn't create a local copy of the archive mailbox on a user's computer, even if it's configured to use Cached Exchange Mode. Users can access an archive mailbox in online mode only.
Exchange ActiveSync	No

NOTE

• In-Place Archiving is a premium feature and requires an Exchange Enterprise client access license (CAL). For details about how to license Exchange, see Exchange licensing FAQs.
• For details about the versions of Outlook that are required to access an archive mailbox, see Outlook license requirements for Exchange features.

Moving messages to the archive mailbox

There are several ways to move messages from a user's primary mailbox to their archive mailbox:

- **Move or copy messages manually**: Users can manually move or copy messages from their primary mailbox or a .pst file to their archive mailbox. The archive mailbox appears as another mailbox Outlook and Outlook on the web or like a mounted .pst file in Outlook.
- **Move or copy messages using Inbox rules**: Users can create Inbox rules in Outlook or Outlook on the web to automatically move messages to a folder in their archive mailbox.
- **Move messages using retention policies**: You can use retention policies to automatically move messages to the archive mailbox. Users can also apply personal tags to move messages to their archive

mailbox. For details about archive and retention policies, see the next section in this topic.

- **Import messages from .pst files**: In Exchange Server, you can use a mailbox import request to import messages from a .pst file to a user's archive or primary mailbox. For details, see Mailbox imports and exports in Exchange Server.

Archiving and retention policies

In Exchange Server, you can apply archive policies to a mailbox to automatically move messages from a user's primary mailbox to the archive mailbox after a specified period. Archive policies are implemented by creating retention tags that use the **Move to Archive** retention action.

Messages are moved to a folder in the archive mailbox that has the same name as the source folder in the primary mailbox. If a folder with the same name doesn't exist in the archive mailbox, it's created when the Managed Folder Assistant moves a message. Re-creating the same folder hierarchy in the archive mailbox allows users to find messages easily.

To learn more about retention policies, retention tags, and the **Move to Archive** retention action, see Retention tags and retention policies in Exchange Server.

Default MRM policy

Exchange Server Setup creates a default archive and retention policy named **Default MRM Policy**. This policy contains retention tags that have the **Move to Archive** action, as shown in the following table.

RETENTION TAG NAME	TAG TYPE	DESCRIPTION
Default 2 year move to archive	Default (DPT)	Messages are automatically moved to the archive mailbox after two years. Applies to items in the entire mailbox that don't have a retention tag applied explicitly or inherited from the folder.
Personal 1 year move to archive	Personal	Messages are automatically moved to the archive mailbox after one year.
Personal 5 year move to archive	Personal	Messages are automatically moved to the archive mailbox after five years.
Personal never move to archive	Personal	Messages are never moved to the archive mailbox.
Recoverable Items 14 days move to archive	Recoverable Items Folder	Messages are moved from the Recoverable Items folder in the user's primary mailbox to the Recoverable Items folder in the archive mailbox. Users attempting to recover deleted items in their archive mailbox must use the Recover Deleted Items tool in the archive mailbox.

If you enable an In-Place Archive for a mailbox user and the mailbox doesn't already have a retention policy assigned, the default archive and retention policy is automatically assigned. After the Managed Folder Assistant processes the mailbox, these tags become available to the user, who can then tag folders or messages to be moved to the archive mailbox. By default, email messages from the entire mailbox are moved to the archive after two years.

Before provisioning archive mailboxes for your users, we recommend that you inform them about the archive policies that will be applied to their mailbox and provide subsequent training or documentation to meet their needs. This should include details about the following:

- Functionality available within the archive, and the default archive retention policies.
- Information about when messages are automatically moved to the archive.
- Information about the folder hierarchy created in the archive mailbox.
- How to apply personal tags (displayed in the Archive policy menu in Outlook and Outlook on the web).

> **NOTE**
>
> If you apply a retention policy to users who have an archive mailbox, the retention policy replaces the default MRM policy. You can create one or more retention tags with the **Move to Archive** action, and then link the tags to the retention policy. You can also add the default **Move to Archive** tags (which are created by Setup and linked to the Default MRM Policy) to any retention policies you create.

Archive quotas

Archive mailboxes are designed so that users can store historical messaging data outside their primary mailbox. Often, users use .pst files due to low mailbox storage quotas and the restrictions imposed when these quotas are exceeded. For example, users can be prevented from sending messages when their mailbox size exceeds the *Prohibit send quota*. Similarly, users can be prevented from sending and receiving messages when their mailbox size exceeds the *Prohibit send and receive quota*.

To eliminate the need for .pst files, you can provide an archive mailbox with storage limits that meet the user's requirements. However, you may still want to retain some control of the storage quotas and growth of archive mailboxes to help monitor costs and expansion.

To help with this control, you can configure archive mailboxes with an *archive warning quota* and an *archive quota*. When an archive mailbox exceeds the specified archive warning quota, a warning event is logged in the Application event log. When an archive mailbox exceeds the specified archive quota, messages are no longer moved to the archive, a warning event is logged in the Application event log, and a quota message is sent to the mailbox user. By default, in Exchange Server, the archive warning quota is set to 90 GB and the archive quota is set to 100 GB.

The following table lists the events logged and warning messages sent when the archive warning quota and archive quota are met.

QUOTA	EVENT ID	TYPE	SOURCE	CATEGORY	MESSAGE
Archive warning quota	10022	Warning	MSExchangeMailboxAssistants	Managed Folder Assistant	`The archive mailbox '<Display Name>:<GUID>:<Mailbox Database>:<Server FQDN>' exceeded the archive warning quota '<Archive warning quota>'. Archive mailbox size is '<Size>' bytes.`

QUOTA	EVENT ID	TYPE	SOURCE	CATEGORY	MESSAGE
Archive quota	8537	Warning	MSExchangeIS	General	`The archive mailbox for <Legacy DN> has exceeded the maximum archive mailbox size. You can't copy or move items into the archive mailbox. All message retention actions that move items to the archive mailbox will fail, and the primary mailbox may contain items with expired retention tags until the archive mailbox is within the maximum size limit. The mailbox owner should be notified about the condition of the archive mailbox.`

In-Place Archiving and other Exchange features

This section explains the functionality between In-Place Archiving and various Exchange features:

- **Exchange Search**: The ability to quickly search messages becomes even more critical with archive mailboxes. For Exchange Search, there's no difference between the primary and archive mailbox. Content in both mailboxes is indexed. Because the archive mailbox isn't cached on a user's computer (even when using Outlook in Cached Exchange Mode), search results for the archive are always provided by Exchange Search. When searching the entire mailbox in Outlook and later and Outlook on the web, search results include the users' primary and archive mailbox.

- **In-Place eDiscovery**: When a discovery manager performs an In-Place eDiscovery search, users' archive mailboxes are also searched. There's no option to exclude archive mailboxes when creating a discovery search from the Exchange admin center (EAC). When using the Exchange Management Shell to create a discovery search, you can exclude the archive by using the *DoNotIncludeArchive* switch. For details, see New-MailboxSearch. To learn more, see In-Place eDiscovery in Exchange Server.

- **In-Place Hold and Litigation Hold**: When you put a mailbox on In-Place Hold or Litigation Hold, the hold is placed on both the primary and the archive mailbox. To learn more, see In-Place Hold and Litigation Hold in Exchange Server.

- **Recoverable Items folder**: The archive mailbox contains its own Recoverable Items folder and is subject to the same Recoverable Items folder quotas as the primary mailbox. To learn more about recoverable items, see Recoverable Items folder in Exchange Server.

- **Archiving Skype for Business content in Exchange**: You can archive instant messaging conversations and shared online meeting documents in the user's primary mailbox. The mailbox must reside on an Exchange Mailbox server and you must have Skype for Business Server 2015 deployed in your organization.

Managing archive mailboxes

In Exchange Server, creating and managing archive mailboxes is integrated with common mailbox management tasks. For step by step procedures, see Manage In-Place Archives in Exchange Server.

- **Creating an archive mailbox**: You can enable an archive mailbox for an existing mailbox or you can create an archive mailbox when creating a new mailbox. .

- **Moving an archive mailbox**: You can move a user's archive mailbox to another mailbox database on the same Mailbox server or to another server, independent of the primary mailbox. To move a user's archive mailbox, you must create a mailbox move request. For details, see Manage on-premises mailbox moves in Exchange Server.

- **Disabling an archive mailbox**: You may want to disable a user's archive mailbox for troubleshooting purposes or if you're moving the primary mailbox to a version of Exchange that doesn't support In-Place Archiving. Disabling an archive is similar to disabling a primary mailbox. In on-premises deployments, a disabled archive mailbox is retained in the mailbox database until the deleted mailbox retention period for that database is reached. During this period, you can reconnect the same disabled archive mailbox to a user's primary mailbox. When the deleted mailbox retention period is reached, the disconnected archive mailbox is purged from the mailbox database.

- **Retrieving mailbox statistics and folder statistics**: You can retrieve mailbox statistics and mailbox folder statistics for a user's archive mailbox by using the *Archive* switch with the Get-MailboxStatistics and Get-MailboxFolderStatistics cmdlets.

- **Test archive connectivity**: In Exchange Server, you can use the Test-ArchiveConnectivity cmdlet to test connectivity to a specified user's on-premises or cloud-based archive mailbox.

Manage In-Place Archives in Exchange Server

8/3/2020 • 8 minutes to read • Edit Online

In-Place Archiving helps you regain control of your organization's messaging data by eliminating the need for personal store (.pst) files and allowing you to meet your organization's message retention and eDiscovery requirements. With archiving enabled, users can store messages in an archive mailbox, which is accessible by using Microsoft Outlook and Outlook on the web.

What do you need to know before you begin?

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "In-Place Archive" entry in the Messaging policy and compliance permissions in Exchange Server topic.
- The procedures in this topic apply to on-premises archive mailboxes. For information about archive mailboxes in Exchange Online, see Enable archive mailboxes in the Security & Compliance Center.
- To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- It's not supported to have a user's primary mailbox reside on a version of Exchange that's older than the user's archive. If the user's primary mailbox is still on Exchange 2010 or Exchange 2013, you need to move it to Exchange 2016 or Exchange 2019 at the same time you move the archive mailbox to Exchange 2016 or Exchange 2019.

Enable an archive mailbox

You can use the Exchange admin center or the Exchange Management Shell to enable archive mailboxes for users that already have a primary mailbox.

Use the EAC to enable an archive mailbox

1. Go to **Recipients** > **Mailboxes**.
2. Select a mailbox.
3. In the details pane, under **In-Place Archive**, click **Enable**.

   Note: You can also bulk-enable archives by selecting multiple mailboxes (use the Shift or Ctrl keys). After selecting multiple mailboxes, in the details pane, click **More options**. Then, under **Archive** click **Enable**.
4. On the **Create In-Place Archive** page, click **OK** to have Exchange automatically select a mailbox database for the archive or click **Browse** to specify one.

Use the Exchange Management Shell to enable an archive mailbox

This example enables the archive mailbox for Tony Smith.

```
Enable-Mailbox "Tony Smith" -Archive
```

This example retrieves mailboxes in database DB01 that don't have an on-premises or cloud-based archive enabled and don't have a name starting with DiscoverySearchMailbox. It pipes the results to the **Enable-Mailbox** cmdlet to enable the archive for all mailboxes on mailbox database DB01.

```
Get-Mailbox -Database DB01 -Filter "ArchiveGuid -Eq `$null -AND ArchiveDomain -eq `$null -AND Name -NotLike
'DiscoverySearchMailbox*'" | Enable-Mailbox -Archive
```

How do you know this worked?

To verify that you've successfully enabled an on-premises archive for an existing mailbox, do one of the following:

- In the EAC, go to **Recipients** > **Mailboxes**, and then select the mailbox from the list. In the details pane, under **In-Place Archive**, confirm that it is set to **Enabled**. Click **View details** to view archive properties, including archive status and the mailbox database in which it is created.

- In the Exchange Management Shell, run the following command to display information about the new archive.

  ```
  Get-Mailbox <MailboxIdentity> | Format-List Name,*Archive*
  ```

- In the Exchange Management Shell, use the **Test-ArchiveConnectivity** cmdlet to test connectivity to the archive. For an example of how to test archive connectivity, see the Examples section in the topic, Test-ArchiveConnectivity.

Enable an archive mailbox when you create a new mailbox

You can also enable an archive mailbox when you first create a new mailbox for a user.

Use the EAC to enable an archive mailbox when you create a new mailbox

1. Go to **Recipients** > **Mailboxes**.

2. Click **New** > **User mailbox**.

3. On the **New user mailbox** page, in the **Alias** box, type an alias for the user.

   Note: If you leave this box blank, the value you type in the **User logon name** box is used for the alias.

4. Select one of the following options:

   - **Existing user that isn't mail-enabled**: Click this button and then click **Browse** to open the **Select User - Entire Forest** dialog box. This dialog box displays a list of Active Directory user accounts in the forest that aren't mail-enabled or don't have Exchange mailboxes. Select the user account you want to mail-enable, and then click **OK**. If you select this option, you don't have to provide user account information because this information already exists in Active Directory.

   - **New user**: Click this button to create a new user account in Active Directory and create a mailbox for the user. If you select this option, you'll have to provide the required user account information.

5. Click **More options** to configure the following settings.

   - **Mailbox database**: Click **Browse** to select a mailbox database in which to store the mailbox. If you don't select a database, Exchange will automatically assign one.

   - **Archive**: Select this check box to create an archive mailbox for the mailbox. If you create an archive mailbox, mailbox items will be moved automatically from the primary mailbox to the archive, based on the default retention policy settings or those you define.

   Click **Browse** to select a database to store the archive mailbox.

6. When you're finished, click **Save** to create the mailbox and its archive.

Use the Exchange Management Shell to enable an archive mailbox when you create a new mailbox

This example creates the user named Chris Ashton in Active Directory, creates the mailbox on mailbox database DB01, and enables and creates an archive mailbox on DB01. To set the initial value of the password, this example creates a variable ($password), prompts you to enter a password, and assigns that password to the variable as a SecureString object.

```
$password = Read-Host "Enter password" -AsSecureString
```

```
New-Mailbox -UserPrincipalName cashton@contoso.com -Alias cashton -Database "DB01" -Archive -Name "Chris
Ashton" -OrganizationalUnit Users -Password $password -FirstName Chris -LastName Ashton
```

How do you know this worked?

To verify that you've successfully created a user mailbox with an on-premises archive, do one of the following:

- In the EAC, go to **Recipients** > **Mailboxes**, and then select the new user mailbox from the list. In the details pane, under **In-Place Archive**, confirm that it is set to **Enabled**. Click **View details** to view archive properties, including archive status and the mailbox database in which it is created.

- In the Exchange Management Shell, run the following command to display information about the new user mailbox and archive.

  ```
  Get-Mailbox <Name> | Format-List Name,RecipientTypeDetails,PrimarySmtpAddress,*Archive*
  ```

- In the Exchange Management Shell, use the **Test-ArchiveConnectivity** cmdlet to test connectivity to the archive. For an example of how to test archive connectivity, see the Examples section in Test-ArchiveConnectivity.

Disable an archive mailbox

You may want to disable a user's archive for troubleshooting purposes or compliance-related reasons. If you disable an archive mailbox, all information in the archive will be kept in the mailbox database until the mailbox retention time expires and the archive is permanently deleted. By default, Exchange keeps deleted mailboxes, including archive mailboxes, for 30 days.

Use the EAC to disable an archive mailbox

1. Go to **Recipients** > **Mailboxes**.

2. Select a mailbox.

3. In the details pane, under **In-Place Archive**, click **Disable**.

   Note: You can also bulk-disable archives by selecting multiple mailboxes (use the Shift or Ctrl keys). After selecting multiple mailboxes, in the details pane, click **More options**. Then, under **Archive** click **Disable**.

Use the Exchange Management Shell to disable an archive mailbox

This example disables the archive mailbox for Chris Ashton's mailbox. It doesn't disable the user's primary mailbox.

```
Disable-Mailbox "Chris Ashton" -Archive
```

How do you know this worked?

To verify that you have successfully disabled an archive mailbox, do the following:

- In the EAC, select the mailbox. In the details pane, check its archive status under **In-Place Archive**.

- In the Exchange Management Shell, run the following command to check the archive properties for the mailbox user.

  ```
  Get-Mailbox "Chris Ashton" | Format-List *Archive*
  ```

If the archive is disabled, the following values are returned for archive-related properties.

PROPERTY	VALUE
ArchiveDatabase (for on-premises archives)	<blank>
ArchiveState	`None`
DisabledArchiveDatabase (for on-premises archives)	*<name of mailbox database>*
DisabledArchiveGuid	*<GUID of disabled archive>*

Re-enable an archive mailbox

When you disable an archive mailbox, it becomes disconnected. A disconnected archive mailbox is retained in the mailbox database for a specified amount of time. By default, Exchange retains disconnected archive mailboxes for 30 days. Within 30 days of disabling an archive mailbox, you can reconnect it to the user's primary mailbox by re-enabling the archive. In this case, the original contents of the archive mailbox are restored. However after 30 days of disabling a mailbox, the contents of the original archive mailbox are permanently deleted (purged from the mailbox database) and can't be recovered. So if you re-enable the archive more than 30 days after disabling it, a new archive mailbox is created when you re-enable it.

Use the EAC to re-enable an archive mailbox

1. Go to **Recipients** > **Mailboxes**.

2. Select the mailbox.

3. In the details pane, under **In-Place Archive**, click **Enable**

4. On the **Create in-place archive** page, click **OK**.

   You can have Exchange automatically select a mailbox database for the re-enabled archive mailbox or you can click **Browse** to specify one.

Use the Exchange Management Shell to re-enable an archive mailbox

Use the **Enable-Mailbox -Archive** command to re-enable an archive mailbox. For example:

```
Enable-Mailbox "Chris Ashton" -Archive
```

How do you know this worked?

To verify that you have successfully connected a disabled archive mailbox to the user's primary mailbox, run the following command to retrieve the mailbox user's archive properties, and verify the values returned for the *ArchiveGuid* and *ArchiveDatabase* properties.

```
Get-Mailbox "Chris Ashton" | Format-List *Archive*
```

As previously stated, if you re-enable an archive mailbox within 30 days of disabling it, the user will be able to access the original contents of their archive mailbox. If you re-enable the archive more than 30 days after disabling it, the new archive mailbox will be empty the first time the user accesses it.

In-Place Hold and Litigation Hold in Exchange Server

8/3/2020 • 14 minutes to read • Edit Online

When a reasonable expectation of litigation exists, organizations are required to preserve electronically stored information (ESI), including email that's relevant to the case. This expectation often exists before the specifics of the case are known, and preservation is often broad. Organizations may need to preserve all email related to a specific topic or all email for certain individuals. Depending on the organization's electronic discovery (eDiscovery) practices, the following measures can be adopted to preserve email:

- End users may be asked to preserve email by not deleting any messages. However, users can still delete email knowingly or inadvertently.

- Automated deletion mechanisms such as messaging records management (MRM) may be suspended. This could result in large volumes of email cluttering the user mailbox, and thus impacting user productivity. Suspending automated deletion also doesn't prevent users from manually deleting email.

- Some organizations copy or move email to an archive to make sure it isn't deleted, altered, or tampered with. This increases costs due to the manual efforts required to copy or move messages to an archive, or third-party products used to collect and store email outside Exchange.

Failure to preserve email can expose an organization to legal and financial risks such as scrutiny of the organization's records retention and discovery processes, adverse legal judgments, sanctions, or fines.

Litigation Hold and In-Place Hold

There are two types of holds available in Exchange Server: Litigation Hold and In-Place Hold. Litigation Hold uses the **LitigationHoldEnabled** property of a mailbox. When Litigation Hold is enabled, all mailbox all items are placed on hold. In contrast, you can use an In-Place Hold to preserve only those items that meet that the criteria of a search query that you define by using the In-Place eDiscovery tool. You can place multiple In-Place Holds on a mailbox, but Litigation Hold is either enabled or disabled for a mailbox. For both types of holds, you can also specify the duration period to hold items. The duration is calculated from the date a mailbox item is received or created. If a duration isn't set, items are held indefinitely or until the hold is removed. If you remove a Litigation Hold from a mailbox, but one or more In-Place Holds are still placed on the mailbox, items matching the In-Place Hold criteria are held for the period specified in the hold settings.

You can use In-Place Hold to place a user on multiple holds. When a user is placed on multiple holds, the search queries from any query-based hold are combined (with **OR** operators). In this case, the maximum number of keywords in all query-based holds placed on a mailbox is 500. If there are more than 500 keywords, then all content in the mailbox is placed on hold (not just that content that matches the search criteria). All content is held until the total number of keywords is reduced to 500 or less.

When you move a mailbox that's on Litigation Hold in Exchange 2010 or Exchange 2013 to a Mailbox server in Exchange 2016, the Litigation Hold setting continues to apply, ensuring that compliance requirements are met during and after the move.

IMPORTANT

When you put a mailbox on Litigation Hold or In-Place Hold, the hold is placed on both the primary and the archive mailbox.

For more information when to use each type of hold, see Place all mailboxes on hold.

Hold goals and features

You can use Litigation Hold and In-Place Hold to accomplish the following goals:

- Place user mailboxes on hold and preserve mailbox items immutably.
- Preserve items indefinitely or for a specific duration.
- Preserve mailbox items deleted by users or automatic processes such as MRM.
- Preserve messages that are forwarded to another mailbox.
- Use query-based In-Place Hold to search for and retain items matching specified criteria (you can also place all items hold by including all mailbox content when you create the hold)
- Place a user on multiple holds for different cases or investigations.
- Keep holds transparent from the user by not having to suspend MRM.
- Use In-Place eDiscovery to search for items that are preserved by being placed on hold

If you're upgrading from Exchange Server 2010, the notion of legal hold is to hold all mailbox data for a user indefinitely or until when hold is removed. In Exchange 2016, In-Place Hold introduced a different model that allows you to specify the following parameters:

- **Query-based hold**: With Litigation Hold, all items in a mailbox are preserved. However, an In-Place Hold allows you to specify which items to hold by using search query parameters such as keywords, senders and recipients, start and end dates, and also specify the message types such as email messages, calendar items, and Skype for Business conversations that you want to place on hold. After you create a query-based In-Place Hold, all existing and future mailbox items (including messages received at a later date) that match the query parameters are preserved. Litigation Hold doesn't support query-based holds.
- **Hold duration**: In both Litigation Hold and In-Place Hold, you can specify how long to hold items. You can either specify an infinite hold duration or a time-based hold duration. The duration is calculated from the date a mailbox item is received or created. For example, if your organization requires that all mailbox items be preserved for 7 years, you can create a time-based hold. So if a mailbox is placed on hold and the hold duration is set to 7 years, and an item in the mailbox is permanently deleted after 2 years from the date it was received, it's held for an 5 years before being purged from the mailbox database.

> **TIP**
>
> You can use a time-based hold together with a retention policy to make sure items are preserved for a specified duration and then permanently removed from Exchange after the retention age and the hold duration expire.

Placing a mailbox on hold

The Legal Hold management role is required to place a mailbox on Litigation Hold or In-Place Hold. But to create a query-based In-Place Hold, you must also be assigned the Mailbox Search role. Users that have been added to the Discovery Management role-based access control (RBAC) role group (or assigned the Legal Hold and Mailbox Search roles) can place users hold and create a query-based In-Place Hold. To learn how to add members to the Discovery Management role group, see Assign eDiscovery permissions in Exchange Server.

You can place a mailbox Litigation Hold on the **Recipients** page in the Exchange admin center or by using the `Set-Mailbox -LitigationHoldEnabled $true` command in the Exchange Management Shell.

The In-Place Hold functionality is integrated with In-Place eDiscovery searches. You can place a mailbox on In-Place Hold by using the **In-Place eDiscovery & Hold** wizard in the EAC or the **New-MailboxSearch** cmdlet in the Exchange Management Shell. To learn how, see:

- Place a mailbox on Litigation Hold
- Create or remove an In-Place Hold

> **NOTE**
>
> If you use Exchange Online Archiving to provision a cloud-based archive for your on-premises mailboxes, you must manage In-Place Holds from your on-premises Exchange Server organization. Hold settings are automatically propagated to the cloud-based archive using DirSync.

Many organizations require that users be informed when they're placed on hold. Additionally, when a mailbox is on hold, any retention policies applicable to the mailbox user don't need to be suspended. Because messages continue to be deleted as expected, users may not notice they're on hold. If your organization requires that users on hold be informed, you can add a notification message to the mailbox user's by populating the **Retention Comment** property and using the **RetentionUrl** property to link to a web page for more information. Outlook 2010 and later versions display the retention comment and URL in the backstage area, which is located on the **Files** ribbon. You can use the **Set-Mailbox** cmdlet to add these properties.

Holds and the Recoverable Items folder

Litigation Hold and In-Place Hold use the Recoverable Items folder to preserve items. The Recoverable Items folder is hidden from the default view of Outlook, Outlook on the web, and other email clients. To learn more about the Recoverable Items folder, see Recoverable Items folder in Exchange Server.

By default, when a user deletes a message from a folder other than the Deleted Items folder, the message is moved to the Deleted Items folder. When a user *soft deletes* an item (by pressing SHIFT+DELETE) or deletes an item from the Deleted Items folder, the message is moved to the Recoverable Items folder, thereby disappearing from the user's view.

Items in the Recoverable Items folder are retained for the deleted item retention period configured on the user's mailbox database. By default, the deleted item retention period is set to 14 days for mailbox databases.

The Recoverable Items folder contains the following subfolders used to store deleted items in various sites and facilitate Litigation Hold and In-Place Hold:

- **Deletions**: Items removed from the Deleted Items folder or soft-deleted from other folders are moved to the Deletions subfolder and are visible to the user when using the Recover Deleted Items feature in Outlook and Outlook on the web. By default, items reside in this folder until the deleted item retention period configured for the mailbox database or the mailbox expires.

- **Purges**: When a user deletes an item from the Recoverable Items folder (by using the Recover Deleted Items tool in Outlook and Outlook on the web, the item is moved to the Purges folder. Items that exceed the deleted item retention period configured on the mailbox database or the mailbox are also moved to the Purges folder. Items in this folder aren't visible to users if they use the Recover Deleted Items tool. When the mailbox assistant processes the mailbox, items in the Purges folder are purged from the mailbox database. When you place the mailbox user on Litigation Hold, the mailbox assistant doesn't purge items in this folder.

- **DiscoveryHolds**: If a user is put on an In-Place Hold, deleted items are moved to this folder. When the mailbox assistant processes the mailbox, it evaluates messages in this folder. Items that match the In-Place Hold query are retained until the hold period specified in the query. If no hold period is specified, items are held indefinitely or until the user is removed from the hold. However, if you put a user who was already on an In-Place Hold on Litigation Hold, the Litigation Hold takes preference. Therefore, deleted items are moved to the Purges folder instead.

- **Versions**: When a user is put on In-Place Hold or Litigation Hold, mailbox items must be protected from tampering or modification by the user or a process. This is done by using a *copy-on-write* process. When a user or a process changes specific properties of a mailbox item, a copy of the original item is saved in the Versions folder before the change is committed. This process is repeated for subsequent changes. Items captured in the Versions folder are also indexed and returned in In-Place eDiscovery searches. After the hold is removed, copies in the Versions folder are removed by the Managed Folder Assistant.

Properties that trigger copy-on-write

ITEM TYPE	PROPERTIES THAT TRIGGER COPY-ON-WRITE
Messages (IPM.Note*) Posts (IPM.Post*)	Subject Body Attachments Senders/Recipients Sent/Received Dates
Items other than messages and posts	Any change to a visible property, except the following: • Item location (when an item is moved between folders) • Item status change (read or unread) • Changes to retention tag applied to an item
Items in the default folder Drafts	None (items in the Drafts folder are exempt from copy on write)

IMPORTANT

Copy-on-write is disabled for calendar items in the organizer's mailbox when meeting responses are received from attendees and the tracking information for the meeting is updated. For calendar items and items that have a reminder set, copy-on-write is disabled for the ReminderTime and ReminderSignalTime properties. Changes to these properties are not captured by copy-on-write. Changes to RSS feeds aren't captured by copy-on-write.

Although the DiscoveryHolds, Purges, and Versions folders aren't visible to the user, all items in the Recoverable Items folder are discoverable by using In-Place eDiscovery. After a mailbox user is removed from In-Place Hold or Litigation Hold, items in the DiscoveryHolds, Purges, and Versions folders are purged by the Managed Folder Assistant.

If a mailbox isn't placed on Litigation Hold or In-Place Hold, items in the Purges folder are permanently deleted from the Recoverable Items folder on a first in, first out basis when the item has resided in the folder for longer than the deleted item retention period.

Holds and mailbox quotas

Items in the Recoverable Items folder aren't calculated toward the user's mailbox quota. In Exchange, the Recoverable Items folder has its own quota. For Exchange, the default values for the *RecoverableItemsWarningQuota* and *RecoverableItemsQuota* mailbox properties are set to 20 GB and 30 GB respectively. To modify these values for a mailbox database for Exchange Server, use the Set-MailboxDatabase cmdlet. To modify them for individual mailboxes, use the Set-Mailbox cmdlet.

When a user's Recoverable Items folder exceeds the warning quota for recoverable items (as specified by the *RecoverableItemsWarningQuota* parameter), an event is logged in the Application event log of the Mailbox server. When the folder exceeds the quota for recoverable items (as specified by the *RecoverableItemsQuota* parameter), users won't be able to empty the Deleted Items folder or permanently delete mailbox items. Also copy-on-write won't be able to create copies of modified items. Therefore, it's critical that you monitor Recoverable Items quotas for mailbox users placed on In-Place Hold.

Holds and email forwarding

Users with mailboxes on Exchange Server can use Outlook and Outlook on the web to set up email forwarding for their mailbox. Email forwarding lets users configure their mailbox to forward email messages sent to their mailbox to another mailbox located inside or outside of their organization. Administrators can also set up mail flow rules (also known as transport rules) to forward messages to another mailbox. In both cases, email forwarding can be configured so that any message sent to the original mailbox isn't copied to that mailbox and is only sent to the forwarding address.

If a mailbox is on hold, additional steps are taken. During the delivery process, the hold settings for the mailbox are checked. If the message meets the hold criteria for the mailbox, a copy of the message is saved to the Inbox folder. That means you can use In-Place eDiscovery to search the original mailbox to find messages that were forwarded to another mailbox.

Preserving archived Skype for Business content

Exchange 2016 and Exchange 2019, Skype for Business, and SharePoint 2016 provide an integrated preservation and eDiscovery experience that allows you to preserve and search for items across the different data stores. Exchange 2016 and 2019 allow you to archive Skype for Business content in Exchange, removing the requirement of having a separate SQL Server database to store archived Lync content. The integrated hold and eDiscovery capability in SharePoint 2016 allows you to preserve and search data across all stores from a single console.

When you place an Exchange Server mailbox on In-Place Hold or Litigation Hold, Skype for Business content (such as instant messaging conversations and files shared in an online meeting) are archived in the mailbox. If you search the mailbox using In-Place eDiscovery, any archived Skype for Business content matching the search query is also returned in search results. You can also restrict the search to Skype for Business content archived in the mailbox.

To enable archiving of Skype for Business content in Exchange Server mailboxes, you must configure Skype for Business Server 2015 integration with Exchange Server. For details, see the following topics:

- Planning for Archiving
- Deploying Archiving

Deleting a mailbox on hold

If you delete a user account that has a mailbox, the Exchange Information store will eventually detect that the mailbox is no longer connected to a user account and mark that mailbox for deletion, even if the mailbox is on hold. If you want to preserve the mailbox, you have to do the following:

1. Instead of deleting the user account, disable the user account.
2. Change the properties of the mailbox to restrict the use and access to the mailbox. For example, set send and receive quotas equal to 1, block who can send messages to the mailbox, and restrict who can access the mailbox.
3. Retain the mailbox until all data has been removed or until preserving the data is no longer required.

Migrating mailboxes on hold from Exchange Server to Microsoft 365 or Office 365

If you have an Exchange hybrid deployment, the following conditions are true when you move (onboard) an on-premises Exchange Server mailbox to Exchange Online in Microsoft 365 or Office 365:

- If the on-premises mailbox is on Litigation Hold or In-Place Hold, the hold settings are preserved after the mailbox is moved to Exchange Online.
- If the on-premises mailbox is on Litigation Hold or In-Place Hold, any content in the Recoverable Items folder is moved to the Exchange Online mailbox.

Hold settings and content in the Recoverable Items folder are also preserved when you move (offboard) an Exchange Online mailbox to your on-premises Exchange Server organization.

> **TIP**
>
> For Exchange Server, an Exchange hybrid deployment is the recommended way to migrate on-premises mailboxes to Microsoft 365 or Office 365.

Create or remove an In-Place Hold

8/3/2020 • 7 minutes to read • Edit Online

An In-Place Hold preserves all mailbox and public folder content, including deleted items and original versions of modified items. All such items can be returned in an In-Place eDiscovery search. When you place an In-Place Hold on a user's mailbox, the contents in the corresponding archive mailbox (if it's enabled) are also placed on hold and returned in a eDiscovery search.

When you create an In-Place Hold, you can place all items in the source mailbox or public folder on hold or you can hold only the items that meet the search criteria specified for the hold. Similarly, you can hold items indefinitely or for a specific amount of time. For more information about In-Place Holds, see In-Place Hold and Litigation Hold in Exchange Server.

You can create In-Place holds in the Exchange admin center (EAC) or in the Exchange Management Shell.

Before you begin

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "In-Place Hold" entry in the Messaging policy and compliance permissions in Exchange Server topic.

- Depending on your Active Directory topology and replication latency, it may take up to an hour for an In-Place Hold to take effect.

- As previously explained, when you place an In-Place Hold on a user's mailbox, the content in the user's archive mailbox is also placed on hold.

- You can only search or place holds on all public folders in your organization. You can't specify individual public folders.

- See the More information section for a description of the In-Place Hold workflow process.

- To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.

Create an In-Place Hold

Use the EAC to create an In-Place Hold

1. Go to **Compliance management** > **In-Place eDiscovery & Hold**, and then click **New** +.

2. In the **New In-Place eDiscovery & Hold** window, on the **Name and description** page, type a name for the hold and an optional description, and then click **Next**.

3. On the **Mailboxes and Public folders** page, select the content sources to search:

   - To exclude mailboxes from the hold (and place a hold on public folders only), click **Don't search any mailboxes**.

   - To include specific mailboxes in the search, click **Specify mailboxes to search**, and then add the mailboxes that you want to search.

   - To place public folders on hold, click **Search all public folders**.

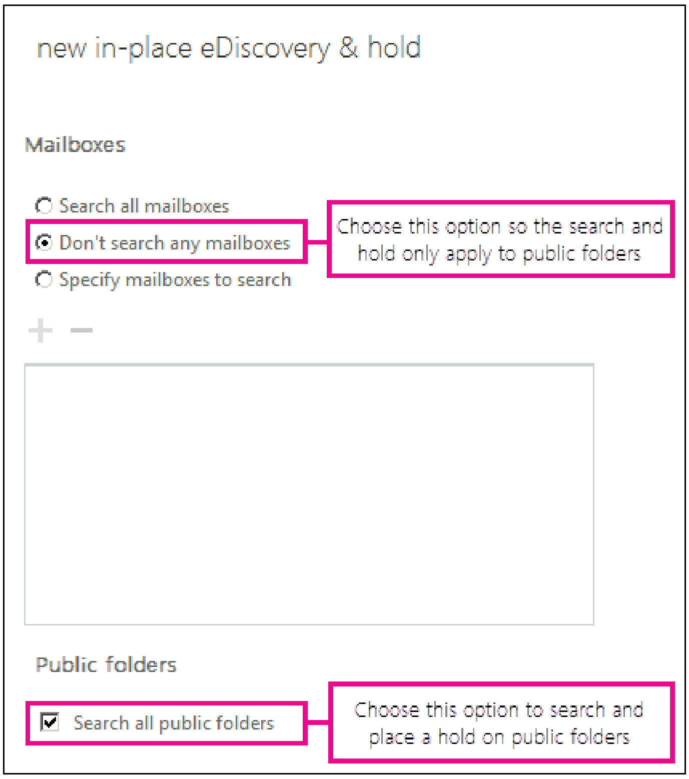

new in-place eDiscovery & hold

Mailboxes

○ Search all mailboxes

● Don't search any mailboxes

○ Specify mailboxes to search

Choose this option so the search and hold only apply to public folders

Public folders

☑ Search all public folders

Choose this option to search and place a hold on public folders

> **IMPORTANT**
>
> You can't select the **Search all mailboxes** option when creating an In-Place Hold. To create an In-Place Hold, you must select the specific mailboxes you want to place on hold.

4. On the **Search query** page, complete the following fields, and then click **Next**.

    - **Include all content**: Select this option to place all content in selected sources on hold.

    - **Filter based on criteria**: Select this option to specify search criteria, including keywords, start and end dates, sender and recipient addresses, and message types.

> **IMPORTANT**
>
> If a user is placed on multiple In-Place Holds, the search queries from any query-based hold are combined (with **OR** operators). In this case, the maximum number of keywords in all query-based holds placed on a mailbox is 500. If there are more than 500 keywords, then all content in the mailbox is placed on hold (not just that content that matches the search criteria). All content is held until the total number of keywords is reduced to 500 or less.

5. On the **In-Place Hold settings** page, click the **Place content matching the search query in selected sources on hold** check box and then select one of the following options:

    - **Hold indefinitely**: Place items returned by the search on an indefinite hold. Items on hold will be preserved until you change the hold duration, remove the mailbox (or public folders) from the search, or remove the search.

    - **Specify number of days to hold items relative to their received date**: Hold items for a specific period. For example, you can use this option if your organization requires that all messages be retained for at least seven years. You can use a *time-based* In-Place Hold along with a retention policy to make sure items are permanently deleted in seven years.

6. Click **Finish** to create the In-Place Hold.

Use the Exchange Management Shell to create an In-Place Hold

This example creates an In-Place Hold named Hold-CaseId012 and adds the mailbox joe@contoso.com to the hold.

> **IMPORTANT**
>
> If you don't specify additional search parameters for an In-Place Hold, all items in the specified source mailboxes are placed on hold. If you don't specify the *ItemHoldPeriod* parameter, items are placed on hold indefinitely or until the mailbox is either removed from hold or the hold is deleted.

```
New-MailboxSearch "Hold-CaseId012" -SourceMailboxes "joe@contoso.com" -InPlaceHoldEnabled $true
```

This example places an In-Place Hold on all public folders in the organization, and holds content for 7 years. The hold doesn't include any mailboxes.

```
New-MailboxSearch -Name "Hold for Public Folders" -AllPublicFolderSources $true -AllSourceMailboxes $false -
ItemHoldPeriod 2555 -InPlaceHoldEnabled $true
```

For detailed syntax and parameter information, see New-MailboxSearch.

How do you know this worked?

To verify that you have successfully created the In-Place Hold, do one of the following:

- Use the EAC to verify that the In-Place Hold is listed in the list view of the **In-Place eDiscovery & Hold** page.

- Use the **Get-MailboxSearch** cmdlet to retrieve the mailbox search and check the hold properties. For example, the following command displays the hold properties for the search named Hold-CaseId012:

  ```
  Get-MailboxSearch "Hold-CaseId012" | Format-List InPlaceHoldEnabled,ItemHoldPeriod,InPlaceHoldIdentity
  ```

- Use the **Get-Mailbox** cmdlet to display In-Place Hold information for specific user mailboxes or public folder mailboxes. For example, the following command displays the GUID for the In-Place Hold:

  ```
  Get-Mailbox "joe@contoso.com" | Format-List InPlaceHolds
  ```

  This example will display the In-Place Hold GUID for all public folder mailboxes in the organization.

  ```
  Get-Mailbox -PublicFolder | Format-List Name,InPlaceHolds
  ```

Remove an In-Place Hold

In Exchange Server, eDiscovery searches are used to hold and search for content in on content sources. You can't remove an In-Place eDiscovery search that's used to place content sources on hold. You must first remove the In-Place Hold by clearing the **Place content matching the search query in selected sources on hold** check box on the **In-Place Hold** page or by setting the *InPlaceHoldEnabled* parameter to `$false` in the Exchange Management Shell. Alternatively, you can remove mailboxes and public folders from an In-Place Hold by changing the value of the *SourceMailboxes* or *AllPublicFolderSources* parameters specified in the search.

Use the EAC to remove an In-Place Hold

1. Go to **Compliance management** > **In-Place eDiscovery & Hold**.

2. In the list view, select the In-Place Hold you want to remove, and then click **Edit** ✏.

3. In **In-Place eDiscovery & Hold** properties, on the **In-Place Hold** page, clear the **Place content matching the search query in selected sources on hold** check box, and then click **Save**.

4. Select the In-Place Hold again from the list view and then click **Delete** 🗑.

5. In warning, click **Yes** to remove the search.

Use the Exchange Management Shell to remove an In-Place Hold

This example first disables In-Place Hold named Hold-CaseId012 and then removes the mailbox search.

```
Set-MailboxSearch "Hold-CaseId012" -InPlaceHoldEnabled $false; Remove-MailboxSearch "Hold-CaseId012"
```

For detailed syntax and parameter information, see Set-Mailboxsearch.

How do you know this worked?

To verify that you have successfully removed an In-Place Hold, do one of the following:

- Use the EAC to verify that the In-Place Hold doesn't appear in the list view of the **In-Place eDiscovery & Hold** tab.

- Use the **Get-MailboxSearch** cmdlet to retrieve all mailbox searches and check that the search you removed is no longer listed.

More information

How does In-Place Hold work?: If a mailbox or public folder is not on hold, an item is moved to the Deletions subfolder in the Recoverable Items folder when it's permanently deleted (Shift + Delete) or deleted from the Deleted Items folder. A deletion policy (how long items are set to be retained) also moves items to the Deletions subfolder when the retention period expires. When a user purges an item in the Recoverable Items folder or when the deleted item retention period expires for an item, the item is moved to the Purges subfolder and marked for permanent deletion. It will be then be purged from Exchange the next time the mailbox is processed by the Managed Folder Assistant (MFA).

When an In-Place Hold is placed on a mailbox or public folder, purged items are not moved to the Purges subfolder but are instead moved to the DiscoveryHolds subfolder and are preserved for the hold duration specified by the In-Place Hold. The hold duration is calculated from the original date an item was received or created, and defines how long items in the DiscoveryHolds subfolder are held. When the hold duration expires for an item in the DiscoveryHolds subfolder, the item it is marked for permanent deletion and will be purged from Exchange the next time the mailbox or public folder is processed by the MFA. If an indefinite hold is placed on a mailbox or public folder, items will never be purged from the DiscoveryHolds subfolder.

The following illustration shows the subfolders in the Recoverable Items folders and the hold workflow process.

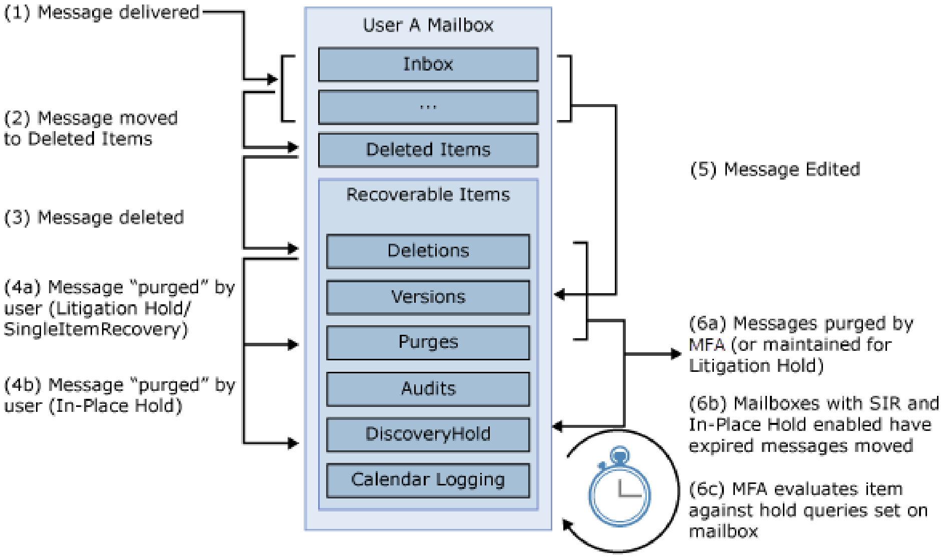
(1) Message delivered
User A Mailbox
Inbox
...
Deleted Items
(2) Message moved to Deleted Items
Recoverable Items
(3) Message deleted
Deletions
(4a) Message "purged" by user (Litigation Hold/ SingleItemRecovery)
Versions
Purges
Audits
(4b) Message "purged" by user (In-Place Hold)
DiscoveryHold
Calendar Logging
(5) Message Edited
(6a) Messages purged by MFA (or maintained for Litigation Hold)
(6b) Mailboxes with SIR and In-Place Hold enabled have expired messages moved
(6c) MFA evaluates item against hold queries set on mailbox

NOTE

If a mailbox is place on Litigation Hold, purged items are moved to the Purges subfolder and preserved for the hold duration configured for the Litigation Hold.

Place a mailbox on Litigation Hold

8/3/2020 • 8 minutes to read • Edit Online

Place a mailbox on Litigation Hold to preserve all mailbox content, including deleted items and original versions of modified items. When you place a mailbox on Litigation Hold, the user's archive mailbox (if it's enabled) is also placed on hold. Deleted and modified items are preserved for a specified period or until you remove the mailbox from Litigation Hold. All such mailbox items are returned in an In-Place eDiscovery in Exchange Server search.

Before you begin

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "In-Place Hold" entry in the Messaging policy and compliance permissions in Exchange Server topic.

- The Litigation Hold setting may take up to 60 minutes to take effect.

- Litigation Hold preserves items in the Recoverable Items folder in the user's mailbox. The default size for this folder is 30 GB. Depending on number and size of items deleted or modified, the size of the Recoverable Items folder of the mailbox may increase quickly. The Recoverable Items folder is configured with a high quota by default. We recommend that you monitor mailboxes that are placed on Litigation Hold on a weekly basis to ensure they don't reach the limits of the Recoverable Items quotas.

- As previously explained, when you place a user's mailbox on Litigation Hold, the user's archive mailbox is also placed on hold.

- Litigation Hold preserves deleted items and also preserves original versions of modified items until the hold is removed. You can optionally specify a hold duration, which preserves a mailbox item for the specified duration period. If you specify a hold duration period, it's calculated from the date a message is received or a mailbox item is created.

- See the More information section for a description of the Litigation Hold workflow process.

Use the EAC to place a mailbox on Litigation Hold

1. Go to **Recipients** > **Mailboxes**.

2. In the list of user mailboxes, click the mailbox that you want to place on Litigation Hold, and then click **Edit** ✏.

3. On the mailbox properties page, click **Mailbox features**.

4. Under **Litigation hold: Disabled**, click **Enable** to place the mailbox on Litigation Hold.

5. On the **Litigation Hold** page, enter the following optional information:

   - **Litigation hold duration (days)**: Use this box to specify how long mailbox items are held when the mailbox is placed on Litigation Hold. The duration is calculated from the date a mailbox item is received or created. If you leave this box blank, items are held indefinitely or until the hold is removed. Use days to specify the duration.

   - **Note**[*]: Use this box to inform the user their mailbox is on Litigation Hold. The note will appear on the **File** tab in Outlook 2010 or later.

   - **URL**[*]: Use this box to direct the user to a website for more information about Litigation Hold. This

URL appears on the **File** tab Outlook 2010 or later.

*If you leave the **Note** and **URL** values blank, the user isn't notified that you placed a litigation hold on their mailbox.

6. Click **Save** on the **Litigation Hold** page, and then click **Save** on the mailbox properties page.

Use the Exchange Management Shell to place a mailbox on Litigation Hold indefinitely

This example places the mailbox bsuneja@contoso.com on Litigation Hold. Items in the mailbox are held indefinitely or until the hold is removed.

```
Set-Mailbox bsuneja@contoso.com -LitigationHoldEnabled $true
```

NOTE

When you place a mailbox on Litigation Hold indefinitely (by not specifying a duration period), the value for the *LitigationHoldDuration* property mailbox is set to `Unlimited`.

Use the Exchange Management Shell to place a mailbox on Litigation Hold and preserve items for a specified duration

This example places the mailbox bsuneja@contoso.com on Litigation Hold and preserves items for 2555 days (approximately 7 years).

```
Set-Mailbox bsuneja@contoso.com -LitigationHoldEnabled $true -LitigationHoldDuration 2555
```

Use the Exchange Management Shell to place all mailboxes on Litigation Hold

Your organization may require that all mailbox data be preserved.

This example places all user mailboxes in the organization on Litigation Hold and sets the hold duration for one year (365 days).

```
Get-Mailbox -ResultSize Unlimited -Filter "RecipientTypeDetails -eq 'UserMailbox'" | Set-Mailbox -
LitigationHoldEnabled $true -LitigationHoldDuration 365
```

The example uses the Get-Mailbox cmdlet to retrieve all mailboxes in the organization, specifies a recipient filter to include all user mailboxes, and then pipes the list of mailboxes to the Set-Mailbox cmdlet to enable the Litigation Hold and set the hold duration.

To place all user mailboxes on an indefinite hold, run the previous command but don't include the *LitigationHoldDuration* parameter.

See the More information section for examples of using other recipient properties in a filter to include or exclude one or more mailboxes.

Use the Exchange Management Shell to remove a mailbox from Litigation Hold

This example removes the mailbox bsuneja@contoso.com from Litigation Hold.

```
Set-Mailbox bsuneja@contoso.com -LitigationHoldEnabled $false
```

How do you know this worked?

To verify that you have successfully placed a mailbox on Litigation Hold, do the one of the following:

- In the EAC:

  1. Go to **Recipients** > **Mailboxes**.

  2. In the list of user mailboxes, click the mailbox that you want to verify Litigation Hold settings for, and then click **Edit** ✏.

  3. On the mailbox properties page, click **Mailbox features**.

  4. Under **Litigation hold**, verify that hold is enabled.

  5. Click **View details** to verify when the mailbox was placed on Litigation Hold and by whom. You can also verify or change the values in the optional **Litigation hold duration (days)**, **Note**, and **URL** boxes.

- In the Exchange Management Shell, run one of the following commands:

  ```
  Get-Mailbox <name of mailbox> | Format-List LitigationHold*
  ```

  or

  ```
  Get-Mailbox -ResultSize Unlimited -Filter "RecipientTypeDetails -eq 'UserMailbox'" | Format-List
  Name,LitigationHold*
  ```

  If a mailbox is placed on Litigation Hold indefinitely, the value for the *LitigationHoldDuration* property mailbox is set to `Unlimited`.

More information

- **How does Litigation Hold work?** In the normal deleted item workflow, a mailbox item is moved to the Deletions subfolder in the Recoverable Items folder when a user permanently deletes it (Shift + Delete) or deletes it from the Deleted Items folder. A deletion policy (which is a retention tag configured with a Delete retention action) also moves items to the Deletions subfolder when the retention period expires. When a user purges an item in the Recoverable Items folder or when the deleted item retention period expires for an item, it's moved to the Purges subfolder in the Recoverable Items folder and marked for permanent deletion. It will be purged from Exchange the next time the mailbox is processed by the Managed Folder Assistant (MFA).

  When a mailbox is placed on Litigation Hold, items in the Purges subfolder are preserved for the hold duration specified by the Litigation Hold. The hold duration is calculated from the original date an item was received or created, and defines how long items in the Purges subfolder are held. When the hold duration expires for an item in the Purges subfolder, the item is marked for permanent deletion and will be purged from Exchange the next time the mailbox is processed by the MFA. If an indefinite hold is placed on a mailbox, items will never be purged from the Purges subfolder.

  The following illustration shows the subfolders in the Recoverable Items folders and the hold workflow process.

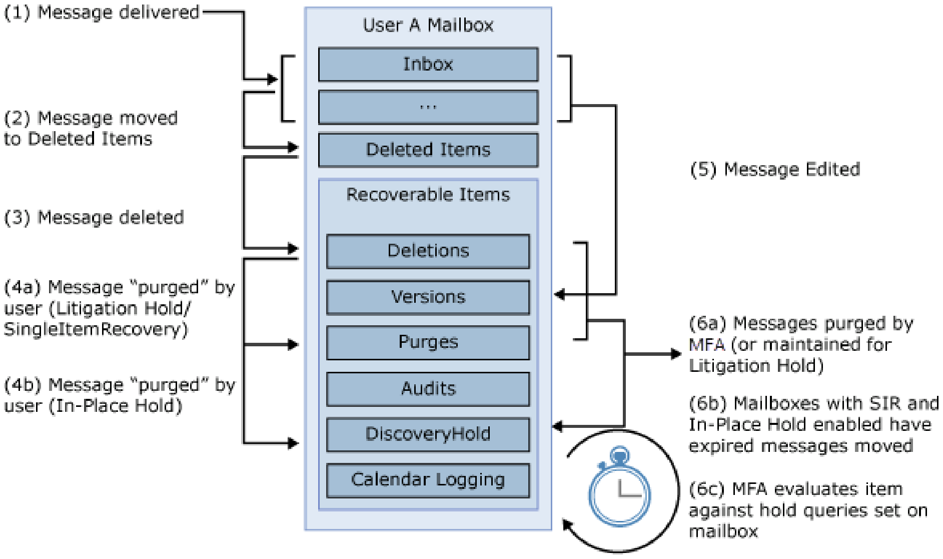
(1) Message delivered
(2) Message moved to Deleted Items
(3) Message deleted
(4a) Message "purged" by user (Litigation Hold/ SingleItemRecovery)
(4b) Message "purged" by user (In-Place Hold)
User A Mailbox
Inbox
...
Deleted Items
Recoverable Items
Deletions
Versions
Purges
Audits
DiscoveryHold
Calendar Logging
(5) Message Edited
(6a) Messages purged by MFA (or maintained for Litigation Hold)
(6b) Mailboxes with SIR and In-Place Hold enabled have expired messages moved
(6c) MFA evaluates item against hold queries set on mailbox

NOTE

If an In-Place Hold is placed on a mailbox, purged items are moved from the Deletions subfolder to the DiscoveryHolds subfolder and are preserved for the hold duration for the In-Place Hold.

- If your organization requires that all mailbox data has to preserved for a specific period of time, consider the following before you place all mailboxes in an organization on Litigation Hold.
  - When you use the previous command to place a hold on all mailboxes in an organization (or a subset of mailboxes matching a specified recipient filter) only mailboxes that exist at the time that you run the command are placed on hold. If you create new mailboxes later, you have to run the command again to place the new mailboxes on hold. If you create new mailboxes often, you can run the command as a scheduled task as frequently as required.
  - Placing all mailboxes on Litigation Hold can significantly impact mailbox sizes. In an Exchange 2016 or Exchange 2019 organization, plan for adequate storage to meet your organization's preservation requirements.
  - The Recoverable Items folder has its own storage limit, so items in the folder don't count towards the mailbox storage limit. As previously explained, preserving mailbox data for a long period of time will result in growth of the Recoverable Items folder in a user's mailbox and archive. We recommend that you periodically monitor the size of this folder by using the **Get-MailboxFolderStatistics** cmdlet to ensure it doesn't reach the limit. For more information, see:
    - Get-MailboxFolderStatistics
    - Clean up or delete items from the Recoverable Items folder.
- The previous command to place a hold on all mailboxes uses a recipient filter that returns all user mailboxes. You can use other recipient properties to return a list of specific mailboxes that you can then pipe to the **Set-Mailbox** cmdlet to place a Litigation Hold on those mailboxes.

  Here are some examples of using the **Get-Mailbox** and **Get-Recipient** cmdlets to return a subset of mailboxes based on common user or mailbox properties. These examples assume that relevant mailbox properties (such as *CustomAttributeN* or *Department*) have been populated.

```
Get-Mailbox -RecipientTypeDetails UserMailbox -ResultSize unlimited -Filter 'CustomAttribute15 -eq
"OneYearLitigationHold"'
```

```
Get-Recipient -RecipientTypeDetails UserMailbox -ResultSize unlimited -Filter 'Department -eq "HR"'
```

```
Get-Recipient -RecipientTypeDetails UserMailbox -ResultSize unlimited -Filter 'PostalCode -eq "98052"'
```

```
Get-Recipient -RecipientTypeDetails UserMailbox -ResultSize unlimited -Filter 'StateOrProvince -eq "WA"'
```

```
Get-Mailbox -ResultSize Unlimited -Filter "RecipientTypeDetails -ne 'DiscoveryMailbox'"
```

You can use other user mailbox properties in a filter to include or exclude mailboxes. For details, see Filterable Properties for the -Filter Parameter.

Place all mailboxes on hold

8/3/2020 • 6 minutes to read • Edit Online

Your organization may require all mailbox data to be preserved for a specific period. You can use Litigation Hold or In-Place Hold to meet this requirement. After you place a mailbox on Litigation Hold or In-Place Hold, mailbox items that are modified or that are permanently deleted are preserved in the Recoverable Items folder for the duration specified by the hold. For more information, see In-Place Hold and Litigation Hold in Exchange Server.

Before you place all mailboxes in an organization on Litigation Hold or In-Place Hold, consider the following:

- When you place mailboxes on hold, content in a user's archive mailbox (if it's enabled) is also placed on hold.
- Placing all mailboxes in an organization on hold can significantly impact mailbox sizes. In an Exchange Server deployment, plan for adequate storage to meet your organization's preservation requirements.
- Preserving mailbox data for a long duration will result in growth of the Recoverable Items folder in a user's primary mailbox and archive mailbox. The Recoverable Items folder has its own storage limit, so items in the folder don't count towards the mailbox storage limit. In Exchange Server, the default storage limit for the Recoverable Items folder is 30 GB. We recommend that you periodically monitor the size of this folder to ensure it doesn't reach the limit. For more information, see Recoverable Items folder in Exchange Server.

Choosing between Litigation Hold and In-Place Hold

Here are some factors to consider when deciding the hold feature you should use to place all mailboxes in your organization on hold.

IF YOU WANT TO...	USE LITIGATION HOLD	USE IN-PLACE HOLD
Use the EAC	Yes For setting a Litigation Hold, the EAC is best suited for quick one-off actions on a few mailboxes. We recommend using the Exchange Management Shell for placing a Litigation Hold for all users in your organization. For details, see Place a mailbox on Litigation Hold.	Yes However, you can't select more than 500 source mailboxes in the EAC. For details, see Create or remove an In-Place Hold.
Use the Exchange Management Shell	Yes	Yes
Place more than 10,000 mailboxes on hold	Yes Litigation Hold is a mailbox property. You can place all mailboxes in an organization on hold by using the **Set-Mailbox** cmdlet.	Yes; use multiple In-Place Holds You can use distribution groups to specify a maximum of 10,000 mailboxes in a single In-Place Hold. To place additional mailboxes on hold, you must create additional In-Place Holds. This will result in additional management overhead. Using Litigation Hold placing large numbers of mailboxes on hold is simpler.

IF YOU WANT TO...	USE LITIGATION HOLD	USE IN-PLACE HOLD
Place many different mailboxes on hold for different periods.	Yes Litigation Hold is a mailbox property. You can place each mailbox (or sets of mailboxes) on hold for a different duration. See the More information section for examples of using recipient properties in a filter so you can place a Litigation Hold on a subset of mailboxes.	Yes If you're placing individual holds on thousands of mailboxes, we recommend using Litigation Hold. However, if you're creating holds for specific events that involve multiple users (such as a legal case), use a single in-Place hold for the group of users. It's not recommended to create separate In-Place Holds for each mailbox as this will create many In-Place Hold queries. This will be more difficult to manage than Litigation Holds. A large number of In-Place Hold objects may also result in slow performance in the EAC when refreshing, creating, or modifying In-Place eDiscovery or In-Place Hold objects.
Automatically place new mailboxes on hold	No You have to place a new mailbox on Litigation Hold after it's created. You can schedule the command or script to run as frequently as required to achieve the same effect.	No You have to add a new mailbox to an existing In-Place Hold, even if you specified a distribution group when you created the In-Place Hold. You can also schedule the command or script to run as frequently as required to achieve the same effect. We recommend that the script check if an existing In-Place Hold has already reached the 10,000 mailbox limit, and then create a new In-Place Hold if required.

Place all mailboxes on Litigation Hold

You can easily and quickly place all mailboxes on hold indefinitely or for a specified hold duration using the Exchange Management Shell. This command places all mailboxes on hold with a hold duration of 2555 days (approximately 7 years).

```
Get-Mailbox -ResultSize Unlimited -Filter "RecipientTypeDetails -eq 'UserMailbox'" | Set-Mailbox -
LitigationHoldEnabled $true -LitigationHoldDuration 2555
```

The example uses the Get-Mailbox cmdlet and a recipient filter to retrieve all user mailboxes in the organization, and then pipes the list of mailboxes to the Set-Mailbox cmdlet to enable the Litigation Hold and specify a hold duration. For more information, see Place a mailbox on Litigation Hold.

Place all mailboxes on In-Place Hold

You can use the EAC to select up to 500 mailboxes and place them on hold. For details, see Create or remove an In-Place Hold.

To place more than 500 users on In-Place Hold, use the Exchange Management Shell. For details, see New-MailboxSearch.

TIP

In hybrid environments, you can use the inactive mailbox feature in Exchange Online to retain mailboxes without consuming a license or requiring an account for the mailbox. The inactive mailbox feature requires an "Exchange Online Plan 2, Office 365 Enterprise E3 and E5 subscriptions" license. If you have an Exchange Online Plan 1 license, you need to assign a separate "Exchange Online Archiving" license to the mailbox. For more information, see Inactive mailbox.

More information

- When you place all mailboxes in your organization on hold, only the mailboxes that exist at the time you run the command are placed on hold. If you create new mailboxes later, run the command again to place them on hold. If you frequently create new mailboxes, you can run the command as a scheduled task as frequently as required.

- Placing mailboxes on hold preserves data by preventing items in the Recoverable Items folder from being deleted until the specified hold duration for an item expires. If a hold is configured to hold items indefinitely, items won't be purged from a mailbox. Also, when a mailbox is on hold the original version of a message is saved before it's modified. Combine Litigation Hold or In-Place Hold with a Retention Policy, which can automatically delete messages (and move them into the Recoverable Items folder) after a specified period, to meet your organization's email retention requirements. See Retention tags and retention policies in Exchange Server for details.

- The Exchange Management Shell command used in this topic to place a Litigation Hold on all mailboxes uses a recipient filter that returns all user mailboxes. You can use other recipient properties to return a list of specific mailboxes that you can then pipe to the **Set-Mailbox** cmdlet to place a Litigation Hold on those mailboxes.

  Here are some examples of using the **Get-Mailbox** and **Get-Recipient** cmdlets to return a subset of mailboxes based on common user or mailbox properties. These examples assume that relevant mailbox properties (such as *CustomAttributeN* or *Department*) have been populated.

  ```
  Get-Mailbox -RecipientTypeDetails UserMailbox -ResultSize unlimited -Filter 'CustomAttribute15 -eq
  "OneYearLitigationHold"'
  ```

  ```
  Get-Recipient -RecipientTypeDetails UserMailbox -ResultSize unlimited -Filter 'Department -eq "HR"'
  ```

  ```
  Get-Recipient -RecipientTypeDetails UserMailbox -ResultSize unlimited -Filter 'PostalCode -eq "98052"'
  ```

  ```
  Get-Recipient -RecipientTypeDetails UserMailbox -ResultSize unlimited -Filter 'StateOrProvince -eq "WA"'
  ```

  ```
  Get-Mailbox -ResultSize Unlimited -Filter "RecipientTypeDetails -ne 'DiscoveryMailbox'"
  ```

  You can use other user mailbox properties in a filter to include or exclude mailboxes. For details, see Filterable Properties for the -Filter Parameter.

Preserve Bcc and expanded distribution group recipients for eDiscovery

8/3/2020 • 4 minutes to read • Edit Online

In-Place Hold and Litigation Hold allow you to preserve mailbox content to meet regulatory compliance and eDiscovery requirements. Information about recipients directly addressed in the To and Cc fields of a message is included in all messages by default, but your organization may require the ability to search for and reproduce details about all recipients of a message. This includes

- **Recipients addressed using the Bcc field of a message**: Bcc recipients are stored in the message in the sender's mailbox, but not included in headers of the message delivered to recipients.
- **Expanded distribution group recipients**: Recipients who receive the message because they're members of a distribution group to which the message was addressed, either in the To, Cc or Bcc fields.

Exchange 2016 and Exchange 2019 preserve information about Bcc and expanded distribution group recipients. You can search for this information by using an In-Place eDiscovery search .

How Bcc recipients and expanded distribution group recipients are preserved

As stated earlier, information about Bcc'ed recipients is stored with the message in the sender's mailbox. This information is indexed and available to In-Place eDiscovery and Hold.

Information about expanded distribution group recipients is stored with the message after you place a mailbox on In-Place Hold or Litigation Hold. Distribution group membership is determined at the time the message is sent. The expanded recipients list stored with the message is not impacted by changes to membership of the group after the message is sent.

INFORMATION ABOUT...	IS STORED IN...	IS STORED BY DEFAULT?	IS ACCESSIBLE TO...
To and Cc recipients	Message properties in the sender and recipients' mailboxes.	Yes	Sender, recipients, and compliance officers
Bcc recipients	Message property in the sender's mailbox.	Yes	Sender and compliance officers
Expanded distribution group recipients	Message properties in the sender's mailbox.	No. Expanded distribution group recipient information is stored after a mailbox is placed on In-Place Hold or Litigation Hold.	Compliance officers

Searching for messages sent to Bcc and expanded distribution group recipients

When searching for messages sent to a recipient, eDiscovery search results now include messages sent to a distribution group that the recipient is a member of. The following table shows the scenarios where messages sent to Bcc and expanded distribution group recipients are returned in eDiscovery searches.

Scenario 1: John is a member of the US-Sales distribution group. This table shows eDiscovery search results when Bob sends a message to John directly or indirectly via a distribution group.

WHEN YOU SEARCH BOB'S MAILBOX FOR MESSAGES SENT...	AND THE MESSAGE IS SENT WITH...	RESULTS INCLUDE MESSAGE?
To:John	John on TO	Yes
To:John	US-Sales on TO	Yes
To:US-Sales	US-Sales on TO	Yes
Cc:John	John on CC	Yes
Cc:John	US-Sales on CC	Yes
Cc:US-Sales	US-Sales on CC	Yes

Scenario 2: Bob sends an email to John (To/Cc) and Jack (Bcc directly, or indirectly via a distribution group). The table below shows eDiscovery search results.

WHEN YOU SEARCH...	FOR MESSAGES SENT...	RESULTS INCLUDE MESSAGE?	NOTES
Bob's mailbox	To/Cc:John	Yes	Presents an indication that Jack was Bcc'ed.
Bob's mailbox	Bcc:Jack	Yes	Presents an indication that Jack was Bcc'ed.
Bob's mailbox	Bcc:Jack (via distribution group)	Yes	List of members of the Bcc'ed distribution group, expanded when the message was sent, is visible in eDiscovery search preview, export and logs.
John's mailbox	To/Cc:John	Yes	No indication of Bcc recipients.
John's mailbox	Bcc:Jack (directly or via distribution group)	No	Bcc information is not stored in the message delivered to recipients. You must search the sender's mailbox.
Jack's mailbox	To/Cc:John (directly or via distribution group)	Yes	To/Cc information is included in message delivered to all recipients.
Jack's mailbox	Bcc:Jack (directly or via distribution group)	No	Bcc information is not stored in the message delivered to recipients. You must search the sender's mailbox.

Frequently asked questions

Q. When and where is Bcc recipient information stored?

A. Bcc recipient information is preserved by default in the original message in sender's mailbox. If the Bcc recipient is a distribution group, distribution group membership is only expanded if the sender's mailbox is on hold.

Q. When and where is the list of expanded distribution group recipients stored?

A. Group membership is expanded at the time the message is sent. The list of expanded distribution group members is stored in the original message in the sender's mailbox. The sender's mailbox must be on In-Place Hold or Litigation Hold.

Q. Can the To/Cc recipients see which recipients were Bcc'ed?

A. No. This information is not included in message headers, and isn't visible to To/Cc recipients. The sender can see the Bcc field stored in the original message stored in their mailbox. Compliance officers can see this information when searching the sender's mailbox.

Q. How can I ensure expanded distribution group recipients are always preserved?

A. To ensure expanded distribution group members are always preserved with a message, Place all mailboxes on hold.

Q. Which types of groups are supported?

A. Distribution groups, mail-enabled security groups, and dynamic distribution groups are supported.

Q. Is there a limit on the number of distribution group recipients that are expanded and stored in the message?

A. Up to 10,000 members of a distribution group is preserved.

Q. Are nested distribution groups supported?

A. Yes, 25 levels of nested distribution groups are expanded.

Q. Where is the Bcc and expanded distribution group recipient information visible?

A. Bcc and expanded distribution group recipients information is visible to Compliance officers when performing an eDiscovery search. Bcc and expanded distribution group recipients are included in search results copied to a Discovery mailbox or exported to a PST file and in the eDiscovery log included in search results. Bcc recipient information is also available in search preview.

Q. What happens if a member of a distribution group is hidden from the organization's global address list (GAL)?

A. There's no impact. If recipients are hidden from the GAL, they're still included in the list of recipients for the expanded distribution group.

In-Place eDiscovery in Exchange Server

8/3/2020 • 25 minutes to read • Edit Online

If your organization adheres to legal discovery requirements (related to organizational policy, compliance, or lawsuits), In-Place eDiscovery in Exchange Server can help you perform discovery searches for relevant content within mailboxes. You can also use In-Place eDiscovery in an Exchange hybrid environment to search on-premises and cloud-based mailboxes in the same search.

> **IMPORTANT**
> In-Place eDiscovery is a powerful feature that allows a user with the correct permissions to potentially gain access to all messaging records stored throughout the Exchange Server organization. It's important to control and monitor discovery activities, including addition of members to the Discovery Management role group, assignment of the Mailbox Search management role, and assignment of mailbox access permission to discovery mailboxes.

How In-Place eDiscovery works

In-Place eDiscovery uses the content indexes created by Exchange Search. Role Based Access Control (RBAC) provides the Discovery Management role group to delegate discovery tasks to non-technical personnel, without the need to provide elevated privileges that may allow a user to make any operational changes to Exchange configuration. The Exchange admin center (EAC) provides an easy-to-use search interface for non-technical personnel such as legal and compliance officers, records managers, and human resources professionals.

Authorized users can perform an In-Place eDiscovery search by selecting the mailboxes, and then specifying search criteria such as keywords, start and end dates, sender and recipient addresses, and message types. After the search is complete, authorized users can then select one of the following actions:

- **Estimate search results** - This option returns an estimate of the total size and number of items that will be returned by the search based on the criteria you specified.
- **Preview search results** - This option provides a preview of the results. Messages returned from each mailbox searched are displayed.
- **Copy search results** - This option lets you copy messages to a discovery mailbox.
- **Export search results** - After search results are copied to a discovery mailbox, you can export them to a PST file.

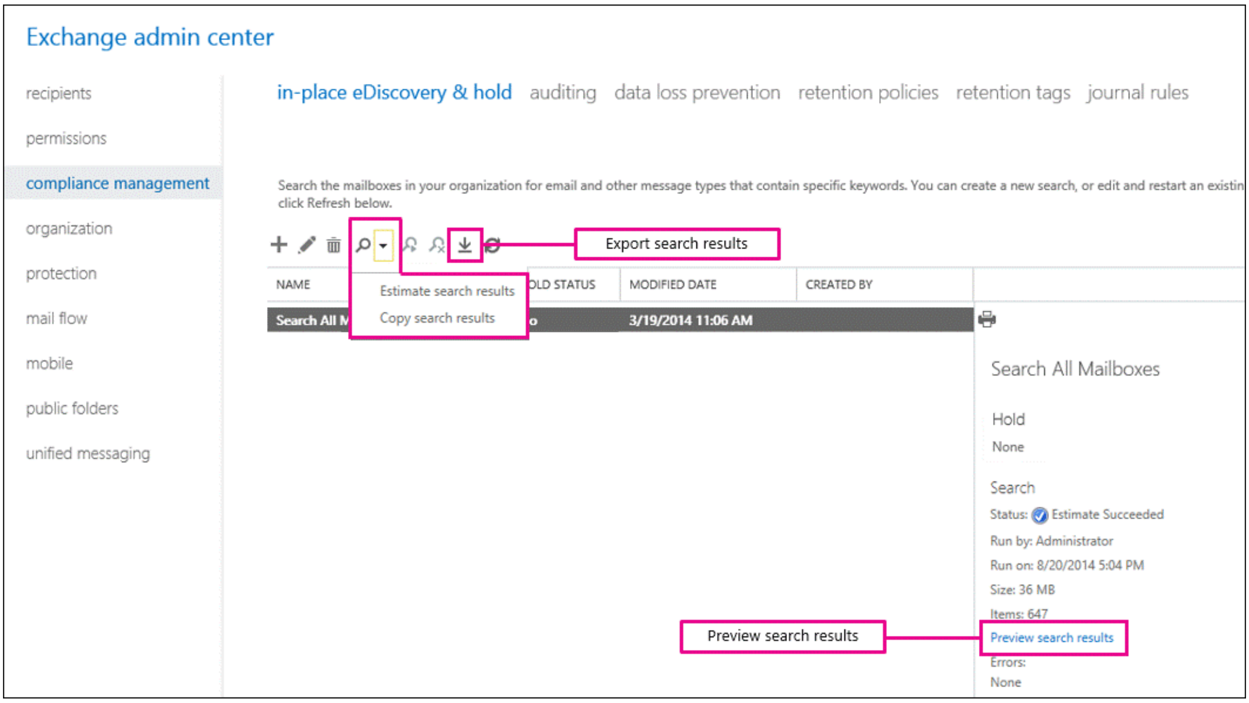

Search the mailboxes in your organization for email and other message types that contain specific keywords. You can create a new search, or edit and restart an existin click Refresh below.

NAME	OLD STATUS	MODIFIED DATE	CREATED BY
Search All M	o	3/19/2014 11:06 AM	

Search All Mailboxes

Hold
None

Search
Status: Estimate Succeeded
Run by: Administrator
Run on: 8/20/2014 5:04 PM
Size: 36 MB
Items: 647
Preview search results
Errors:
None

In-Place eDiscovery uses Keyword Query Language (KQL). Users familiar with KQL can construct powerful search queries to search content indexes. For more information about KQL, see Keyword Query Language syntax reference.

In-Place eDiscovery permissions

For authorized users to perform In-Place eDiscovery searches, you need to add them to the Discovery Management role group. This role group consists of two management roles: the Mailbox Search Role, which allows a user to perform an In-Place eDiscovery search, and the Legal Hold Role, which allows a user to place a mailbox on In-Place Hold and Litigation Hold.

By default, permissions to perform In-Place eDiscovery-related tasks aren't assigned to any user or Exchange administrators. Exchange administrators who are members of the Organization Management role group can add users to the Discovery Management role group and create custom role groups to narrow the scope of a discovery manager to a subset of users. To learn more about adding users to the Discovery Management role group, see Assign eDiscovery permissions in Exchange Server.

IMPORTANT

If a user isn't added to the Discovery Management role group or isn't assigned the Mailbox Search role, the **In-Place eDiscovery & Hold** user interface isn't displayed in the EAC, and the In-Place eDiscovery (***MailboxSearch**) cmdlets aren't available in the Exchange Management Shell.

Auditing of RBAC role changes, which is enabled by default, makes sure that adequate records are kept to track assignment of the Discovery Management role group. You can use the administrator role group report to search for changes made to administrator role groups. For more information, see Search the role group changes or administrator audit logs.

Using In-Place eDiscovery

Users who have been added to the Discovery Management role group can perform In-Place eDiscovery searches. You can perform a search using the web-based interface in the EAC. This makes it easier for non-technical users such as records managers, compliance officers, or legal and HR professionals to use In-Place eDiscovery. You can also use the Exchange Management Shell to perform a search. For more information, see Create an In-Place eDiscovery search in Exchange Server

The **In-Place eDiscovery & Hold** wizard in the EAC allows you to create an In-Place eDiscovery search and also use In-Place Hold to place search results on hold. When you create an In-Place eDiscovery search, a search object is created in the In-Place eDiscovery system mailbox. This object can be manipulated to start, stop, modify, and remove the search. After you create the search, you can choose to get an estimate of search results, which includes keyword statistics that help you determine query effectiveness. You can also do a live preview of items returned in the search, allowing you to view message content, the number of messages returned from each source mailbox and the total number of messages. You can use this information to further fine-tune your query if required.

When satisfied with the search results, you can copy them to a discovery mailbox. You can also use the EAC or Outlook to export a discovery mailbox or some of its content to a PST file.

When creating an In-Place eDiscovery search, you must specify the following parameters:

- **Name** - The search name is used to identify the search. When you copy search results to a discovery mailbox, a folder is created in the discovery mailbox using the search name and the timestamp to uniquely identify search results in a discovery mailbox.

- **Sources** - You can choose to search all mailboxes in your Exchange Server organization or specify the mailboxes to search. You can also choose to search all public folders. If you also want to use the same search to place items on hold, you must specify the mailboxes. You can also place all public folders on In-Place Hold. You can specify a distribution group to include mailbox users who are members of that group. Membership of the group is calculated once when creating the search and subsequent changes to group membership aren't automatically reflected in the search. A user's primary and archive mailboxes are included in the search.

- **Search query** - You can either include all mailbox content from the specified mailboxes or use a search query to return items that are more relevant to the case or investigation. You can specify the following parameters in a search query:

  - **Keywords** - You can specify keywords and phrases to search message content. You can also use the logical operators **AND**, **OR**, and **NOT**. Additionally, Exchange Server also supports the **NEAR** operator, allowing you to search for a word or phrase that's in proximity to another word or phrase.

    To search for an exact match of a multiple word phrase, you must enclose the phrase in quotation marks. For example, searching for the phrase "plan and competition" returns messages that contain an exact match of the phrase, whereas specifying **plan AND competition** returns messages that contain the words **plan** and **competition** anywhere in the message.

    Exchange Server also supports the Keyword Query Language (KQL) syntax for In-Place eDiscovery searches. For more information about KQL, see Keyword Query Language syntax reference.

    > **NOTE**
    >
    > In-Place eDiscovery does not support regular expressions.

    You must capitalize logical operators such as **AND** and **OR** for them to be treated as operators instead of keywords. We recommend that you use explicit parenthesis for any query that mixes multiple logical operators to avoid mistakes or misinterpretations. For example, if you want to search for messages that contain either WordA or WordB AND either WordC or WordD, you must use **(WordA OR WordB) AND (WordC OR WordD)**.

  - **Start and End dates** - By default, In-Place eDiscovery doesn't limit searches by a date range. To search messages sent during a specific date range, you can narrow the search by specifying the start and end dates. If you don't specify an end date, the search will return the latest results every time you restart it.

- **Senders and recipients** - To narrow down the search, you can specify the senders or recipients of messages. You can use email addresses, display names, or the name of a domain to search for items sent to or from everyone in the domain. For example, to find email sent by or sent to anyone at Contoso, Ltd, specify **@contoso.com** in the **From** or the **To/cc** field in the EAC. You can also specify **@contoso.com** in the *Senders* or *Recipients* parameters in the Exchange Management Shell

- **Message types** - By default, all message types are searched. You can restrict the search by selecting specific message types such as email, contacts, documents, journal, meetings, notes and Lync content.

The following screenshot shows an example of a search query in the EAC.

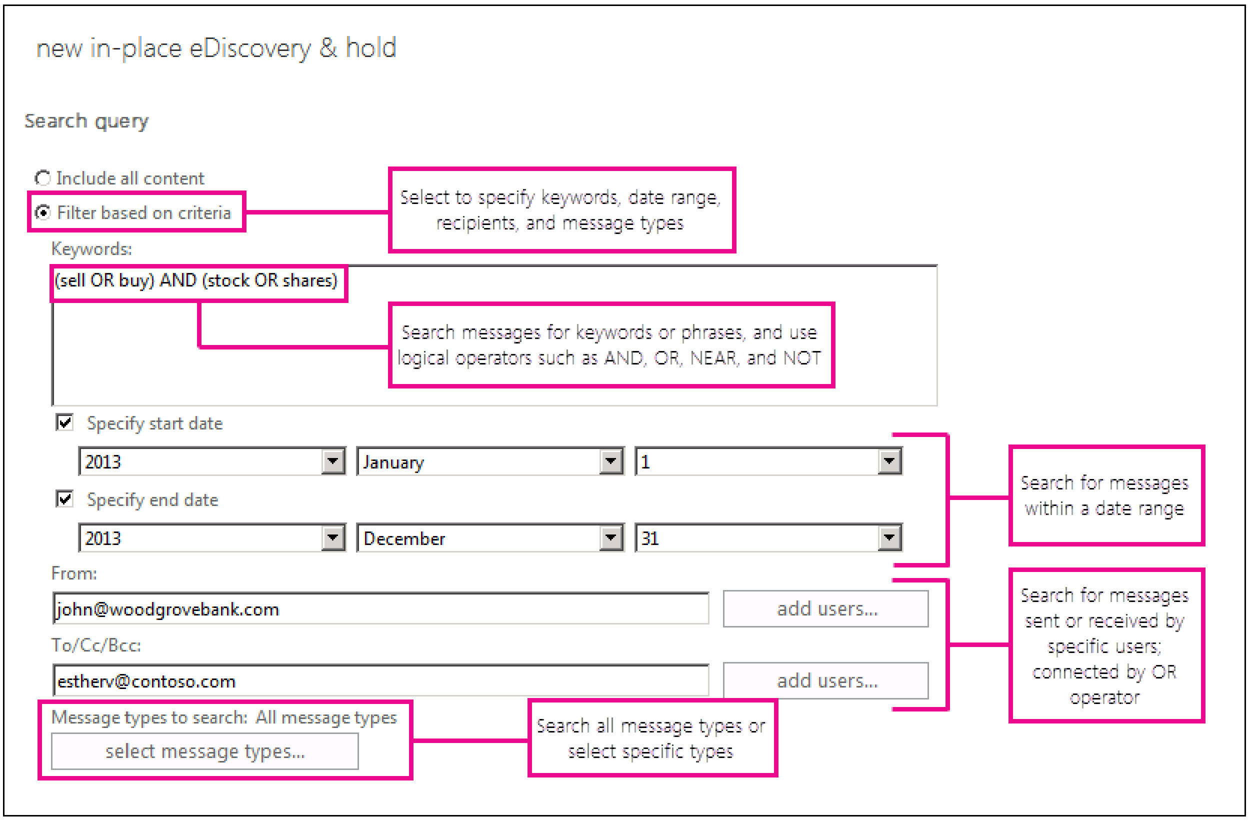

new in-place eDiscovery & hold

Search query

○ Include all content
● Filter based on criteria

Select to specify keywords, date range, recipients, and message types

Keywords:
(sell OR buy) AND (stock OR shares)

Search messages for keywords or phrases, and use logical operators such as AND, OR, NEAR, and NOT

☑ Specify start date
2013 | January | 1

☑ Specify end date
2013 | December | 31

Search for messages within a date range

From:
john@woodgrovebank.com

To/Cc/Bcc:
estherv@contoso.com

Search for messages sent or received by specific users; connected by OR operator

Message types to search: All message types

Search all message types or select specific types

When using In-Place eDiscovery, also consider the following:

- **Attachments** - In-Place eDiscovery searches attachments supported by Exchange Search. For details, see Default Filters for Exchange Search. In on-premises deployments, you can add support for additional file types by installing search filters (also known as an iFilter) for the file type on Mailbox servers.

- **Unsearchable items** - Unsearchable items are mailbox items that can't be indexed by Exchange Search. Reasons they can't be indexed include the lack of an installed search filter for an attached file, a filter error, and encrypted messages. For a successful eDiscovery search, your organization may be required to include such items for review. When copying search results to a discovery mailbox or exporting them to a PST file, you can include unsearchable items. For more information, see Unsearchable Items in Exchange eDiscovery.

- **Encrypted items** - Because messages encrypted using S/MIME aren't indexed by Exchange Search, In-Place eDiscovery doesn't search these messages. If you select the option to include unsearchable items in search results, these S/MIME encrypted messages are copied to the discovery mailbox.

- **De-duplication** - When copying search results to a discovery mailbox or exporting search results to a PST file, you can enable *de-duplication* of search results to copy only one instance of a unique message to the discovery mailbox. De-duplication has the following benefits:

  - Lower storage requirement and smaller discovery mailbox size due to reduced number of messages copied.

  - Reduced workload for discovery managers, legal counsel, or others involved in reviewing search results.

  - Reduced cost of eDiscovery, depending on the number of duplicate items excluded from search results.

- **IRM-protected items** - Messages protected using Information Rights Management (IRM) are indexed by Exchange Search and therefore included in the search results if they match query parameters. Messages must be protected by using an Active Directory Rights Management Services (AD RMS) cluster in the same Active Directory forest as the Mailbox server. For more information, see Information Rights Management in Exchange Server.

  Important:

  - When Exchange Search fails to index an IRM-protected message, either due to a decryption failure or because IRM is disabled, the protected message isn't added to the list of failed items. If you select the option to include unsearchable items in search results, the results may not include IRM-protected messages that could not be decrypted.

  - To include IRM-protected messages in a search, you can create another search to include messages with .rpmsg attachments. You can use the query string `attachment:rpmsg` to search all IRM-protected messages in the specified mailboxes, whether successfully indexed or not. This may result in some duplication of search results in scenarios where one search returns messages that match the search criteria, including IRM-protected messages that have been indexed successfully. The search doesn't return IRM-protected messages that couldn't be indexed.

  - Performing a second search for all IRM-protected messages also includes the IRM-protected messages that were successfully indexed and returned in the first search. Additionally, the IRM-protected messages returned by the second search may not match the search criteria such as keywords used for the first search.

Estimate, preview, and copy search results

After an In-Place eDiscovery search is completed, you can view search result estimates in the Details pane in the EAC. The estimate includes number of items returned and total size of those items. You can also view keyword statistics, which returns details about number of items returned for each keyword used in the search query. This information is helpful in determining query effectiveness. If the query is too broad, it may return a much bigger data set, which could require more resources to review and raise eDiscovery costs. If the query is too narrow, it may significantly reduce the number of records returned or return no records at all. You can use the estimates and keyword statistics to fine-tune the query to meet your requirements.

> **NOTE**
>
> In Exchange Server, keyword statistics also include statistics for non-keyword properties such as dates, message types, and senders/recipients specified in a search query.

You can also preview the search results to further ensure that messages returned contain the content you're searching for and further fine-tune the query if required. eDiscovery Search Preview displays the number of messages returned from each mailbox searched and the total number of messages returned by the search. The preview is generated quickly without requiring you to copy messages to a discovery mailbox.

After you're satisfied with the quantity and quality of search results, you can copy them to a discovery mailbox. When copying messages, you have the following options:

- **Include unsearchable items** - For details about the types of items that are considered unsearchable, see the eDiscovery search considerations in the previous section.

- **Enable de-duplication** - De-duplication reduces the dataset by only including a single instance of a unique record if multiple instances are found in one or more mailboxes searched.

- **Enable full logging** - By default, only basic logging is enabled when copying items. You can select full logging to include information about all records returned by the search.

- **Send me mail when the copy is completed** - An In-Place eDiscovery search can potentially return a large number of records. Copying the messages returned to a discovery mailbox can take a long time. Use this option to get an email notification when the copying process is completed. For easier access using Outlook on the web, the notification includes a link to the location in a discovery mailbox where the messages are copied.

For more information, see Copy eDiscovery search results to a discovery mailbox.

Export search results to a PST file

After search results are copied to a discovery mailbox, you can export the search results to a PST file.

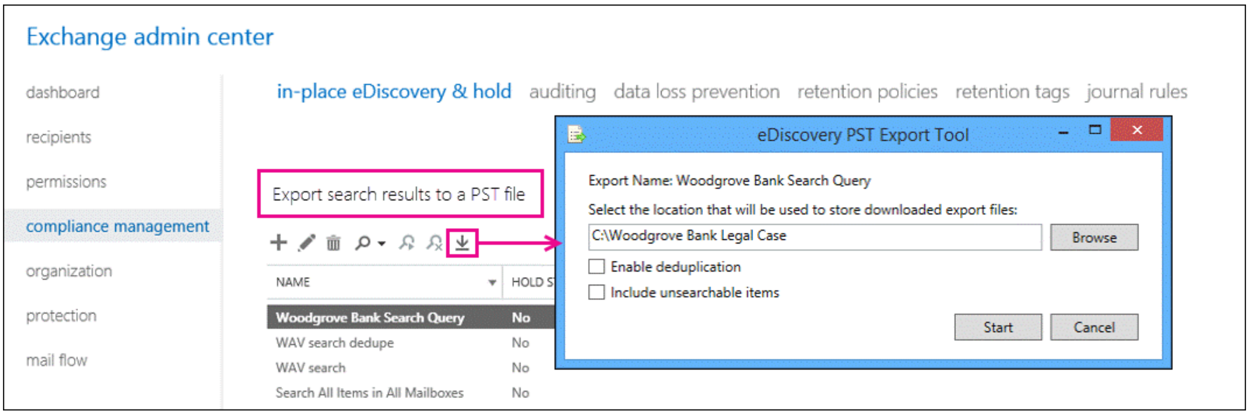

Export Name: Woodgrove Bank Search Query

Select the location that will be used to store downloaded export files:

C:\Woodgrove Bank Legal Case

□ Enable deduplication

□ Include unsearchable items

NAME	HOLD S
Woodgrove Bank Search Query	No
WAV search dedupe	No
WAV search	No
Search All Items in All Mailboxes	No

After search results are exported to a PST file, you or other users can open them in Outlook to review or print messages returned in the search results. For more information, see Export eDiscovery search results to a PST file.

Logging for In-Place eDiscovery searches

There are two types of logging available for In-Place eDiscovery searches.

- **Basic logging** - Basic logging is enabled by default for all In-Place eDiscovery searches. It includes information about the search and who performed it. Information captured about basic logging appears in the body of the email message sent to the mailbox where the search results are stored. The message is located in the folder created to store search results.

- **Full logging** - Full logging includes information about all messages returned by the search. This information is provided in a comma-separated value (.csv) file attached to the email message that contains the basic logging information. The name of the search is used for the .csv file name. This information may be required for compliance or record-keeping purposes. To enable full logging, you must select the **Enable full logging** option when copying search results to a discovery mailbox in the EAC. If you're using the Exchange Management Shell, specify the full logging option using the *LogLevel* parameter.

NOTE

When using the Exchange Management Shell to create or modify an In-Place eDiscovery search, you can also disable logging.

Besides the search log included when copying search results to a discovery mailbox, Exchange also logs cmdlets used by the EAC or the Exchange Management Shell to create, modify or remove In-Place eDiscovery searches. This information is logged in the admin audit log entries. For details, see Administrator audit logging in Exchange Server.

Discovery mailboxes

After you create an In-Place eDiscovery search, you can copy the search results to a target mailbox. The EAC allows you to select a discovery mailbox as the target mailbox. A discovery mailbox is a special type of mailbox that provides the following functionality:

- **Easier and secure target mailbox selection** - When you use the EAC to copy In-Place eDiscovery search results, only discovery mailboxes are made available as a repository in which to store search results. This eliminates the possibility of a discovery manager accidentally selecting another user's mailbox or an unsecured mailbox in which to store potentially sensitive messages.

- **Large mailbox storage quota** - The target mailbox should be able to store a large amount of message data that may be returned by an In-Place eDiscovery search. By default, discovery mailboxes have a mailbox storage quota of 50 GB. This storage quota can't be increased.

- **More secure by default** - Like all mailbox types, a discovery mailbox has an associated Active Directory user account. However, this account is disabled by default. Only users explicitly authorized to access a discovery mailbox have access to it. Members of the Discovery Management role group are assigned Full Access permissions to the default discovery mailbox. Any additional discovery mailboxes you create don't have mailbox access permissions assigned to any user.

- **Email delivery disabled** - Users can't send email to a discovery mailbox. Email delivery to discovery mailboxes is prohibited by using delivery restrictions. This preserves the integrity of search results copied to a discovery mailbox. By default, discovery mailboxes aren't displayed in your organization's global address list.

Exchange Server Setup creates one discovery mailbox with the display name **Discovery Search Mailbox**. You can use the Exchange Management Shell to create additional discovery mailboxes. By default, the discovery mailboxes you create won't have any mailbox access permissions assigned. You can assign Full Access permissions for a discovery manager to access messages copied to a discovery mailbox. For details, see Create a Discovery Mailbox .

In-Place eDiscovery also uses a system mailbox with the display name **SystemMailbox{e0dc1c29-89c3-4034-b678-e6c29d823ed9}** to hold In-Place eDiscovery metadata. System mailboxes aren't visible in the EAC or in Exchange address lists. Before removing a mailbox database where the In-Place eDiscovery system mailbox is located, you must move the mailbox to another mailbox database. If the mailbox is removed or corrupted, your discovery managers are unable to perform eDiscovery searches until you re-create the mailbox. For details, see Re-Create the Discovery System Mailbox.

In-Place eDiscovery and In-Place Hold

As part of eDiscovery requests, you may be required to preserve mailbox content until a lawsuit or investigation is disposed. Messages deleted or altered by the mailbox user or any processes must also be preserved. In Exchange Server, this is accomplished by using In-Place Hold. For details, see In-Place Hold and Litigation Hold in

Exchange Server.

You can use the **In-Place eDiscovery & Hold** wizard to search items and preserve them for as long as they're required for eDiscovery or to meet other business requirements. When using the same search for both In-Place eDiscovery and In-Place Hold, be aware of the following:

- You can't use the option to place a hold on all mailboxes in your organization. You must select the mailboxes or distribution groups. However, you can place all public fol ders in your organization on hold.
- You can't remove an In-Place eDiscovery search if the search is also used for In-Place Hold. You must first disable the In-Place Hold option in a search and then remove the search.

Preserving mailboxes for In-Place eDiscovery

When an employee leaves an organization, it's a common practice to disable or remove the mailbox. After you disable a mailbox, it is disconnected from the user account but remains in the mailbox for a certain period, 30 days by default. The Managed Folder Assistant does not process disconnected mailboxes and any retention policies are not applied during this period. You can't search content of a disconnected mailbox. Upon reaching the deleted mailbox retention period configured for the mailbox database, the mailbox is purged from the mailbox database.

If your organization requires that retention settings be applied to messages of employees who are no longer in the organization or if you may need to retain an ex-employee's mailbox for an ongoing or future eDiscovery search, do not disable or remove the mailbox. You can take the following steps to ensure the mailbox can't be accessed and no new messages are delivered to it.

1. Disable the Active Directory user account using **Active Directory Users & Computers** or other Active Directory or account provisioning tools or scripts. This prevents mailbox logon using the associated user account.

   IMPORTANT

   Users with Full Access mailbox permission will still be able to access the mailbox. To prevent access by others, you must remove their Full Access permission from the mailbox. For information about how to remove Full Access mailbox permissions on a mailbox, see Manage permissions for recipients.

2. Set the message size limit for messages that can be sent from or received by the mailbox user to a very low value, 1 KB for example. This prevents delivery of new mail to and from the mailbox. For details, see Configure message size limits for a mailbox.

3. Configure delivery restrictions for the mailbox so nobody can send messages to it. For details, see Configure message delivery restrictions for a mailbox.

IMPORTANT

You must take the above steps along with any other account management processes required by your organization, but without disabling or removing the mailbox or removing the associated user account.

When planning to implement mailbox retention for messaging retention management (MRM) or In-Place eDiscovery, you must take employee turnover into consideration. Long-term retention of ex-employee mailboxes will require additional storage on Mailbox servers and also result in an increase in Active Directory database because it requires that the associated user account be retained for the same duration. Additionally, it may also require changes to your organization's account provisioning and management processes.

Different search results

Because In-Place eDiscovery performs searches on live data, it's possible that two searches of the same content sources and that use the same search query can return different results. Estimated search results can also be different from the actual search results that are copied to a discovery mailbox. This can happen even when rerunning the same search within a short amount of time. There are several factors that can affect the consistency of search results.

- The continual indexing of incoming email because Exchange Search continuously crawls and indexes your organization's mailbox databases and transport pipeline.
- Deletion of email by users or automated processes.
- Bulk importing large amounts of email, which takes time to index.

If you do experience dissimilar results for the same search, consider placing mailboxes on hold to preserve content, running searches during off-peak hours, and allowing time for indexing after importing large amounts of email.

Custom management scopes for In-Place eDiscovery

You can use a custom management scope to let specific people or groups use In-Place eDiscovery to search a subset of mailboxes in your Exchange Server organization. For example, you might want to let a discovery manager search only the mailboxes of users in a specific location or department. You do this by creating a custom management scope that uses a custom recipient filter to control which mailboxes can be searched. Recipient filter scopes use filters to target specific recipients based on recipient type or other recipient properties.

For In-Place eDiscovery, the only property on a user mailbox that you can use to create a recipient filter for a custom scope is distribution group membership. If you use other properties, such as *CustomAttributeN*, *Department*, or *PostalCode*, the search fails when it's run by a member of the role group that's assigned the custom scope. For more information, see Create a custom management scope for In-Place eDiscovery searches.

In-Place eDiscovery and Exchange Search

In-Place eDiscovery uses the content indexes created by Exchange Search. Exchange Search has been retooled to use Microsoft Search Foundation, a rich search platform that comes with significantly improved indexing and querying performance and improved search functionality. Because the Microsoft Search Foundation is also used by other Office products, including SharePoint Server, it offers greater interoperability and similar query syntax across these products.

With a single content indexing engine, no additional resources are used to crawl and index mailbox databases for In-Place eDiscovery when eDiscovery requests are received by IT departments.

For more information about the file formats indexed by Exchange Search, see File Formats Indexed By Exchange Search.

eDiscovery in an Exchange hybrid deployment

In a hybrid deployment, which is an environment where some mailboxes exist on your on-premises Mailbox servers and some mailboxes exist in a cloud-based organization, you can perform In-Place eDiscovery searches of your cloud-based mailboxes using the EAC in your on-premises organization. If you intend to copy messages to a discovery mailbox, you must select an on-premises discovery mailbox. Messages from cloud-based mailboxes that are returned in search results are copied to the specified on-premises discovery mailbox. To learn more about hybrid deployments, see Exchange Server Hybrid Deployments.

To successfully perform cross-premises eDiscovery searches in an Exchange Server hybrid organization, you will have to configure OAuth (Open Authorization) authentication between your Exchange on-premises and Exchange Online organizations so that you can use In-Place eDiscovery to search on-premises and cloud-based mailboxes. OAuth authentication is a server-to-server authentication protocol that allows applications to authenticate to each other.

OAuth authentication supports the following eDiscovery scenarios in an Exchange hybrid deployment:

- Search on-premises mailboxes that use Exchange Online Archiving for cloud-based archive mailboxes.
- Search on-premises and cloud-based mailboxes in the same eDiscovery search.

For more information about the eDiscovery scenarios that require OAuth authentication to be configured in an Exchange hybrid deployment, see Using Oauth Authentication to Support eDiscovery in an Exchange Hybrid Deployment. For step-by-step instructions for configuring OAuth authentication to support eDiscovery, see Configure OAuth Authentication Between Exchange and Exchange Online Organizations.

For step-by-step instructions for configuring OAuth authentication to support eDiscovery, see Configure OAuth Authentication Between Exchange and Exchange Online Organizations.

Integration with SharePoint Server

Exchange Server offers integration with SharePoint Server, allowing a discovery manager to use the eDiscovery Center in SharePoint Server to perform the following tasks:

- **Search and preserve content from a single location** - An authorized discovery manager can search and preserve content across SharePoint and Exchange, including Lync content such as instant messaging conversations and shared meeting documents archived in Exchange mailboxes.
- **Case management** - eDiscovery Center uses a case management approach to eDiscovery, allowing you to create cases and search and preserve content across different content repositories for each case.
- **Export search results** - A discovery manager can use eDiscovery Center to export search results. Mailbox content included in search results is exported to a PST file.

SharePoint Server also uses Microsoft Search Foundation for content indexing and querying. Regardless of whether a discovery manager uses the EAC or the eDiscovery Center to search Exchange content, the same mailbox content is returned.

Before you can use eDiscovery Center in SharePoint Server to search Exchange mailboxes, you must establish trust between the two applications. In Exchange and SharePoint, this is done using OAuth authentication. For details, see Configure Exchange for SharePoint eDiscovery Center. eDiscovery searches performed from SharePoint are authorized by Exchange using RBAC. For a SharePoint user to be able to perform an eDiscovery search of Exchange mailboxes, they must be assigned delegated Discovery Management permission in Exchange. To be able to preview mailbox content returned in an eDiscovery search performed using SharePoint eDiscovery Center, the discovery manager must have a mailbox in the same Exchange organization.

In-Place eDiscovery limits and throttling policies

In Exchange Server, the resources In-Place eDiscovery uses are controlled with throttling policies.

The default throttling policy contains the following parameters. You can change the default values to meet your organization's requirements by creating a new throttling policy with an Organization scope and name it as "DiscoveryThrottlingPolicy" only.

PARAMETER	DESCRIPTION	DEFAULT VALUE
DiscoveryMaxConcurrency	The maximum number of In-Place eDiscovery searches a user can perform concurrently.	2
DiscoveryMaxMailboxes	The maximum number of mailboxes that can be searched in a single In-Place eDiscovery search. Public folder mailboxes are also counted against the source mailbox limit.	10,000[1]
DiscoveryMaxStatsSearchMailboxes	The maximum number of mailboxes that can be searched in a single In-Place eDiscovery search that still allows you to view keyword statistics.	100 **Note**: After you run an eDiscovery search estimate, you can view keyword statistics. These statistics show details about the number of items returned for each keyword used in the search query. If more than 100 source mailboxes are included in the search, an error will be returned if you try to view keyword statistics.
DiscoveryMaxKeywords	The maximum number of keywords that can be specified in a single In-Place eDiscovery search.	500
DiscoveryPreviewSearchResultsPageSize	The maximum number of items displayed on a single page when previewing In-Place eDiscovery search results.	200
DiscoverySearchTimeoutPeriod	The number of minutes that an In-Place eDiscovery search will run before it times out.	10 minutes

[1] Archive mailboxes are counted against the source mailbox limit. That means you can search a maximum of 5,000 mailboxes if the corresponding archive mailbox is enabled for all 5,000 mailboxes.

In-Place eDiscovery documentation

The following table contains links to Exchange Server topics that will help you learn about and manage In-Place eDiscovery.

TOPIC	DESCRIPTION
Assign eDiscovery permissions in Exchange Server	Learn how to give a user access to use In-Place eDiscovery in the EAC (and by using the corresponding cmdlets) to search Exchange mailboxes.
Create an In-Place eDiscovery search in Exchange Server	Learn how to create an In-Place eDiscovery search, and how to estimate and preview eDiscovery search results.
Copy eDiscovery search results to a discovery mailbox	Learn how to copy the results of an eDiscovery search to a discovery mailbox.

TOPIC	DESCRIPTION
Export eDiscovery search results to a PST file	Learn how to export the results of an eDiscovery search to a PST file.
Message properties and search operators for In-Place eDiscovery in Exchange Server	Learn which email message properties can be searched using In-Place eDiscovery. The topic provides syntax examples for each property, information about search operators such as **AND** and **OR**, and information about other search query techniques such as using double quotation marks (" ") and prefix wildcards.
Search and place a hold on public folders using In-Place eDiscovery	Learn how to use In-Place eDiscovery to search and place a hold on all public folders in your organization.

Assign eDiscovery permissions in Exchange Server

8/3/2020 • 2 minutes to read • Edit Online

If you want users to be able to use Exchange Server In-Place eDiscovery, you first need to add them to the Discovery Management role group. Members of the Discovery Management role group have Full Access mailbox permissions to the default discovery mailbox, which is called **Discovery Search Mailbox**.

Caution

Members of the Discovery Management role group can access sensitive message content. Specifically, these members can use In-Place eDiscovery to search all mailboxes in your Exchange organization, preview the search results (and other mailbox items), copy them to a discovery mailbox, and export the search results to a .pst file. In most organizations, this permission is assigned to legal, compliance, or Human Resources personnel.

What do you need to know before you begin?

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Role groups" entry in the Role management permissions topic.
- By default, the Discovery Management role group doesn't contain any members. Therefore, administrators that have the Organization Management role assigned can't create or manage discovery searches without being added to the Discovery Management role group.
- In Exchange Server, members of the Organization Management role group can create an In-Place Hold to place all mailbox content on hold. However, to create a query-based In-Place Hold, the user must be a member of the Discovery Management role group or have the Mailbox Search role assigned.
- You can only add *security principals* to the Discovery Management role group (users or groups that can be assigned permissions). For example:
  - User mailboxes
  - Mail users
  - Security groups
  - Other role groups
- To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.

Use the EAC to add a user to the Discovery Management role group

1. In the EAC, go to **Permissions** > **Admin roles**, select the **Discovery Management** role group, and then click **Edit** ✎.
2. On the resulting **Role Group** page, in the **Members** section, click **Add** +.
3. In the resulting **Select Members** dialog, select an available user or group, and then click **Add**. Repeat this step as many times as necessary. When you're finished, click **OK**.
4. Back on the **Role Group** page, click **Save**.

Use the Exchange Management Shell to add a user to the Discovery Management role group

To add a user to the Discovery Management role group, use the following syntax:

```
Add-RoleGroupMember -Identity "Discovery Management" -Member <Identity>
```

This example adds the user Bsuneja to the Discovery Management role group.

```
Add-RoleGroupMember -Identity "Discovery Management" -Member Bsuneja
```

This example add the members of the mail-enabled security group named Contoso Compliance Management.

```
Add-RoleGroupMember -Identity "Discovery Management" -Member "Contoso Compliance Management"
```

For more information, see Add-RoleGroupMember.

How do you know this worked?

To verify that you've added the user to the Discovery Management role group, use either of the following procedures:

- In the EAC, go to **Permissions** > **Admin roles**, and select the **Discovery Management** role group. In the details pane, verify that the user is listed in the **Members** section.

- In the Exchange Management Shell, run the following command to view the members of the Discovery Management role group.

  ```
  Get-RoleGroupMember -Identity "Discovery Management"
  ```

Create an In-Place eDiscovery search in Exchange Server

8/3/2020 • 9 minutes to read • Edit Online

Use an In-Place eDiscovery search to search for content across all mailboxes and public folders in your Exchange Server organization. This includes searching permanently deleted items and original versions of modified items (in the Recoverable Items folder) for users placed on Litigation Hold or In-Place Hold. For more information about these searches, see In-Place eDiscovery in Exchange Server.

Before you begin

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "In-Place eDiscovery" entry in the Messaging policy and compliance permissions in Exchange Server topic.

- To create eDiscovery searches, you have to have an SMTP address in the organization that you're creating the searches in. In an Exchange hybrid organization, your on-premises Exchange mailbox must have a corresponding mail user account in your Microsoft 365 or Office 365 organization, such as the tenant administrator account, that account must be assigned an Exchange Online license. For more information about the Microsoft 365 or Office 365 licensing requirements for in-place eDiscovery searches, see Exchange Online Service Description.

- Exchange Server Setup creates a Discovery mailbox called **Discovery Search Mailbox** to copy search results. You can create additional Discovery mailboxes. For details, see Create a discovery mailbox.

- When you create a search, messages returned in search results aren't copied automatically to a discovery mailbox. After you create the search, you can use the Exchange admin center (EAC) to estimate and preview search results or copy them to a discovery mailbox. You can also export the search results to a .pst file. For details, see:

  - Use the EAC to estimate or preview search results (later in this topic)

  - Copy eDiscovery search results to a discovery mailbox

  - Export eDiscovery search results to a PST file

- To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.

Use the EAC to create a search

As previously explained, to create eDiscovery searches, you have to sign in to a user account that has an SMTP address in your organization.

1. Go to **Compliance management** > **In-place eDiscovery & Hold**, and then click **New** +.

2. In the **New In-Place eDiscovery & Hold** window, on the **Name and description** page, type a name for the search, add an optional description, and then click **Next**.

3. On the **Mailboxes and Public folders** page, select the content sources to search:

   - To include all mailboxes in the search, click **Search all mailboxes**. If you select this option, you won't be able to enable an In-Place Hold for the search.

- To exclude mailboxes from the search (and search only public folders), click **Don't search any mailboxes**.

- To include specific mailboxes in the search, click **Specify mailboxes to search**, and then add that mailboxes that you want to search.

- To include public folders in the search (or to place public folders on hold), click **Search all public folders**. For more information about searching public folders, see Search and place a hold on public folders using In-Place eDiscovery.

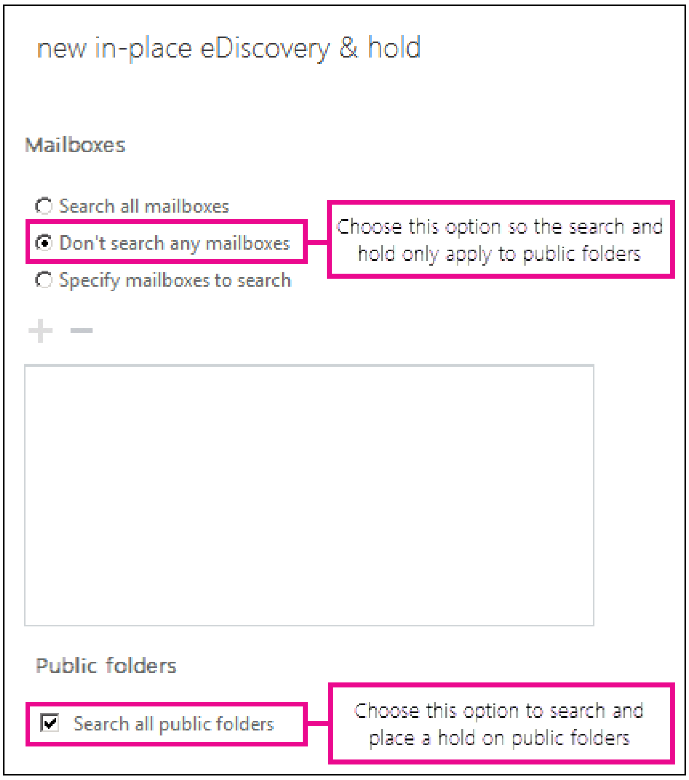

new in-place eDiscovery & hold

Mailboxes

○ Search all mailboxes

● Don't search any mailboxes

○ Specify mailboxes to search

Choose this option so the search and hold only apply to public folders

Public folders

☑ Search all public folders

Choose this option to search and place a hold on public folders

4. On the **Search query** page, complete the following fields:

   - **Include all content**: Select this option to include all content in the search results. If you select this option, you can't specify additional search criteria.

   - **Filter based on criteria**: Select this option to specify search criteria, including keywords, start and end dates, sender and recipient addresses, and message types. For more information about search queries, see Message properties and search operators for In-Place eDiscovery in Exchange Server.

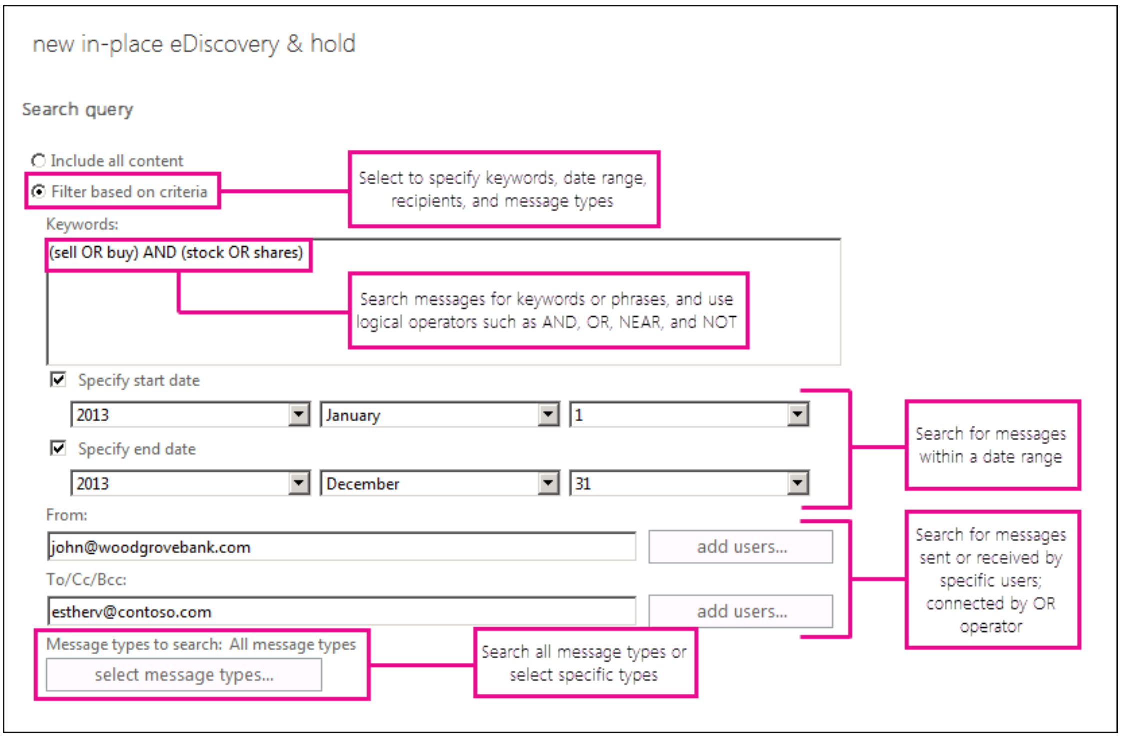

new in-place eDiscovery & hold

Search query

○ Include all content

● Filter based on criteria

Select to specify keywords, date range, recipients, and message types

Keywords:

(sell OR buy) AND (stock OR shares)

Search messages for keywords or phrases, and use logical operators such as AND, OR, NEAR, and NOT

☑ Specify start date

2013 | January | 1

☑ Specify end date

2013 | December | 31

Search for messages within a date range

From:

john@woodgrovebank.com

add users...

To/Cc/Bcc:

estherv@contoso.com

add users...

Search for messages sent or received by specific users; connected by OR operator

Message types to search: All message types

select message types...

Search all message types or select specific types

> **NOTE**
>
> The **From:** and **To/Cc/Bcc:** fields are connected by an **OR** operator in the search query that's created when you run the search. That means any message sent or received by any of the specified users (and matches the other search criteria) is included in the search results. The dates are connected by an **AND** operator.

5. On the **In-Place Hold settings** page, you can select the **Place content matching the search query in selected sources on hold** check box, and then select one of the following options to place items on In-Place Hold:

   - **Hold indefinitely**: Select this option to place the returned items on an indefinite hold. Items on hold will be preserved until you remove the content source from the search or if you delete the search.

   - **Specify number of days to hold items relative to their received date** Use this option to hold items for a specific period. For example, you can use this option if your organization requires that all messages be retained for at least seven years. You can use a *time-based* In-Place Hold along with a retention policy to make sure items are deleted in seven years.

   > **IMPORTANT**
   >
   > When placing content sources or specific items on In-Place Hold for legal purposes, it's generally recommended to hold items indefinitely and remove the hold when the case or investigation is completed.

6. Click **Finish** to save the search and return an estimate of the total size and number of items that will be returned by the search based on the criteria you specified. Estimates are displayed in the details pane. Click **Refresh** 🔄 to update the information displayed in the details pane.

Use the Exchange Management Shell to create a search

Here are four examples of using the Exchange Management Shell to search and place a hold on content in mailboxes and public folders. For detailed syntax and parameter information about using the Exchange Management Shell to create eDiscovery searches, see New-MailboxSearch

Example 1

This example creates the search Discovery-CaseId012 for items containing the keywords **Contoso** and **ProjectA**. The search results are place on In-Place hold, with an unlimited hold duration. The search also includes the following criteria:

- Start date: 1/1/2013
- End date: 12/31/2015
- Source mailbox: DG-Finance
- Target mailbox: Discovery Search Mailbox
- Message types: Email
- Log level: Full

> **IMPORTANT**
>
> If you don't specify a search query, a date range, or a message type, all items in the source mailboxes or public folders are returned in the results. The results would be similar to selecting **Include all content** on the **Search query** page in the EAC.

```
New-MailboxSearch "Discovery-CaseId012" -StartDate "01/01/2013" -EndDate "12/31/2015" -SourceMailboxes "DG-
Finance" -TargetMailbox "Discovery Search Mailbox" -SearchQuery '"Contoso" AND "Project A"' -MessageTypes
Email -IncludeUnsearchableItems -LogLevel Full -InPlaceHoldEnabled $true
```

```
Start-MailboxSearch "Discovery-CaseId012"
```

After using the Exchange Management Shell to create an In-Place eDiscovery search, you have to start the search by using the **Start-MailboxSearch** cmdlet to copy messages to the discovery mailbox specified in the *TargetMailbox* parameter. For details, see Copy eDiscovery search results to a discovery mailbox.

> **NOTE**
>
> When using the *StartDate* and *EndDate* parameters, you have to use the date format of mm/dd/yyyy, even if your local machine settings are configured to use a different date format, such as dd/mm/yyyy. For example, to search for messages sent between April 1, 2015 and July 1, 2015, you would use **04/01/2015** and **07/01/2015** for the start and end dates.

Example 2

This example creates an In-Place eDiscovery search named HRCase090116 that searches for email messages sent by Alex Darrow to Sara Davis in 2015.

```
New-MailboxSearch "HRCase090116" -StartDate "01/01/2015" -EndDate "12/31/2015" -SourceMailboxes alexd,sarad -
SearchQuery 'From:alexd@contoso.com AND To:sarad@contoso.com' -MessageTypes Email -TargetMailbox "Discovery
Search Mailbox" -IncludeUnsearchableItems -LogLevel Full
```

```
Start-MailboxSearch "HRCase090116"
```

Example 3

This example creates an estimate-only search that searches all public folders in the organization for items sent between January 1, 2015 and June 30, 2015, and that contain the phrase "patent infringement". The search doesn't include any mailboxes. The **Start-MailboxSearch** cmdlet is used to start the estimate-only search.

```
New-MailboxSearch -Name "Northwind Subpoena-All PFs" -AllPublicFolderSources $true -AllSourceMailboxes $false
-SearchQuery "patent infringement" -StartDate "01/01/2015" -EndDate "06/30/2015" -TargetMailbox "Discovery
Search Mailbox" -EstimateOnly
```

```
Start-MailboxSearch "Northwind Subpoena-All PFs"
```

Example 4

This example searches all mailboxes and public folders for any content that contains the words "price list" and "Contoso" and that was sent after January 1, 2015. The **Start-MailboxSearch** cmdlet is use to run the search and copy the search results to the discovery mailbox.

```
New-MailboxSearch -Name "Contoso Litigation" -AllSourceMailboxes $true -AllPublicFolderSources $true -
SearchQuery '"price list" AND "contoso"' -StartDate "01/01/2015" -TargetMailbox "Discovery Search Mailbox"
```

```
Start-MailboxSearch "Contoso Litigation"
```

Use the EAC to estimate or preview search results

After you create an eDiscovery search, you can use the EAC to get an estimate and preview of the search results. If you created a new search using the **New-MailboxSearch** cmdlet, you can use the Exchange Management Shell to start the search to get an estimate of the search results.

1. Go to **Compliance management** > **In-Place eDiscovery & Hold**.

2. In the list view, select the search, and then do one of the following:

   - Click **Search** 🔍 > **Estimate search results** to return an estimate of the total size and number of items that will be returned by the search based on the criteria you specified. Selecting this option restarts the search and performs an estimate.

     Search estimates are displayed in the details pane. Click **Refresh** ⟳ to update the information displayed in the details pane.

   - Click **Preview search results** in the details pane to preview the results after the search estimate is completed. Selecting this option opens the **eDiscovery search preview** window. All messages returned from the mailboxes or public folders that were searched are displayed.

     NOTE

     The mailboxes or public folders that were searched are listed in the right pane in the **eDiscovery search preview** window. For each source, the number of items returned and the total size of these items is also displayed. All items returned by the search are listed in the right pane, and can be sorted by newest or oldest date. Items from each mailbox or public folder can't be displayed in the right pane by clicking a source in the left pane. To view the items returned from a specific mailbox or public folder, you can copy the search results and view the items in the discovery mailbox.

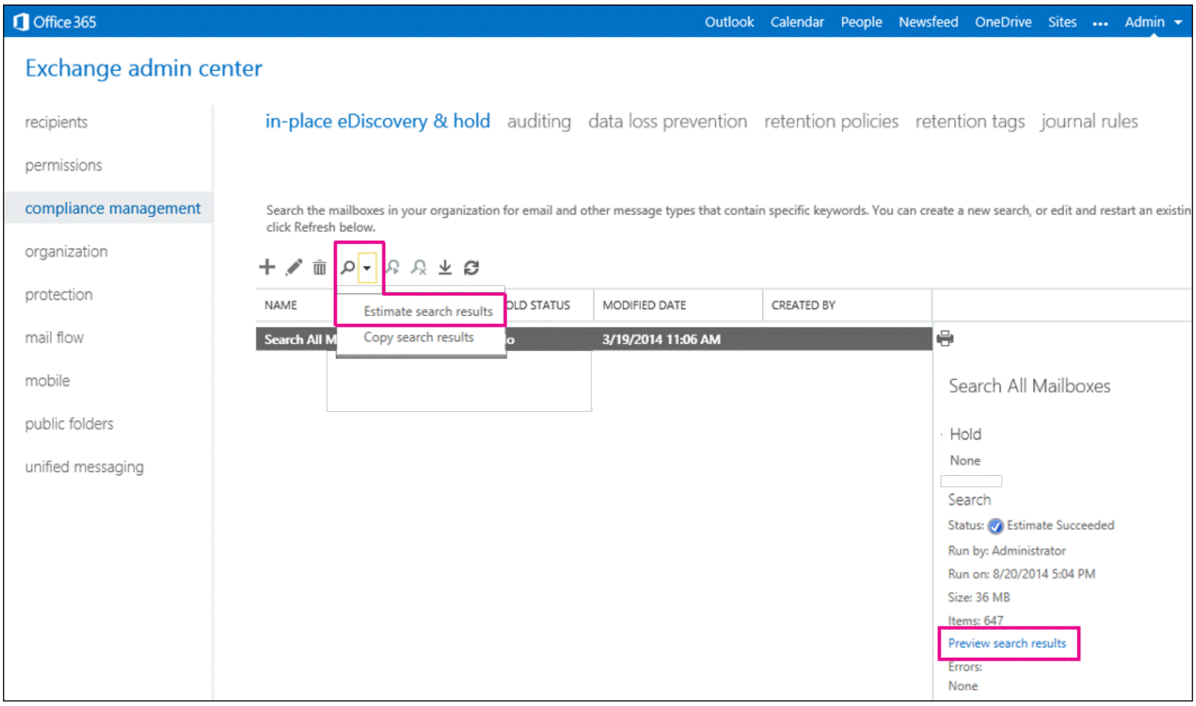

Search the mailboxes in your organization for email and other message types that contain specific keywords. You can create a new search, or edit and restart an existin click Refresh below.

NAME	OLD STATUS	MODIFIED DATE	CREATED BY
Search All M	o	3/19/2014 11:06 AM	

Search All Mailboxes

Hold
None

Search
Status: Estimate Succeeded
Run by: Administrator
Run on: 8/20/2014 5:04 PM
Size: 36 MB
Items: 647
Errors:
None

Use the Exchange Management Shell to estimate search results

You can use the *EstimateOnly* switch to get an estimate of the search results and not copy the results to a discovery mailbox. You have to start an estimate-only search with the **Start-MailboxSearch** cmdlet. Then you can retrieve the estimated search results by using the **Get-MailboxSearch** cmdlet. You can't use the Exchange Management Shell to preview messages returned in search results.

For example, you would run the following commands to create a new search and then display an estimate of the search results:

```
New-MailboxSearch "FY15 Q2 Financial Results" -StartDate "04/01/2015" -EndDate "06/30/2015" -SourceMailboxes
"DG-Finance" -SearchQuery '"Financial" AND "Fabrikam"' -EstimateOnly -IncludeKeywordStatistics
```

```
Start-MailboxSearch "FY15 Q2 Financial Results"
```

```
Get-MailboxSearch "FY15 Q2 Financial Results"
```

To display specific information about the estimated search results from the previous example, you could run the following command:

```
Get-MailboxSearch "FY15 Q2 Financial Results" | Format-List
Name,Status,LastRunBy,LastStartTime,LastEndTime,Sources,SearchQuery,ResultSizeEstimate,ResultNumberEstimate,Er
rors,KeywordHits
```

More information

- After you create a new eDiscovery search, you can copy search results to the discovery mailbox and export those search results to a PST file. For more information, see:
  - Copy eDiscovery search results to a discovery mailbox
  - Export eDiscovery search results to a PST file

- After you run an eDiscovery search estimate (that includes keywords in the search criteria), you can view keyword statistics by clicking **View keyword statistics** in the details pane for the selected search. These statistics show details about the number of items returned for each keyword used in the search query. However, if more than 100 source mailboxes are included in the search, an error will be returned if you try to view keyword statistics. To view keyword statistics, no more than 100 source mailboxes can be included in the search.

- If you use **Get-MailboxSearch** in Exchange Online to retrieve information about an eDiscovery search, you have to specify the name of a search to return a complete list of the search properties; for example, `Get-MailboxSearch "Contoso Legal Case"`. If you run the **Get-MailboxSearch** cmdlet without using any parameters, the following properties aren't returned:

  - `SourceMailboxes`
  - `Sources`
  - `PublicFolderSources`
  - `SearchQuery`
  - `ResultsLink`
  - `PreviewResultsLink`
  - `Errors`

  The reason is that it requires a lot of resources to return these properties for all eDiscovery searches in your organization.

Copy eDiscovery search results to a discovery mailbox

8/3/2020 • 4 minutes to read • Edit Online

After you create an In-Place eDiscovery search in Exchange Server, you can use the Exchange admin center (EAC) to copy the results to a discovery mailbox. You can also use the Exchange Management Shell to start an eDiscovery search that was created using the **New-MailboxSearch** cmdlet, which will copy the results to the discovery mailbox that was specified when you created the search.

What do you need to know before you begin?

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "In-Place eDiscovery" entry in the Messaging policy and compliance permissions in Exchange Server topic.

- An eDiscovery search has to be created, by using the EAC or the Exchange Management Shell, before you can copy the search results. For details, see Create an In-Place eDiscovery search in Exchange Server.

- Exchange Server Setup creates a discovery mailbox called **Discovery Search Mailbox** to copy search results. You can create additional discovery mailboxes. For details, see Create a Discovery Mailbox.

- It might take 5 minutes or longer to copy search results to a discovery mailbox, depending on the number of mailbox items returned in the results.

- To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.

Use the EAC to copy search results

1. In the EAC, go to **Compliance management** > **In-Place eDiscovery & Hold**.

2. In the list view, select an eDiscovery search.

3. Click **Search** 🔎, and then click **Copy search results** from the drop-down list.

4. In **Copy Search Results**, select from the following options:

   - **Include unsearchable items**: Select this check box to include mailbox items that couldn't be searched (for example, messages with attachments of file types that couldn't be indexed by Exchange Search). For more information, see Unsearchable Items in Exchange eDiscovery.

   - **Enable de-duplication**: Select this check box to exclude duplicate messages. Only a single instance of a message will be copied to the discovery mailbox.

   - **Enable full logging**: Select this check box to include a full log in search results.

   - **Send me mail when the copy is completed**: Select this check box to get an email notification when the search is completed.

   - **Copy results to this discovery mailbox**: Click **Browse** to select the discovery mailbox where you want the search results copied to.

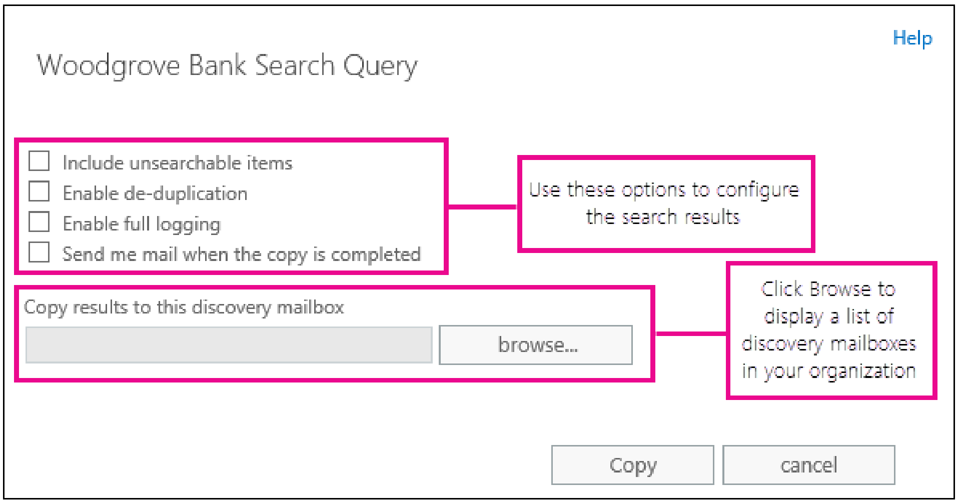
Help
Woodgrove Bank Search Query
Copy results to this discovery mailbox
browse...
Copy
cancel

□ Include unsearchable items
□ Enable de-duplication
□ Enable full logging
□ Send me mail when the copy is completed

Use these options to configure the search results

Click Browse to display a list of discovery mailboxes in your organization

5. Click **Copy** to start the process to copy the search results to the specified discovery mailbox.

6. Click **Refresh** 🔄 to update the information about the copying status that is displayed in the details pane.

7. When copying is complete, click **Open** to open the discovery mailbox to view the search results.

Use the Exchange Management Shell to copy search results

After using the **New-MailboxSearch** cmdlet to create an In-Place eDiscovery search, you need to start the search to copy messages to the discovery mailbox you specified in the *TargetMailbox* parameter. For information about creating eDiscovery searches by using the Exchange Management Shell, see:

- Create an In-Place eDiscovery search in Exchange Server

- New-MailboxSearch

In the following example, you would run the following command to start an eDiscovery search named *Fabrikam Investigation* to copy the search results to the discovery mailbox that was specified by the *TargetMailbox* parameter when the search was created.

```
Start-MailboxSearch "Fabrikam Investigation"
```

If you used the *EstimateOnly* switch to get an estimate of the search results, you have to remove the switch before you can copy the search results. You also have to specify a discovery mailbox to copy to search results to. For example, say you created an estimate-only search by using the following command:

```
New-MailboxSearch "FY15 Q2 Financial Results" -StartDate "04/01/2015" -EndDate "06/30/2015" -SourceMailboxes
"DG-Finance" -SearchQuery '"Financial" AND "Fabrikam"' -EstimateOnly -IncludeUnsearchableItems
```

To copy the results of this search to a discovery mailbox, you would run the following commands:

```
Set-MailboxSearch "FY15 Q2 Financial Results" -EstimateOnly $false -TargetMailbox "Discovery Search Mailbox"
```

```
Start-MailboxSearch "FY15 Q2 Financial Results"
```

For more information about these cmdlets, see the following topics:

- Set-Mailboxsearch

- Start-MailboxSearch

More information

- After you copy search results to the discovery mailbox, you can also export those search results to a PST file. For more information, see Export eDiscovery search results to a PST file. Note that you can export search results without having to copy them to a discovery mailbox. You can create an estimate-only search, start it, and then export the search results.

- For more information about unsearchable items, see Unsearchable Items in Exchange eDiscovery.

- If you're copying all mailbox content within a specific date range (by not specifying any keywords in the search criteria), then all unsearchable items within that date range will be automatically included in the search results. Therefore, don't select the **Include unsearchable items** checkbox when copying search results. Otherwise, a duplicate copy of all unsearchable items will be copied to the discovery mailbox.

- In addition to copying the search results to a discovery mailbox, you can also estimate or preview the search results for a selected search.

  - **Estimate search results**: This option returns an estimate of the total size and number of items that will be returned by the search based on the criteria you specified. Estimates are displayed in the details pane in the EAC.

  - **Preview search results**: This option lets you preview the search results returned by the search instead of having to copy them to a discovery mailbox to view. This lets you quickly determine whether the search results are relevant. After you preview the results, you can revise your search query to narrow the search results and rerun the search. Items in the preview page are read-only versions of the actual search results, so you can't move, edit, delete or forward on the preview page.

    For more information, see Use the EAC to estimate or preview search results.

Export eDiscovery search results to a PST file

8/3/2020 • 4 minutes to read • Edit Online

You can use the eDiscovery Export tool in the Exchange admin center (EAC) to export the results of an In-Place eDiscovery search to an Outlook Data File, which is also called a PST file. Search results will contain items from mailboxes and public folders, depending on the content sources from the eDiscovery search. This lets you distribute search results to other people within your organization, such as a human resources manager or records manager, or to opposing counsel in a legal case. After search results are exported to a PST file, you or other users can open them in Outlook to review or print messages returned in the search results. PST files can also be opened in third-party eDiscovery and reporting applications.

What do you need to know before you begin?

- The amount of time it takes to export search results will vary based on the amount and size of the search results that will be exported.

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "In-Place eDiscovery" entry in the Messaging policy and compliance permissions in Exchange Server topic.

- You'll need an active mail account attached to the account you wish to export.

- The computer you use to export search results to a PST file needs to meet the following system requirements:

  - 32- and 64-bit versions of Windows 7 and later versions

  - Microsoft .NET Framework 4.7

  - A supported browser:

  - Internet Explorer 8 and later versions

    OR

  - Mozilla Firefox or Google Chrome, with the ClickOnce add-in installed

Use the EAC to export In-Place eDiscovery search results to a PST file

1. In the EAC, go to **Compliance management** > **In-Place eDiscovery & Hold**.

2. In the list view, select the eDiscovery search you want to export the results of, and then click **Export to a PST file**.

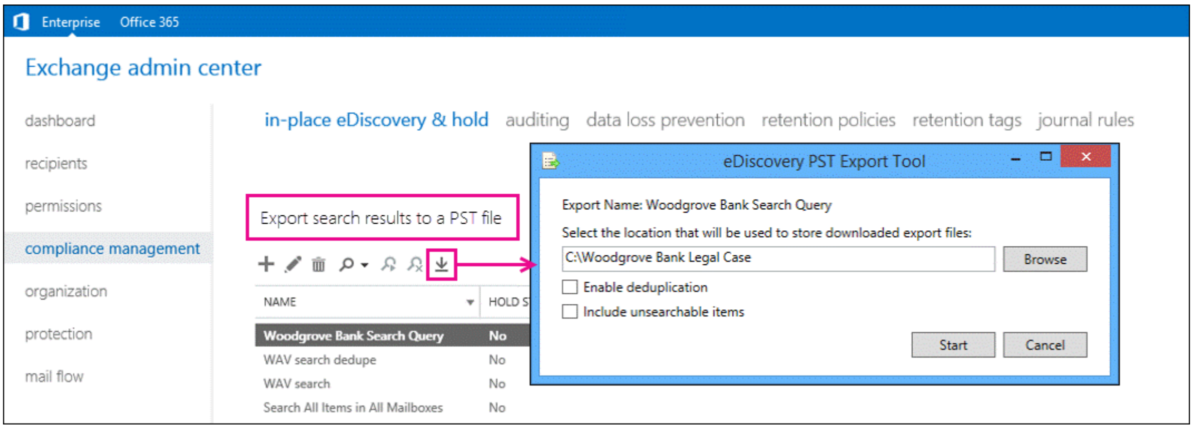

NAME	HOLD S
Woodgrove Bank Search Query	No
WAV search dedupe	No
WAV search	No
Search All Items in All Mailboxes	No

eDiscovery PST Export Tool

Export Name: Woodgrove Bank Search Query

Select the location that will be used to store downloaded export files:

C:\Woodgrove Bank Legal Case

□ Enable deduplication

□ Include unsearchable items

3. In the **eDiscovery PST Export Tool** window, do the following:

   - Click **Browse** to specify the location where you want to download the PST file.

   - Click the **Enable deduplication** checkbox to exclude duplicate messages. Only a single instance of a message will be included in the PST file.

   - Click the **Include unsearchable items** checkbox to include items that couldn't be searched (for example, messages with attachments of file types that couldn't be indexed by Exchange Search). Unsearchable items are exported to a separate PST file.

   Note: Including unsearchable items when you export eDiscovery search results takes longer when mailboxes or public folders contain a lot of unsearchable items. To reduce the time it takes to export search results and prevent large PST export files, consider the following recommendations:

   - Create multiple eDiscovery searches that each search a fewer number of source mailboxes.

   - Create an eDiscovery search that only includes public folders.

   - If you're exporting all mailbox or public folder content within a specific date range (by not specifying any keywords in the search criteria), then all unsearchable items within that date range will be automatically included in the search results. Therefore, don't select the **Include unsearchable items** checkbox.

4. Click **Start** to export the search results to a PST file.

   A window is displayed that contains status information about the export process.

More information

- Another way to reduce the size of PST export files is to export only the unsearchable items for an eDiscovery search. To do this, create or edit a search, specify a start date in the future, and then remove any keywords from the **Keywords** box. This will result in no search results being returned. When you copy or export the search results and select the **Include unsearchable items** checkbox, only the unsearchable items will be copied to the discovery mailbox or exported to a PST file.

- If you enable deduplication, all search results are exported in a single PST file. If you don't enable deduplication, a separate PST file is exported for each mailbox (including public folder mailboxes if the search includes public folders) that contains search results. And as previously stated, unsearchable items are exported to a separate PST file.

- In addition to the PST files that contain the search results, two other files are also exported:

   - A configuration file (.txt file format) that contains information about the PST export request, such as the name of the eDiscovery search that was exported, the date and time of the export, whether de-duplication and unsearchable items were enabled, the search query, and the content sources that were searched.

  - A search results log (.csv file format) that contains an entry for each message returned in the search results. Each entry identifies the content source where the message is located. If you've enabled de-duplication, this helps you identify all mailboxes or public folders that contain a duplicate message.

- The name of the search is the first part of the filename for each file that is exported. Also, the date and time of the export request is appended to the filename of each PST file and the results log.

- You can't use the PST export tool with accounts that require mult-factor authentication (MFA). Instead, you need to create an app password for the PST export tool. For instructions, see Create an app password for Microsoft 365.

Message properties and search operators for In-Place eDiscovery in Exchange Server

8/3/2020 • 9 minutes to read • Edit Online

This topic describes the properties of Exchange email messages that you can search by using In-Place eDiscovery & Hold in Exchange Server 2016 or Exchange Server 2019. The topic also describes Boolean search operators and other search query techniques that you can use to refine eDiscovery search results.

In-Place eDiscovery uses Keyword Query Language (KQL). For more details, see Keyword Query Language syntax reference.

Searchable properties in Exchange

The following table lists email message properties that can be searched using an In-Place eDiscovery search or by using the **New-MailboxSearch** or the **Set-MailboxSearch** cmdlet. The table includes an example of the *property:value* syntax for each property and a description of the search results returned by the examples.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLES	SEARCH RESULTS RETURNED BY THE EXAMPLES
Attachment	The names of files attached to an email message.	attachment:annualreport.ppt attachment:annual*	Messages that have an attached file with a name matching annualreport.ppt, e.g. "annualreport.ppt" or "2017 annualreport.ppt". In the second example, using the wildcard returns messages with the word "annual" in the file name of an attachment.
Bcc	The BCC field of an email message.[1]	bcc:pilarp@contoso.com bcc:pilarp bcc:"Pilar Pinilla"	All examples return messages with Pilar Pinilla included in the Bcc field.
Category	The categories to search. Categories can be defined by users by using Outlook or Outlook on the web (formerly known as Outlook Web App). Valid values are: • blue • green • orange • purple • red • yellow	category:"Red Category"	Messages that have been assigned the red category in the source mailboxes.
Cc	The CC field of an email message.[1]	cc:pilarp@contoso.com cc:"Pilar Pinilla"	In both examples, messages with Pilar Pinilla specified in the CC field.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLES	SEARCH RESULTS RETURNED BY THE EXAMPLES

From	The sender of an email message.[1]	from:pilarp@contoso.com
		from:contoso.com
			Messages sent by the specified user or sent from a specified domain.

Importance	The importance of an email message, which a sender can specify when sending a message. By default, messages are sent with normal importance, unless the sender sets the importance as **high** or **low**.	importance:high importance:medium importance:low	Messages that are marked as high importance, medium importance, or low importance.
Kind	The message type to search. Valid values are: • contacts • docs • email • faxes • im • journals • meetings • notes • posts • rssfeeds • tasks • voicemail	kind:email kind:email OR kind:im OR kind:voicemail	Email messages that meet the search criteria. The second example returns email messages, instant messaging conversations, and voice messages that meet the search criteria.
Participants	All the people fields in an email message; these fields are From, To, CC, and BCC.[1]	participants:garthf@contoso.com participants:contoso.com	Messages sent by or sent to garthf@contoso.com. The second example returns all messages sent by or sent to a user in the contoso.com domain.
Received	The date that an email message was received by a recipient.	received:04/15/2015 received>=01/01/2015 AND received<=03/31/2015	Messages that were received on April 15, 2014. The second example returns all messages received between January 1, 2014 and March 31, 2014.
Recipients	All recipient fields in an email message; these fields are To, CC, and BCC.[1]	recipients:garthf@contoso.com recipients:contoso.com	Messages sent to garthf@contoso.com. The second example returns messages sent to any recipient in the contoso.com domain.
Sent	The date that an email message was sent by the sender.	sent:07/01/2015 sent>=06/01/2015 AND sent<=07/01/2015	Messages that were sent on the specified date or sent within the specified date range.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLES	SEARCH RESULTS RETURNED BY THE EXAMPLES
Size	The size of an item, in bytes.	size>26214400 size:1..1048576	Messages larger than 25 MB. The second example returns messages from 1 through 1,048,576 bytes (1 MB) in size.
Subject	The text in the subject line of an email message.	subject:"Quarterly Financials" subject:northwind	Messages that contain the exact phrase "Quarterly Financials" anywhere in the text of the subject line. The second example returns all messages that contain the word northwind in the subject line.
To	The To field of an email message.[1]	to:annb@contoso.com to:annb to:"Ann Beebe"	All examples return messages where Ann Beebe is specified in the To: line.

[1] For the value of a recipient property, you can use the SMTP address, display name, or alias to specify a user. For example, you can use annb@contoso.com, annb, or "Ann Beebe" to specify the user Ann Beebe.

Supported search operators

Boolean search operators, such as **AND**, **OR**, and **NOT**, help you define more-precise mailbox searches by including or excluding specific words in the search query. Other techniques, such as using property operators (such as >= or ..), quotation marks, parentheses, and wildcards, help you refine eDiscovery search queries. The following table lists the operators that you can use to narrow or broaden search results.

OPERATOR	USAGE	DESCRIPTION
AND	keyword1 AND keyword2	Returns messages that include all of the specified keywords or `property:value` expressions.
+	keyword1 +keyword2 +keyword3	Returns items that contain *either* `keyword2` or `keyword3` *and* that also contain `keyword1`. Therefore, this example is equivalent to the query `(keyword2 OR keyword3) AND keyword1` . Note that the query `keyword1 + keyword2` (with a space after the **+** symbol) isn't the same as using the **AND** operator. This query would be equivalent to `"keyword1 + keyword2"` and return items with the exact phase `"keyword1 + keyword2"`.

OPERATOR	USAGE	DESCRIPTION
OR	keyword1 OR keyword2	Returns messages that include one or more of the specified keywords or `property:value` expressions.
NOT	keyword1 NOT keyword2 NOT from:"Ann Beebe"	Excludes messages specified by a keyword or a `property:value` expression. For example, `NOT from:"Ann Beebe"` excludes messages sent by Ann Beebe.
-	keyword1 -keyword2	The same as the **NOT** operator. This query returns items that contain `keyword1` and excludes items that contain `keyword2`.
NEAR	keyword1 NEAR(n) keyword2	Returns messages with words that are near each other, where n equals the number of words apart. For example, `best NEAR(5) worst` returns messages where the word "worst" is within five words of "best". If no number is specified, the default distance is eight words.
:	property:value	The colon (:) in the `property:value` syntax specifies that the property value being searched for equals the specified value. For example, `recipients:garthf@contoso.com` returns any message sent to garthf@contoso.com.
<	property<value	Denotes that the property being searched is less than the specified value. [1]
>	property>value	Denotes that the property being searched is greater than the specified value.[1]
<=	property<=value	Denotes that the property being searched is less than or equal to a specific value.[1]
>=	property>=value	Denotes that the property being searched is greater than or equal to a specific value.[1]
..	property:value1..value2	Denotes that the property being searched is greater than or equal to value1 and less than or equal to value2.[1]

OPERATOR	USAGE	DESCRIPTION
" "	"fair value" subject:"Quarterly Financials"	Use double quotation marks (" ") to search for an exact phrase or term in keyword and `property:value` search queries.
*	cat* subject:set*	Prefix wildcard searches (where the asterisk is placed at the end of a word) match for zero or more characters in keywords or `property:value` queries. For example, `subject:set*` returns messages that contain the word set, setup, and setting (and other words that start with "set") in the subject line.
( )	(fair OR free) AND from:contoso.com (IPO OR initial) AND (stock OR shares) (quarterly financials)	Parentheses group together Boolean phrases, `property:value` items, and keywords. For example, `(quarterly financials)` returns items that contain the words quarterly and financials.

[1] Use this operator for properties that have date or numeric values.

[2] Boolean search operators must be uppercase; for example, **AND**. Using lowercase operators in search queries will return an error.

Unsupported characters in search queries

Unsupported characters in a search query typically cause a search error or return unintended results. Unsupported characters are often hidden and they're typically added to a query when you copy the query or parts of the query from other applications (such as Microsoft Word or Microsoft Excel) and copy them to the keyword box on the query page of In-Place eDiscovery search.

Here's a list of the unsupported characters for an In-Place eDiscovery search query.

- **Smart quotation marks**: Smart single and double quotation marks (also called *curly quotes*) aren't supported. Only straight quotation marks can be used in a search query.
- **Non-printable and control characters**: Non-printable and control characters don't represent a written symbol, such as a alpha-numeric character. Examples of non-printable and control characters include characters that format text or separate lines of text.
- **Left-to-right and right-to-left marks**: These are control characters used to indicate text direction for left-to-right languages (such as English and Spanish) and right-to-left languages (such as Arabic and Hebrew).
- **Lowercase Boolean operators**: As previous explained, you have to use uppercase Boolean operators, such as **AND** and **OR**, in a search query. Note that the query syntax will often indicate that a Boolean operator is being used even though lowercase operators might be used; for example, `(WordA or WordB) and (WordC or WordD)`.

How to prevent unsupported characters in your search queries? The best way to prevent unsupported characters is to just type the query in the keyword box. Alternatively, you can copy a query from Word or Excel and then paste it to file in a plain text editor, such as Microsoft Notepad. Then save the text file and select **ANSI** in the **Encoding** drop-down list. This will remove any formatting and unsupported characters. Then you can copy and paste the query from the text file to the keyword query box.

Search tips and tricks

- Keyword searches are not case sensitive. For example, **cat** and **CAT** return the same results.

- A space between two keywords or two `property:value` expressions is the same as using **AND**. For example, `from:"Sara Davis" subject:reorganization` returns all messages sent by Sara Davis that contain the word **reorganization** in the subject line.

- Use syntax that matches the `property:value` format. Values are not case-sensitive, and they can't have a space after the operator. If there is a space, your intended value will just be full-text searched. For example, **to: pilarp** searches for "pilarp" as a keyword, rather than for messages that were sent to pilarp.

- When searching a recipient property, such as To, From, Cc, or Recipients, you can use an SMTP address, alias, or display name to denote a recipient. For example, you can use pilarp@contoso.com, pilarp, or "Pilar Pinilla".

- You can use only prefix wildcard searches (for example, **cat*** or **set***). Suffix wildcard searches (*cat) or substring wildcard searches (*cat*) aren't supported.

- When searching a property, use double quotation marks (" ") if the search value consists of multiple words. For example **subject:budget Q1** returns messages that contain **budget** in the in the subject line and that contain **Q1** anywhere in the message or in any of the message properties. Using **subject:"budget Q1"** returns all messages that contain **budget Q1** anywhere in the subject line.

- To exclude content marked with a certain property value from your search results, place a minus sign (-) before the name of the property. For example, **-from:"Sara Davis"** will exclude any messages sent by Sara Davis.

Search and place a hold on public folders using In-Place eDiscovery

8/3/2020 • 4 minutes to read • Edit Online

You can use In-Place eDiscovery to search for content in public folders and place content in public folders on In-Place Hold. Like content in user mailboxes, content in public folders might be relevant if your organization has to respond to legal requests such as lawsuits or regulatory investigations.

Before you begin

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "In-Place eDiscovery" entry in the Messaging policy and compliance permissions in Exchange Server topic.
- You can include mailboxes and public folders in the same eDiscovery search.
- You can use an In-Place Hold to place content in public folders on hold. But if you select the option to search all mailboxes in your organization, you can't use the search to place a hold on any of the content sources of the search.

Use the EAC to search and place a hold on public folders

1. Go to **Compliance management** > **In-place eDiscovery & hold**.
2. Click **New** +.
3. On the **Name and description** page, type a name for the search, add an optional description, and then click **Next**.
4. On the **Mailboxes and Public folders** page, under **Public folders**, click **Search all public folders**. Additionally, you can configure whether to include mailboxes in the search:
   - To exclude mailboxes from the search, click **Don't search any mailboxes**.
   - To include specific mailboxes in the search, click **Specify mailboxes to search**, and then add that mailboxes that you want to search.

   > **NOTE**
   >
   > As previously explained, if you select the **Search all mailboxes**option, you won't be able to enable an In-Place Hold for the search.

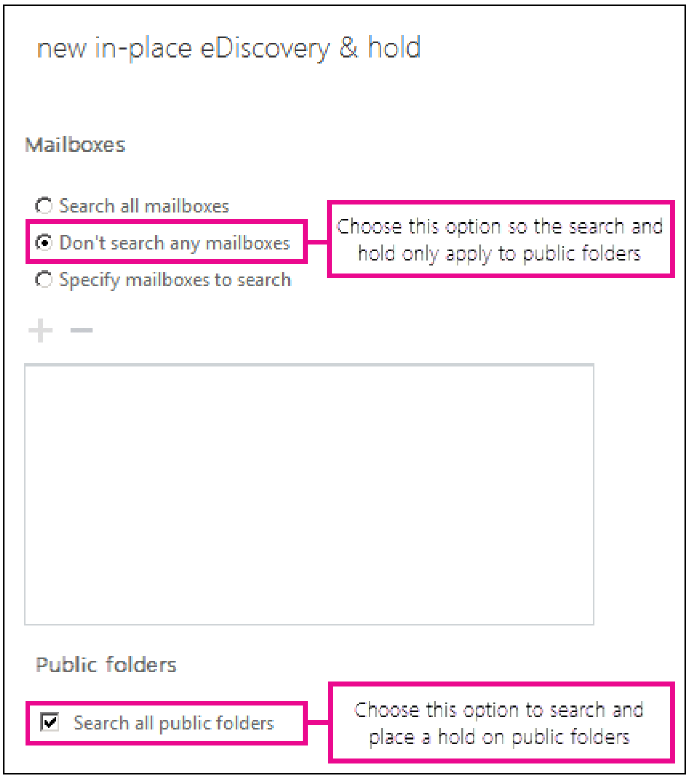

new in-place eDiscovery & hold

Mailboxes

○ Search all mailboxes

● Don't search any mailboxes

○ Specify mailboxes to search

Choose this option so the search and hold only apply to public folders

Public folders

☑ Search all public folders

Choose this option to search and place a hold on public folders

5. On the **Search query** page, complete the following fields:

   - **Include all content**: Select this option to include all content in the selected sources in the search results. If you select this option, you can't specify additional search criteria.

   - **Filter based on criteria**: Select this option to specify search criteria, including keywords, start and end dates, sender and recipient addresses, and message types.

6. On the **In-Place Hold settings** page, you can select the **Place content matching the search query in selected mailboxes on hold** to place an In-Place Hold on all public folders in your organization. Leave the check box unselected to not place content on hold. If you place content on hold, select one of the following options for the hold duration:

   - **Hold indefinitely**: Click this button to place items returned by the search on an indefinite hold. Items on hold will be preserved until you remove public folders from the search or remove the search.

   - **Specify number of days to hold items relative to their received date**: Click this button to hold items in public folders for a specific period. For example, you can use this option if your organization requires that public folder content be retained for at least seven years.

7. Click **Finish** to save the search and return an estimate of the total size and number of items that will be returned by the search or placed on hold based on the criteria you specified.

   Estimates are displayed in the details pane on the **In-Place eDiscovery & Hold** page. Select a search and then click **Refresh** ⟳ to update the information about the search that's displayed in the details pane.

Use the Exchange Management Shell to search and place a hold on public folders

Here are three examples of using the Exchange Management Shell to search and place a hold on public folders.

Example 1

This example creates an estimate-only search that searches all public folders in the organization for items sent between January 1, 2015 and June 30, 2015 and that contain the phrase "patent infringement". The search doesn't include any mailboxes. The **Start-MailboxSearch** cmdlet is used to start the estimate-only search.

```
New-MailboxSearch -Name "Northwind Subpoena-All PFs" -AllPublicFolderSources $true -AllSourceMailboxes $false
-SearchQuery "patent infringement" -StartDate "01/01/2015" -EndDate "06/30/2015" -TargetMailbox "Discovery
Search Mailbox" -EstimateOnly
```

```
Start-MailboxSearch "Northwind Subpoena-All PFs"
```

Example 2

This example places all content in all public folders on In-Place hold, with an unlimited hold duration. The **Start-MailboxSearch** cmdlet is use to run the search and place the content on hold.

```
New-MailboxSearch -Name "Hold for all PFs" -AllPublicFolderSources $true -AllSourceMailboxes $false -
EstimateOnly -InPlaceHoldEnabled $true
```

```
Start-MailboxSearch "Hold for all PFs"
```

Example 3

This example searches all mailboxes and public folders for any content that contains the words "price list" and "Contoso" and that was sent after January 1, 2015. The **Start-MailboxSearch** cmdlet is use to run the search and copy the search results to the discovery mailbox.

```
New-MailboxSearch -Name "Contoso Litigation" -AllSourceMailboxes $true -AllPublicFolderSources $true -
SearchQuery '"price list" AND "contoso"' -StartDate "01/01/2015" -TargetMailbox "Discovery Search Mailbox"
```

```
Start-MailboxSearch "Contoso Litigation"
```

More information

- You can only search or place holds on all public folders in your organization. You can't select specific public folders to search.

- Moving public folders to a different public folder mailbox doesn't affect searching or placing holds on public folders that have been moved.

- Public folder mailboxes are counted against the source mailbox limit for the eDiscovery search.

- You can't delete public folders that are on In-Place Hold. You will have to remove the hold before you can delete any public folder.

- Mail-enabling a public folder doesn't impact using In-Place eDiscovery to search or place holds on public folders. Mail-enabled and non-mail enabled public folders can be searched and placed on hold.

Use Compliance Search to search all mailboxes in Exchange Server

8/3/2020 • 10 minutes to read • Edit Online

The Compliance Search feature in Exchange Server allows you to search all mailboxes in your organization. Unlike In-Place eDiscovery where you can search up to 10,000 mailboxes, there are no limits for the number of target mailboxes in a single search. For scenarios that require you to perform organization-wide searches, you can use the **New-ComplianceSearch** cmdlet to search all mailboxes. Then you can use the workflow features of In-Place eDiscovery to perform other eDiscovery-related tasks, such as placing mailboxes on hold and exporting search results. For example, let's say you have to search all mailboxes to identify specific custodians that are responsive to a legal case. You can use the **New-ComplianceSearch** cmdlet to search all mailboxes in your organization to identify those that are responsive to the case. Then you can use that list of custodian mailboxes as the source mailboxes for an In-Place eDiscovery. Using In-Place eDiscovery also allows you to put a hold on those source mailboxes, copy search results to a discovery mailbox, and export the search results.

This topic includes a script that you can run to create an In-Place eDiscovery search by using the list of source mailboxes and search query from a compliance search that is created by running the **New-ComplianceSearch** cmdlet.

Step 1: Run the New-ComplianceSearch cmdlet to search all mailboxes

The first step is to use the Exchange Management Shell to create a compliance search that searches all mailboxes in your organization. There's no limit for the number of mailboxes for a single compliance search. Specify an appropriate keyword query (or a query for sensitive information types) so that the search returns only those source mailboxes that are relevant to your investigation. If necessary, refine the search query to narrow the scope of search results and source mailboxes that are returned.

> **NOTE**
>
> If the source compliance search doesn't return any results, an In-Place eDiscovery won't be created when you run the script in Step 3. You may have to revise the search query and then rerun the compliance search to return search results.

Here's an example of using the **New-ComplianceSearch** cmdlet to search all mailboxes in your organization. The search query returns all messages sent between October 1, 2015 and October 31, 2015 and that contain the phrase "financial report" in the subject line. The first command creates the search, and the second command runs the search.

```
New-ComplianceSearch -Name "Search All-Financial Report" -ExchangeLocation all -ContentMatchQuery
'sent>=01/01/2015 AND sent<=06/30/2015 AND subject:"financial report"'
```

```
Start-ComplianceSearch -Identity "Search All-Financial Report"
```

For more information, see New-ComplianceSearch.

IMPORTANT

When you create a compliance search by using the **New-ComplianceSearch** cmdlet, a shadow In-Place eDiscovery search is created (but not started) and displayed on the **In-Place eDiscovery & Hold** page in the Exchange admin center (EAC). It's also returned by using the **Get-MailboxSearch** cmdlet. This mailbox search is named **ComplianceSearchName - shadow**. We recommend that you delete the shadow In-Place eDiscovery search, and use the script in Step 3 to create the In-Place eDiscovery search. The functionality of creating a shadow search will be removed in a cumulative update for Exchange 2016.

(Optional) Step 2: Verify the number of source mailboxes in the compliance search

A compliance search will return a maximum of 500 source mailboxes that contain search results. If there are more than 500 mailboxes that contain content that matches the search query, only the top 500 mailboxes with the most search results are included in the compliance search that you created in the previous step. So if more than 500 mailboxes contain search results, some of those mailboxes won't be included in the list of source mailboxes copied to the new In-Place eDiscovery search created in Step 3.

To help you create a compliance search with no more than 500 source mailboxes, follow these steps to run a script that displays the number of source mailboxes (that contain search results) returned by the compliance search you created in Step 1.

1. Save the following text to a Windows PowerShell script file by using a filename suffix of .ps1. For example, you could save it to a file named SourceMailboxes.ps1.

```
[CmdletBinding()]
Param(
     [Parameter(Mandatory=$True,Position=1)]
     [string]$SearchName
)
$search = Get-ComplianceSearch $SearchName
if ($search.Status -ne "Completed")
{
                "Please wait until the search finishes.";
                break;
}
$results = $search.SuccessResults;
if (($search.Items -le 0) -or ([string]::IsNullOrWhiteSpace($results)))
{
                "The compliance search " + $SearchName + " didn't return any useful results.";
                break;
}
$mailboxes = @();
$lines = $results -split '[\r\n]+';
foreach ($line in $lines)
{
    if ($line -match 'Location: (\S+),.+Item count: (\d+)' -and $matches[2] -gt 0)
    {
        $mailboxes += $matches[1];
    }
}
"Number of mailboxes that have search hits: " + $mailboxes.Count
```

2. In the Exchange Management Shell, go to the folder where the script you created in the previous step is located, and then run the script; for example:

```
.\SourceMailboxes.ps1
```

3. When prompted by the script, type the name of the compliance search that you created in Step 1.

   The script displays the number of source mailboxes that contain search results.

If there are more than 500 source mailboxes, try creating two (or more) compliance searches. For example, search half of your organization's mailboxes in one compliance search and the other half in another compliance search. You could also change the search criteria to reduce the number of mailboxes that contain search results. For example, you could specify a date range or refine the keyword query.

Step 3: Run the script to create an In-Place eDiscovery search from the Compliance Search

The next step is to run a script that will convert an existing compliance search to an In-Place eDiscovery search. Here's what the script does:

- Prompts you for the name of the compliance search to convert.
- Verifies that the compliance search has completed running. If the compliance search doesn't return any results, and In-Place eDiscovery won't be created.
- Saves a list of the source mailboxes from the compliance search that contain search results to a variable.
- Creates a new In-Place eDiscovery search, with the following properties. Note that the new search isn't started. You'll start it in step 4.
  - **Name**: The name of the new search uses this format: *<Name of compliance search>* _MBSearch1. If you run the script again and use the same source compliance search, the search will be named *<Name of compliance search>* _MBSearch2.
  - **Source mailboxes**: All mailboxes from the compliance search that contain search results.
  - **Search query**: The new search uses the search query from the compliance search. If the compliance search includes all content (where the search query is blank) the new search will also have a blank search query and will include all content found in the source mailboxes.
  - **Estimate only search**: The new search is marked as an estimate-only search. It won't copy search results to a discovery mailbox after you start it.

1. Save the following text to a Windows PowerShell script file by using a filename suffix of ps1. For example, you could save it to a file named MBSearchFromComplianceSearch.ps1.

```
[CmdletBinding()]
Param(
    [Parameter(Mandatory=$True,Position=1)]
    [string]$SearchName,
    [switch]$original,
    [switch]$restoreOriginal
)
$search = Get-ComplianceSearch $SearchName
if ($search.Status -ne "Completed")
{
 "Please wait until the search finishes";
 break;
}
$results = $search.SuccessResults;
if (($search.Items -le 0) -or ([string]::IsNullOrWhiteSpace($results)))
{
 "The compliance search " + $SearchName + " didn't return any useful results";
 "A mailbox search object wasn't created";
 break;
}
$mailboxes = @();
$lines = $results -split '[\r\n]+';
foreach ($line in $lines)
{
    if ($line -match 'Location: (\S+),.+Item count: (\d+)' -and $matches[2] -gt 0)
    {
        $mailboxes += $matches[1];
    }
}
$msPrefix = $SearchName + "_MBSearch";
$I = 1;
$mbSearches = Get-MailboxSearch;
while ($true)
{
    $found = $false;
    $mbsName = "$msPrefix$I";
    foreach ($mbs in $mbSearches)
    {
        if ($mbs.Name -eq $mbsName)
        {
            $found = $true;
            break;
        }
    }
    if (!$found)
    {
        break;
    }
    $I++;
}
$query = $search.KeywordQuery;
if ([string]::IsNullOrWhiteSpace($query))
{
    $query = $search.ContentMatchQuery;
}
if ([string]::IsNullOrWhiteSpace($query))
{
 New-MailboxSearch "$msPrefix$i" -SourceMailboxes $mailboxes -EstimateOnly;
}
else
{
 New-MailboxSearch "$msPrefix$i" -SourceMailboxes $mailboxes -SearchQuery $query -EstimateOnly;
}
```

2. In the Exchange Management Shell, go to the folder where the script that you created in the previous step is located, and then run the script; for example:

```
.\MBSearchFromComplianceSearch.ps1
```

3. When prompted by the script, type the name of the compliance search that you want to covert to an In-Place eDiscovery search (for example, the search that you created in Step 1) , and then press **Enter**.

   If the script is successful, a new In-Place eDiscovery search is created with a status of **NotStarted**. Run the command `Get-MailboxSearch <Name of compliance search>_MBSearch1 | FL` to display the properties of the new search.

Step 4: Start the In-Place eDiscovery search

The script that you run in Step 3 creates a new In-Place eDiscovery search, but doesn't start it. The next step is to start the search so you can get an estimate of the search results.

1. In the Exchange admin center (EAC), go to **Compliance management** > **In-Place eDiscovery & Hold**.

2. In the list view, select the In-Place eDiscovery search that you created in Step 3.

3. Click **Search** (🔍) > **Estimate search results** to start the search and return an estimate of the total size and number of items returned by the search.

   The estimates are displayed in the details pane. Click **Refresh** (🔄) to update the information displayed in the details pane.

4. To preview the results after the search is completed, click **Preview search results** in the details pane.

> **TIP**
> Alternatively, you can use the Exchange Management Shell to start the In-Place eDiscovery search; for example `Start-MailboxSearch -Identity <Name of compliance search>_MBSearch1`.

Next steps after creating and running the In-Place eDiscovery search

After you create and start the In-Place eDiscovery search that was created by the script in Step 3, you can use the normal In-Place eDiscovery workflow to perform different eDiscovery actions on the search results.

Create an In-Place Hold

1. In the EAC, go to **Compliance management** > **In-Place eDiscovery & Hold**.

2. In the list view, select the In-Place eDiscovery search that you created in Step 3, and then click **Edit** (✎).

3. On the **In-Place Hold** page, select the **Place content matching the search query in selected mailboxes on hold** check box and then select one of the following options:

   - **Hold indefinitely**: Choose this option to place items returned by the search on an indefinite hold. Items on hold will be preserved until you remove the mailbox from the search or remove the search.

   - **Specify number of days to hold items relative to their received date**: Choose this option to hold items for a specific period. The duration is calculated from the date a mailbox item is received or created.

4. Click **Save** to create the In-Place Hold and restart the search.

Copy the search results

1. In the EAC, go to **Compliance management** > **In-Place eDiscovery & Hold**.

2. In the list view, select the In-Place eDiscovery search that you created in Step 3.

3. Click **Search** (🔍), and then click **Copy search results** from the drop-down list.

4. In **Copy Search Results**, select from the following options:

   - **Include unsearchable items**: Select this check box to include mailbox items that couldn't be searched (for example, messages with attachments of file types that couldn't be indexed by Exchange Search).

   - **Enable de-duplication**: Select this check box to exclude duplicate messages. Only a single instance of a message will be copied to the discovery mailbox.

   - **Enable full logging**: Select this check box to include a full log in search results.

   - **Send me mail when the copy is completed**: Select this check box to get an email notification when the search is completed.

   - **Copy results to this discovery mailbox**: Click **Browse** to select the discovery mailbox where you want the search results copied to.

5. Click **Copy** to start the process to copy the search results to the specified discovery mailbox.

6. Click **Refresh** (🔄) to update the information about the copying status that is displayed in the details pane.

7. When copying is complete, click **Open** to open the discovery mailbox to view the search results.

Export the search results

1. In the EAC, go to **Compliance management** > **In-Place eDiscovery & Hold**.

2. In the list view, select the In-Place eDiscovery search that you created in Step 3, and then click **Export to a PST file**.

3. In the list view, select the In-Place eDiscovery search you want to export the results of, and then click **Export to a PST file**.

4. In the **eDiscovery PST Export Tool** window, do the following:

   - Click **Browse** to specify the location where you want to download the PST file.

   - Click the **Enable deduplication** checkbox to exclude duplicate messages. Only a single instance of a message will be included in the PST file.

   - Click the **Include unsearchable items** checkbox to include mailbox items that couldn't be searched (for example, messages with attachments of file types that couldn't be indexed by Exchange Search). Unsearchable items are exported to a separate PST file.

5. Click **Start** to export the search results to a PST file.

   A window is displayed that contains status information about the export process.

Search for and delete messages in Exchange Server

8/3/2020 • 6 minutes to read • Edit Online

You can use the **New-ComplianceSearch** and **New-ComplianceSearchAction** cmdlets to search for and delete an email message from all mailboxes in your organization. This can help you find and remove potentially harmful or high-risk email, such as:

- Messages that contain dangerous attachment or virus
- Phishing messages
- Messages that contain sensitive data

Why use the New-ComplianceSearch and New-ComplianceSearchAction cmdlets instead of using the Search-Mailbox cmdlet to delete messages? In previous versions of Exchange, you could run the `Search-Mailbox -DeleteContent` command to search for and delete email messages. You can still do that in Exchange Server, but you can only search a maximum of 10,000 mailboxes in a single search by using the **Search-Mailbox** cmdlet. For **New-ComplianceSearch**, there are no limits for the number of mailboxes in a single search. This lets large organizations perform organization-wide search and delete operations.

Here's the workflow for the search and delete process:

Step 1: Create and run a Compliance Search to find the message to delete

Step 2: Delete the message

See the More information section for description of what happens to deleted messages and how to get the status of a search and delete operation.

Caution

Search and delete is a powerful feature that allows anyone that is assigned the necessary permissions to delete email messages from mailboxes in your organization.

Before you begin

- To use the **New-ComplianceSearch** and **Start-ComplianceSearchAction** cmdlets to create and run a Compliance Search, and to use the **New-ComplianceSearchAction** cmdlet to delete messages, you have to be assigned the Mailbox Search management role. Administrators aren't assigned this role by default. To assign yourself this role so that you can search mailboxes and delete messages, add yourself as a member of the Discovery Management role group. See Assign eDiscovery permissions in Exchange Server.
- A maximum of 10 items per mailbox can be removed at once. Because the capability to search for and remove messages is intended to be an incident-response tool, this limit helps ensure that messages are quickly removed from mailboxes. This feature isn't intended to clean up user mailboxes.

Step 1: Create and run a Compliance Search to find the message to delete

The first step is to create and run a Compliance Search to find the message that you want to remove from mailboxes in your organization. You can create the search by running the **New-ComplianceSearch** and **Start-ComplianceSearch** cmdlets. The messages that match the query for this search will be deleted by running the **New-ComplianceSearchAction** cmdlet in Step 2.

In this example, the commands will create and start a search of all mailboxes in the organization for a message that contains the words "Update your account information" in the subject line.

1. Open the Exchange Management Shell.

2. Run the following commands.

```
New-ComplianceSearch -Name "Remove Phishing Message" -ExchangeLocation all -ContentMatchQuery
'subject:"Update your account information"'
```

```
Start-ComplianceSearch -Identity "Remove Phishing Message"
```

For information about creating a Compliance Search and configuring search queries, see the following topics:

- New-ComplianceSearch
- Start-ComplianceSearch
- Message properties and search operators for In-Place eDiscovery in Exchange Server

Tips for finding messages to remove

The goal of the search query is to narrow the results of the search to only the message or messages that you want to remove. Here are some tips:

- If you know the exact text or phrase used in the subject line of the message, use the **Subject** property in the search query.
- If you know that exact date (or date range) of the message, include the **Received** property in the search query.
- If you know who sent the message, include the **From** property in the search query.
- Preview the search results to verify that the search returned only the message (or messages) that you want to delete.
- Use the search estimate statistics (by running the Get-ComplianceSearch cmdlet) to get a count of the total number of search results.

Here are two examples of queries to find suspicious email messages.

- This query returns messages that were received by users between April 13, 2016 and April 14, 2016 and that contain the words "action" and "required" in the subject line.

  ```
  (Received:4/13/2016..4/14/2016) AND (Subject:'Action required')
  ```

- This query returns messages that were sent by chatsuwloginsset12345@outlook.com and that contain the exact phrase "Update your account information" in the subject line.

  ```
  (From:chatsuwloginsset12345@outlook.com) AND (Subject:"Update your account information")
  ```

Step 2: Delete the message

After you've created and refined a Compliance Search to return the message that you want to remove, the final step is to run the **New-ComplianceSearchAction** cmdlet to delete the message. Deleted messages are moved to a user's Recoverable Items folder.

In this example, the command will delete the search results returned by a Compliance Search named "Remove Phishing Message".

1. Open the Exchange Management Shell.

2. Run the following command.

```
New-ComplianceSearchAction -SearchName "Remove Phishing Message" -Purge -PurgeType SoftDelete
```

More information

- **What happens after you delete a message?**: A message that is deleted by using the `New-ComplianceSearchAction -Purge -PurgeType SoftDelete` command is moved to the Deletions folder in the user's Recoverable Items folder. It isn't immediately purged from the Exchange database. The user can recover messages in the Deleted Items folder for the duration based on the deleted item retention period configured for the mailbox. After this retention period expires (or if user purges the message before it expires), the message is moved to the Purges folder and can no longer be accessed by the user. Once in the Purges folder, the message is again retained for the duration based on the deleted item retention period configured for the mailbox if single items recovery is enabled for the mailbox. (In Exchange, single item recovery is enabled by default when a new mailbox is created.) After the deleted item retention period expires, the message is marked from permanent deletion and will be purged from the Exchange database the next time that the mailbox is processed by the Managed Folder assistant.

- **How do you know that messages are deleted and moved to the users' Recoverable Items folder?**: If you run the same Compliance Search after you delete a message, you will still see the same number of search results (and might assume that the message wasn't deleted from user mailboxes). This is because a Compliance Search searches the Recoverable Items folder, which is where the deleted message is moved to after you run the `New-ComplianceSearchAction -Purge -PurgeType SoftDelete` command. To verify that messages where moved to the Recoverable Items folder, you can run an In-Place eDiscovery search (using the same source mailboxes and search criteria as the Compliance Search created in Step 1) and the copy the search results to discovery mailbox. Then you can view the search results in the discovery mailbox and verify that the messages was moved to the Recoverable Items folder. See Use Compliance Search to search all mailboxes in Exchange Server for details about creating an In-Place eDiscovery search that uses the list of source mailboxes and search query from a Compliance Search.

- **What happens if a message is deleted from a mailbox that has been placed on In-Place Hold or Litigation Hold?**: After the message is purged (either by the user or after the deleted item retention period expires), the message is retained until the hold duration expires. If the hold duration is unlimited, then items are retained until the hold is removed or the hold duration is changed.

- **How to get status on the search and delete operation?** Run the **Get-ComplianceSearchAction**: to get the status on the delete operation. Note that the object that is created when you run the **New-ComplianceSearchAction** cmdlet is named by using this format: `<name of Compliance Search>_Purge`.

Messaging records management in Exchange Server

8/3/2020 • 6 minutes to read • Edit Online

Users send and receive email every day. If left unmanaged, the volume of email generated and received each day can inundate users, impact user productivity, and expose your organization to risks. As a result, email lifecycle management is a critical component for most organizations.

Messaging records management (MRM) is the records management technology in Exchange Server that helps organizations manage email lifecycle and reduce the legal risks associated with email. Deploying MRM can help your organization in several ways:

- **Meet business requirements**: Depending on your organization's messaging policies, you may need to retain important email messages for a certain period. For example, a user's mailbox may contain critical messages related to business strategy, transactions, product development, or customer interactions.

- **Meet legal and regulatory requirements**: Many organizations have a legal or regulatory requirement to store messages for a designated period and remove messages older than that period. Storing messages longer than necessary may increase your organization's legal or financial risks.

- **Increase user productivity**: If left unmanaged, the ever-increasing volume of email in your users' mailboxes can also impact their productivity. For example, although newsletter subscriptions and automated notifications may have informational value when they're received, users may not remove them after reading (often they're never read). Many of these types of messages don't have a retention value beyond a few days. Using MRM to remove such messages can help reduce information clutter in users' mailboxes, thereby increasing productivity.

- **Improve storage management**: Due to expectations driven by free consumer email services, many users keep old messages for a long period or never remove them. Maintaining large mailboxes is increasingly becoming a standard practice, and users shouldn't be forced to change their work habits based on restrictive mailbox quotas. However, retaining messages beyond the period that's necessary for business, legal, or regulatory reasons also increases storage costs.

MRM provides the flexibility to implement the records management policy that best meets your organization's requirements. With a good understanding of MRM, In-Place Archiving, and In-Place Hold, you can help meet your goals of managing mailbox storage and meeting regulatory retention requirements.

MRM in Exchange Server

In Exchange Server, MRM is accomplished through the use of retention tags and retention policies. Retention tags are used to apply retention settings to an entire mailbox and default mailbox folders such as Inbox and Deleted Items. You can also create and deploy retention tags that Outlook 2010 and later and Outlook on the web users can use to apply to folders or individual messages. After they're created, you add retention tags to a retention policy and then apply the policy to users. The Managed Folder Assistant processes mailboxes and applies retention settings in the user's retention policy. To learn more about retention policies, see Retention tags and retention policies in Exchange Server.

When a message reaches its retention age specified in the applicable retention tag, the Managed Folder Assistant takes the retention action specified by the tag. Messages can then be deleted permanently or deleted with the ability to recover them. If an archive has been provisioned for the user, you can also use retention tags to move items to the user's In-Place Archive.

MRM strategies

You can use retention policies to enforce basic message retention for an entire mailbox or for specific default folders. Although there are several strategies for deploying MRM, here are some of the most common:

Remove all messages after a specified period: In this strategy, you implement a single MRM policy that removes all messages after a certain period. In this strategy, there's no classification of messages. You can implement this policy by creating a single default policy tag (DPT) for the mailbox. However, this doesn't ensure that messages are retained for the specified period. Users can still delete messages before retention period is reached.

Move messages to archive mailboxes: In this strategy, you implement MRM policies that move items to the user's archive mailbox. An archive mailbox provides additional storage for users to maintain old and infrequently accessed content. Retention tags that move items are also known as *archive policies*. Within the same retention policy, you can combine a DPT and personal tags to move items, and a DPT, RPTs, and personal tags to delete items. To learn more about archiving policies, see In-Place Archiving in Exchange Server.

Remove messages based on folder location: In this strategy, you implement MRM policies based on email location. For example, you can specify that messages in the Inbox are retained for one year and messages in the Junk Email folder are retained for 60 days. You can implement this policy by using a combination of retention policy tags (RPTs) for each default folder you want to configure and a DPT for the entire mailbox. The DPT applies to all custom folders and all default folders that don't have an RPT applied.

> **NOTE**
>
> In Exchange Server, you can create RPTs for the Calendar and Tasks folders. If you don't want items in these folders or other default folders to expire, you can create a disabled retention tag for that default folder.

Allow users to classify messages: In this strategy, you implement MRM policies that include a baseline retention setting for all messages but allow users to classify messages based on business or regulatory requirements. In this case, users become an important part of your records management strategy - often they have the best understanding of a message's retention value.

Users can apply different retention settings to messages that need to be retained for a longer or shorter period. You can implement this policy using a combination of the following:

- A DPT for the mailbox
- Personal tags that users can apply to custom folders or individual messages
- (Optional) Additional RPTs to expire items in specific default folders

For example, you can use a retention policy with personal tags that have a shorter retention period (such as two days, one week, or one month), as well as personal tags that have a longer retention period (such as one, two, or five years). Users can apply personal tags with the shorter retention periods for items such as newsletter subscriptions that may lose their value within days of receiving them, and apply the tags with longer periods to preserve items that have a high business value. They can also automate the process by using Inbox rules in Outlook to apply a personal tag to messages that match rule conditions.

Retain messages for eDiscovery purposes: In this strategy, you implement MRM policies that remove messages from mailboxes after a specified period but also retain them in the Recoverable Items folder for In-Place eDiscovery in Exchange Server purposes, even if the messages were deleted by the user or another process.

You can meet this requirement by using a combination of retention policies and In-Place Hold and Litigation Hold in Exchange Server or Litigation Hold. Retention policies remove messages from the mailbox after the specified period. A time-based In-Place Hold or Litigation Hold preserves messages that were deleted or modified before that period. For example, to retain messages for seven years, you can create a retention policy with a DPT that deletes messages in seven years and Litigation Hold to hold messages for seven years. Messages that aren't removed by users will be deleted after seven years; messages deleted by users before the seven year period will be retained in the Recoverable Items folder for seven years. To learn more about this folder, see Recoverable Items folder in Exchange Server.

Optionally, you can use RPTs and personal tags to allow users to clean up their mailboxes. However, In-Place Hold and Litigation Hold continues to retain the deleted messages until the hold period expires.

NOTE

A time-based In-Place Hold or Litigation Hold is similar to what was informally referred to as a *rolling legal hold* in Exchange 2010. Rolling legal hold was implemented by configuring the deleted item retention period for a mailbox database or individual mailbox. However, deleted item retention retains deleted and modified items based on the date deleted. In-Place Hold and Litigation Hold preserves items based on the date they're received or created. This ensures that messages are preserved for at least the specified period.

Retention tags and retention policies in Exchange Server

8/3/2020 • 14 minutes to read • Edit Online

Messaging records management (MRM) helps organizations to manage email lifecycle and reduce legal risks associated with email and other communications. MRM makes it easier to keep messages needed to comply with company policy, government regulations, or legal needs, and to remove content that has no legal or business value.

Messaging records management strategy

MRM in Exchange Server is accomplished by using *retention tags* and *retention policies*. Before discussing the details about each of these retention features, let's learn how the features are used in the overall MRM strategy:

- Assigning *retention policy tags* (RPTs) to default folders, such as the Inbox and Deleted Items.
- Applying *default policy tags* (DPTs) to mailboxes to manage the retention of all untagged items.
- Allowing the user to assign *personal tags* to custom folders and individual items.
- Separating MRM functionality from users' Inbox management and filing habits. Users aren't required to file messages in managed folders based on retention requirements. Individual messages can have a different retention tag than the one applied to the folder in which they're located.

The following figure illustrates the tasks involved in implementing this strategy.

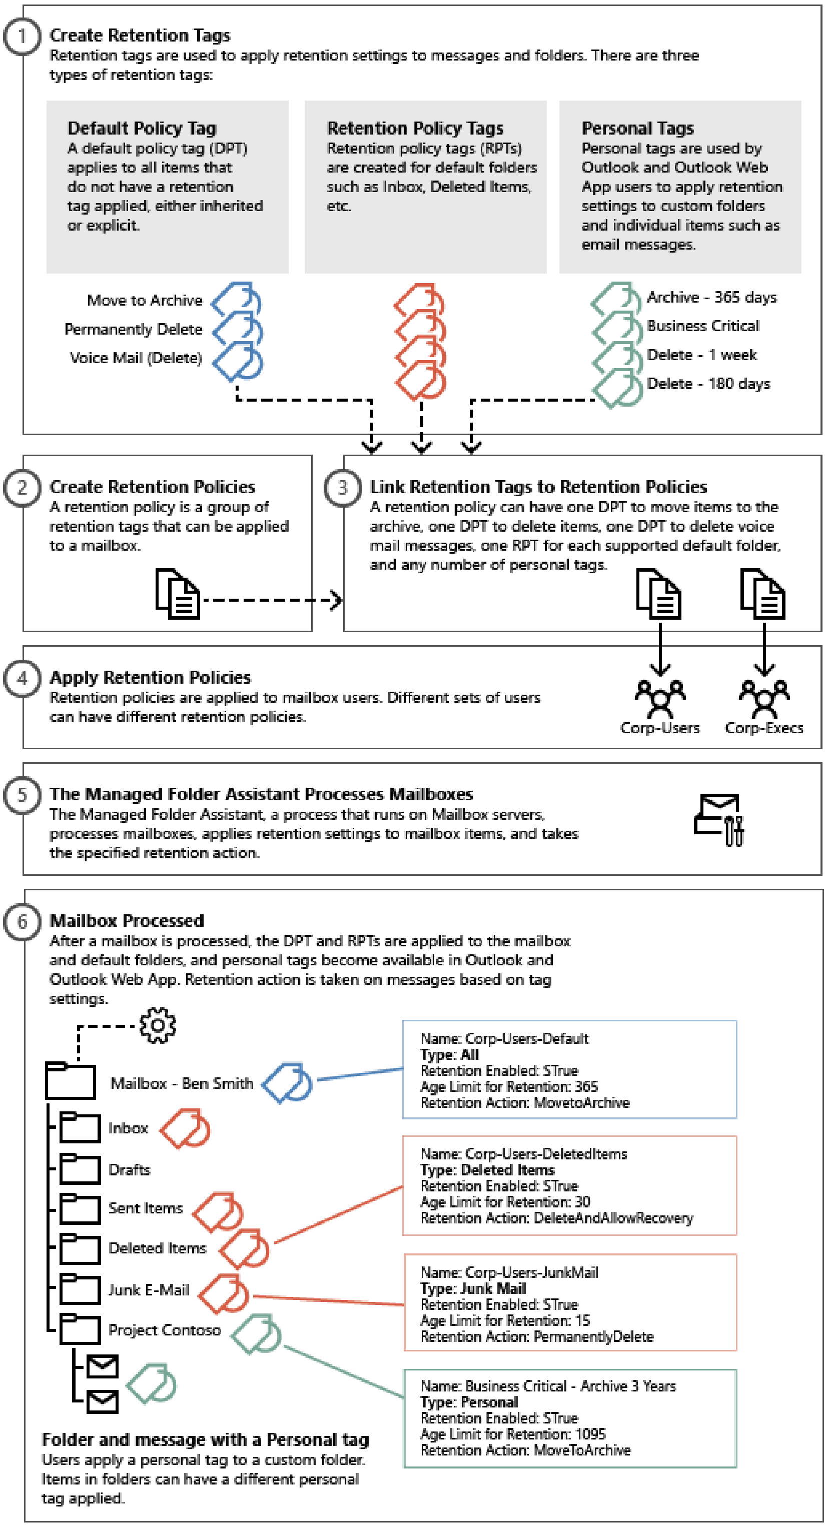

1 Create Retention Tags

Retention tags are used to apply retention settings to messages and folders. There are three types of retention tags:

Default Policy Tag

A default policy tag (DPT) applies to all items that do not have a retention tag applied, either inherited or explicit.

Retention Policy Tags

Retention policy tags (RPTs) are created for default folders such as Inbox, Deleted Items, etc.

Personal Tags

Personal tags are used by Outlook and Outlook Web App users to apply retention settings to custom folders and individual items such as email messages.

2 Create Retention Policies

A retention policy is a group of retention tags that can be applied to a mailbox.

3 Link Retention Tags to Retention Policies

A retention policy can have one DPT to move items to the archive, one DPT to delete items, one DPT to delete voice mail messages, one RPT for each supported default folder, and any number of personal tags.

4 Apply Retention Policies

Retention policies are applied to mailbox users. Different sets of users can have different retention policies.

5 The Managed Folder Assistant Processes Mailboxes

The Managed Folder Assistant, a process that runs on Mailbox servers, processes mailboxes, applies retention settings to mailbox items, and takes the specified retention action.

6 Mailbox Processed

After a mailbox is processed, the DPT and RPTs are applied to the mailbox and default folders, and personal tags become available in Outlook and Outlook Web App. Retention action is taken on messages based on tag settings.

Name: Corp-Users-Default
Type: All
Retention Enabled: STrue
Age Limit for Retention: 365
Retention Action: MovetoArchive

Name: Corp-Users-DeletedItems
Type: Deleted Items
Retention Enabled: STrue
Age Limit for Retention: 30
Retention Action: DeleteAndAllowRecovery

Name: Corp-Users-JunkMail
Type: Junk Mail
Retention Enabled: STrue
Age Limit for Retention: 15
Retention Action: PermanentlyDelete

Name: Business Critical - Archive 3 Years
Type: Personal
Retention Enabled: STrue
Age Limit for Retention: 1095
Retention Action: MoveToArchive

Folder and message with a Personal tag

Users apply a personal tag to a custom folder. Items in folders can have a different personal tag applied.

Retention tags

As you can see, retention tags are used to apply retention settings to folders and individual items such as email messages and voice mail. These settings specify how long a message remains in a mailbox and the action to take when the message reaches the specified retention age. When a message reaches its retention age, it's moved to the user's In-Place Archive or deleted.

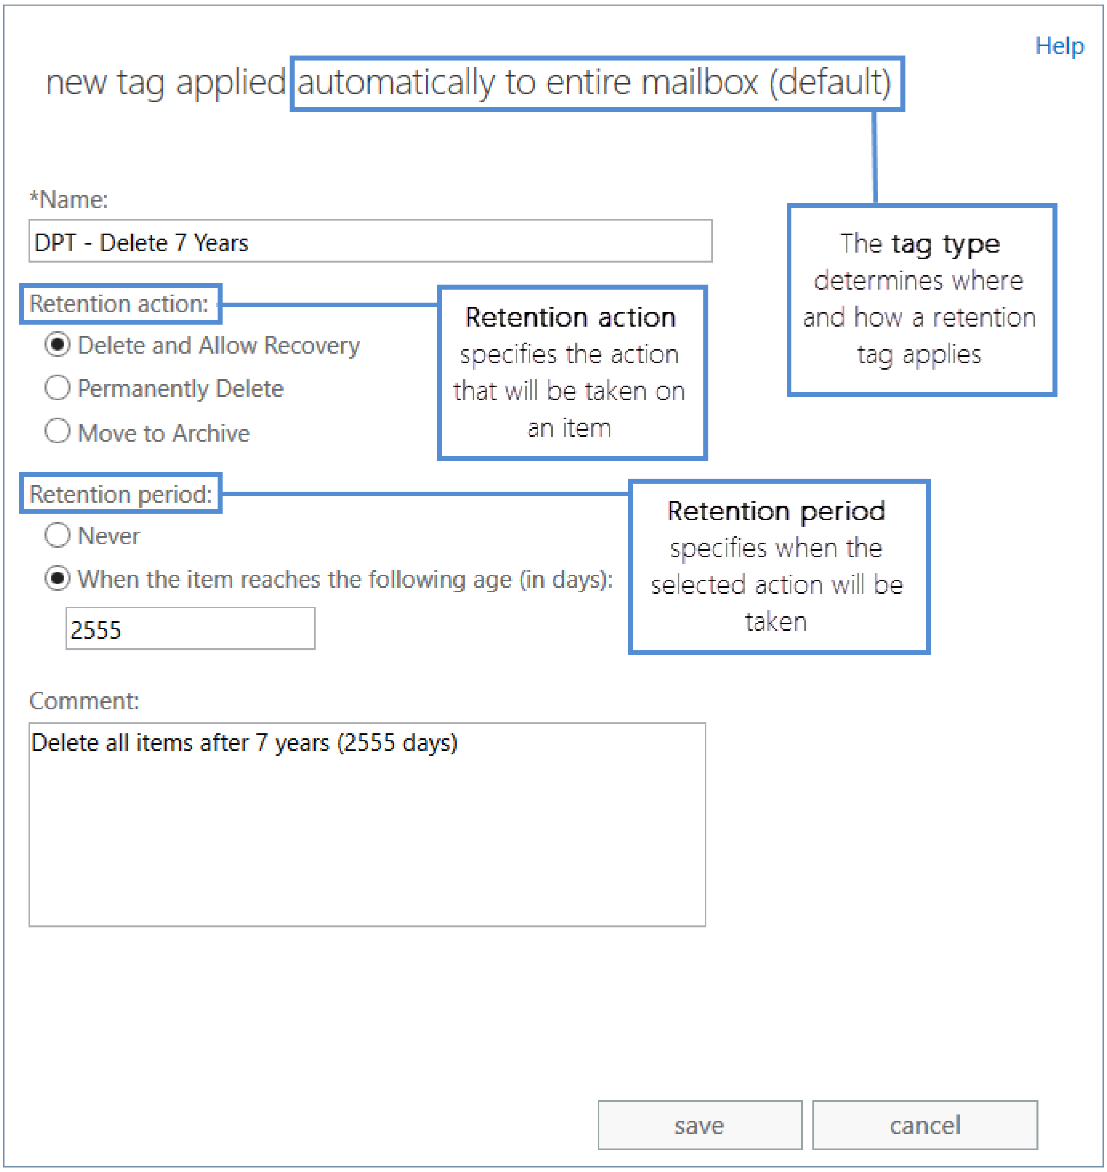

new tag applied automatically to entire mailbox (default)

*Name:
DPT - Delete 7 Years

Retention action:
● Delete and Allow Recovery
○ Permanently Delete
○ Move to Archive

Retention period:
○ Never
● When the item reaches the following age (in days):
2555

Comment:
Delete all items after 7 years (2555 days)

The **tag type** determines where and how a retention tag applies

Retention action specifies the action that will be taken on an item

Retention period specifies when the selected action will be taken

Retention tags allow users to tag their own mailbox folders and individual items for retention. Users no longer have to file items in managed folders provisioned by an administrator based on message retention requirements.

Types of retention tags

Retention tags are classified into the following three types based on who can apply them and where in a mailbox they can be applied.

TYPE OF RETENTION TAG	APPLIED...	APPLIED BY...	AVAILABLE ACTIONS...	DETAILS
Default policy tag (DPT)	Automatically to entire mailbox A DPT applies to *untagged* items, which are mailbox items that don't have a retention tag applied directly or by inheritance from the folder.	Administrator	Move to archive Delete and allow recovery Permanently delete	Users can't change DPTs applied to a mailbox.

TYPE OF RETENTION TAG	APPLIED...	APPLIED BY...	AVAILABLE ACTIONS...	DETAILS
Retention policy tag (RPT)	Automatically to a default folder Default folders are folders created automatically in all mailboxes, for example: **Inbox**, **Deleted Items**, and **Sent Items**. See the list of supported default folders in Default folders that support Retention Policy Tags.	Administrator	Delete and allow recovery Permanently delete	Users can't change the RPT applied to a default folder.
Personal tag	Manually to items and folders Users can automate tagging by using Inbox rules to either move a message to a folder that has a particular tag or to apply a personal tag to the message.	Users	Move to archive Delete and allow recovery Permanently delete	Personal tags allow your users to determine how long an item should be retained. For example, the mailbox can have a DPT to delete items in seven years, but a user can create an exception for items such as newsletters and automated notifications by applying a personal tag to delete them in three days.

More about personal tags

Personal tags are available to Outlook and Outlook on the web users as part of their retention policy. In Outlook and Outlook on the web, personal tags with the **Move to Archive** action appear as **Archive Policy**, and personal tags with the **Delete and Allow Recovery** or **Permanently Delete** actions appear as **Retention Policy**, as shown here:

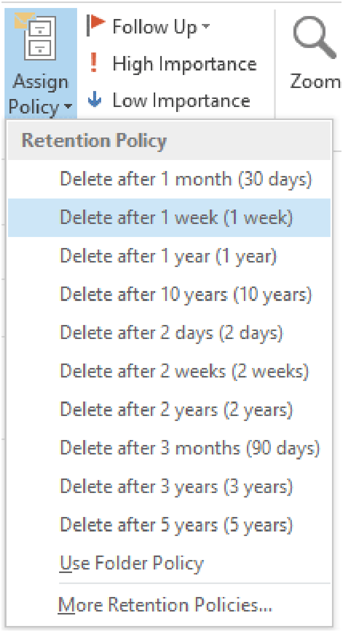
Assign Policy
Follow Up
High Importance
Low Importance
Zoom
Retention Policy
Delete after 1 month (30 days)
Delete after 1 week (1 week)
Delete after 1 year (1 year)
Delete after 10 years (10 years)
Delete after 2 days (2 days)
Delete after 2 weeks (2 weeks)
Delete after 2 years (2 years)
Delete after 3 months (90 days)
Delete after 3 years (3 years)
Delete after 5 years (5 years)
Use Folder Policy
More Retention Policies...

Outlook

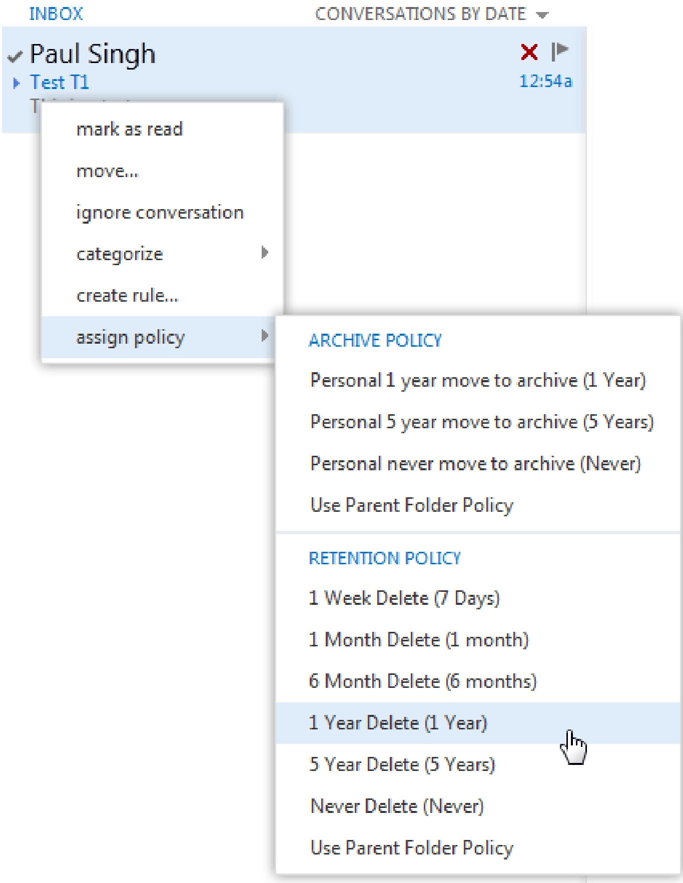

Outlook on the web

Users can apply personal tags to folders they create or to individual items. Messages that have a personal tag applied are always processed based on the personal tag's settings. Users can apply a personal tag to a message so that it's moved or deleted sooner or later than the settings specified in the DPT or RPTs applied to that user's mailbox. You can also create personal tags with retention disabled. This allows users to tag items so they're never moved to an archive or never expire.

NOTE

Users can apply archive policies to default folders, user-created folders or subfolders, and individual items. Users can apply a retention policy to user-created folders or subfolders and individual items (including subfolders and items in a default folder), but not to default folders.

Users can also use the Exchange admin center (EAC) to select additional personal tags that aren't linked to their retention policy. The selected tags then become available in Outlook and Outlook on the web. To enable users to select additional tags from the EAC, you must add the MyRetentionPolicies Role to the user's role assignment policy. To learn more about role assignment policies for users, see Understanding Management Role Assignment Policies. If you allow users to select additional personal tags, all personal tags in your Exchange organization become available to them.

NOTE

Personal tags are a premium feature. Mailboxes with policies that contain these tags (or as a result of users adding the tags to their mailbox) require an Exchange Enterprise client access license (CAL).

Retention age

When you enable a retention tag, you must specify a retention age for the tag. This age indicates the number of days to retain a message after it arrives in the user's mailbox.

The retention age for non-recurring items (such as email messages) is calculated differently than items that have an end date or recurring items (such as meetings and tasks). To learn how retention age is calculated for different types of items, see How retention age is calculated in Exchange Server.

You can also create retention tags with retention disabled or disable tags after they're created. Because messages that have a disabled tag applied aren't processed, no retention action is taken. As a result, users can use a disabled personal tag as a **Never Move** tag or a **Never Delete** tag to override a DPT or RPT that would otherwise apply to the message.

Retention actions

When creating or configuring a retention tag, you can select one of the following retention actions to be taken when an item reaches its retention age:

RETENTION ACTION	ACTION TAKEN...	EXCEPT...
Move to archive	Moves the message to the user's archive mailbox Only available for DPTs and personal tags For details about archiving, see In-Place Archiving in Exchange Server.	If the user doesn't have an archive mailbox, no action is taken.
Delete and allow recovery:	Emulates the behavior when the user empties the Deleted Items folder. Items are moved to the Recoverable Items folder in Exchange Server in the mailbox and preserved until the *deleted item retention* period. Provides the user a second chance to recover the item using the **Recover Deleted Items** dialog box in Outlook or Outlook on the web	If you've set the deleted item retention period to zero days, items are permanently deleted. For details, see Configure Deleted Item retention and Recoverable Items quotas.
Permanently delete	Permanently deletes messages. You can't recover messages after they're permanently deleted.	If mailbox is placed on In-Place Hold and Litigation Hold in Exchange Server or Litigation Hold, items are preserved in the Recoverable Items folder based on hold parameters. In-Place eDiscovery in Exchange Server will still return these items in search results.
Mark as past retention limit	Marks a message as expired. In Outlook, and Outlook on the web, expired items are displayed with the notification stating 'This item has expired' and 'This item will expire in 0 days'.	N. A.

> **NOTE**
>
> Default Policy tag (DPT) with **Move to Archive** action always overwrites the Retention Policy tag (RPT) or the Personal tag (PT), when the age limit for retention of DPT is lower than RPT or PT.

For details about how to create retention tags, see Create a retention policy in Exchange Server.

Retention policies

To apply one or more retention tags to a mailbox, you need to add them to a retention policy and then apply the policy to mailboxes. A mailbox can't have more than one retention policy. Retention tags can be linked to or unlinked from a retention policy at any time, and the changes automatically take effect for all mailboxes that have the policy applied.

A retention policy can have the following retention tags:

RETENTION TAG TYPE	TAGS IN A POLICY
Default policy tag (DPT)	One DPT with the **Move to archive** action One DPT with the **Delete and allow Recovery** or **Permanently delete** actions One DPT for voice mail messages with the **Delete and allow recovery** or **Permanently delete** action
Retention policy tags (RPTs)	One RPT for each supported default folder **Note**: You can't link more than one RPT for a particular default folder (such as **Deleted Items**) to the same retention policy.
Personal tags	Any number of personal tags **Note**: Many personal tags in a policy can confuse users. We recommend adding no more than 10 personal tags to a retention policy.

NOTE

Although a retention policy doesn't need to have any retention tags linked to it, we don't recommend using this scenario. If mailboxes with retention policies don't have retention tags linked to them, this may cause mailbox items to never expire.

A retention policy can contain both archive tags (tags that move items to the personal archive mailbox) and deletion tags (tags that delete items). A mailbox item can also have both types of tags applied. Archive mailboxes don't have a separate retention policy. The same retention policy is applied to the primary and archive mailbox.

When planning to create retention policies, you must consider whether they'll include both archive and deletion tags. As mentioned earlier, a retention policy can have one DPT that uses the **Move to archive** action and one DPT that uses either the **Delete and allow recovery** or **Permanently delete** action. The DPT with the **Move to archive** action must have a lower retention age than the DPT with a deletion action. For example, you can use a DPT with the **Move to archive** action to move items to the archive mailbox in two years, and a DPT with a deletion action to remove items from the mailbox in seven years. Items in both primary and archive mailboxes will be deleted after seven years.

Default retention policy

Exchange Setup creates the retention policy **Default MRM Policy**. The policy is applied automatically if you create an archive for the new user and don't specify a retention policy

You can modify tags included in the Default MRM Policy, for example by changing the retention age or retention action, disable a tag or modify the policy by adding or removing tags from it. The updated policy is applied to mailboxes the next time they're processed by the Managed Folder Assistant.

For more details, including a list of retention tags linked to the policy, see Default Retention Policy.

Managed Folder Assistant

The Managed Folder Assistant, a mailbox assistant that runs on Mailbox servers, processes mailboxes that have a retention policy applied.

The Managed Folder Assistant applies the retention policy by inspecting items in the mailbox and determining whether they're subject to retention. It then stamps items subject to retention with the appropriate retention tags and takes the specified retention action on items past their retention age.

The Managed Folder Assistant is a throttle-based assistant. Throttle-based assistants are always running and don't need to be scheduled. The system resources they can consume are throttled. You can configure the Managed Folder Assistant to process all mailboxes on a Mailbox server within a certain period (known as a *work cycle*). Additionally, at a specified interval (known as the *work cycle checkpoint*), the assistant refreshes the list of mailboxes to be processed. During the refresh, the assistant adds newly created or moved mailboxes to the queue. It also reprioritizes existing mailboxes that haven't been processed successfully due to failures and moves them higher in the queue so they can be processed during the same work cycle.

You can also use the Start-ManagedFolderAssistant cmdlet to manually trigger the assistant to process a specified mailbox. To learn more, see Configure and run the Managed Folder Assistant in Exchange Server.

> **NOTE**
>
> The Managed Folder Assistant doesn't take any action on messages that aren't subject to retention, specified by disabling the retention tag. You can also disable a retention tag to temporarily suspend items with that tag from being processed.

Moving items between folders

A mailbox item moved from one folder to another inherits any tags applied to the folder to which it's moved. If an item is moved to a folder that doesn't have a tag assigned, the DPT is applied to it. If the item has a tag explicitly assigned to it, the tag always takes precedence over any folder-level tags or the default tag.

Applying a retention tag to a folder in the archive

When the user applies a personal tag to a folder in the archive, if a folder with the same name exists in the primary mailbox and has a different tag, the tag on that folder in the archive changes to match the one in the primary mailbox. This is by design to avoid any confusion about items in a folder in the archive having a different expiry behavior than the same folder in the user's primary mailbox. For example, the user has a folder named Project Contoso in the primary mailbox with a *Delete - 3 years* tag and a Project Contoso folder also exists in the archive mailbox. If the user applies a *Delete - 1 year* personal tag to delete items in the folder after 1 year. When the mailbox is processed again, the folder reverts to the Delete - 3 Years tag.

Removing or deleting a retention tag from a retention policy

When a retention tag is removed from the retention policy applied to a mailbox, the tag is no longer available to the user and can't be applied to items in the mailbox.

Existing items that have been stamped with that tag continue to be processed by the Managed Folder Assistant based on those settings and any retention action specified in the tag is applied to those messages.

However, if you delete the tag, the tag definition stored in Active Directory is removed. This causes the Managed Folder Assistant to process all items in a mailbox and restamp the ones that have the removed tag applied. Depending on the number of mailboxes and messages, this process may significantly consume resources on all Mailbox servers that contain mailboxes with retention policies that include the removed tag.

> **IMPORTANT**
>
> If a retention tag is removed from a retention policy, any existing mailbox items with the tag applied will continue to expire based on the tag's settings. To prevent the tag's settings from being applied to any items, you should delete the tag. Deleting a tag removes it from any retention policies where it's included.

Disabling a retention tag

If you disable a retention tag, the Managed Folder Assistant ignores items that have that tag applied. Items that have a retention tag for which retention is disabled are either never moved or never deleted, depending on the specified retention action. Because these items are still considered tagged items, the DPT doesn't apply to them. For example, if you want to troubleshoot retention tag settings, you can temporarily disable a retention tag to stop the Managed Folder Assistant from processing messages with that tag.

> **NOTE**
>
> The retention period for a disabled retention tag is displayed to the user as **Never**. If a user tags an item believing it will never be deleted, enabling the tag later may result in unintentional deletion of items the user didn't want to delete. The same is true for tags with the **Move to archive** action.

Retention hold

When users are temporarily away from work and don't have access to their email, retention settings can be applied to new messages before they return to work or access their email. Depending on the retention policy, messages may be deleted or moved to the user's personal archive. You can temporarily suspend retention policies from processing a mailbox for a specified period by placing the mailbox on retention hold. When you place a mailbox on retention hold, you can also specify a retention comment that informs the mailbox user (or another user authorized to access the mailbox) about the retention hold, including when the hold is scheduled to begin and end. Retention comments are displayed in supported Outlook clients. You can also localize the retention hold comment in the user's preferred language.

> **NOTE**
>
> Placing a mailbox on retention hold doesn't affect how mailbox storage quotas are processed. Depending on the mailbox usage and applicable mailbox quotas, consider temporarily increasing the mailbox storage quota for users when they're on vacation or don't have access to email for an extended period. For more information about mailbox storage quotas, see Configure storage quotas for a mailbox.

During long absences from work, users may accrue a large amount of email. Depending on the volume of email and the length of absence, it may take these users several weeks to sort through their messages. In these cases, consider the additional time it may take the users to catch up on their mail before removing them from retention hold.

If your organization has never implemented MRM, and your users aren't familiar with its features, you can also use retention holds during the initial *warm up and training* phase of your MRM deployment. You can create and deploy retention policies and educate users about the policies without the risk of having items moved or deleted before users can tag them. A few days before the warm up and training period ends, you should remind users of the warm-up deadline. After the deadline, you can remove the retention hold from user mailboxes, allowing the Managed Folder Assistant to process mailbox items and take the specified retention action.

How retention age is calculated in Exchange Server

8/3/2020 • 4 minutes to read • Edit Online

The Managed Folder Assistant (MFA) is one of many mailbox assistant processes that runs on mailbox servers. Its job is to process mailboxes that have a Retention Policy applied, add the Retention Tags included in the policy to the mailbox, and process items in the mailbox. If the items have a retention tag, the assistant tests the age of those items. If an item has exceeded its retention age, it takes the specified retention action. Retention actions include moving an item to the user's archive, deleting the item and allowing recovery, or deleting the item permanently.

See Retention tags and retention policies in Exchange Server for more information.

Determining the age of different types of items

The retention age of mailbox items is calculated from the date of delivery or the date of creation for items such as drafts that are not delivered but created by the user. When the Managed Folder Assistant processes items in a mailbox, it stamps a start date and an expiration date for all items that have retention tags with the **Delete and Allow Recovery** or **Permanently Delete** retention action. Items that have an archive tag are also stamped with a move date.

Items in the Deleted Items folder and items which may have a start and end date, such as calendar items (meetings and appointments) and tasks, are handled differently as shown in this table.

IF THE ITEM TYPE IS...	AND THE ITEM IS...	THE RETENTION AGE IS CALCULATED BASED ON...
Email message Document Fax Journal item Meeting request, response, or cancellation Missed call Notes	Not in the Deleted Items folder	Delivery date or date of creation
Email message Document Fax Journal item Meeting request, response, or cancellation Missed call Notes	In the Deleted Items folder	Date of delivery or creation unless the item was deleted from a folder that does not have an inherited or implicit retention tag. If an item is in a folder that doesn't have an inherited or implicit retention tag applied, the item isn't processed by the MFA and therefore doesn't have a start date stamped by it. When the user deletes such an item, and the MFA processes it for the first time in the Deleted Items folder, it stamps the current date as the start date.

IF THE ITEM TYPE IS...	AND THE ITEM IS...	THE RETENTION AGE IS CALCULATED BASED ON...
Calendar	Not in the Deleted Items folder	Non-recurring calendar items expire according to their end date. Recurring calendar items expire according to the end date of their last occurrence. Recurring calendar items with no end date don't expire.
Calendar	In the Deleted Items folder	A calendar item expires according to its `message-received date`, if one exists. If a calendar item doesn't have a `message-received date`, it expires according to its `message-creation date`. If a calendar item has neither a `message-received date` nor a `message-creation date`, it doesn't expire.
Task	Not in the Deleted Items folder	Non-recurring tasks: • A non-recurring task expires according to its `message-received date`, if one exists. • If a non-recurring task doesn't have a `message-received date`, it expires according to its `message-creation date`. • If a non-recurring task has neither a `message-received date` nor a `message-creation date`, it doesn't expire. A recurring task expires according to the `end date` of its last occurrence. If a recurring task doesn't have an `end date`, it doesn't expire. A regenerating task (which is a recurring task that regenerates a specified time after the preceding instance of the task is completed) doesn't expire.
Task	In the Deleted Items folder	A task expires according to its `message-received date`, if one exists. If a task doesn't have a `message-received date`, it expires according to its `message-creation date`. If a task has neither a `message-received date` nor a `message-creation date`, it doesn't expire.

IF THE ITEM TYPE IS...	AND THE ITEM IS...	THE RETENTION AGE IS CALCULATED BASED ON...
Contact	In any folder	Contacts aren't stamped with a start date or an expiration date, so they're skipped by the Managed Folder Assistant and don't expire.
Corrupted	In any folder	Corrupted items are skipped by the Managed Folder Assistant and don't expire.

Examples

IF THE USER..	THE RETENTION TAGS ON FOLDER...	THE MANAGED FOLDER ASSISTANT...
Receives a message in the Inbox on 01/26/2016. Deletes the message on 2/27/2016.	Inbox: Delete in 365 days Deleted Items: Delete in 30 days	Processes the message in the Inbox on 1/26/2016, stamps it with a start date of 01/26/2016 and an expiration date of 01/26/2017. Processes the message again in the Deleted Items folder on 2/27/2016. It recalculates the expiration date based on the same start date (01/26/2016). Because the item is older than 30 days, it is expired immediately.
Receives a message in the Inbox on 01/26/2016. Deletes the message on 2/27/2016.	Inbox: None (inherited or implicit) Deleted Items: Delete in 30 days	Processes the message in the Deleted Items folder on 02/27/2016 and determines the item doesn't have a start date. It stamps the current date as the start date, and 03/27/2016 as the expiration date. The item is expired on 3/27/2016, which is 30 days after the user deleted or moved it to the Deleted Items folder.

More information

Items in mailboxes placed on Retention Hold aren't removed until the hold is removed.

If a mailbox is placed on In-Place Hold or Litigation Hold, expiring items are removed from the Inbox but preserved in the Recoverable Items folder until the mailbox is removed from In-Place Hold and Litigation Hold in Exchange Server.

In hybrid deployments, the same retention tags and retention policies must exist in your on-premises and Exchange Online organizations in order to consistently move and expire items across both organizations. See Export and import retention tags for more information.

Create a retention policy in Exchange Server

8/3/2020 • 6 minutes to read • Edit Online

Learn how to use retention policies to manage an email lifecycle in Exchange 2016 and Exchange 2019. Retention policies are applied by creating retention tags, adding them to a retention policy, and applying the policy to mailbox users.

What do you need to know before you begin?

- Estimated time to complete this task: 30 minutes.
- Procedures in this topic require specific permissions. See each procedure for its permissions information.
- Mailboxes to which you apply retention policies must reside on servers running Exchange Server 2010 or later.

Step 1: Create a retention tag

You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Messaging records management" entry in the Messaging policy and compliance permissions in Exchange Server topic.

Use the Exchange admin center (EAC) to create a retention tag

1. Go to **Compliance management** > **Retention tags**, and click **Add** +.
2. Select one of the following options:
   - **Applied automatically to entire mailbox (default)**: Creates a default policy tag (DPT). You can use DPTs to create a default deletion policy and a default archive policy, which applies to all items in the mailbox.

     > **NOTE**
     >
     > You can't use the EAC to create a DPT to delete voice mail items. For details about how to create a DPT to delete voice mail items, see the Exchange Management Shell example below.

   - **Applied automatically to a default folder**: Creates a retention policy tag (RPT) for a default folder such as **Inbox** or **Deleted Items**.

     > **NOTE**
     >
     > You can only create RPTs with the **Delete and allow recovery** or **Permanently delete** actions.

   - **Applied by users to items and folders (personal)**: Creates personal tags. These tags allow Outlook and Outlook on the web users to apply archive or deletion settings to a message or folders that are different from the settings applied to the parent folder or the entire mailbox.
3. The **New retention tag** page title and options vary depending on the type of tag you select. Complete the following fields:
   - **Name**: Enter a name for the retention tag. Retention tag names are displayed to users in Outlook and

Outlook on the web along with the retention period.

- **Apply this tag to the following default folder**: Available only if you selected this option in Step 2.
- **Retention action**: Select one of the following actions to take after the item reaches its retention period:
- **Delete and Allow Recovery**: Deletes items but allow users to recover them using the **Recover Deleted Items** option in Outlook or Outlook on the web. Items are retained until the deleted item retention period configured for the mailbox database or the mailbox user is reached.
- **Permanently Delete**: Permanently deletes the item from the mailbox database.

> **IMPORTANT**
>
> Mailboxes or items subject to In-Place Hold or litigation hold will be retained and returned in In-Place eDiscovery searches. To learn more, see In-Place Hold and Litigation Hold in Exchange Server.

- **Move to Archive**: Available only if you're creating a DPT or a personal tag. Select this action to move items to the user's In-Place Archive.
- **Retention period**: Select one of the following options:
- **Never**: Specifies that items should never be deleted or moved to the archive.
- **When the item reaches the following age (in days)**: Specifies the number of days to retain items before they're moved or deleted. The retention age for all supported items except Calendar and Tasks is calculated from the date an item is received or created. Retention age for Calendar and Tasks items is calculated from the end date.
- **Comment**: Optional field used for administrative notes or comments. The field isn't displayed to users.

Use the Exchange Management Shell to create a retention tag

Use the **New-RetentionPolicyTag** cmdlet to create a retention tag. Different options available in the cmdlet allow you to create different types of retention tags. Use the *Type* parameter to create a DPT (`All`), RPT (specify a default folder type, such as `Inbox`) or a personal tag (`Personal`).

This example creates a DPT to delete all messages in the mailbox after 7 years (2,556 days).

```
New-RetentionPolicyTag -Name "DPT-Corp-Delete" -Type All -AgeLimitForRetention 2556 -RetentionAction
DeleteAndAllowRecovery
```

This example creates a DPT to move all messages to the In-Place Archive in 2 years (730 days).

```
New-RetentionPolicyTag -Name "DPT-Corp-Move" -Type All -AgeLimitForRetention 730 -RetentionAction
MoveToArchive
```

This example creates a DPT to delete voice mail messages after 20 days.

```
New-RetentionPolicyTag -Name "DPT-Corp-Voicemail" -Type All -MessageClass Voicemail -AgeLimitForRetention 20 -
RetentionAction DeleteAndAllowRecovery
```

This example creates a RPT to permanently delete messages in the Junk EMail folder after 30 days.

```
New-RetentionPolicyTag -Name "RPT-Corp-JunkMail" -Type JunkEmail -AgeLimitForRetention 30 -RetentionAction
PermanentlyDelete
```

This example creates a personal tag to never delete a message.

```
New-RetentionPolicyTag -Name "Never Delete" -Type Personal -RetentionAction DeleteAndAllowRecovery -
RetentionEnabled $false
```

Step 2: Create a retention policy

You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Messaging records management" entry in the Messaging policy and compliance permissions in Exchange Server topic.

Use the EAC to create a retention policy

1. Go to **Compliance management** > **Retention policies**, and click **Add** +.

2. In **New Retention Policy**, complete the following fields:

   - **Name**: Enter a name for the retention policy.

   - **Retention tags**: Click **Add** + to select the tags you want to add to this retention policy.

     A retention policy can contain the following tags:

     - One DPT with the **Move to Archive** action

     - One DPT with the **Delete and Allow Recovery** or **Permanently Delete** actions

     - One DPT for voice mail messages with the **Delete and Allow Recovery** or **Permanently Delete** actions

     - One RPT per default folder such as **Inbox** to delete items

     - Any number of personal tags

     > **NOTE**
     >
     > Although you can add any number of personal tags to a retention policy, having many personal tags with different retention settings can confuse users. We recommend linking no more than ten personal tags to a retention policy.

   You can create a retention policy without adding any retention tags to it, but items in the mailbox to which the policy is applied won't be moved or deleted. You can also add and remove retention tags from a retention policy after you create it.

Use the Exchange Management Shell to create a retention policy

This example creates the retention policy RetentionPolicy-Corp and uses the *RetentionPolicyTagLinks* parameter to associate five tags to the policy.

```
New-RetentionPolicy "RetentionPolicy-Corp" -RetentionPolicyTagLinks "DPT-Corp-Delete","DPT-Corp-Move","DPT-
Corp-Voicemail","RPT-Corp-JunkMail","Never Delete"
```

For detailed syntax and parameter information, see New-RetentionPolicy.

Step 3: Apply a retention policy to mailbox users

After you create a retention policy, you must apply it to mailbox users. You can apply different retention policies to different set of users. For detailed instructions, see Apply a retention policy to mailboxes in Exchange Server.

How do you know this task worked?

After you create retention tags, add them to a retention policy, and apply the policy to a mailbox user, the next time the MRM mailbox assistant processes the mailbox, messages are moved or deleted based on settings you configured in the retention tags.

To verify that you have applied the retention policy, do the following:

1. Run the following Exchange Management Shell command to run the MRM assistant manually against a single mailbox.

   ```
   Start-ManagedFolderAssistant -Identity <mailbox identity>
   ```

2. Log on to the mailbox using Outlook or Outlook on the web and verify that messages are deleted or moved to an archive in accordance with the policy configuration.

Retention tags and retention policies in Exchange Server

8/3/2020 • 14 minutes to read • Edit Online

Messaging records management (MRM) helps organizations to manage email lifecycle and reduce legal risks associated with email and other communications. MRM makes it easier to keep messages needed to comply with company policy, government regulations, or legal needs, and to remove content that has no legal or business value.

Messaging records management strategy

MRM in Exchange Server is accomplished by using *retention tags* and *retention policies*. Before discussing the details about each of these retention features, let's learn how the features are used in the overall MRM strategy:

- Assigning *retention policy tags* (RPTs) to default folders, such as the Inbox and Deleted Items.
- Applying *default policy tags* (DPTs) to mailboxes to manage the retention of all untagged items.
- Allowing the user to assign *personal tags* to custom folders and individual items.
- Separating MRM functionality from users' Inbox management and filing habits. Users aren't required to file messages in managed folders based on retention requirements. Individual messages can have a different retention tag than the one applied to the folder in which they're located.

The following figure illustrates the tasks involved in implementing this strategy.

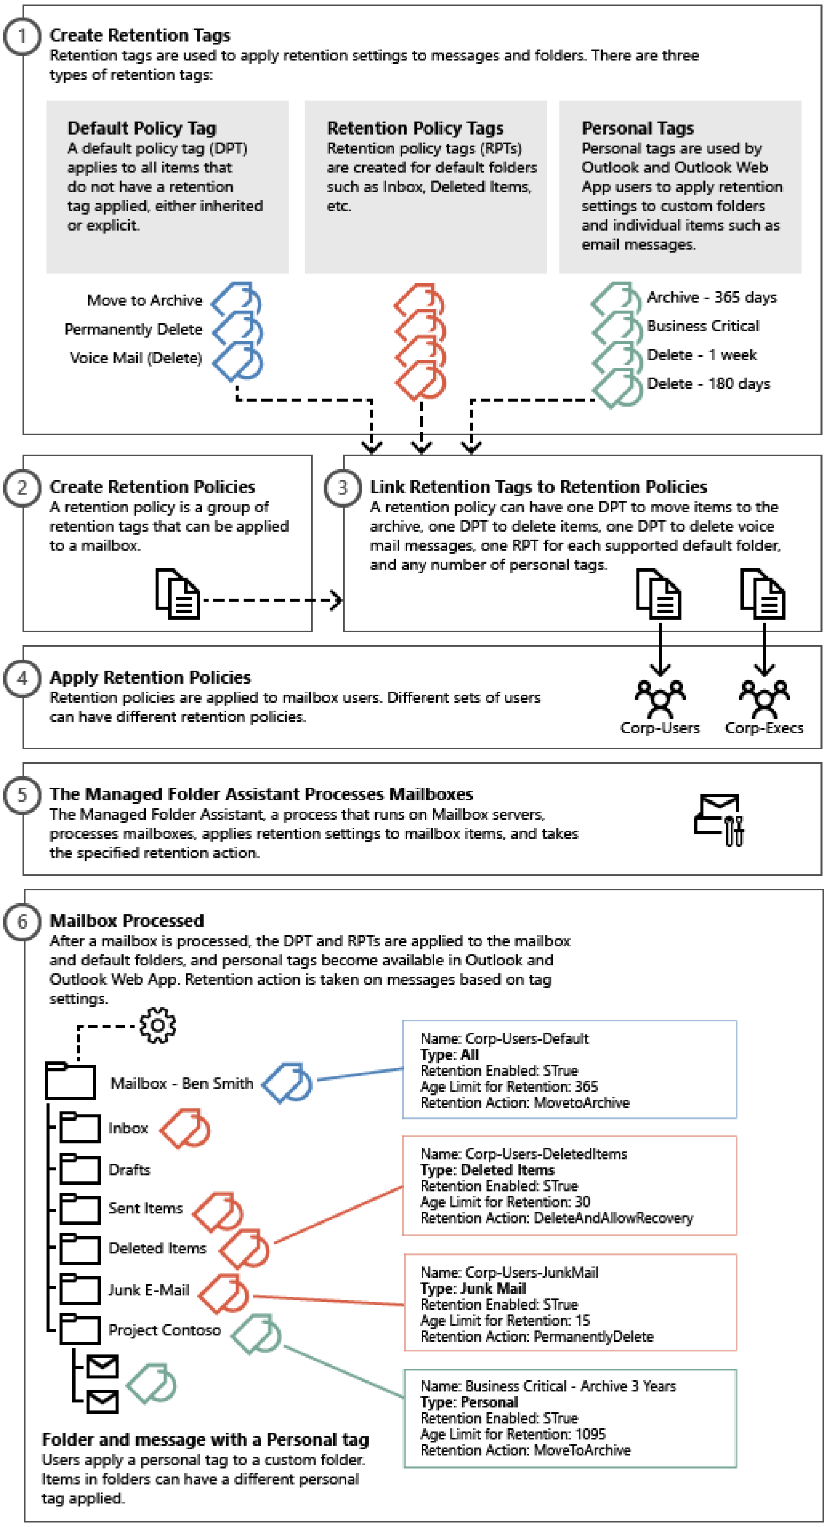

Create Retention Tags
Retention tags are used to apply retention settings to messages and folders. There are three types of retention tags:

Default Policy Tag
A default policy tag (DPT) applies to all items that do not have a retention tag applied, either inherited or explicit.

Retention Policy Tags
Retention policy tags (RPTs) are created for default folders such as Inbox, Deleted Items, etc.

Personal Tags
Personal tags are used by Outlook and Outlook Web App users to apply retention settings to custom folders and individual items such as email messages.

Create Retention Policies
A retention policy is a group of retention tags that can be applied to a mailbox.

Link Retention Tags to Retention Policies
A retention policy can have one DPT to move items to the archive, one DPT to delete items, one DPT to delete voice mail messages, one RPT for each supported default folder, and any number of personal tags.

Apply Retention Policies
Retention policies are applied to mailbox users. Different sets of users can have different retention policies.

The Managed Folder Assistant Processes Mailboxes
The Managed Folder Assistant, a process that runs on Mailbox servers, processes mailboxes, applies retention settings to mailbox items, and takes the specified retention action.

Mailbox Processed
After a mailbox is processed, the DPT and RPTs are applied to the mailbox and default folders, and personal tags become available in Outlook and Outlook Web App. Retention action is taken on messages based on tag settings.

Name: Corp-Users-Default
Type: All
Retention Enabled: STrue
Age Limit for Retention: 365
Retention Action: MovetoArchive

Name: Corp-Users-DeletedItems
Type: Deleted Items
Retention Enabled: STrue
Age Limit for Retention: 30
Retention Action: DeleteAndAllowRecovery

Name: Corp-Users-JunkMail
Type: Junk Mail
Retention Enabled: STrue
Age Limit for Retention: 15
Retention Action: PermanentlyDelete

Name: Business Critical - Archive 3 Years
Type: Personal
Retention Enabled: STrue
Age Limit for Retention: 1095
Retention Action: MoveToArchive

Folder and message with a Personal tag
Users apply a personal tag to a custom folder. Items in folders can have a different personal tag applied.

Retention tags

As you can see, retention tags are used to apply retention settings to folders and individual items such as email messages and voice mail. These settings specify how long a message remains in a mailbox and the action to take when the message reaches the specified retention age. When a message reaches its retention age, it's moved to the user's In-Place Archive or deleted.

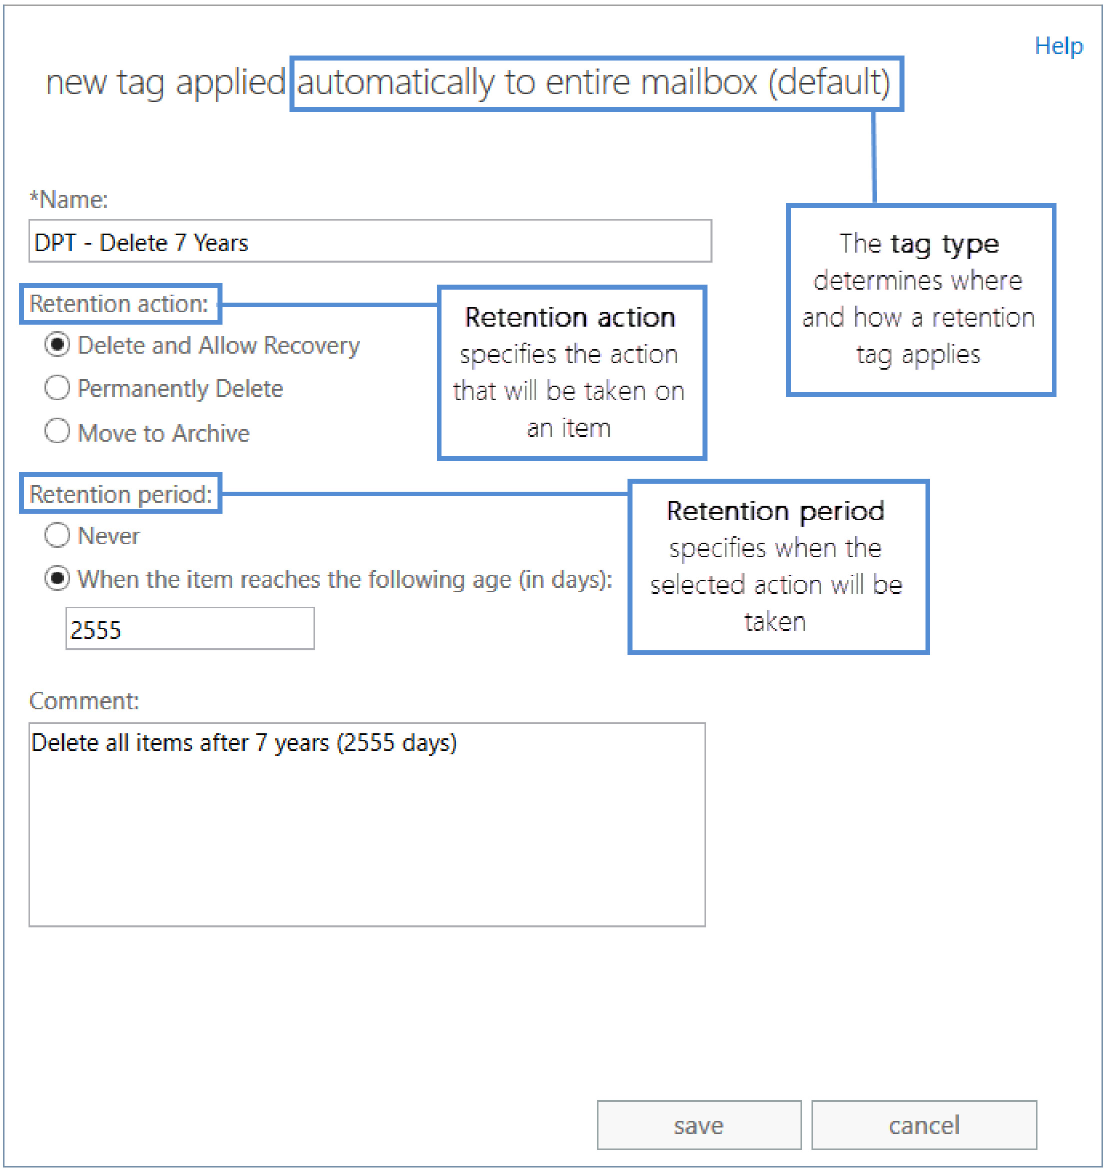

Help
new tag applied automatically to entire mailbox (default)
The tag type determines where and how a retention tag applies
Retention action specifies the action that will be taken on an item
Retention period specifies when the selected action will be taken
save
cancel

*Name:
DPT - Delete 7 Years

Retention action:
● Delete and Allow Recovery
○ Permanently Delete
○ Move to Archive

Retention period:
○ Never
● When the item reaches the following age (in days):
2555

Comment:
Delete all items after 7 years (2555 days)

Retention tags allow users to tag their own mailbox folders and individual items for retention. Users no longer have to file items in managed folders provisioned by an administrator based on message retention requirements.

Types of retention tags

Retention tags are classified into the following three types based on who can apply them and where in a mailbox they can be applied.

TYPE OF RETENTION TAG	APPLIED...	APPLIED BY...	AVAILABLE ACTIONS...	DETAILS
Default policy tag (DPT)	Automatically to entire mailbox A DPT applies to *untagged* items, which are mailbox items that don't have a retention tag applied directly or by inheritance from the folder.	Administrator	Move to archive Delete and allow recovery Permanently delete	Users can't change DPTs applied to a mailbox.

TYPE OF RETENTION TAG	APPLIED...	APPLIED BY...	AVAILABLE ACTIONS...	DETAILS
Retention policy tag (RPT)	Automatically to a default folder Default folders are folders created automatically in all mailboxes, for example: **Inbox**, **Deleted Items**, and **Sent Items**. See the list of supported default folders in Default folders that support Retention Policy Tags.	Administrator	Delete and allow recovery Permanently delete	Users can't change the RPT applied to a default folder.
Personal tag	Manually to items and folders Users can automate tagging by using Inbox rules to either move a message to a folder that has a particular tag or to apply a personal tag to the message.	Users	Move to archive Delete and allow recovery Permanently delete	Personal tags allow your users to determine how long an item should be retained. For example, the mailbox can have a DPT to delete items in seven years, but a user can create an exception for items such as newsletters and automated notifications by applying a personal tag to delete them in three days.

More about personal tags

Personal tags are available to Outlook and Outlook on the web users as part of their retention policy. In Outlook and Outlook on the web, personal tags with the **Move to Archive** action appear as **Archive Policy**, and personal tags with the **Delete and Allow Recovery** or **Permanently Delete** actions appear as **Retention Policy**, as shown here:

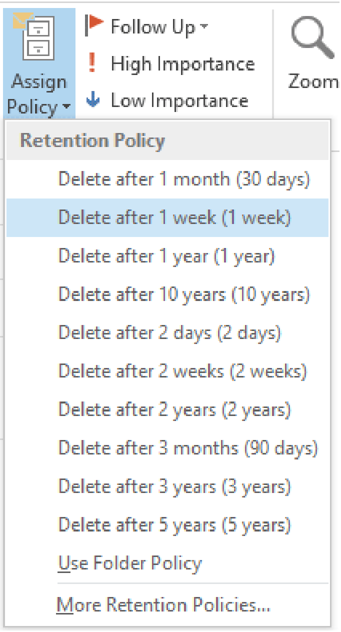
Assign Policy
Follow Up
High Importance
Low Importance
Zoom
Retention Policy
Delete after 1 month (30 days)
Delete after 1 week (1 week)
Delete after 1 year (1 year)
Delete after 10 years (10 years)
Delete after 2 days (2 days)
Delete after 2 weeks (2 weeks)
Delete after 2 years (2 years)
Delete after 3 months (90 days)
Delete after 3 years (3 years)
Delete after 5 years (5 years)
Use Folder Policy
More Retention Policies...

Outlook

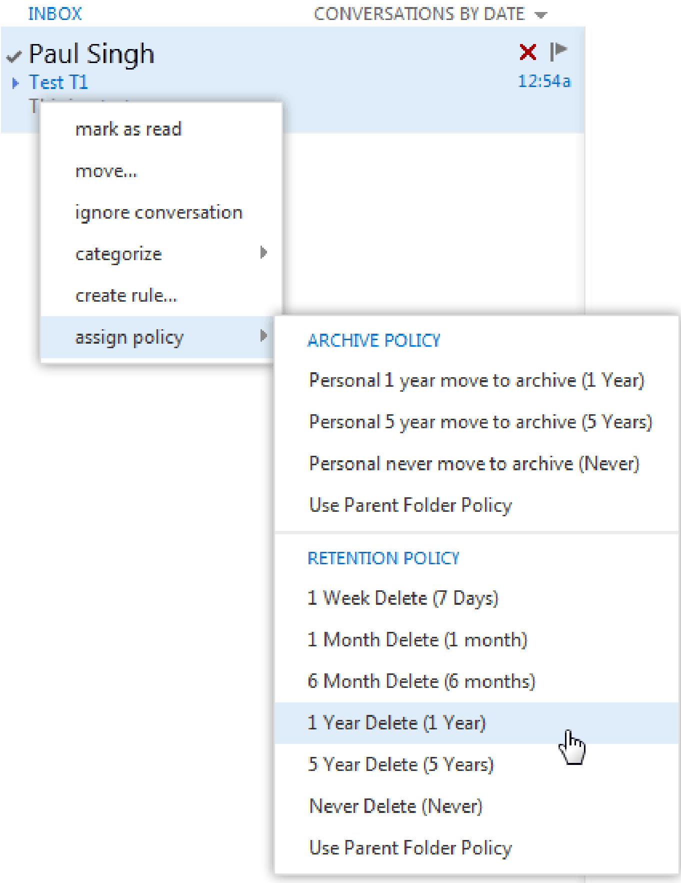

Outlook on the web

Users can apply personal tags to folders they create or to individual items. Messages that have a personal tag applied are always processed based on the personal tag's settings. Users can apply a personal tag to a message so that it's moved or deleted sooner or later than the settings specified in the DPT or RPTs applied to that user's mailbox. You can also create personal tags with retention disabled. This allows users to tag items so they're never moved to an archive or never expire.

NOTE

Users can apply archive policies to default folders, user-created folders or subfolders, and individual items. Users can apply a retention policy to user-created folders or subfolders and individual items (including subfolders and items in a default folder), but not to default folders.

Users can also use the Exchange admin center (EAC) to select additional personal tags that aren't linked to their retention policy. The selected tags then become available in Outlook and Outlook on the web. To enable users to select additional tags from the EAC, you must add the MyRetentionPolicies Role to the user's role assignment policy. To learn more about role assignment policies for users, see Understanding Management Role Assignment Policies. If you allow users to select additional personal tags, all personal tags in your Exchange organization become available to them.

NOTE

Personal tags are a premium feature. Mailboxes with policies that contain these tags (or as a result of users adding the tags to their mailbox) require an Exchange Enterprise client access license (CAL).

Retention age

When you enable a retention tag, you must specify a retention age for the tag. This age indicates the number of days to retain a message after it arrives in the user's mailbox.

The retention age for non-recurring items (such as email messages) is calculated differently than items that have an end date or recurring items (such as meetings and tasks). To learn how retention age is calculated for different types of items, see How retention age is calculated in Exchange Server.

You can also create retention tags with retention disabled or disable tags after they're created. Because messages that have a disabled tag applied aren't processed, no retention action is taken. As a result, users can use a disabled personal tag as a **Never Move** tag or a **Never Delete** tag to override a DPT or RPT that would otherwise apply to the message.

Retention actions

When creating or configuring a retention tag, you can select one of the following retention actions to be taken when an item reaches its retention age:

RETENTION ACTION	ACTION TAKEN...	EXCEPT...
Move to archive	Moves the message to the user's archive mailbox Only available for DPTs and personal tags For details about archiving, see In-Place Archiving in Exchange Server.	If the user doesn't have an archive mailbox, no action is taken.
Delete and allow recovery:	Emulates the behavior when the user empties the Deleted Items folder. Items are moved to the Recoverable Items folder in Exchange Server in the mailbox and preserved until the *deleted item retention* period. Provides the user a second chance to recover the item using the **Recover Deleted Items** dialog box in Outlook or Outlook on the web	If you've set the deleted item retention period to zero days, items are permanently deleted. For details, see Configure Deleted Item retention and Recoverable Items quotas.
Permanently delete	Permanently deletes messages. You can't recover messages after they're permanently deleted.	If mailbox is placed on In-Place Hold and Litigation Hold in Exchange Server or Litigation Hold, items are preserved in the Recoverable Items folder based on hold parameters. In-Place eDiscovery in Exchange Server will still return these items in search results.
Mark as past retention limit	Marks a message as expired. In Outlook, and Outlook on the web, expired items are displayed with the notification stating 'This item has expired' and 'This item will expire in 0 days'.	N. A.

For details about how to create retention tags, see Create a retention policy in Exchange Server.

Retention policies

To apply one or more retention tags to a mailbox, you need to add them to a retention policy and then apply the policy to mailboxes. A mailbox can't have more than one retention policy. Retention tags can be linked to or unlinked from a retention policy at any time, and the changes automatically take effect for all mailboxes that have the policy applied.

A retention policy can have the following retention tags:

RETENTION TAG TYPE	TAGS IN A POLICY
Default policy tag (DPT)	One DPT with the **Move to archive** action One DPT with the **Delete and allow Recovery** or **Permanently delete** actions One DPT for voice mail messages with the **Delete and allow recovery** or **Permanently delete** action
Retention policy tags (RPTs)	One RPT for each supported default folder **Note**: You can't link more than one RPT for a particular default folder (such as **Deleted Items**) to the same retention policy.
Personal tags	Any number of personal tags **Note**: Many personal tags in a policy can confuse users. We recommend adding no more than 10 personal tags to a retention policy.

NOTE

Although a retention policy doesn't need to have any retention tags linked to it, we don't recommend using this scenario. If mailboxes with retention policies don't have retention tags linked to them, this may cause mailbox items to never expire.

A retention policy can contain both archive tags (tags that move items to the personal archive mailbox) and deletion tags (tags that delete items). A mailbox item can also have both types of tags applied. Archive mailboxes don't have a separate retention policy. The same retention policy is applied to the primary and archive mailbox.

When planning to create retention policies, you must consider whether they'll include both archive and deletion tags. As mentioned earlier, a retention policy can have one DPT that uses the **Move to archive** action and one DPT that uses either the **Delete and allow recovery** or **Permanently delete** action. The DPT with the **Move to archive** action must have a lower retention age than the DPT with a deletion action. For example, you can use a DPT with the **Move to archive** action to move items to the archive mailbox in two years, and a DPT with a deletion action to remove items from the mailbox in seven years. Items in both primary and archive mailboxes will be deleted after seven years.

Default retention policy

Exchange Setup creates the retention policy **Default MRM Policy**. The policy is applied automatically if you create an archive for the new user and don't specify a retention policy

You can modify tags included in the Default MRM Policy, for example by changing the retention age or retention action, disable a tag or modify the policy by adding or removing tags from it. The updated policy is applied to mailboxes the next time they're processed by the Managed Folder Assistant.

For more details, including a list of retention tags linked to the policy, see Default Retention Policy.

Managed Folder Assistant

The Managed Folder Assistant, a mailbox assistant that runs on Mailbox servers, processes mailboxes that have a retention policy applied.

The Managed Folder Assistant applies the retention policy by inspecting items in the mailbox and determining whether they're subject to retention. It then stamps items subject to retention with the appropriate retention tags and takes the specified retention action on items past their retention age.

The Managed Folder Assistant is a throttle-based assistant. Throttle-based assistants are always running and don't need to be scheduled. The system resources they can consume are throttled. You can configure the Managed Folder Assistant to process all mailboxes on a Mailbox server within a certain period (known as a *work cycle*). Additionally, at a specified interval (known as the *work cycle checkpoint*), the assistant refreshes the list of mailboxes to be processed. During the refresh, the assistant adds newly created or moved mailboxes to the queue. It also reprioritizes existing mailboxes that haven't been processed successfully due to failures and moves them higher in the queue so they can be processed during the same work cycle.

You can also use the Start-ManagedFolderAssistant cmdlet to manually trigger the assistant to process a specified mailbox. To learn more, see Configure and run the Managed Folder Assistant in Exchange Server.

> **NOTE**
>
> The Managed Folder Assistant doesn't take any action on messages that aren't subject to retention, specified by disabling the retention tag. You can also disable a retention tag to temporarily suspend items with that tag from being processed.

Moving items between folders

A mailbox item moved from one folder to another inherits any tags applied to the folder to which it's moved. If an item is moved to a folder that doesn't have a tag assigned, the DPT is applied to it. If the item has a tag explicitly assigned to it, the tag always takes precedence over any folder-level tags or the default tag.

Applying a retention tag to a folder in the archive

When the user applies a personal tag to a folder in the archive, if a folder with the same name exists in the primary mailbox and has a different tag, the tag on that folder in the archive changes to match the one in the primary mailbox. This is by design to avoid any confusion about items in a folder in the archive having a different expiry behavior than the same folder in the user's primary mailbox. For example, the user has a folder named Project Contoso in the primary mailbox with a *Delete - 3 years* tag and a Project Contoso folder also exists in the archive mailbox. If the user applies a *Delete - 1 year* personal tag to delete items in the folder after 1 year. When the mailbox is processed again, the folder reverts to the Delete - 3 Years tag.

Removing or deleting a retention tag from a retention policy

When a retention tag is removed from the retention policy applied to a mailbox, the tag is no longer available to the user and can't be applied to items in the mailbox.

Existing items that have been stamped with that tag continue to be processed by the Managed Folder Assistant based on those settings and any retention action specified in the tag is applied to those messages.

However, if you delete the tag, the tag definition stored in Active Directory is removed. This causes the Managed Folder Assistant to process all items in a mailbox and restamp the ones that have the removed tag applied. Depending on the number of mailboxes and messages, this process may significantly consume resources on all Mailbox servers that contain mailboxes with retention policies that include the removed tag.

> **IMPORTANT**
>
> If a retention tag is removed from a retention policy, any existing mailbox items with the tag applied will continue to expire based on the tag's settings. To prevent the tag's settings from being applied to any items, you should delete the tag. Deleting a tag removes it from any retention policies where it's included.

Disabling a retention tag

If you disable a retention tag, the Managed Folder Assistant ignores items that have that tag applied. Items that have a retention tag for which retention is disabled are either never moved or never deleted, depending on the specified retention action. Because these items are still considered tagged items, the DPT doesn't apply to them. For example, if you want to troubleshoot retention tag settings, you can temporarily disable a retention tag to stop the Managed Folder Assistant from processing messages with that tag.

NOTE

The retention period for a disabled retention tag is displayed to the user as **Never**. If a user tags an item believing it will never be deleted, enabling the tag later may result in unintentional deletion of items the user didn't want to delete. The same is true for tags with the **Move to archive** action.

Retention hold

When users are temporarily away from work and don't have access to their email, retention settings can be applied to new messages before they return to work or access their email. Depending on the retention policy, messages may be deleted or moved to the user's personal archive. You can temporarily suspend retention policies from processing a mailbox for a specified period by placing the mailbox on retention hold. When you place a mailbox on retention hold, you can also specify a retention comment that informs the mailbox user (or another user authorized to access the mailbox) about the retention hold, including when the hold is scheduled to begin and end. Retention comments are displayed in supported Outlook clients. You can also localize the retention hold comment in the user's preferred language.

NOTE

Placing a mailbox on retention hold doesn't affect how mailbox storage quotas are processed. Depending on the mailbox usage and applicable mailbox quotas, consider temporarily increasing the mailbox storage quota for users when they're on vacation or don't have access to email for an extended period. For more information about mailbox storage quotas, see Configure storage quotas for a mailbox.

During long absences from work, users may accrue a large amount of email. Depending on the volume of email and the length of absence, it may take these users several weeks to sort through their messages. In these cases, consider the additional time it may take the users to catch up on their mail before removing them from retention hold.

If your organization has never implemented MRM, and your users aren't familiar with its features, you can also use retention holds during the initial *warm up and training* phase of your MRM deployment. You can create and deploy retention policies and educate users about the policies without the risk of having items moved or deleted before users can tag them. A few days before the warm up and training period ends, you should remind users of the warm-up deadline. After the deadline, you can remove the retention hold from user mailboxes, allowing the Managed Folder Assistant to process mailbox items and take the specified retention action.

Apply a retention policy to mailboxes in Exchange Server

8/3/2020 • 2 minutes to read • Edit Online

You can use retention policies to group one or more retention tags and apply them to mailboxes to enforce message retention settings. A mailbox can't have more than one retention policy.

Caution

Messages are expired based on settings defined in the retention tags linked to the policy. These settings include actions such moving messages to the archive or permanently deleting them. Before applying a retention policy to one or more mailboxes, we recommended that you test the policy and inspect each retention tag associated with it.

What do you need to know before you begin?

- Estimated time to complete: 5 minutes.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Applying retention policies" entry in the Messaging policy and compliance permissions in Exchange Server topic.

Use the Exchange admin center to apply a retention policy to a single mailbox

1. Go to **Recipients** > **Mailboxes**.
2. In the list view, select the mailbox to which you want to apply the retention policy, and then click **Edit** ✎.
3. In **User Mailbox**, click **Mailbox features**.
4. In the **Retention policy** list, select the policy you want to apply to the mailbox, and then click **Save**.

Use the Exchange admin center to apply a retention policy to multiple mailboxes

1. Go to **Recipients** > **Mailboxes**.
2. In the list view, use the Shift or Ctrl keys to select multiple mailboxes.
3. In the details pane, click **More options**.
4. Under **Retention Policy**, click **Update**.
5. In **Bulk Assign Retention Policy**, select the retention policy you want to apply to the mailboxes, and then click **Save**.

Use the Exchange Management Shell to apply a retention policy to a single mailbox

This example applies the retention policy RP-Finance to Morris's mailbox.

```
Set-Mailbox "Morris" -RetentionPolicy "RP-Finance"
```

For detailed syntax and parameter information, see Set-Mailbox.

Use the Exchange Management Shell to apply a retention policy to multiple mailboxes

This example applies the new retention policy New-Retention-Policy to all mailboxes that have the old policy Old-Retention-Policy.

```
$OldPolicy=(Get-RetentionPolicy "Old-Retention-Policy").distinguishedName
Get-Mailbox -Filter "RetentionPolicy -eq '$OldPolicy'" -Resultsize Unlimited | Set-Mailbox -RetentionPolicy
"New-Retention-Policy"
```

This example applies the retention policy RetentionPolicy-Corp to all mailboxes in the Exchange organization.

~~~~
Get-Mailbox -ResultSize unlimited | Set-Mailbox -RetentionPolicy "RetentionPolicy-Corp"
```PowerShell

This example applies the retention policy RetentionPolicy-Finance to all mailboxes in the Finance
organizational unit.

```PowerShell
Get-Mailbox -OrganizationalUnit "Finance" -ResultSize Unlimited | Set-Mailbox -RetentionPolicy
"RetentionPolicy-Finance"
~~~~

For detailed syntax and parameter information, see Get-Mailbox and Set-Mailbox.

How do you know this worked?

To verify that you have applied the retention policy, run the Get-Mailbox cmdlet to retrieve the retention policy for the mailbox or mailboxes.

This example retrieves the retention policy for Morris's mailbox.

```
Get-Mailbox Morris | Select RetentionPolicy
```

This command retrieves all mailboxes that have the retention policy RP-Finance applied.

```
Get-Mailbox -ResultSize unlimited | Where-Object {$_.RetentionPolicy -eq "RP-Finance"} | Format-Table
Name,RetentionPolicy -Auto
```

Configure and run the Managed Folder Assistant in Exchange Server

8/3/2020 • 4 minutes to read • Edit Online

The *Managed Folder Assistant* (MFA) is an Exchange Mailbox Assistant that applies and processes the message retention settings that are configured in retention policies.

As in Exchange 2013, the Managed Folder Assistant in Exchange 2016 and Exchange 2019 is a throttle-based assistant that's always running. The MFA doesn't need to be scheduled, and the system resources that are consumed by the MFA can be throttled. You can configure the Managed Folder Assistant to process all mailboxes on a Mailbox server within a certain time period that's known as a *work cycle*. By default, the work cycle for the MFA is one day (all mailboxes on the server are processed by the MFA every day).

You can also force the MFA to immediately process a specified mailbox.

What do you need to know before you begin?

- You can only use PowerShell to perform this procedure. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.

- Although the *ManagedFolderAssistantSchedule* parameter is available in Exchange Server, it doesn't work on Exchange 2016 or Exchange 2019 servers. It's only used for coexistence with previous versions of Exchange.

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Messaging records management" entry in the Messaging policy and compliance permissions in Exchange Server topic.

Configure the Managed Folder Assistant

Configuring the interval for when the MFA processes mailboxes is a two-step process:

1. Configure the work cycle for the MFA.

2. Apply the new work cycle value for the MFA.

Step 1: Use the Exchange Management Shell to configure the work cycle for the Managed Folder Assistant

To configure the work cycle for the MFA, use this syntax:

```
New-SettingOverride -Name "<UniqueOverrideName>" -Component TimeBasedAssistants -Section ELCAssistant -
Parameters @("WorkCycle=<Timespan>") -Reason "<DescriptiveReason>" [-Server <ServerName>]
```

Notes:

- To specify a *<TimeSpan>* value, use the syntax `d.hh:mm:ss`, where *d* = days, *hh* = hours, *mm* = minutes, and *ss* = seconds.

- To configure the same work cycle for the MFA on all Exchange 2016 and Exchange 2019 Mailbox servers in the Active Directory forest, don't use the *Server* parameter.

- To configure the work cycle for the MFA on a specific Exchange 2016 and Exchange 2019 Mailbox server, use the *Server* parameter and the name (not the fully qualified domain name or FQDN) of the server. This method is useful when you need to specify different work cycle values for the MFA on different Exchange servers.

This example configures the work cycle for the MFA to two days (the MFA processes mailboxes every two days). Because we aren't using the *Server* parameter, the setting is applied to all Exchange 2016 and Exchange 2019 Mailbox servers in the organization.

- **Setting override name**: "MFA WorkCycle Override" (must be unique)
- **WorkCycle**: `2.00:00:00` (2 days; note the value `2` also works)
- **Override reason**: Process mailboxes every 2 days

```
New-SettingOverride -Name "MFA WorkCycle Override" -Component TimeBasedAssistants -Section ELCAssistant -
Parameters @("WorkCycle=2.00:00:00") -Reason "Process mailboxes every 2 days"
```

This example specifies the same 2 day work cycle for the MFA, but only on the server named Mailbox01.

```
New-SettingOverride -Name "Mailbox01 MFA WorkCycle Override" -Component TimeBasedAssistants -Section
ELCAssistant -Parameters @("WorkCycle=2.00:00:00") -Reason "Process mailboxes every 2 days" -Server Mailbox01
```

Step 2: Use the Exchange Management Shell to apply the new the work cycle value for the Managed Folder Assistant

To apply the new the work cycle value for the MFA, use this syntax:

```
Get-ExchangeDiagnosticInfo -Process Microsoft.Exchange.Directory.TopologyService -Component
VariantConfiguration -Argument Refresh [-Server <ServerName>]
```

Notes:

- If you didn't use the *Server* parameter in Step 1, don't use it here. If you used the *Server* parameter in Step 1, use the same server name here.
- If you delete the custom work cycle value for the MFA by using the **Remove-SettingOverride** cmdlet, you still need to run this command to change the work cycle back to the default value of one day.

This example applies the new work cycle value for the MFA on all Exchange 2016 and Exchange 2019 Mailbox servers in the organization.

```
Get-ExchangeDiagnosticInfo -Process Microsoft.Exchange.Directory.TopologyService -Component
VariantConfiguration -Argument Refresh
```

This example applies the new work cycle value for the MFA on the server named Mailbox01.

```
Get-ExchangeDiagnosticInfo -Process Microsoft.Exchange.Directory.TopologyService -Component
VariantConfiguration -Argument Refresh -Server Mailbox01
```

How do you know this worked?

To verify that you've successfully configured the work cycle for the Managed Folder Assistant on one or more servers, replace *<ServerName>* with the name of the server (not the FQDN), and run the following command to verify the value of the **WorkCycle** property:

```
[xml]$diag=Get-ExchangeDiagnosticInfo -Server <ServerName> -Process MSExchangeMailboxAssistants -Component
VariantConfiguration -Argument "Config,Component=TimeBasedAssistants"
$diag.Diagnostics.Components.VariantConfiguration.Configuration.TimeBasedAssistants.ElcAssistant
```

Use the Exchange Management Shell to start the Managed Folder Assistant on a specific mailbox

To trigger the MFA to immediately process a mailbox, use this syntax:

```
Start-ManagedFolderAssistant -Identity <MailboxIdentity>
```

This example triggers the Managed Folder Assistant to immediately process Morris Cornejo's mailbox.

```
Start-ManagedFolderAssistant -Identity morris.cornejo@contoso.com
```

For detailed syntax and parameter information, see Start-ManagedFolderAssistant.

Administrator audit logging in Exchange Server

8/3/2020 • 13 minutes to read • Edit Online

You can use administrator audit logging in Exchange Server to log when a user or administrator makes a change in your organization. By keeping a log of the changes, you can trace changes to the person who made the change, augment your change logs with detailed records of the change as it was implemented, comply with regulatory requirements and requests for discovery, and more.

By default, administrator audit logging is enabled in new installations of Exchange Server.

What gets audited

Cmdlets that are run directly in the Exchange Management Shell are audited. In addition, operations performed using the Exchange admin center (EAC) are also logged because those operations run cmdlets in the background.

Cmdlets, regardless of where they're run, are audited if a cmdlet is on the cmdlet auditing list and one or more parameters on that cmdlet are on the parameter auditing list. Audit logging is intended to show what actions have been taken to modify objects in an Exchange organization rather than what objects have been viewed.

Notes:

- A cmdlet might not be logged if an error occurs before the cmdlet calls the Admin Audit Log cmdlet extension agent. If an error occurs after the Admin Audit Log agent is called, the cmdlet is logged along with the associated error. For more information, see the Admin Audit Log agent section later in this topic.
- Changes to the audit log configuration are refreshed every 60 minutes on computers that have the Exchange Management Shell open at the time a configuration change is made. If you want to apply the changes immediately, close and then open the Exchange Management Shell again on each computer.
- A command may take up to 15 minutes after it's run to appear in audit log search results. This is because audit log entries must be indexed before they can be searched. If a command doesn't appear in the administrator audit log, wait a few minutes and run the search again.

Admin audit logging configuration

By default, when admin audit logging is enabled, a log entry is created every time any cmdlet is run. If you don't want to audit every cmdlet that's run, you can configure audit logging to audit only the cmdlets and parameters you're interested in. You configure audit logging with the **Set-AdminAuditLogConfig** cmdlet. The parameters referenced in the following sections are used with this cmdlet.

> **IMPORTANT**
> Changes to the administrator audit log configuration are always logged, regardless of whether the **Set-AdminAuditLogConfig** cmdlet is included in the list of cmdlets being audited or whether audit logging is enabled or disabled.

When a command is run, Exchange inspects the cmdlet that was used. If the cmdlet that was run matches any of the cmdlets provided with the *AdminAuditLogCmdlets* parameter, Exchange then checks the parameters specified in the *AdminAuditLogParameters* parameter. If at least one or more parameters from the parameters list are matched, Exchange logs the cmdlet that was run. The following sections contain more information about each aspect of the audit logging configuration.

For more information about managing audit logging configuration, see Manage administrator audit logging.

Cmdlets

You can control which cmdlets are audited by providing a list of cmdlets, and their parameters, that you want to log. When you configure audit logging, you can specify to audit every cmdlet, or you can specify the cmdlets you want to audit by using the *AdminAuditLogCmdlets* parameter. You can specify full cmdlet names, such as **New-Mailbox**, or you can specify partial cmdlet names and enclose those names in wildcard characters, such as an asterisk (`*`). For example, if you want to log when any cmdlet that contains the string `Transport` runs, you can specify a value of `*Transport*`. You can use a mix of full cmdlet names and partial cmdlet names at the same time to tailor the audit logging configuration to your needs.

To audit all cmdlets, specify only the wildcard character (*). This is the default setting.

Parameters

In addition to specifying which cmdlets you want to log, you can also indicate that cmdlets should only be logged if certain parameters on those cmdlets are used. Use the *AdminAuditLogParameters* parameter to specify which parameters should be logged. As with cmdlets, you can specify full parameter names, such as `Database`, or partial parameter names enclosed in wildcard characters (`*`), such as `*Address*`, or a combination of both.

To audit all parameters, specify only the wildcard character (*). This is the default setting.

Admin audit log age limit

By default, admin audit logging is configured to store audit log entries for 90 days. After 90 days, the audit log entry is deleted. You can change the audit log age limit using the *AdminAuditLogAgeLimit* parameter. For example, to change the age limit to 180 days, use the command `Set-AdminAuditLogConfig -AdminAuditLogAgeLimit 180`. You can also specify the number of days, hours, minutes, and seconds that audit log entries should be kept. To specify a value, use the format `dd.hh:mm:ss` where the following applies:

- **dd**: The number of days to keep the audit log entry.
- **hh**: The number of hours to keep the audit log entry.
- **mm**: The number of minutes to keep the audit log entry.
- **ss**: The number of seconds to keep the audit log entry.

You need to specify multiple years by using the `dd` field. For example, 365 days equals one year; 730 days equals two years; 913 days equals two years and six months. For example, to set the audit log age limit to two years and six months, use the value `913`.

Notes:

- You can set the admin audit log age limit to a value that's less than the current age limit. If you do this, any audit log entry whose age exceeds the new age limit is deleted.
- If you set the age limit to 0, Exchange deletes all the entries in the audit log.
- We recommend that you assign permissions to configure the audit log age limit only to highly trusted users.

Verbose logging

By default, the admin audit log records only the cmdlet name, cmdlet parameters (and values specified), the object that was modified, who ran the cmdlet, when the cmdlet was run, and on what server the cmdlet was run. The admin audit log doesn't log what properties were modified on the object. If you want the admin audit log to also include the properties of the object that were modified, you can enable verbose logging by setting the *LogLevel* parameter to `Verbose`. When you enable verbose logging, in addition to the information logged by default, the properties modified on an object, including their old and new values, are included in the admin audit log.

Test cmdlets

Cmdlets that begin with the verb **Test** aren't logged by default. You can indicate that **Test** cmdlets should be logged by setting the *TestCmdletLoggingEnabled* parameter to `$true`. Although you can enable logging of test cmdlets, we recommend that you do this only for short periods of time because test cmdlets can produce a large number of audit log entries.

Admin audit log

Each time a cmdlet is logged, an admin audit log entry is created. The audit log entries are stored in the admin audit log, which is stored in a hidden, dedicated arbitration mailbox that can only be accessed by using the EAC, the **Search-AdminAuditLog** cmdlet, or the **New-AdminAuditLogSearch** cmdlet. The following sections provide information about:

- What's included in the admin audit log.
- Reports available on the EAC **Auditing** page.
- Admin audit log search cmdlets.

Audit log contents

Each audit log entry contains the information described in the following table. The audit log contains one or more audit log entries. The number of audit log entries is controlled by the audit log age limit specified using the `Set-AdminAuditLogConfig -AdminAuditLogAgeLimit` command. Any audit log entry that exceeds the age limit is deleted.

Audit log entry fields

FIELD	DESCRIPTION
`RunspaceId`	This field is used internally by Exchange.
`ObjectModified`	This field contains the object that was modified by the cmdlet specified in the `CmdletName` field.
`CmdletName`	This field contains the name of the cmdlet that was run by the user in the `Caller` field.
`CmdletParameters`	This field contains the parameters that were specified when the cmdlet in the `CmdletName` field was run. Also stored in this field, but not visible in the default output, is the value specified with the parameter, if any.
`ModifiedProperties`	This field contains the properties that were modified on the object in the `ObjectModified` field. Also stored in this field, but not visible in the default output, are the old value of the property and the new value that was stored. **Important**: This field is only populated if the *LogLevel* parameter on the **Set-AdminAuditLogConfig** cmdlet is set to `verbose`.
`Caller`	This field contains the user account of the user who ran the cmdlet in the `CmdletName` field.
`Succeeded`	This field specifies whether the cmdlet in the `CmdletName` field ran successfully. The value is either `True` or `False`.

FIELD	DESCRIPTION
`Error`	This field contains the error message generated if the cmdlet in the `CmdletName` field failed to complete successfully.
`RunDate`	This field contains the date and time when the cmdlet in the `CmdletName` field was run. The date and time are stored in Coordinated Universal Time (UTC) format.
`OriginatingServer`	This field indicates the server on which the cmdlet specified in the `CmdletName` field was run.
`Identity`	This field is used internally by Exchange.
`IsValid`	This field is used internally by Exchange.
`ObjectState`	This field is used internally by Exchange.

EAC auditing reports

The **Auditing** page in the EAC has several reports that provide information about various types of compliance and administrative configuration changes. The following reports provide information about configuration changes in your organization:

- **Administrator role group report**: This report enables you to search for changes to management role groups that you specify within a specified timeframe. The results that are returned include the role groups that have been changed, who changed them and when, and what changes were made. A maximum of 3,000 entries can be returned. If your search might return more than 3,000 entries, use the **Administrator audit log** report or the **Search-AdminAuditLog** cmdlet.

- **Admin audit log report**: This report enables you to view entries in the admin audit log recorded within a specified time frame. You can also export admin audit log entries to a XML file and then send the file via email to a recipient you specify. For more information about the contents of the XML file, see Administrator audit log structure.

For information about how to use these reports, see Search the role group changes or administrator audit logs.

Search-AdminAuditLog cmdlet

When you run the **Search-AdminAuditLog** cmdlet, all the audit log entries that match your search criteria are returned. You can specify the following search criteria:

- **Cmdlets**: Specifies the cmdlets you want to search for in the admin audit log.

- **Parameters**: Specifies the parameters, separated by commas, you want to search for in the admin audit log. You can only search for parameters if you specify a cmdlet to search for.

- **End date**: Scopes the admin audit log results to log entries that occurred on or before the specified date.

- **Start date**: Scopes the admin audit log results to log entries that occurred on or after the specified date.

- **Object IDs**: Specifies that only admin audit log entries that contain the specified changed objects should be returned

- **User IDs**: Specifies that only the admin audit log entries that contain the specified IDs of the user who ran the cmdlet should be returned.

- **Successful completion**: Specifies whether only admin audit log entries that indicated a success or failure should be returned.

Each audit log entry contains the information described in the table in Audit log contents. By default, only the first 1,000 log entries that match the search criteria are returned. However, you can override this default and return more or fewer entries using the *ResultSize* parameter. You can specify a value of `Unlimited` with the *ResultSize* parameter to return all log entries that match the specified criteria.

For information about how to use the **Search-AdminAuditLog** cmdlet, see Search the role group changes or administrator audit logs.

New-AdminAuditLogSearch cmdlet

The **New-AdminAuditLogSearch** cmdlet searches the admin audit log just like the **Search-AdminAuditLog** cmdlet. However, instead of displaying the results of the search in the Exchange Management Shell, the **New-AdminAuditLogSearch** cmdlet performs the search and then sends the results to a recipient you specify via an email message. The results are included as an XML attachment to the email message.

You can use the same search criteria with the **New-AdminAuditLogSearch** cmdlet that's used on the **Search-AdminAuditLog** cmdlet. For a list of the search criteria, see Search-AdminAuditLog cmdlet.

After you run the **New-AdminAuditLogSearch** cmdlet, Exchange may take up to 15 minutes to deliver the report to the specified recipient. The XML file attached report can be a maximum of 10 MB. The XML file contains the same information described in the table in Audit log contents. For more information about the structure of the XML file, see Administrator audit log structure.

> **NOTE**
>
> Outlook Web App doesn't allow you to open XML attachments by default. You can either configure Exchange to allow XML attachments to be viewed using Outlook Web App, or you can use another email client, such as Microsoft Outlook, to view the attachment. For information about how to configure Outlook Web App to allow you to view an XML attachment, see View or configure Outlook on the web virtual directories in Exchange Server.

For information about how to use the **New-AdminAuditLogSearch** cmdlet, see Search the role group changes or administrator audit logs.

Manual admin audit log entries

In addition to logging Exchange cmdlets when they're run, Exchange Server enables you to manually write log entries to the audit log. Exchange Server supports this using the **Write-AdminAuditLog** cmdlet. Situations where you might want to add a manual log entry include the following:

- Custom script entry and exit
- Change control information
- Maintenance start and end times

With the **Write-AdminAuditLog** cmdlet, you specify a string of text to include in the audit log using the *Comment* parameter. The *Comment* parameter accepts an alphanumeric string up to 500 characters. Included in the manual audit log entry along with the comment string is all of the same information captured when an Exchange cmdlet is logged. For a description of each field included in the audit log, see the table in Audit log contents.

You can retrieve manual audit log entries the same way as any other log entry, using the EAC **Auditing** page or using the **Search-AdminAuditLog** or **New-AdminAuditLogSearch** cmdlets.

To view the contents of the *Comment* parameter on the **Write-AdminAuditLog** cmdlet in a manual audit log entry, see Search the role group changes or administrator audit logs.

Active Directory replication

Administrator audit logging relies on Active Directory replication to replicate the configuration settings you specify to the domain controllers in your organization. Depending on your replication settings, the changes you make may not be immediately applied to all servers running Exchange in your organization.

Admin Audit Log agent

The Admin Audit Log built-in cmdlet extension agent performs admin audit logging of cmdlet operations in Exchange Server. This agent reads the audit log configuration and then performs an evaluation of each cmdlet run in your organization. If the criteria you've specified in the admin audit log configuration matches the cmdlet that's being run, the agent generates an audit log entry.

The Admin Audit Log agent is enabled by default, which is required for admin audit logging to function. It can't be disabled, and its priority can't be changed. For more information about cmdlet extension agents, see Cmdlet Extension Agents.

Administrator audit log structure

8/3/2020 • 3 minutes to read • Edit Online

Administrator audit logs contain a record of all the cmdlets and parameters that have been run in the Exchange Management Shell and by the Exchange admin center (EAC). They're created on-demand when you run the admin audit log report in the EAC, or when you run the **New-AdminAuditLogSearch** cmdlet in the Exchange Management Shell. For more information about audit logs, see Administrator audit logging in Exchange Server.

Audit log XML tags and attributes

The audit logs are XML files and can contain multiple audit log entries. The following table describes each XML tag and its associated attributes.

ELEMENT	ATTRIBUTE	DESCRIPTION
`<?xml version="1.0" encoding="utf-8"?>`	N/A	This is the XML document declaration tag. It's included in every audit log XML file and contains the XML version number and the character encoding value.
`SearchResults`	N/A	This tag contains all the audit log entries in the XML file. The `Event` tag is a child of this tag. There is only one `SearchResults` tag per XML file.
`Event`		This tag contains the audit log entry for an individual cmdlet. This tag contains the `Caller`, `Cmdlet`, `ObjectModified`, `RunDate`, `Succeeded`, `Error`, and `OriginatingServer` attributes. The `CmdletParameters` and `ModifiedProperties` tags are children of this tag. There is one `Event` tag per audit log entry.
	`Caller`	This attribute contains the user account of the user who ran the cmdlet in the `Cmdlet` attribute.
	`Cmdlet`	This attribute contains the name of the cmdlet that was run by the user in the `Caller` attribute.
	`ObjectModified`	This attribute contains the object that was modified by the cmdlet specified in the `Cmdlet` attribute. The `ModifiedProperties` tag shows which properties were modified on this object.

ELEMENT	ATTRIBUTE	DESCRIPTION
	`RunDate`	This attribute contains the date and time when the cmdlet in the `Cmdlet` attribute was run.
	`Succeeded`	This attribute specifies whether the cmdlet in the `Cmdlet` attribute ran successfully. The value is either `True` or `False`.
	`Error`	This attribute contains the error message generated if the cmdlet in the `Cmdlet` attribute failed to complete successfully. If no error was encountered, the value is set to `None`.
	`OriginatingServer`	This attribute contains the server on which the cmdlet specified in the `Cmdlet` attribute was run.
`CmdletParameters`	N/A	This tag contains all of the parameters specified when the cmdlet was run. The `Parameter` tag is a child of this tag. There is one `CmdletParameters` tag per `Event` tag.
`Parameter`		This tag contains an individual parameter that was specified when the cmdlet was run. This tag contains the `Name` and `Value` attributes. There can be multiple `Parameter` tags per `CmdletParameters` tag.
	`Name`	This attribute contains the name of the parameter that was specified on the cmdlet that was run.
	`Value`	This attribute contains the value that was provided on the parameter specified in the `Name` attribute.
`ModifiedProperties`	N/A	This tag contains all of the properties that were modified by the cmdlet that was run. The `Property` tag is a child of this tag. There is one `ModifiedProperties` tag per `Event` tag. **Important**: This tag is only populated if the *LogLevel* parameter on the **Set-AdminAuditLogConfig** cmdlet is set to `Verbose`.

ELEMENT	ATTRIBUTE	DESCRIPTION
`Property`		This tag contains an individual property that was specified when the cmdlet was run. This tag contains the `Name` , `OldValue` , and `NewValue` attributes. There can be multiple `Property` tags per `ModifiedProperties` tag.
	`Name`	This attribute contains the name of the property that was modified when the cmdlet was run.
	`OldValue`	This attribute contains the value that was contained in the property specified in the `Name` attribute before it was changed.
	`NewValue`	This attribute contains the value that the property in the `Name` attribute was changed to.

Example of an admin audit log entry

The following is an example of a typical log entry in the admin audit log.

```
<?xml version="1.0" encoding="utf-8"?>
<SearchResults>
  <Event Caller="corp.e16.contoso.com/Users/Administrator" Cmdlet="Set-Mailbox"
ObjectModified="corp.e16.contoso.com/Users/david" RunDate="2015-10-18T15:48:15-07:00" Succeeded="true"
Error="None" OriginatingServer="WIN8MBX (15.01.0396.030)">
    <CmdletParameters>
      <Parameter Name="Identity" Value="david" />
      <Parameter Name="ProhibitSendReceiveQuota" Value="10 GB (10,737,418,240 bytes)" />
    </CmdletParameters>
    <ModifiedProperties>
      <Property Name="ProhibitSendReceiveQuota" OldValue="35 GB (37,580,963,840 bytes)" NewValue="10 GB
(10,737,418,240 bytes)" />
    </ModifiedProperties>
  </Event>
</SearchResults>
```

Based on the information in this log entry, we know the following occurred:

- On 10/18/2017 at 3:48 P.M. Pacific Daylight Time (UTC-7), the user `Administrator` ran the cmdlet **Set-Mailbox**.
- The two following parameters were provided when the **Set-Mailbox** cmdlet was run:
  - *Identity* with a value of `david`
  - *ProhibitSendReceiveQuota* with a value of `10GB`
- The *ProhibitSendReceiveQuota* property on the object `david` was modified with a new value of `10GB` , which replaced the old value of `35GB` .

NOTE

The modified properties are saved to the audit log because the *LogLevel* parameter on the `Set-AdminAuditLogConfig` cmdlet was set to `Verbose` in this example.

- The operation completed successfully without any errors.

Manage administrator audit logging

8/3/2020 • 5 minutes to read • Edit Online

Administrator audit logging in Exchange Server enables you to create a log entry each time a specified cmdlet is run. Log entries provide you with information about what cmdlet was run, which parameters were used, who ran the cmdlet, and what objects were affected. For more information about administrator audit logging, see Administrator audit logging in Exchange Server.

What do you need to know before you begin?

- Estimated time to complete each procedure: less than 5 minutes
- You can only use PowerShell to perform this procedure. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Administrator audit logging" entry in the Exchange infrastructure and PowerShell permissions topic.
- Admin audit logging relies on Active Directory replication to replicate the configuration settings you specify to the domain controllers in your organization. Depending on your replication settings, the changes you make may not be immediately applied to all Exchange 2016 and Exchange 2019 servers in your organization.
- Changes to the audit log configuration are refreshed every 60 minutes on computers that have the Exchange Management Shell open at the time a configuration change is made. If you want to apply the changes immediately, close and then open the Exchange Management Shell again on each computer.
- A command may take up to 15 minutes after it's run to appear in audit log search results. This is because audit log entries must be indexed before they can be searched. If a command doesn't appear in the administrator audit log, wait a few minutes and run the search again.

Specify the cmdlets to be audited

By default, audit logging creates a log entry for every cmdlet that's run. If you're enabling audit logging for the first time and want this behavior, you don't have to change the cmdlet audit list. If you've previously specified cmdlets to audit and now want to audit all cmdlets, you can audit all cmdlets by specifying the asterisk (*) wildcard character with the *AdminAuditLogCmdlets* parameter on the **Set-AdminAuditLogConfig** cmdlet, as shown in the following command.

```
Set-AdminAuditLogConfig -AdminAuditLogCmdlets *
```

You can specify which cmdlets to audit by providing a list of cmdlets using the *AdminAuditLogCmdlets* parameter. When you provide the list of cmdlets to audit, you can provide single cmdlets, cmdlets with the asterisk (*) wildcard characters, or a mix of both. Each entry in the list is separated by commas. The following values are all valid:

- `New-Mailbox`
- `*TransportRule`
- `*Management*`

- `Set-Transport*`

This example audits the cmdlets specified in the preceding list.

```
Set-AdminAuditLogConfig -AdminAuditLogCmdlets New-Mailbox, *TransportRule, *Management*, Set-Transport*
```

For detailed syntax and parameter information, see Set-AdminAuditLogConfig.

Specify the parameters to be audited

By default, audit logging creates a log entry for every cmdlet that's run, regardless of the parameters specified. If you're enabling audit logging for the first time and want this behavior, you don't have to change the parameter audit list. If you've previously specified parameters to audit and now want to audit all parameters, you can do so by specifying the asterisk (*) wildcard character with the *AdminAuditLogParameters* parameter on the **Set-AdminAuditLogConfig** cmdlet, as shown in the following command.

```
Set-AdminAuditLogConfig -AdminAuditLogParameters *
```

You can specify which parameters you want to audit by using the *AdminAuditLogParameters* parameter. When you provide the list of parameters to audit, you can provide single parameters, parameters with the asterisk (*) wildcard characters, or a mix of both. Each entry in the list is separated by commas. The following values are all valid:

- `Database`
- `*Address*`
- `Custom*`
- `*Region`

> **NOTE**
> For an audit log entry to be created when a command is run, the command must include at least one or more parameters that exist on at least one or more cmdlets specified with the *AdminAuditLogCmdlets* parameter.

This example audits the parameters specified in the preceding list.

```
Set-AdminAuditLogConfig -AdminAuditLogParameters Database, *Address*, Custom*, *Region
```

For detailed syntax and parameter information, see Set-AdminAuditLogConfig.

Specify the admin audit log age limit

The audit log age limit determines how long audit log entries will be retained. When a log entry exceeds the age limit, it's deleted. The default is 90 days.

You can specifiy the age limit in days. Or you can specify the number of days, hours, minutes, and seconds that audit log entries should be kept. To specify a value more specific than days, use the format dd.hh.mm:ss where the following applies:

- **dd**: Number of days to keep the audit log entry
- **hh**: Number of hours to keep the audit log entry

- **mm**: Number of minutes to keep the audit log entry
- **ss**: Number of seconds to keep the audit log entry

Caution

You can set the audit log age limit to a value that's less than the current age limit. If you do this, any audit log entry whose age exceeds the new age limit will be deleted. > If you set the age limit to 0, Exchange deletes all the entries in the audit log. > We recommend that you assign permissions to configure the audit log age limit only to highly trusted users.

This example specifies an age limit of two years and six months.

```
Set-AdminAuditLogConfig -AdminAuditLogAgeLimit 913
```

For detailed syntax and parameter information, see Set-AdminAuditLogConfig.

Enable or disable logging of Test cmdlets

Cmdlets that start with the verb **Test** aren't logged by default. This is because **Test** cmdlets can generate a significant amount of data in a short time. Only enable the logging of **Test** cmdlets for short periods of time.

This command enables the logging of **Test** cmdlets.

```
Set-AdminAuditLogConfig -TestCmdletLoggingEnabled $true
```

This command disables the logging of **Test** cmdlets.

```
Set-AdminAuditLogConfig -TestCmdletLoggingEnabled $false
```

For detailed syntax and parameter information, see Set-AdminAuditLogConfig.

Disable admin audit logging

To disable admin audit logging, use the following command.

```
Set-AdminAuditLogConfig -AdminAuditLogEnabled $false
```

Enable admin audit logging

To enable admin audit logging, use the following command.

```
Set-AdminAuditLogConfig -AdminAuditLogEnabled $true
```

View admin audit logging settings

To view the admin audit logging settings that you've configured for your organization, use the following command.

```
Get-AdminAuditLogConfig
```

Mailbox audit logging in Exchange Server

8/3/2020 • 6 minutes to read • Edit Online

Because mailboxes can contain sensitive, high business impact (HBI) information and personally identifiable information (PII), it's important that you track who logs on to the mailboxes in your organization and what actions are taken. It's especially important to track access to mailboxes by users other than the mailbox owner. These users are referred to as *delegate users*.

By using *mailbox audit logging*, you can log mailbox access by mailbox owners, delegates (including administrators with full access permissions to mailboxes), and administrators.

When you enable audit logging for a mailbox, you can specify which user actions (for example, accessing, moving, or deleting a message) will be logged for a logon type (administrator, delegate user, or owner). Audit log entries also include important information such as the client IP address, host name, and process or client used to access the mailbox. For items that are moved, the entry includes the name of the destination folder.

Mailbox audit logs

Mailbox audit logs are generated for each mailbox that has mailbox audit logging enabled. Log entries are stored in the Recoverable Items folder in the audited mailbox, in the Audits subfolder. This ensures that all audit log entries are available from a single location, regardless of which client access method was used to access the mailbox or which server or computer an administrator uses to access the mailbox audit log. If you move a mailbox to another Mailbox server, the mailbox audit logs for that mailbox are also moved because they're located in the mailbox.

By default, mailbox audit log entries are retained in the mailbox for 90 days and then deleted. You can modify this retention period by using the *AuditLogAgeLimit* parameter with the Set-Mailbox cmdlet. If a mailbox is on In-Place Hold or Litigation Hold, audit log entries are only retained until the audit log retention period for the mailbox is reached. To retain audit log entries longer, you have to increase the retention period by changing the value for the *AuditLogAgeLimit* parameter. You can also export audit log entries before the retention period is reached. For more information, see:

- Export Mailbox Audit Logs
- Create a Mailbox Audit Log Search

Enabling mailbox audit logging

Mailbox audit logging is enabled per mailbox. Use the **Set-Mailbox** cmdlet to enable or disable mailbox audit logging. For details, see Enable or disable mailbox audit logging for a mailbox.

When you enable mailbox audit logging for a mailbox, access to the mailbox and certain administrator and delegate actions are logged by default. To log actions taken by the mailbox owner, you must specify which owner actions should be audited.

Mailbox actions logged by mailbox audit logging

The following table lists the actions logged by mailbox audit logging, including the logon types for which the action can be logged. Note that an administrator who has been assigned the Full Access permission to a user's mailbox is considered a delegate user.

If you no longer require certain types of mailbox actions to be audited, you should modify the mailbox's audit logging configuration to disable those actions. Existing log entries aren't purged until the age limit for audit log entries is reached.

ACTION	DESCRIPTION	ADMIN	DELEGATE	OWNER
Copy	An item is copied to another folder.	Yes	No	No
Create	An item is created in the Calendar, Contacts, Notes, or Tasks folder in the mailbox; for example, a new meeting request is created. Note that message or folder creation isn't audited.	Yes[1]	Yes[1]	Yes
FolderBind	A mailbox folder is accessed.	Yes[1]	Yes[2]	No
HardDelete	An item is deleted permanently from the Recoverable Items folder.	Yes[1]	Yes[1]	Yes
MailboxLogin	The user signed in to their mailbox.	No	No	Yes[3]
MessageBind	An item is accessed in the reading pane or opened.	Yes	No	No
Move	An item is moved to another folder.	Yes[1]	Yes	Yes
MoveToDeletedItems	An item is moved to the Deleted Items folder.	Yes[1]	Yes	Yes
SendAs	A message is sent using Send As permissions.	Yes[1]	Yes[1]	No
SendOnBehalf	A message is sent using Send on Behalf permissions.	Yes[1]	Yes	No
SoftDelete	An item is deleted from the Deleted Items folder.	Yes[1]	Yes[1]	Yes
Update	An item's properties are updated.	Yes[1]	Yes[1]	Yes

[1] Audited by default if auditing is enabled for a mailbox.

[2] Entries for folder bind actions performed by delegates are consolidated. One log entry is generated for individual

folder access within a time span of 24 hours.

[3] Auditing for owner logins to a mailbox works only for POP3, IMAP4, or OAuth logins. It doesn't work for NTLM or Kerberos logins to the mailbox.

Searching the mailbox audit log

You can use the following methods to search mailbox audit log entries:

- **Synchronously search a single mailbox**: You can use the Search-MailboxAuditLog cmdlet to synchronously search mailbox audit log entries for a single mailbox. The cmdlet displays search results in the Exchange Management Shell window. For details, see Search Mailbox Audit Log for a Mailbox.
- **Asynchronously search one or more mailboxes**: You can create a mailbox audit log search to asynchronously search mailbox audit logs for one or more mailboxes, and then have the search results sent to a specified email address. The search results are sent as an XML attachment. To create the search, use the New-MailboxAuditLogSearch cmdlet. For details, see Create a Mailbox Audit Log Search.
- **Use auditing reports in the Exchange admin center (EAC)**: You can use the **Auditing** tab in the EAC to run a non-owner mailbox access report (contains entries for admin and delete actions) or export non-owner entries from the mailbox audit log. For details, see:
  - Run a non-owner mailbox access report
  - Export Mailbox Audit Logs

Mailbox audit log entries

The following table describes the fields logged in a mailbox audit log entry.

FIELD	POPULATED WITH
Operation	One of the following actions: Copy Create FolderBind HardDelete MailboxLogin MessageBind Move MoveToDeletedItems SendAs SendOnBehalf SoftDelete Update
OperationResult	One of the following results: Failed PartiallySucceeded Succeeded
LogonType	Logon type of the user who performed the operation. Logon types include: Owner Delegate Admin
DestFolderId	Destination folder GUID for move operations.

FIELD	POPULATED WITH
DestFolderPathName	Destination folder path for move operations.
FolderId	Folder GUID.
FolderPathName	Folder path.
ClientInfoString	Details that identify which client or Exchange component performed the operation.
ClientIPAddress	Client computer IP address.
ClientMachineName	Client computer name.
ClientProcessName	Name of the client application process.
ClientVersion	Client application version.
InternalLogonType	The type of internal user (a person in your organization) who performed the operation. The possible values for this field are the same ones as the **LogonType** field.
MailboxOwnerUPN	Mailbox owner user principal name (UPN).
MailboxOwnerSid	Mailbox owner security identifier (SID).
DestMailboxOwnerUPN	Destination mailbox owner UPN, logged for cross-mailbox operations.
DestMailboxOwnerSid	Destination mailbox owner SID, logged for cross-mailbox operations.
DestMailboxOwnerGuid	Destination mailbox owner GUID.
CrossMailboxOperation	Information about whether the operation logged is a cross-mailbox operation (for example, copying or moving messages between mailboxes).
LogonUserDisplayName	Display name of user who is logged on.
DelegateUserDisplayName	Delegate user display name.
LogonUserSid	SID of user who is logged on.
SourceItems	ItemID of mailbox items on which the logged action is performed (for example, move or delete). For operations performed on a number of items, this field is returned as a collection of items.
SourceFolders	Source folder GUID.
ItemId	Item ID.

FIELD	POPULATED WITH
ItemSubject	Item subject.
MailboxGuid	Mailbox GUID.
MailboxResolvedOwnerName	Mailbox user resolved name in the format *DOMAIN\SamAccountName*.
LastAccessed	Time when the operation was performed.
Identity	Audit log entry ID.

More information

- **Administrator access to mailboxes**: Mailboxes are considered to be accessed by an administrator only in the following scenarios:
    - In-Place eDiscovery is used to search a mailbox.
    - The New-MailboxExportRequest cmdlet is used to export a mailbox.
    - Microsoft Exchange Server MAPI Client and Collaboration Data Objects is used to access the mailbox.
- **Bypassing mailbox auditing logging**: Mailbox access by authorized automated processes such as accounts used by third-party tools or accounts used for lawful monitoring can create a large number of mailbox audit log entries and may not be of interest to your organization. You can configure such accounts to bypass mailbox audit logging. For details, see Bypass a User Account From Mailbox Audit Logging.
- **Logging mailbox owner actions**: For mailboxes such as the Discovery Search Mailbox, which may contain more sensitive information, consider enabling mailbox audit logging for mailbox owner actions such as message deletion.

Enable or disable mailbox audit logging for a mailbox

8/3/2020 • 4 minutes to read • Edit Online

With mailbox audit logging in Exchange Server, you can track logons to a mailbox as well as what actions are taken while the user is logged on. When you enable mailbox audit logging for a mailbox, some actions performed by administrators and delegates are logged by default. None of the actions performed by the mailbox owner are logged by default. To learn more about mailbox audit logging and what actions can be logged, see Mailbox audit logging in Exchange Server.

Caution

Auditing of mailbox owner actions can generate a large number of mailbox audit log entries and is therefore disabled by default. We recommend that you only enable auditing of specific owner actions needed to meet business or compliance requirements.

What do you need to know before you begin?

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Mailbox audit logging" entry in the Messaging policy and compliance permissions in Exchange Server topic.

- Entries in the mailbox audit log are retained for 90 days, by default. See the More information section change how long entries are retained.

- You can't use the Exchange admin center (EAC) to enable or disable mailbox audit logging. You have to use the Exchange Management Shell. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.

- An administrator who has been assigned the Full Access permission to a user's mailbox is considered a delegate user.

- Mailboxes are considered to be accessed by an administrator only in the following scenarios:

  - In-Place eDiscovery is used to search a mailbox.

  - The **New-MailboxExportRequest** cmdlet is used to export a mailbox.

  - Microsoft Exchange Server MAPI Editor is used to access the mailbox.

Enable or disable mailbox audit logging

You can use the Exchange Management Shell to enable or disable mailbox audit logging for a mailbox. This enables or disables logging of all operations specified for administrator, delegates, and the mailbox owner.

This example enables mailbox audit logging for Ben Smith's mailbox.

```
Set-Mailbox -Identity "Ben Smith" -AuditEnabled $true
```

This example enables mailbox audit logging for all user mailboxes in your organization.

```
Get-Mailbox -ResultSize Unlimited -Filter "RecipientTypeDetails -eq 'UserMailbox'" | Select PrimarySmtpAddress
| ForEach {Set-Mailbox -Identity $_.PrimarySmtpAddress -AuditEnabled $true}
```

This example disables mailbox audit logging for Ben Smith's mailbox.

```
Set-Mailbox -Identity "Ben Smith" -AuditEnabled $false
```

For detailed syntax and parameter information, see Set-Mailbox.

Configure mailbox audit logging settings for administrator, delegate, and owner access

When mailbox audit logging is enabled for a mailbox, only the administrator, delegate, and owner actions specified in the audit logging configuration for the mailbox are logged.

This example specifies that the `MessageBind` and `FolderBind` actions performed by administrators will be logged for Ben Smith's mailbox.

```
Set-Mailbox -Identity "Ben Smith" -AuditAdmin MessageBind,FolderBind -AuditEnabled $true
```

This example specifies that the `SendAs` or `SendOnBehalf` actions performed by delegate users will be logged for Ben Smith's mailbox.

```
Set-Mailbox -Identity "Ben Smith" -AuditDelegate SendAs,SendOnBehalf -AuditEnabled $true
```

This example specifies that the `HardDelete` action performed by the mailbox owner will be logged for Ben Smith's mailbox.

```
Set-Mailbox -Identity "Ben Smith" -AuditOwner HardDelete -AuditEnabled $true
```

For detailed syntax and parameter information, see Set-Mailbox.

How do you know this worked?

To verify that you have successfully enabled mailbox audit logging for a mailbox and specified the correct logging settings for administrator, delegate, or owner access, use the Get-Mailbox cmdlet to retrieve the mailbox audit logging settings for that mailbox.

This example retrieves Ben Smith's mailbox settings and pipes the specified audit settings, including the audit log age limit, to the **Format-List** cmdlet.

```
Get-Mailbox "Ben Smith" | Format-List Audit*
```

A value of `True` for the **AuditEnabled** property verifies that audit logging is enabled.

This example retrieves the auditing settings for all user mailboxes in your organization.

```
Get-Mailbox -ResultSize Unlimited -Filter "RecipientTypeDetails -eq 'UserMailbox'" | Format-List Name,Audit*
```

More information

The actions that are audited for each type of user may not all be displayed when you run the **Get-Mailbox** cmdlet. But you can run the following commands to display all the audited actions for a specific user logon type.

```
Get-Mailbox <identity of mailbox> | Select-Object -ExpandProperty AuditAdmin
```

```
Get-Mailbox <identity of mailbox> | Select-Object -ExpandProperty AuditDelegate
```

```
Get-Mailbox <identity of mailbox> | Select-Object -ExpandProperty AuditOwner
```

By default, entries in the mailbox audit log are kept for 90 days. When an entry is older than 90 days, it's deleted. You can use the **Set-Mailbox** cmdlet to change this setting so items are kept for a longer (or shorter) period of time.

This example increases the age limit for mailbox audit log entries in Pilar Pinilla's mailbox to 180 days.

```
Set-Mailbox -Identity "Pilar Pinilla" -AuditLogAgeLimit 180
```

This example decreases the age limit for mailbox audit log entries for all user mailboxes in your organization to 60 days.

```
Get-Mailbox -ResultSize Unlimited -Filter "RecipientTypeDetails -eq 'UserMailbox'" | Set-Mailbox -
AuditLogAgeLimit 60
```

Run a non-owner mailbox access report

8/3/2020 • 4 minutes to read • Edit Online

The **Non-Owner Mailbox Access Report** in the Exchange admin center (EAC) lists the mailboxes that have been accessed by someone other than the person who owns the mailbox. When a non-owner accesses a mailbox, Exchange logs information about this action. Exchange stores this mailbox audit log as an email message in a hidden folder in the audited mailbox. The report displays entries from this log as search results and includes any mailboxes accessed by a non-owner, who accessed each mailbox and when, the actions performed by non-owners, and whether or not the actions were successful.

Exchange logs specific actions by non-owners, which includes administrators and users who have been assigned permissions to a mailbox (who are called *delegated users*). You can also narrow the search to users inside or outside your organization. By default, Exchange retains entries in the mailbox audit log for 90 days.

You enable mailbox audit logging in the Exchange Management Shell.

What do you need to know before you begin?

- Estimated time to complete: 5 minutes.
- To open the EAC, see Exchange admin center in Exchange Server. To open the Exchange Management Shell, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Mailbox audit logging" entry in the Messaging policy and compliance permissions in Exchange Server topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Step 1: Use the Exchange Management Shell to enable mailbox audit logging

You have to enable mailbox audit logging for each mailbox that you want to include in a non-owner mailbox access report. If you don't enable mailbox audit logging, you won't get any results when you run a report.

To enable mailbox audit logging for a single mailbox, run the following command in the Exchange Management Shell:

```
Set-Mailbox <Identity> -AuditEnabled $true
```

For example, to enable mailbox auditing for a user named Florence Flipo, run the following command.

```
Set-Mailbox "Florence Flipo" -AuditEnabled $true
```

To enable mailbox auditing for all user mailboxes in your organization, run the following commands:

```
$UserMailboxes = Get-mailbox -Filter "RecipientTypeDetails -eq 'UserMailbox'"
```

```
$UserMailboxes | ForEach {Set-Mailbox $_.Identity -AuditEnabled $true}
```

How do you know this worked?

Run the following command to verify that you've successfully configured mailbox audit logging.

```
Get-Mailbox | Format-List Name,AuditEnabled
```

A value of `True` for the *AuditEnabled* property verifies that audit logging is enabled.

Step 2: Use the EAC to run a non-owner mailbox access report

1. In the EAC, navigate to **Compliance Management** > **Auditing**.

2. Click **Run a non-owner mailbox access report**.

   By default, Exchange runs the report for non-owner access to any mailboxes in the organization over the past two weeks. Audit logging was enabled for the mailboxes listed in the search results.

3. To view non-owner access for a specific mailbox, select the mailbox from the list of mailboxes. View the search results in the details pane.

Notes:

- Want to narrow the search results? Select the start date, end date, or both, and select specific mailboxes to search. Click **Search** to re-run the report.

- You can also specify that you want to search for the non-owner access type, also called the logon type. Here are your options:

  - **All non-owners**: Search for access by administrators and delegated users inside your organization. Also includes access user outside of your organization.

  - **External users**: Search for access by users outside of your organization.

  - **Administrators and delegated users**: Search for access by administrators and delegated users inside your organization.

  - **Administrators**: Search for access by administrators in your organization.

How do you know this worked?

To verify that you've successfully run a non-owner mailbox access report, check the search results pane. The results pane displays the mailboxes that you ran the report for, whether an individual user or a group of mailboxes. If there are no results for a specific mailbox, it's possible there was no non-owner access or there was no non-owner access in the specified date range. As we previously recommended, be sure to verify that you enabled audit logging for the mailboxes you want to search for access by non-owners.

What gets logged in the mailbox audit log?

When you run a non-owner mailbox access report, the EAC search results displays entries from the mailbox audit log. Each report entry contains this information:

- Who accessed the mailbox and when.
- The actions performed by the non-owner.
- The affected message and its folder location.
- Whether the action was successful.

The following table describes the types of actions logged, and whether these actions are logged by default for access by administrators and for access by delegated users. If you want to track actions that aren't logged by default, you have to use the Exchange Management Shell to enable logging of those actions.

ACTION	DESCRIPTION	ADMINISTRATORS	DELEGATED USERS
Update	A message was changed.	Yes	Yes
Copy	A message was copied to another folder.	No	No
Move	A message was moved to another folder.	Yes	No
Move To Deleted Items	A message was moved to the Deleted Items folder.	Yes	No
Soft-delete	A message was deleted from the Deleted Items folder.	Yes	Yes
Hard-delete	A message was purged from the Recoverable Items folder.	Yes	Yes
FolderBind	A mailbox folder was accessed.	Yes	No
Send as	A message was sent using SendAs permission. This means another user sent the message as though it came from the mailbox owner.	Yes	Yes
Send on behalf of	A message was sent using SendOnBehalf permission. This means another user sent the message on behalf of the mailbox owner. The message will indicate to the recipient who the message was sent on behalf of and who actually sent the message.	Yes	No
MessageBind	A message was viewed in the preview pane or opened.	No	No

Data loss prevention in Exchange Server

8/3/2020 • 6 minutes to read • Edit Online

Data loss prevention (DLP) is important in Exchange Server because business critical email communication often includes sensitive data. DLP features make managing sensitive data in email messages easier than ever before by balancing compliance requirements without unnecessarily hindering the productivity of workers. For a conceptual overview of DLP, watch the following video.

DLP policies are simple packages that are collections of mail flow rules (also known as transport rules) that contain specific conditions, actions, and exceptions that filter messages and attachments based on their content. You can create a DLP policy, yet choose to not activate it. This allows you to test your policies without affecting mail flow. For more information, see Test a mail flow rule.

DLP policies can use the full power of mail flow rules to detect and then act on messages in transit. For example, a mail flow rule can perform deep content analysis through keyword matches, dictionary matches, text pattern matches through regular expressions, and other content examination techniques to detect content that violates your organization's DLP policies. Document fingerprinting is also available to help you detect sensitive information in standard forms. For more information, see the following topics:

- Document fingerprinting
- Mail flow rules in Exchange Server
- Integrating classification rules with mail flow rules

In addition to the customizable DLP policies themselves, you can also inform email senders when they're about to violate one of your policies, even before they send a message that contains sensitive information. You do this by configuring Policy Tips. Policy Tips present a brief note about the possible policy violations in Outlook 2013 or later, Outlook on the web (formerly known as Outlook Web App), and Outlook on the web for devices. For more information, see Policy Tips.

Notes:

- DLP is a premium feature that requires an Exchange Enterprise Client Access License (CAL). For more information about CALs and server licensing, see Exchange licensing FAQs.
- In hybrid environments where some mailboxes are in on-premises Exchange and some are in Exchange Online, DLP policies are only applied in Exchange Online. Messages that are sent between on-premises users don't have DLP policies applied, because the messages don't leave the on-premises environment.

Looking for management tasks related to Data Loss Prevention? See DLP Procedures.

Establish policies to protect sensitive data

The data loss prevention features can help you identify and monitor many categories of sensitive information that you have defined within the conditions of your policies, such as private identification numbers or credit card numbers. You have the option of defining your own custom policies and mail flow rules, or you can use the DLP policy templates that are included in Exchange to get started quickly. A *policy template* is a model that includes a range of conditions, rules, and actions that you can choose from to create and save an actual DLP policy that will help you inspect messages. For more information about the included policy templates, see DLP Policy Templates Supplied in Exchange.

There are three different methods that you can use to implement DLP:

- **Apply an out-of-the-box template supplied in Exchange**: The quickest way to start using DLP policies is to create and implement a new policy by using a template. This saves you the effort of building a new set of rules from nothing. You need to know what type of data you want to check for or which compliance regulation you're attempting to address. You also need to know your organization's expectations for processing this data. For more information, see DLP Policy Templates Supplied in Exchange and Create a DLP Policy From a Template.

- **Import a pre-built policy file from outside your organization**: You can import policies that were created by independent software vendors. In this way, you can extend the DLP solution to meet your business requirements. For more information, see Define Your Own DLP Templates and Information Types and Import a DLP Policy From a File.

- **Create a custom policy without any pre-existing conditions**: Your enterprise may have its own requirements for monitoring certain types of data that's known to exist within a messaging system. You can create a custom policy entirely on your own to find and act on your own unique message data. You need to know the requirements and constraints of the environment where the DLP policy will be enforced to create effective custom policies. For more information, see Create a Custom DLP Policy.

After you add a policy, you can review and change its rules, deactivate the policy, or remove it completely. For more information, see Manage DLP Policies.

Sensitive information types in DLP policies

When you create or change DLP policies, you can include rules that look for sensitive information. The sensitive information types that are listed in the topic Sensitive information types in Exchange Server are available for you to use in your policies. You can customize the conditions within a policy, such as how many times something has to be found before an action is taken, or the action to take. For more information about creating DLP policies see, Create a Custom DLP Policy. For more information about mail flow rules, see Mail flow rules in Exchange Server.

To make it easy for you to use rules that look for sensitive information, Exchange comes with policy templates that already include some of the sensitive information types. You can't add conditions for all of the sensitive information types, because the templates are designed to help you focus on the most common types of compliance-related data within your organization. For more information about the pre-built templates, see DLP Policy Templates Supplied in Exchange.

You can create many DLP policies for your organization, and enable them all so that many different types of information are looked for. You can also create a DLP policy that isn't based on an existing template. To create such a policy, see Create a Custom DLP Policy. For more information about the available sensitive information types, see Sensitive information types in Exchange Server.

Detecting sensitive form data with Document Fingerprinting

Exchange lets you use Document Fingerprinting to easily create a sensitive information type that's based on a standard form.

Policy Tips notify users about sensitive content expectations

You can use Policy Tip notification messages to inform email senders about possible compliance issues while they are composing an email message. When you configure a Policy Tip in a DLP policy, the notification message will only show up if something in the sender's email message matches the conditions described in your policy. Policy Tips are similar to MailTips that were introduced in Exchange 2010. For more information, see Policy Tips.

Detecting sensitive information along with traditional message classification

A key factor in the strength of a DLP solution is the ability to correctly identify confidential or sensitive content that may be unique to your organization, regulatory needs, geography, or other business needs. The Exchange DLP architecture uses deep content analysis coupled with detection criteria that you establish through rules in your DLP policies. Helping to prevent data loss in Exchange requires you to configure the appropriate set of sensitive information rules that provide a high degree of protection while minimizing disruptions to mail flow that are caused by false positives and negatives. These types of rules (referred to throughout the DLP information as *sensitive information detection*) function within the framework of mail flow rules to enable DLP capabilities. To learn more about these features, see Integrating sensitive information rules with mail flow rules.

You can still apply traditional message classifications to messages, and you can combine these classifications with sensitive information detection. You can use these features together within a single DLP policy, or operate them independently (concurrently). To learn more about the traditional Exchange 2010 message classifications, see Understanding Message Classifications.

Information about DLP-processed messages

To see information about messages that contain DLP policy detections in your environment, see View DLP policy detection reports and Create incident reports for DLP policy detections. Data related to DLP detections is highly integrated in the delivery reports.

For more information

- Messaging policy and compliance in Exchange Server
- DLP Procedures
- View DLP policy detection reports)
- Document Fingerprinting

Sensitive information types in Exchange Server

8/3/2020 • 69 minutes to read • Edit Online

Data loss prevention (DLP) includes 80 sensitive information types that are ready for you to use in your DLP policies. This topic lists all of these sensitive information types and shows what a DLP policy looks for when it detects each type. A sensitive information type is defined by a pattern that can be identified by a regular expression or a function. In addition, corroborative evidence such as keywords and checksums can be used to identify a sensitive information type. Confidence level and proximity are also used in the evaluation process.

ABA Routing Number

Format: 9 digits which may be in a formatted or unformatted pattern.

Pattern:

Formatted:

- Four digits beginning with 0, 1, 2, 3, 6, 7, or 8
- A hyphen
- Four digits
- A hyphen
- A digit

Unformatted: 9 consecutive digits beginning with 0, 1, 2, 3, 6, 7, or 8

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_aba_routing` finds content that matches the pattern.
- A keyword from `Keyword_ABA_Routing` is found.

```
<!-- ABA Routing Number -->
<Entity id="cb353f78-2b72-4c3c-8827-92ebe4f69fdf" patternsProximity="300" recommendedConfidence="75">
      <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_aba_routing" />
        <Match idRef="Keyword_ABA_Routing" />
      </Pattern>
 </Entity>
```

Keywords:

KEYWORD_ABA_ROUTING
aba aba # aba routing # aba routing number aba# abarouting# aba number abaroutingnumber american bank association routing # american bank association routing number americanbankassociationrouting# americanbankassociationroutingnumber bank routing number bankrouting# bankroutingnumber routing transit number RTN

Argentina National Identity (DNI) Number

Format: Eight digits separated by periods

Pattern: Eight digits:

- Two digits
- A period
- Three digits
- A period
- Three digits

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_argentina_national_id` finds content that matches the pattern.
- A keyword from `Keyword_argentina_national_id` is found.

```
<!-- Argentina National Identity (DNI) Number -->
<Entity id="eefbb00e-8282-433c-8620-8f1da3bffdb2" recommendedConfidence="75" patternsProximity="300">
   <Pattern confidenceLevel="75">
      <IdMatch idRef="Regex_argentina_national_id"/>
      <Match idRef="Keyword_argentina_national_id"/>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_ARGENTINA_NATIONAL_ID
Argentina National Identity number Identity Identification National Identity Card DNI NIC National Registry of Persons Documento Nacional de Identidad Registro Nacional de las Personas Identidad Identificación

Australia Bank Account Number

Format: 6-10 digits with or without a bank state branch number

Pattern: Account number is 6-10 digits. Australia bank state branch number:

- Three digits
- A hyphen
- Three digits

Checksum: No

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_australia_bank_account_number` finds content that matches the pattern..
- A keyword from `Keyword_australia_bank_account_number` is found.
- The regular expression `Regex_australia_bank_account_number_bsb` finds content that matches the pattern.

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_australia_bank_account_number` finds content that matches the pattern..
- A keyword from `Keyword_australia_bank_account_number` is found.

```
<!-- Australia Bank Account Number -->
<Entity id="74a54de9-2a30-4aa0-a8aa-3d9327fc07c7" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
        <IdMatch idRef="Regex_australia_bank_account_number" />
        <Match idRef="Keyword_australia_bank_account_number" />
        <Match idRef="Regex_australia_bank_account_number_bsb" />
  </Pattern>
  <Pattern confidenceLevel="75">
        <IdMatch idRef="Regex_australia_bank_account_number" />
        <Match idRef="Keyword_australia_bank_account_number" />
  </Pattern>
 </Entity>
```

Keywords:

KEYWORD_AUSTRALIA_BANK_ACCOUNT_NUMBER
swift bank code correspondent bank base currency usa account holder address bank address information account fund transfers bank charges bank details banking information full names iaea

Australia Driver's License Number

Format: Nine letters and digits

Pattern: Nine letters and digits:

- Two digits or letters (not case sensitive)
- Two digits
- Five digits or letters (not case sensitive)

  OR

- 1-2 optional letters (not case sensitive)
- 4-9 digits

  OR

- Nine digits or letters (not case sensitive)

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_australia_drivers_license_number` finds content that matches the pattern.
- A keyword from `Keyword_australia_drivers_license_number` is found.
- No keyword from `Keyword_australia_drivers_license_number_exclusions` is found.

```
<!-- Australia Drivers License Number -->
<Entity id="1cbbc8f5-9216-4392-9eb5-5ac2298d1356" patternsProximity="300" recommendedConfidence="75">
   <Pattern confidenceLevel="75">
        <IdMatch idRef="Regex_australia_drivers_license_number" />
        <Match idRef="Keyword_australia_drivers_license_number" />
        <Any minMatches="0" maxMatches="0">
          <Match idRef="Keyword_australia_drivers_license_number_exclusions" />
        </Any>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_AUSTRALIA_DRIVERS_LICENSE_NUMBER	KEYWORD_AUSTRALIA_DRIVERS_LICENSE_NUMBER_EXCLUSIONS
international driving permits	aaa
australian automobile association	DriverLicense
sydney nsw	DriverLicenses
international driving permit	Driver License
DriverLicence	Driver Licenses
DriverLicences	DriversLicense
Driver Lic	DriversLicenses
Driver Licence	Drivers License
Driver Licences	Drivers Licenses
DriversLic	Driver'License
DriversLicence	Driver'Licenses
DriversLicences	Driver' License
Drivers Lic	Driver' Licenses
Drivers Lics	Driver'sLicense
Drivers Licence	Driver'sLicenses
Drivers Licences	Driver's License
Driver'Lic	Driver's Licenses
Driver'Lics	DriverLicense#
Driver'Licence	DriverLicenses#
Driver'Licences	Driver License#
Driver' Lic	Driver Licenses#
Driver' Lics	DriversLicense#
Driver' Licence	DriversLicenses#
Driver' Licences	Drivers License#
Driver'sLic	Drivers Licenses#
Driver'sLics	Driver'License#
Driver'sLicence	Driver'Licenses#
Driver'sLicences	Driver' License#
Driver's Lic	Driver' Licenses#
Driver's Lics	Driver'sLicense#
Driver's Licence	Driver'sLicenses#
Driver's Licences	Driver's License#
DriverLic#	Driver's Licenses#
DriverLics#	
DriverLicence#	
DriverLicences#	
Driver Lic#	
Driver Lics#	
Driver Licence#	
Driver Licences#	
DriversLic#	
DriversLics#	
DriversLicence#	
DriversLicences#	
Drivers Lic#	
Drivers Lics#	
Drivers Licence#	
Drivers Licences#	
Driver'Lic#	
Driver'Lics#	
Driver'Licence#	
Driver'Licences#	
Driver' Lic#	
Driver' Lics#	
Driver' Licence#	
Driver' Licences#	
Driver'sLic#	
Driver'sLics#	
Driver'sLicence#	
Driver'sLicences#	

KEYWORD_AUSTRALIA_DRIVERS_LICENSE_NUMBER	KEYWORD_AUSTRALIA_DRIVERS_LICENSE_NUMBER_EXCLUSIONS
Driver's Lic# Driver's Lics# Driver's Licence# Driver's Licences#	

Australia Medical Account Number

Format: 10-11 digits

Pattern: 10-11 digits:

- First digit is in the range 2-6
- Ninth digit is a check digit
- Tenth digit is the issue digit
- Eleventh digit (optional) is the individual number

Checksum: Yes

Definition:

A DLP policy is 95% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_australian_medical_account_number` finds content that matches the pattern.
- A keyword from `Keyword_Australia_Medical_Account_Number` is found.
- The checksum passes.

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_australian_medical_account_number` finds content that matches the pattern.
- The checksum passes.

```
<!-- Australia Medical Account Number -->
<Entity id="104a99a0-3d3b-4542-a40d-ab0b9e1efe63" recommendedConfidence="85" patternsProximity="300">
    <Pattern confidenceLevel="95">
     <IdMatch idRef="Func_australian_medical_account_number"/>
     <Any minMatches="1">
     <Match idRef="Keyword_Australia_Medical_Account_Number"/>
     </Any>
  </Pattern>
<Pattern confidenceLevel="85">
     <IdMatch idRef="Func_australian_medical_account_number"/>
     <Any minMatches="0" maxMatches="0">
  <Match idRef="Keyword_Australia_Medical_Account_Number"/>
     </Any>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_AUSTRALIA_MEDICAL_ACCOUNT_NUMBER
bank account details medicare payments mortgage account bank payments information branch credit card loan department of human services local service medicare

Australia Passport Number

Format: A letter followed by seven digits

Pattern: A letter (not case sensitive) followed by seven digits

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_australia_passport_number` finds content that matches the pattern.
- A keyword from `Keyword_passport` or `Keyword_australia_passport_number` is found.

```
<!-- Australia Passport Number -->
<Entity id="29869db6-602d-4853-ab93-3484f905df50" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
        <IdMatch idRef="Regex_australia_passport_number" />
        <Any minMatches="1">
          <Match idRef="Keyword_passport" />
          <Match idRef="Keyword_australia_passport_number" />
        </Any>
   </Pattern>
</Entity>
```

Keywords:

KEYWORD_PASSPORT	KEYWORD_AUSTRALIA_PASSPORT_NUMBER
Passport Number Passport No Passport # Passport# PassportID Passportno passportnumber パスポート パスポート番号 パスポートのNum パスポート＃ Numéro de passeport Passeport n ° Passeport Non Passeport # Passeport# PasseportNon Passeportn °	passport passport details immigration and citizenship commonwealth of australia department of immigration residential address department of immigration and citizenship visa national identity card passport number travel document issuing authority

Australia Tax File Number

Format: 8-9 digits

Pattern: 8-9 digits typically presented with spaces as follows:

- Three digits
- An optional space
- Three digits
- An optional space
- 2-3 digits where the last digit is a check digit

Checksum: Yes

Definition:

A DLP policy is 95% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_australian_tax_file_number` finds content that matches the pattern.
- A keyword from `Keyword_Australia_Tax_File_Number` is found.
- No keyword from `Keyword_number_exclusions` is found.
- The checksum passes.

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_australian_tax_file_number` finds content that matches the pattern.
- No keyword from `Keyword_Australia_Tax_File_Number` or `Keyword_number_exclusions` is found.
- The checksum passes.

```
<!-- Australia Tax File Number -->
<Entity id="e29bc95f-ff70-4a37-aa01-04d17360a4c5" patternsProximity="300" recommendedConfidence="85">
    <Pattern confidenceLevel="95">
        <IdMatch idRef="Func_australian_tax_file_number" />
        <Any minMatches="1">
          <Match idRef="Keyword_Australia_Tax_File_Number" />
        </Any>
        <Any minMatches="0" maxMatches="0">
          <Match idRef="Keyword_number_exclusions" />
        </Any>
  </Pattern>
  <Pattern confidenceLevel="85">
        <IdMatch idRef="Func_australian_tax_file_number" />
        <Any minMatches="0" maxMatches="0">
          <Match idRef="Keyword_Australia_Tax_File_Number" />
          <Match idRef="Keyword_number_exclusions" />
        </Any>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_AUSTRALIA_TAX_FILE_NUMBER	KEYWORD_NUMBER_EXCLUSIONS
australian business number marginal tax rate medicare levy portfolio number service veterans withholding tax individual tax return tax file number	00000000 11111111 22222222 33333333 44444444 55555555 66666666 77777777 88888888 99999999 000000000 111111111 222222222 333333333 444444444 555555555 666666666 777777777 888888888 999999999 0000000000 1111111111 2222222222 3333333333 4444444444 5555555555 6666666666 7777777777 8888888888 9999999999

Belgium National Number

Format: 11 digits plus delimiters

Pattern: 11 digits plus delimiters:

- Six digits and two periods in the format YY.MM.DD for date of birth

- A hyphen
- Three sequential digits (odd for males, even for females)
- A period
- Two digits that are a check digit

Checksum: Yes

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_belgium_national_number` finds content that matches the pattern.
- A keyword from `Keyword_belgium_national_number` is found.
- The checksum passes.

```
<!-- Belgium National Number -->
  <Entity id="fb969c9e-0fd1-4b18-8091-a2123c5e6a54" recommendedConfidence="75" patternsProximity="300">
   <Pattern confidenceLevel="75">
     <IdMatch idRef="Func_belgium_national_number"/>
     <Match idRef="Keyword_belgium_national_number"/>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_BELGIUM_NATIONAL_NUMBER
Identity Registration Identification ID Identiteitskaart Registratie nummer Identificatie nummer Identiteit Registratie Identificatie Carte d'identité numéro d'immatriculation numéro d'identification identité inscription Identifikation Identifizierung Identifikationsnummer Personalausweis Registrierung Registrationsnummer

Brazil Legal Entity Number (CNPJ)

Format: 14 digits that include a registration number, branch number, and check digits, plus delimiters

Pattern: 14 digits, plus delimiters:

- Two digits
- A period
- Three digits
- A period
- Three digits (these first eight digits are the registration number)
- A forward slash
- Four-digit branch number
- A hyphen
- Two digits which are check digits

Checksum: Yes

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_brazil_cnpj` finds content that matches the pattern.
- A keyword from `Keyword_brazil_cnpj` is found.
- The checksum passes.

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_brazil_cnpj` finds content that matches the pattern.
- The checksum passes.

```
<!-- Brazil Legal Entity Number (CNPJ) -->
<Entity id="9b58b5cd-5e90-4df6-b34f-1ebcc88ceae4" recommendedConfidence="85" patternsProximity="300">
   <Pattern confidenceLevel="85">
     <IdMatch idRef="Func_brazil_cnpj"/>
     <Match idRef="Keyword_brazil_cnpj"/>
  </Pattern>
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Func_brazil_cnpj"/>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_BRAZIL_CNPJ

CNPJ
CNPJ/MF
CNPJ-MF
National Registry of Legal Entities
Taxpayers Registry
Legal entity
Legal entities
Registration Status
Business
Company
CNPJ
Cadastro Nacional da Pessoa Jurídica
Cadastro Geral de Contribuintes
CGC
Pessoa jurídica
Pessoas jurídicas
Situação cadastral
Inscrição
Empresa

Brazil CPF Number

Format: 11 digits that include a check digit and can be formatted or unformatted

Pattern:

Formatted:

- Three digits
- A period
- Three digits
- A period
- Three digits
- A hyphen
- Two digits which are check digits

Unformatted: 11 digits where the last two digits are check digits

Checksum: Yes

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_brazil_cpf` finds content that matches the pattern.
- A keyword from `Keyword_brazil_cpf` is found.
- The checksum passes.

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_brazil_cpf` finds content that matches the pattern.

- The checksum passes.

```
<!-- Brazil CPF Number -->
<Entity id="78e09124-f2c3-4656-b32a-c1a132cd2711" recommendedConfidence="85" patternsProximity="300">
  <Pattern confidenceLevel="85">
     <IdMatch idRef="Func_brazil_cpf"/>
     <Match idRef="Keyword_brazil_cpf"/>
  </Pattern>
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Func_brazil_cpf"/>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_BRAZIL_CPF
CPF Identification Registration Revenue Cadastro de Pessoas Físicas Imposto Identificação Inscrição Receita

Brazil National ID Card (RG)

Format:

- Registro Geral (old format): Nine digits plus delimiters
- Registro de Identidade (RIC) (new format): 11 digits plus a hyphen

Pattern:

Registro Geral (old format):

- Two digits
- A period
- Three digits
- A period
- Three digits
- A hyphen
- One digit which is a check digit

Registro de Identidade (RIC) (new format)

- 10 digits
- A hyphen
- One digit which is a check digit

Checksum: Yes

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_brazil_rg` finds content that matches the pattern.
- A keyword from `Keyword_brazil_rg` is found.
- The checksum passes.

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_brazil_rg` finds content that matches the pattern.
- The checksum passes.

```
<!-- Brazil National ID Card (RG) -->
<Entity id="486de900-db70-41b3-a886-abdf25af119c" recommendedConfidence="85" patternsProximity="300">
  <Pattern confidenceLevel="85">
     <IdMatch idRef="Func_brazil_rg"/>
     <Match idRef="Keyword_brazil_rg"/>
  </Pattern>
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Func_brazil_rg"/>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_BRAZIL_RG
National ID Registration Cédula de identidade Registro Geral RG Registro de Identidade RIC Número de registo Registro

Canada Bank Account Number

Format: Seven or twelve digits

Pattern: A Canada Bank Account Number is seven or twelve digits. A Canada bank account transit number is:

- Five digits
- A hyphen
- Three digits

  OR

- A zero "0"
- Eight digits

Checksum: No

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_canada_bank_account_number` finds content that matches the pattern.
- A keyword from `Keyword_canada_bank_account_number` is found.
- The regular expression `Regex_canada_bank_account_transit_number` finds content that matches the pattern.

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_canada_bank_account_number` finds content that matches the pattern.
- A keyword from `Keyword_canada_bank_account_number` is found.

```
<!-- Canada Bank Account Number -->
<Entity id="552e814c-cb50-4d94-bbaa-bb1d1ffb34de" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
        <IdMatch idRef="Regex_canada_bank_account_number" />
        <Match idRef="Keyword_canada_bank_account_number" />
        <Match idRef="Regex_canada_bank_account_transit_number" />
   </Pattern>
   <Pattern confidenceLevel="75">
        <IdMatch idRef="Regex_canada_bank_account_number" />
        <Match idRef="Keyword_canada_bank_account_number" />
   </Pattern>
</Entity>
```

Keywords:

KEYWORD_CANADA_BANK_ACCOUNT_NUMBER
canada savings bonds canada revenue agency canadian financial institution direct deposit form canadian citizen legal representative notary public commissioner for oaths child care benefit universal child care canada child tax benefit income tax benefit harmonized sales tax social insurance number income tax refund child tax benefit territorial payments institution number deposit request banking information direct deposit

Canada Driver's License Number

Format: Varies by province

Pattern: Various patterns covering Alberta, British Columbia, Manitoba, New Brunswick, Newfoundland/Labrador, Nova Scotia, Ontario, Prince Edward Island, Quebec, and Saskatchewan

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_[province_name]_drivers_license_number` finds content that matches the pattern.
- A keyword from `Keyword_[province_name]_drivers_license_name` is found.
- A keyword from `Keyword_canada_drivers_license` is found.

```
<!-- Canada Driver's License Number -->
    <Entity id="37186abb-8e48-4800-ad3c-e3d1610b3db0" patternsProximity="300" recommendedConfidence="75">
      <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_alberta_drivers_license_number" />
        <Match idRef="Keyword_alberta_drivers_license_name" />
        <Match idRef="Keyword_canada_drivers_license" />
      </Pattern>
      <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_british_columbia_drivers_license_number" />
        <Match idRef="Keyword_british_columbia_drivers_license_name" />
        <Match idRef="Keyword_canada_drivers_license" />
      </Pattern>
      <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_manitoba_drivers_license_number" />
        <Match idRef="Keyword_manitoba_drivers_license_name" />
        <Match idRef="Keyword_canada_drivers_license" />
      </Pattern>
      <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_new_brunswick_drivers_license_number" />
        <Match idRef="Keyword_new_brunswick_drivers_license_name" />
        <Match idRef="Keyword_canada_drivers_license" />
      </Pattern>
      <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_newfoundland_labrador_drivers_license_number" />
        <Match idRef="Keyword_newfoundland_labrador_drivers_license_name" />
        <Match idRef="Keyword_canada_drivers_license" />
      </Pattern>
      <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_nova_scotia_drivers_license_number" />
        <Match idRef="Keyword_nova_scotia_drivers_license_name" />
        <Match idRef="Keyword_canada_drivers_license" />
      </Pattern>
      <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_ontario_drivers_license_number" />
        <Match idRef="Keyword_ontario_drivers_license_name" />
        <Match idRef="Keyword_canada_drivers_license" />
      </Pattern>
      <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_prince_edward_island_drivers_license_number" />
        <Match idRef="Keyword_prince_edward_island_drivers_license_name" />
        <Match idRef="Keyword_canada_drivers_license" />
      </Pattern>
      <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_quebec_drivers_license_number" />
        <Match idRef="Keyword_quebec_drivers_license_name" />
        <Match idRef="Keyword_canada_drivers_license" />
      </Pattern>
      <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_saskatchewan_drivers_license_number" />
        <Match idRef="Keyword_saskatchewan_drivers_license_name" />
        <Match idRef="Keyword_canada_drivers_license" />
      </Pattern>
    </Entity>
```

Keywords:

KEYWORD_[PROVINCE_NAME]_DRIVERS_LICENSE_NAME	KEYWORD_CANADA_DRIVERS_LICENSE
The province abbreviation, for example AB The province name, for example Alberta	DL DLS CDL CDLS DriverLic DriverLics DriverLicense DriverLicenses

KEYWORD_[PROVINCE_NAME]_DRIVERS_LICENSE_NAME	KEYWORD_CANADA_DRIVERS_LICENSE
	DriverLicenses DriverLicence DriverLicences Driver Lic Driver Lics Driver License Driver Licenses Driver Licence Driver Licences DriversLic DriversLics DriversLicence DriversLicences DriversLicense DriversLicenses Drivers Lic Drivers Lics Drivers License Drivers Licenses Drivers Licence Drivers Licences Driver'Lic Driver'Lics Driver'License Driver'Licenses Driver'Licence Driver'Licences Driver' Lic Driver' Lics Driver' License Driver' Licenses Driver' Licence Driver' Licences Driver'sLic Driver'sLics Driver'sLicense Driver'sLicenses Driver'sLicence Driver'sLicences Driver's Lic Driver's Lics Driver's License Driver's Licenses Driver's Licence Driver's Licences Permis de Conduire id ids idcard number idcard numbers idcard # idcard #s idcard card idcard cards idcard identification number identification numbers identification # identification #s identification card identification cards identification DL# DLS# CDL# CDLS#

KEYWORD_[PROVINCE_NAME]_DRIVERS_LICENSE_NAME	KEYWORD_CANADA_DRIVERS_LICENSE
	DriverLic# DriverLics# DriverLicense# DriverLicenses# DriverLicence# DriverLicences# Driver Lic# Driver Lics# Driver License# Driver Licenses# Driver License# Driver Licences# DriversLic# DriversLics# DriversLicense# DriversLicenses# DriversLicence# DriversLicences# Drivers Lic# Drivers Lics# Drivers License# Drivers Licenses# Drivers Licence# Drivers Licences# Driver'Lic# Driver'Lics# Driver'License# Driver'Licenses# Driver'Licence# Driver'Licences# Driver' Lic# Driver' Lics# Driver' License# Driver' Licenses# Driver' Licence# Driver' Licences# Driver'sLic# Driver'sLics# Driver'sLicense# Driver'sLicenses# Driver'sLicence# Driver'sLicences# Driver's Lic# Driver's Lics# Driver's License# Driver's Licenses# Driver's Licence# Driver's Licences# Permis de Conduire# id# ids# idcard card# idcard cards# idcard# identification card# identification cards# identification#

Canada Health Service Number

Format: 10 digits

Pattern: 10 digits

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_canada_health_service_number` finds content that matches the pattern.
- A keyword from `Keyword_canada_health_service_number` is found.

```
<!-- Canada Health Service Number -->
<Entity id="59c0bf39-7fab-482c-af25-00faa4384c94" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
        <IdMatch idRef="Regex_canada_health_service_number" />
        <Any minMatches="1">
          <Match idRef="Keyword_canada_health_service_number" />
        </Any>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_CANADA_HEALTH_SERVICE_NUMBER
personal health number patient information health services speciality services automobile accident patient hospital psychiatrist workers compensation disability

Canada Passport Number

Format: Two uppercase letters followed by six digits

Pattern: Two uppercase letters followed by six digits

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_canada_passport_number` finds content that matches the pattern.
- A keyword from `Keyword_canada_passport_number` or `Keyword_passport` is found.

```
<!-- Canada Passport Number -->
<Entity id="14d0db8b-498a-43ed-9fca-f6097ae687eb" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
        <IdMatch idRef="Regex_canada_passport_number" />
        <Any minMatches="1">
          <Match idRef="Keyword_canada_passport_number" />
          <Match idRef="Keyword_passport" />
        </Any>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_CANADA_PASSPORT_NUMBER	KEYWORD_PASSPORT
canadian citizenship canadian passport passport application passport photos certified translator canadian citizens processing times renewal application	Passport Number Passport No Passport # Passport# PassportID Passportno passportnumber パスポート パスポート番号 パスポートのNum パスポート＃ Numéro de passeport Passeport n ° Passeport Non Passeport # Passeport# PasseportNon Passeportn °

Canada Personal Health Identification Number (PHIN)

Format: Nine digits

Pattern: Nine digits

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_canada_phin` finds content that matches the pattern.
- At least two keywords from `Keyword_canada_phin` or `Keyword_canada_provinces` are found..

```
<!-- Canada PHIN -->
<Entity id="722e12ac-c89a-4ec8-a1b7-fea3469f89db" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
        <IdMatch idRef="Regex_canada_phin" />
        <Any minMatches="2">
          <Match idRef="Keyword_canada_phin" />
          <Match idRef="Keyword_canada_provinces" />
        </Any>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_CANADA_PHIN	KEYWORD_CANADA_PROVINCES
social insurance number health information act income tax information manitoba health health registration prescription purchases benefit eligibility personal health power of attorney registration number personal health number practitioner referral wellness professional patient referral health and wellness	Nunavut Quebec Northwest Territories Ontario British Columbia Alberta Saskatchewan Manitoba Yukon Newfoundland and Labrador New Brunswick Nova Scotia Prince Edward Island Canada

Canada Social Insurance Number

Format: Nine digits with optional hyphens or spaces

Pattern:

Formatted:

- Three digits
- A hyphen or space
- Three digits
- A hyphen or space
- Three digits

Unformatted: Nine digits

Checksum: Yes

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_canadian_sin` finds content that matches the pattern.
- At least two of any combination of the following:

  - A keyword from `Keyword_sin` is found.
  - A keyword from `Keyword_sin_collaborative` is found.
  - The function `Func_eu_date` finds a date in the right date format.
- The checksum passes.

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_unformatted_canadian_sin` finds content that matches the pattern.
- A keyword from `Keyword_sin` is found.
- The checksum passes.

```
<!-- Canada Social Insurance Number -->
<Entity id="a2f29c85-ecb8-4514-a610-364790c0773e" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
        <IdMatch idRef="Func_canadian_sin" />
        <Any minMatches="2">
          <Match idRef="Keyword_sin" />
          <Match idRef="Keyword_sin_collaborative" />
          <Match idRef="Func_eu_date" />
        </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_unformatted_canadian_sin" />
        <Match idRef="Keyword_sin" />
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_SIN	KEYWORD_SIN_COLLABORATIVE
sin social insurance numero d'assurance sociale sins ssn ssns social security numero d'assurance social national identification number national id sin# soc ins social ins	driver's license drivers license driver's licence drivers licence DOB Birthdate Birthday Date of Birth

Chile Identity Card Number

Format: 7-8 digits plus delimiters a check digit or letter

Pattern: 7-8 digits plus delimiters:

- 1-2 digits
- A period
- Three digits

- A period
- Three digits
- A dash
- One digit or letter (not case sensitive) which is a check digit

Checksum: Yes

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_chile_id_card` finds content that matches the pattern.
- A keyword from `Keyword_chile_id_card` is found.
- The checksum passes.

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_chile_id_card` finds content that matches the pattern.
- The checksum passes.

```
<!-- Chile Identity Card Number -->
<Entity id="4e979794-49a0-407e-a0b9-2c536937b925" recommendedConfidence="85" patternsProximity="300">
  <Pattern confidenceLevel="85">
     <IdMatch idRef="Func_chile_id_card"/>
     <Match idRef="Keyword_chile_id_card"/>
  </Pattern>
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Func_chile_id_card"/>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_CHILE_ID_CARD
National Identification Number Identity card ID Identification Rol Único Nacional RUN Rol Único Tributario RUT Cédula de Identidad Número De Identificación Nacional Tarjeta de identificación Identificación

China Resident Identity Card (PRC) Number

Format: 18 digits

Pattern: 18 digits:

- Six digits which are an address code
- Eight digits in the form YYYYMMDD which are the date of birth
- Three digits which are an order code
- One digit which is a check digit

Checksum: Yes

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_china_resident_id` finds content that matches the pattern.
- A keyword from `Keyword_china_resident_id` is found.
- The checksum passes.

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_china_resident_id` finds content that matches the pattern.
- The checksum passes.

```
<!-- China Resident Identity Card (PRC) Number -->
<Entity id="c92daa86-2d16-4871-901f-816b3f554fc1" recommendedConfidence="85" patternsProximity="300">
  <Pattern confidenceLevel="85">
     <IdMatch idRef="Func_china_resident_id"/>
     <Match idRef="Keyword_china_resident_id"/>
  </Pattern>
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Func_china_resident_id"/>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_CHINA_RESIDENT_ID
Resident Identity Card PRC National Identification Card 身份证 居民 身份证 居民身份证 鉴定 身分證 居民 身份證 鑑定

Credit Card Number

Format: 14 digits which can be formatted or unformatted (dddddddddddddd) and must pass the Luhn test.

Pattern: Very complex and robust pattern that detects cards from all major brands worldwide, including Visa, MasterCard, Discover Card, JCB, American Express, gift cards, and diner cards.

Checksum: Yes, the Luhn checksum

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_credit_card` finds content that matches the pattern.
- One of the following is true:
  - A keyword from `Keyword_cc_verification` is found.
  - A keyword from `Keyword_cc_name` is found.
  - The function `Func_expiration_date` finds a date in the right date format.
- The checksum passes.

A DLP policy is 65% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_credit_card` finds content that matches the pattern.
- The checksum passes.

```xml
<!-- Credit Card Number -->
<Entity id="50842eb7-edc8-4019-85dd-5a5c1f2bb085" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
        <IdMatch idRef="Func_credit_card" />
        <Any minMatches="1">
          <Match idRef="Keyword_cc_verification" />
          <Match idRef="Keyword_cc_name" />
          <Match idRef="Func_expiration_date" />
        </Any>
  </Pattern>
  <Pattern confidenceLevel="65">
        <IdMatch idRef="Func_credit_card" />
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_CC_VERIFICATION	KEYWORD_CC_NAME
card verification card identification number cvn cid cvc2 cvv2 pin block security code security number security no issue number issue no cryptogramme numéro de sécurité numero de securite kreditkartenprüfnummer kreditkartenprufnummer prüfziffer	amex american express americanexpress Visa mastercard master card mc mastercards master cards diner's Club diners club dinersclub discover card discovercard discover cards JCB japanese card bureau carte blanche

prufziffer
sicherheits-Kode
sicherheitscode
sicherheitsnummer
verfalldatum
codice di verifica
cod. sicurezza
cod sicurezza
n autorizzazione
código
codigo
cod. seg
cod seg
código de segurança
codigo de seguranca
codigo de segurança
código de seguranca
cód. segurança
cod. seguranca cod. segurança
cód. seguranca
cód segurança
cod seguranca cod segurança
cód seguranca
número de verificação
numero de verificacao
ablauf
gültig bis
gültigkeitsdatum
gultig bis
gultigkeitsdatum
scadenza
data scad
fecha de expiracion
fecha de venc
vencimiento
válido hasta
valido hasta
vto
data de expiração
data de expiracao
data em que expira
validade
valor
vencimento
Venc

carteblanche
credit card
cc#
cc#:
expiration date
exp date
expiry date
date d'expiration
date d'exp
date expiration
bank card
bankcard
card number
card num
cardnumber
cardnumbers
card numbers
creditcard
credit cards
creditcards
ccn
card holder
cardholder
card holders
cardholders
check card
checkcard
check cards
checkcards
debit card
debitcard
debit cards
debitcards
atm card
atmcard
atm cards
atmcards
enroute
en route
card type
carte bancaire
carte de crédit
carte de credit
numéro de carte
numero de carte
nº de la carte
nº de carte
kreditkarte
karte
karteninhaber
karteninhabers
kreditkarteninhaber
kreditkarteninstitut
kreditkartentyp
eigentümername
kartennr
kartennummer
kreditkartennummer
kreditkarten-nummer
carta di credito
carta credito
n. carta
n carta
nr. carta
nr carta
numero carta

KEYWORD_CC_VERIFICATION	KEYWORD_CC_NAME
	numero carta numero della carta numero di carta tarjeta credito tarjeta de credito tarjeta crédito tarjeta de crédito tarjeta de atm tarjeta atm tarjeta debito tarjeta de debito tarjeta débito tarjeta de débito nº de tarjeta no. de tarjeta no de tarjeta numero de tarjeta número de tarjeta tarjeta no tarjetahabiente cartão de crédito cartão de credito cartao de crédito cartao de credito cartão de débito cartao de débito cartão de debito cartao de debito débito automático debito automatico número do cartão numero do cartão número do cartao numero do cartao número de cartão numero de cartão número de cartao numero de cartao nº do cartão nº do cartao nº. do cartão no do cartão no do cartao no. do cartão no. do cartao

Croatia Identity Card Number

Format: Nine digits

Pattern: Nine consecutive digits

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_croatia_id_card` finds content that matches the pattern.
- A keyword from `Keyword_croatia_id_card` is found.

```
<!--Croatia Identity Card Number-->
<Entity id="ff12f884-c20a-4189-b185-34c8e7258d47" recommendedConfidence="75" patternsProximity="300">
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Func_croatia_id_card"/>
     <Match idRef="Keyword_croatia_id_card"/>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_CROATIA_ID_CARD
Croatian identity card Osobna iskaznica

Croatia Personal Identification (OIB) Number

Format: 10 digits

Pattern: 10 digits:

- Six digits in the form DDMMYY which are the date of birth
- Four digits where the final digit is a check digit

Checksum: Yes

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_croatia_oib_number` finds content that matches the pattern.
- A keyword from `Keyword_croatia_oib_number` is found.
- The checksum passes.

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_croatia_oib_number` finds content that matches the pattern.
- The checksum passes.

```
<!-- Croatia Personal Identification (OIB) Number -->
<Entity id="31983b6d-db95-4eb2-a630-b44bd091968d" recommendedConfidence="85" patternsProximity="300">
  <Pattern confidenceLevel="85">
     <IdMatch idRef="Func_croatia_oib_number"/>
     <Match idRef="Keyword_croatia_oib_number"/>
  </Pattern>
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Func_croatia_oib_number"/>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_CROATIA_OIB_NUMBER
Personal Identification Number Osobni identifikacijski broj OIB

Czech National Identity Card Number

Format: 10 digits containing a forward slash

Pattern: 10 digits:

- Six digits which are the date of birth
- A forward slash
- Four digits where the final digit is a check digit

Checksum: Yes

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_czech_id_card` finds content that matches the pattern.
- A keyword from `Keyword_czech_id_card` is found.
- The checksum passes.

```
<!-- Czech National Identity Card Number -->
<Entity id="60c0725a-4eb6-455b-9dda-05d8a7396497" recommendedConfidence="85" patternsProximity="300">
  <Pattern confidenceLevel="85">
     <IdMatch idRef="Func_czech_id_card"/>
     <Match idRef="Keyword_czech_id_card"/>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_CZECH_ID_CARD
Czech national identity card Občanský průka

Denmark Personal Identification Number

Format: 10 digits containing a hyphen

Pattern: 10 digits:

- Six digits in the format DDMMYY which are the date of birth
- A hyphen
- Four digits where the final digit is a check digit

Checksum: Yes

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_denmark_id` finds content that matches the pattern.
- A keyword from `Keyword_denmark_id` is found.
- The checksum passes.

```
<!-- Denmark Personal Identification Number -->
<Entity id="6c4f2fef-56e1-4c00-8093-88d7a01cf460" recommendedConfidence="75" patternsProximity="300">
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Regex_denmark_id"/>
     <Match idRef="Keyword_denmark_id"/>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_DENMARK_ID
Personal Identification Number CPR Det Centrale Personregister Personnummer

Drug Enforcement Agency (DEA) Number

Format: Two letters followed by seven digits

Pattern: Pattern must include all of the following:

- One letter (not case sensitive) from this set of possible letters: abcdefghjklmnprstux, which is a registrant code
- One letter (not case sensitive), which is the first letter of the registrant's last name
- Seven digits, the last of which is the check digit

Checksum: Yes

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_dea_number` finds content that matches the pattern.
- The checksum passes.

```
<!-- DEA Number -->
<Entity id="9a5445ad-406e-43eb-8bd7-cac17ab6d0e4" recommendedConfidence="85" patternsProximity="300">
  <Pattern confidenceLevel="85">
     <IdMatch idRef="Func_dea_number"/>
  </Pattern>
</Entity>
```

Keywords: None

EU Debit Card Number

Format: 16 digits

Pattern: Very complex and robust pattern

Checksum: Yes

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_eu_debit_card` finds content that matches the pattern.
- At least one of the following is true:
  - A keyword from `Keyword_eu_debit_card` is found.
  - A keyword from `Keyword_card_terms_dict` is found.
  - A keyword from `Keyword_card_security_terms_dict` is found.
  - A keyword from `Keyword_card_expiration_terms_dict` is found.
  - The function `Func_eu_date1` finds a date in the right date format.
  - The function `Func_eu_date2` finds a date in the right date format.
- The checksum passes.

```
<!-- EU Debit Card Number -->
<Entity id="0e9b3178-9678-47dd-a509-37222ca96b42" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
        <IdMatch idRef="Func_eu_debit_card" />
        <Any minMatches="1">
          <Match idRef="Keyword_eu_debit_card" />
          <Match idRef="Keyword_card_terms_dict" />
          <Match idRef="Keyword_card_security_terms_dict" />
          <Match idRef="Keyword_card_expiration_terms_dict" />
          <Match idRef="Func_expiration_date" />
          <Match idRef="Func_eu_date" />
          <Match idRef="Func_eu_date1" />
          <Match idRef="Func_eu_date2" />
        </Any>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_EU_DEBIT_CARD	KEYWORD_CARD_TERMS_DICT	KEYWORD_CARD_SECURITY_TERMS_DICT	KEYWORD_CARD_EXPIRATION_TERMS_DICT
account number	acct nbr	card identification number	ablauf
card number	acct num	card verification	data de expiracao
card no.	acct no	cardi la verifica	data de expiração
security number	american express	cid	data del exp
cc#	americanexpress	cod seg	data di exp
	americano espresso	cod seguranca	data di scadenza
	amex	cod segurança	data em que expira
	atm card	cod sicurezza	data scad
	atm cards	cod. seg	data scadenza
	atm kaart	cod. seguranca	date de validité
	atmcard	cod. seguranca	datum afloop

KEYWORD_EU_DEBIT_CARD	KEYWORD_CARD_TERMS_DICT	KEYWORD_CARD_SECURITY_TERMS_DICT	KEYWORD_CARD_EXPIRATION_TERMS_DICT
	atmcard atmcards atmkaart atmkaarten bancontact bank card bankkaart card holder card holders card num card number card numbers card type cardano numerico cardholder cardholders cardnumber cardnumbers carta bianca carta credito carta di credito cartao de credito cartao de crédito cartao de debito cartao de débito carte bancaire carte blanche carte bleue carte de credit carte de crédit carte di credito carteblanche cartão de credito cartão de crédito cartão de debito cartão de débito cb ccn check card check cards checkcard checkcards chequekaart cirrus cirrus-edc-maestro controlekaart controlekaarten credit card credit cards creditcard creditcards debetkaart debetkaarten debit card debit cards debitcard debitcards debito automatico diners club dinersclub discover discover card discover cards discovercard discovercards débito automático	cod. segurança cod sicurezza codice di sicurezza codice di verifica codigo codigo de seguranca codigo de segurança crittogramma cryptogram cryptogramme cv2 cvc cvc2 cvn cvv cvv2 cód seguranca cód segurança cód. seguranca cód. segurança código código de seguranca código de segurança de kaart controle geeft nr uit issue no issue number kaartidentificatienummer kreditkartenprufnummer kreditkartenprüfnummer kwestieaantal no. dell'edizione no. di sicurezza numero de securite numero de verificacao numero dell'edizione numero di identificazione della scheda numero di sicurezza numero van veiligheid numéro de sécurité nº autorizzazione número de verificação perno il blocco pin block prufziffer prüfziffer security code security no security number sicherheits kode sicherheitscode sicherheitsnummer speldblok veiligheid nr veiligheidsaantal veiligheidscode veiligheidsnummer verfalldatum	datum afloop datum van exp de afloop espira espira exp date exp datum expiration expire expires expiry fecha de expiracion fecha de venc gultig bis gultigkeitsdatum gültig bis gültigkeitsdatum la scadenza scadenza valable validade valido hasta valor venc vencimento vencimiento verloopt vervaldag vervaldatum vto válido hasta

KEYWORD_EU_DEBIT_CARD	KEYWORD_CARD_TERMS_DICT	KEYWORD_CARD_SECURITY_TERMS_DICT	KEYWORD_CARD_EXPIRATION_TERMS_DICT
	edc eigentümername european debit card hoofdkaart hoofdkaarten in viaggio japanese card bureau japanse kaartdienst jcb kaart kaart num kaartaantal kaartaantallen kaarthouder kaarthouders karte karteninhaber karteninhabers kartennr kartennummer kreditkarte kreditkarten-nummer kreditkarteninhaber kreditkarteninstitut kreditkartennummer kreditkartentyp maestro master card master cards mastercard mastercards mc mister cash n carta n. carta no de tarjeta no do cartao no do cartão no. de tarjeta no. do cartao no. do cartão nr carta nr. carta numeri di scheda numero carta numero de cartao numero de carte numero de cartão numero de tarjeta numero della carta numero di carta numero di scheda numero do cartao numero do cartão numéro de carte nº carta nº de carte nº de la carte nº de tarjeta nº do cartao nº do cartão nº. do cartão número de cartao número de cartão número de tarjeta número do cartao		

KEYWORD_EU_DEBIT_CARD	KEYWORD_CARD_TERMS_DICT	KEYWORD_CARD_SECURITY_TERMS_DICT	KEYWORD_CARD_EXPIRATION_TERMS_DICT
	numero do cartao scheda dell'assegno scheda dell'atmosfera scheda dell'atmosfera scheda della banca scheda di controllo scheda di debito scheda matrice schede dell'atmosfera schede di controllo schede di debito schede matrici scoprono la scheda scoprono le schede solo supporti di scheda supporto di scheda switch tarjeta atm tarjeta credito tarjeta de atm tarjeta de credito tarjeta de debito tarjeta debito tarjeta no tarjetahabiente tipo della scheda ufficio giapponese della scheda v pay v-pay visa visa plus visa electron visto visum vpay		

Finland National ID

Format: Six digits plus a character indicating a century plus three digits plus a check digit

Pattern: Pattern must include all of the following:

- Six digits in the format DDMMYY which are a date of birth
- Century marker (either '-', '+' or 'a')
- Three-digit personal identification number
- A digit or letter (case insensitive) which is a check digit

Checksum: Yes

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_finnish_national_id` finds content that matches the pattern.
- A keyword from `Keyword_finnish_national_id` is found.
- The checksum passes.

```
<!-- Finnish National ID-->
<Entity id="338FD995-4CB5-4F87-AD35-79BD1DD926C1" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
          <IdMatch idRef="Func_finnish_national_id" />
          <Match idRef="Keyword_finnish_national_id" />
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_FINNISH_NATIONAL_ID
Sosiaaliturvatunnus SOTU Henkilötunnus HETU Personbeteckning Personnummer

Finland Passport Number

Format: Combination of nine letters and digits

Pattern: Combination of nine letters and digits:

- Two letters (not case sensitive)
- Seven digits

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_finland_passport_number` finds content that matches the pattern.
- A keyword from `Keyword_finland_passport_number` is found.

```
<!-- Finland Passport Number -->
<Entity id="d1685ac3-1d3a-40f8-8198-32ef5669c7a5" recommendedConfidence="75" patternsProximity="300">
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Regex_finland_passport_number"/>
     <Match idRef="Keyword_finland_passport_number"/>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_FINLAND_PASSPORT_NUMBER
Passport Passi

France Driver's License Number

Format: 12 digits

Pattern: 12 digits with validation to discount similar patterns such as French telephone numbers

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters::

- The function `Func_french_drivers_license` finds content that matches the pattern.
- At least one of the following is true:
  - A keyword from `Keyword_french_drivers_license` is found.
  - The function `Func_eu_date` finds a date in the right date format.

```
<!-- France Driver's License Number -->
<Entity id="18e55a36-a01b-4b0f-943d-dc10282a1824" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_french_drivers_license" />
        <Any minMatches="1">
          <Match idRef="Keyword_french_drivers_license" />
          <Match idRef="Func_eu_date" />
        </Any>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_FRENCH_DRIVERS_LICENSE
drivers licence drivers license driving licence driving license permis de conduire licence number license number licence numbers license numbers

France National ID Card (CNI)

Format: 12 digits

Pattern: 12 digits

Checksum: No

Definition:

A DLP policy is 65% confident that it's detected this type of sensitive information if, within a proximity of 300 characters: The regular expression `Regex_france_cni` finds content that matches the pattern.

```
<!-- France CNI -->
<Entity id="f741ac74-1bc0-4665-b69b-f0c7f927c0c4" patternsProximity="300" recommendedConfidence="65">
  <Pattern confidenceLevel="65">
        <IdMatch idRef="Regex_france_cni" />
  </Pattern>
</Entity>
```

Keywords: None

France Passport Number

Format: Nine digits and letters

Pattern: Nine digits and letters:

- Two digits
- Two letters (not case sensitive)
- Five digits

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_fr_passport` finds content that matches the pattern.
- A keyword from `Keyword_passport` is found..

```
<!-- France Passport Number -->
<Entity id="3008b884-8c8c-4cd8-a289-99f34fc7ff5d" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_fr_passport" />
        <Match idRef="Keyword_passport" />
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_PASSPORT
Passport Number Passport No Passport # Passport# PassportID Passportno passportnumber パスポート パスポート番号 パスポートのNum パスポート # Numéro de passeport Passeport n ° Passeport Non Passeport # Passeport# PasseportNon Passeportn °

France Social Security Number (INSEE)

Format: 15 digits

Pattern:

Must match one of two patterns:

- 13 digits followed by a space followed by two digits, or
- 15 consecutive digits

Checksum: Yes

Definition:

A DLP policy is 95% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_french_insee` or `Func_fr_insee` finds content that matches the pattern.
- A keyword from `Keyword_fr_insee` is found.
- The checksum passes.

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_french_insee` or `Func_fr_insee` finds content that matches the pattern.
- No keyword from `Keyword_fr_insee` is found.
- The checksum passes.

```
<!-- France INSEE -->
<Entity id="71f62b97-efe0-4aa1-aa49-e14de253619d" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="95">
        <IdMatch idRef="Func_french_insee" />
        <Match idRef="Func_fr_insee" />
        <Any minMatches="1">
          <Match idRef="Keyword_fr_insee" />
        </Any>
  </Pattern>
  <Pattern confidenceLevel="85">
        <IdMatch idRef="Func_french_insee" />
        <Match idRef="Func_fr_insee" />
        <Any minMatches="0" maxMatches="0">
          <Match idRef="Keyword_fr_insee" />
        </Any>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_FR_INSEE
insee securité sociale securite sociale national id national identification numéro d'identité no d'identité no. d'identité numero d'identite no d'identite no. d'identite social security number social security code social insurance number le numéro d'identification nationale d'identité nationale numéro de sécurité sociale le code de la sécurité sociale numéro d'assurance sociale numéro de sécu code sécu

German Driver's License Number

Format: Combination of 11 digits and letters

Pattern: 11 digits and letters (not case sensitive):

- A digit or letter
- Two digits
- Six digits or letters
- A digit
- A digit or letter

Checksum: Yes

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_german_drivers_license` finds content that matches the pattern.
- At least one of the following is true:
  - A keyword from `Keyword_german_drivers_license_number` is found.
  - A keyword from `Keyword_german_drivers_license_collaborative` is found.
  - A keyword from `Keyword_german_drivers_license` is found.
- The checksum passes.

```
<!-- German Driver's License Number -->
<Entity id="91da9335-1edb-45b7-a95f-5fe41a16c63c" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_german_drivers_license" />
        <Any minMatches="1">
          <Match idRef="Keyword_german_drivers_license_number" />
          <Match idRef="Keyword_german_drivers_license_collaborative" />
          <Match idRef="Keyword_german_drivers_license" />
        </Any>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_GERMAN_DRIVERS_LICENSE_NUMBER	KEYWORD_GERMAN_DRIVERS_LICENSE_COLLABORATIVE	KEYWORD_GERMAN_DRIVERS_LICENSE
Führerschein Fuhrerschein Fuehrerschein Führerscheinnummer Fuhrerscheinnummer Fuehrerscheinnummer Führerschein- Fuhrerschein- Fuehrerschein- FührerscheinnummerNr FuhrerscheinnummerNr FuehrerscheinnummerNr FührerscheinnummerKlasse FuhrerscheinnummerKlasse FuehrerscheinnummerKlasse Führerschein- Nr Fuhrerschein- Nr Fuehrerschein- Nr Führerschein- Klasse Fuhrerschein- Klasse Fuehrerschein- Klasse FührerscheinnummerNr FuhrerscheinnummerNr FuehrerscheinnummerNr FührerscheinnummerKlasse FuhrerscheinnummerKlasse FuehrerscheinnummerKlasse Führerschein- Nr Fuhrerschein- Nr Fuehrerschein- Nr Führerschein- Klasse Fuhrerschein- Klasse Fuehrerschein- Klasse DL DLS Driv Lic Driv Licen Driv License Driv Licenses Driv Licence Driv Licences Driv Lic Driver Licen Driver License Driver Licenses Driver Licence Driver Licences	Nr-Führerschein Nr-Fuhrerschein Nr-Fuehrerschein No-Führerschein No-Fuhrerschein No-Fuehrerschein N-Führerschein N-Fuhrerschein N-Fuehrerschein Nr-Führerschein Nr-Fuhrerschein Nr-Fuehrerschein No-Führerschein No-Fuhrerschein No-Fuehrerschein N-Führerschein N-Fuhrerschein N-Fuehrerschein	ausstellungsdatum ausstellungsort ausstellende behöde ausstellende behorde ausstellende behoerde

KEYWORD_GERMAN_DRIVERS_LICENSE_NUMBER	KEYWORD_GERMAN_DRIVERS_LICENSE_COLLABORATIVE	KEYWORD_GERMAN_DRIVERS_LICENSE
Driver Licences Drivers Lic Drivers Licen Drivers License Drivers Licenses Drivers Licence Drivers Licences Driver's Lic Driver's Licen Driver's License Driver's Licenses Driver's Licence Driver's Licences Driving Lic Driving Licen Driving License Driving Licenses Driving Licence Driving Licences		

German Identity Card Number

Format:

- Since 1 November 2010: Nine letters and digits
- From 1 April 1987 until 31 October 2010: 10 digits

Pattern:

Since 1 November 2010:

- One letter (not case sensitive)
- Eight digits

From 1 April 1987 until 31 October 2010: 10 digits

Checksum: No

Definition:

A DLP policy is 65% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_germany_id_card` finds content that matches the pattern.
- A keyword from `Keyword_germany_id_card` is found.

```
<!-- Germany Identity Card Number -->
<Entity id="e577372f-c42e-47a0-9d85-bebed1c237d4" recommendedConfidence="65" patternsProximity="300">
  <Pattern confidenceLevel="65">
     <IdMatch idRef="Regex_germany_id_card"/>
     <Match idRef="Keyword_germany_id_card"/>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_GERMANY_ID_CARD
Identity Card ID Identification Personalausweis Identifizierungsnummer Ausweis Identifikation

German Passport Number

Format: 10 digits or letters

Pattern: Pattern must include all of the following:

- First character is a digit or a letter from this set (C, F, G, H, J, K)
- Three digits
- Five digits or letters from this set (C, -H, J-N, P, R, T, V-Z)
- A digit

Checksum: Yes

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_german_passport` finds content that matches the pattern.
- A keyword from any of the five keyword lists is found.
- The checksum passes.

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_german_passport_data` finds content that matches the pattern.
- A keyword from any of the five keyword lists is found.
- The checksum passes.

```
<!-- German Passport Number -->
<Entity id="2e3da144-d42b-47ed-b123-fbf78604e52c" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
        <IdMatch idRef="Func_german_passport" />
        <Any minMatches="1">
          <Match idRef="Keyword_german_passport" />
          <Match idRef="Keyword_german_passport_collaborative" />
          <Match idRef="Keyword_german_passport_number" />
          <Match idRef="Keyword_german_passport1" />
          <Match idRef="Keyword_german_passport2" />
        </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_german_passport_data" />
        <Any minMatches="1">
          <Match idRef="Keyword_german_passport" />
          <Match idRef="Keyword_german_passport_collaborative" />
          <Match idRef="Keyword_german_passport_number" />
          <Match idRef="Keyword_german_passport1" />
          <Match idRef="Keyword_german_passport2" />
        </Any>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_GERMAN_PASSPORT	KEYWORD_GERMAN_PASSPORT_COLLABORATIVE	KEYWORD_GERMAN_PASSPORT_NUMBER	KEYWORD_GERMAN_PASSPORT1	KEYWORD_GERMAN_PASSPORT2
reisepass reisepasse reisepassnummer passport passports	geburtsdatum ausstellungsdatum ausstellungsort	No-Reisepass Nr-Reisepass	Reisepass-Nr	bnationalit.t

Greece National ID Card

Format: Combination of 7-8 letters and numbers plus a dash

Pattern:

Seven letters and numbers (old format):

- One letter (any letter of the Greek alphabet)
- A dash
- Six digits

Eight letters and numbers (new format):

- Two letters whose uppercase character occurs in both the Greek and Latin alphabets (ABEZHIKMNOPTYX)
- A dash
- Six digits

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_greece_id_card` finds content that matches the pattern.
- A keyword from `Keyword_greece_id_card` is found.

```
<!-- Greece National ID Card -->
<Entity id="82568215-1da1-46d3-874a-d2294d81b5ac" recommendedConfidence="85" patternsProximity="300">
  <Pattern confidenceLevel="85">
     <IdMatch idRef="Regex_greece_id_card"/>
     <Match idRef="Keyword_greece_id_card"/>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_GREECE_ID_CARD
Greek identity Card Tautotita Δελτίο αστυνομικής ταυτότητας Ταυτότητα

Hong Kong Identity Card (HKID) Number

Format: Combination of 8-9 letters and numbers plus optional parentheses around the final character

Pattern: Combination of 8-9 letters:

- 1-2 letters (not case sensitive)
- Six digits
- The final character (any digit or the letter A), which is the check digit and is optionally enclosed in parentheses.

Checksum: Yes

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_hong_kong_id_card` finds content that matches the pattern.
- A keyword from `Keyword_hong_kong_id_card` is found.
- The checksum passes.

A DLP policy is 65% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_hong_kong_id_card` finds content that matches the pattern.
- The checksum passes.

```
<!-- Hong Kong Identity Card (HKID) number -->
<Entity id="e63c28a7-ad29-4c17-a41a-3d2a0b70fd9c" recommendedConfidence="75" patternsProximity="300">
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Func_hong_kong_id_card"/>
     <Match idRef="Keyword_hong_kong_id_card"/>
  </Pattern>
  <Pattern confidenceLevel="65">
     <IdMatch idRef="Func_hong_kong_id_card"/>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_HONG_KONG_ID_CARD
Hong Kong Identity Card HKID ID card 香港身份證 香港永久性居民身份證

India Permanent Account Number

Format: 10 letters or digits

Pattern: 10 letters or digits:

- Five letters (not case sensitive)
- Four digits
- A letter which is an alphabetic check digit

Checksum: Yes

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_india_permanent_account_number` finds content that matches the pattern.
- A keyword from `Keyword_india_permanent_account_number` is found.
- The checksum passes.

```
<!-- India Permanent Account Number -->
<Entity id="2602bfee-9bb0-47a5-a7a6-2bf3053e2804" recommendedConfidence="85" patternsProximity="300">
  <Pattern confidenceLevel="85">
     <IdMatch idRef="Regex_india_permanent_account_number"/>
     <Match idRef="Keyword_india_permanent_account_number"/>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_INDIA_PERMANENT_ACCOUNT_NUMBER
Permanent Account Number PAN

India Unique Identification (Aadhaar) Number

Format: 12 digits containing optional spaces or dashes

Pattern: 12 digits:

- Four digits
- An optional space or dash
- Four digits
- An optional space or dash
- The final digit which is the check digit

Checksum: Yes

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_india_aadhaar` finds content that matches the pattern.
- A keyword from `Keyword_india_aadhar` is found.
- The checksum passes.

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_india_aadhaar` finds content that matches the pattern.
- The checksum passes.

```
<!-- India Unique Identification (Aadhaar) number -->
<Entity id="1ca46b29-76f5-4f46-9383-cfa15e91048f" recommendedConfidence="85" patternsProximity="300">
  <Pattern confidenceLevel="85">
     <IdMatch idRef="Func_india_aadhaar"/>
     <Match idRef="Keyword_india_aadhar"/>
  </Pattern>
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Func_india_aadhaar"/>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_INDIA_AADHAR
Aadhar Aadhaar UID ■■■■

Indonesia Identity Card (KTP) Number

Format: 16 digits containing optional periods

Pattern: 16 digits:

- Two-digit province code
- A period (optional)
- Two-digit regency or city code
- Two-digit subdistrict code
- A period (optional)
- Six digits in the format DDMMYY which are the date of birth
- A period (optional)
- Four digits

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_indonesia_id_card` finds content that matches the pattern.
- A keyword from `Keyword_indonesia_id_card` is found.

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters: The regular expression `Regex_indonesia_id_card` finds content that matches the pattern.

```
<!-- Indonesia Identity Card (KTP) Number -->
<Entity id="da68fdb0-f383-4981-8c86-82689d3b7d55" recommendedConfidence="85" patternsProximity="300">
  <Pattern confidenceLevel="85">
     <IdMatch idRef="Regex_indonesia_id_card"/>
     <Match idRef="Keyword_indonesia_id_card"/>
  </Pattern>
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Regex_indonesia_id_card"/>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_INDONESIA_ID_CARD
KTP Kartu Tanda Penduduk Nomor Induk Kependudukan

International Banking Account Number (IBAN)

Format: Country code (two letters) plus check digits (two digits) plus bban number (up to 30 characters)

Pattern:

Pattern must include all of the following:

- Two-letter country code
- Two check digits (followed by an optional space)
- 1-7 groups of four letters or digits (can be separated by spaces)
- 1-3 letters or digits

The format for each country is slightly different. The IBAN sensitive information type covers these 60 countries: ad, ae, al, at, az, ba, be, bg, bh, ch, cr, cy, cz, de, dk, do, ee, es, fi, fo, fr, gb, ge, gi, gl, gr, hr, hu, ie, il, is, it, kw, kz, lb, li, lt, lu, lv, mc, md, me, mk, mr, mt, mu, nl, no, pl, pt, ro, rs, sa, se, si, sk, sm, tn, tr, vg

Checksum: Yes

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_iban` finds content that matches the pattern.
- The checksum passes.

```
<Entity id="e7dc4711-11b7-4cb0-b88b-2c394a771f0e" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
        <IdMatch idRef="Func_iban" />
  </Pattern>
</Entity>
```

Keywords: None

IP Address

Format: IPv4 or IPv6 address

Pattern:

- IPv4: Complex pattern which accounts for formatted (periods) and unformatted (no periods) versions of the IPv4 addresses.
- IPv6: Complex pattern which accounts for formatted IPv6 numbers (which include colons).

Checksum: No

Definition:

For IPv4, a DLP policy is 95% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_ipv4_address` finds content that matches the pattern.
- A keyword from `Keyword_ipaddress` is found.

For IPv6, a DLP policy is 95% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_ipv6_address` finds content that matches the pattern.
- No keyword from `Keyword_ipaddress` is found.

For IPv4, a DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_ipv4_address` finds content that matches the pattern.
- No keyword from `Keyword_ipaddress` is found.

For IPv6, a DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_ipv6_address` finds content that matches the pattern.
- No keyword from `Keyword_ipaddress` is found.

```xml
<Entity id="1daa4ad5-e2dd-4ca4-a788-54722c09efb2" patternsProximity="300" recommendedConfidence="85">
    <Pattern confidenceLevel="95">
        <IdMatch idRef="Regex_ipv4_address" />
        <Any minMatches="1">
          <Match idRef="Keyword_ipaddress" />
        </Any>
    </Pattern>
    <Pattern confidenceLevel="95">
        <IdMatch idRef="Regex_ipv6_address" />
        <Any minMatches="1">
          <Match idRef="Keyword_ipaddress" />
        </Any>
    </Pattern>
    <Pattern confidenceLevel="85">
        <IdMatch idRef="Regex_ipv4_address" />
        <Any minMatches="0" maxMatches="0">
          <Match idRef="Keyword_ipaddress" />
        </Any>
    </Pattern>
    <Pattern confidenceLevel="85">
        <IdMatch idRef="Regex_ipv6_address" />
        <Any minMatches="0" maxMatches="0">
          <Match idRef="Keyword_ipaddress" />
        </Any>
    </Pattern>
</Entity>
```

Keywords:

KEYWORD_IPADDRESS
ip address internet protocol כתובת ה-IP

Ireland Personal Public Service (PPS) Number

Format:

- New format (1 Jan 2013 and later): Seven digits followed by two letters
- Old format (31 Dec 2012 and earlier): Seven digits followed by 1-2 letters

Pattern:

New format (1 Jan 2013 and later)

- Seven digits

- A letter (not case sensitive) which is an alphabetic check digit
- The letter "A" or "H" (not case sensitive)

Old format (31 Dec 2012 and earlier)

- Seven digits
- 1-2 letters (not case sensitive)

Checksum: Yes

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_ireland_pps` finds content that matches the pattern.
- One of the following is true:
    - A keyword from `Keyword_ireland_pps` is found.
    - The function `Func_eu_date` finds a date in the right date format.
- The checksum passes.

A DLP policy is 65% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_ireland_pps` finds content that matches the pattern.
- The checksum passes.

```
<!-- Ireland Personal Public Service (PPS) Number -->
<Entity id="1cdb674d-c19a-4fcf-9f4b-7f56cc87345a" recommendedConfidence="85" patternsProximity="300">
  <Pattern confidenceLevel="85">
     <IdMatch idRef="Func_ireland_pps"/>
     <Any minMatches="1">
  <Match idRef="Keyword_ireland_pps"/>
  <Match idRef="Func_eu_date"/>
     </Any>
  </Pattern>
  <Pattern confidenceLevel="65">
     <IdMatch idRef="Func_ireland_pps"/>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_IRELAND_PPS
Personal Public Service Number PPS Number PPS Num PPS No. PPS # PPS# PPSN Public Services Card Uimhir Phearsanta Seirbhíse Poiblí Uimh. PSP PSP

Israel Bank Account Number

Format: 13 digits

Pattern:

Formatted:

- Two digits
- A dash
- Three digits
- A dash
- Eight digits

Unformatted: 13 consecutive digits

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_israel_bank_account_number` finds content that matches the pattern.
- A keyword from `Keyword_israel_bank_account_number` is found.

```
<!-- Israel Bank Account Number -->
<Entity id="7d08b2ff-a0b9-437f-957c-aeddbf9b2b25" patternsProximity="300" recommendedConfidence="75">
    <Pattern confidenceLevel="75">
        <IdMatch idRef="Regex_israel_bank_account_number" />
        <Any minMatches="1">
          <Match idRef="Keyword_israel_bank_account_number" />
        </Any>
    </Pattern>
</Entity>
```

Keywords:

KEYWORD_ISRAEL_BANK_ACCOUNT_NUMBER
Bank Account Number Bank Account Account Number מספר חשבון בנק

Israel National ID

Format: Nine digits

Pattern: Nine consecutive digits

Checksum: Yes

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_israeli_national_id_number` finds content that matches the pattern.
- A keyword from `Keyword_Israel_National_ID` is found.
- The checksum passes.

```
<!-- Israel National ID Number -->
<Entity id="e05881f5-1db1-418c-89aa-a3ac5c5277ee" patternsProximity="300" recommendedConfidence="75">
    <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_israeli_national_id_number" />
        <Any minMatches="1">
          <Match idRef="Keyword_Israel_National_ID" />
        </Any>
    </Pattern>
</Entity>
```

Keywords:

KEYWORD_ISRAEL_NATIONAL_ID
מספר זהות National ID Number

Italy Driver's License Number

Format: A combination of 10 letters and digits

Pattern: A combination of 10 letters and digits:

- One letter (not case sensitive)
- The letter "A" or "V" (not case sensitive)
- Seven letters (not case sensitive), digits, or the underscore character
- One letter (not case sensitive)

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_italy_drivers_license_number` finds content that matches the pattern.
- A keyword from `Keyword_italy_drivers_license_number` is found.

```
<!-- Italy Driver's license Number -->
<Entity id="97d6244f-9157-41bd-8e0c-9d669a5c4d71" patternsProximity="300" recommendedConfidence="75">
    <Pattern confidenceLevel="75">
        <IdMatch idRef="Regex_italy_drivers_license_number" />
        <Any minMatches="1">
          <Match idRef="Keyword_italy_drivers_license_number" />
        </Any>
    </Pattern>
</Entity>
```

Keywords:

KEYWORD_ITALY_DRIVERS_LICENSE_NUMBER
numero di patente di guida patente di guida

Japan Bank Account Number

Format: Seven or eight digits

Pattern:

Bank account number: Seven or eight digits

Bank account branch code:

- Four digits
- A space or dash (optional)
- Three digits

Checksum: No

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_jp_bank_account` finds content that matches the pattern.
- A keyword from `Keyword_jp_bank_account` is found.
- One of the following is true:
  - The function `Func_jp_bank_account_branch_code` finds content that matches the pattern.
  - A keyword from `Keyword_jp_bank_branch_code` is found.

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_jp_bank_account` finds content that matches the pattern.
- A keyword from `Keyword_jp_bank_account` is found.

```
<!-- Japan Bank Account Number -->
<Entity id="d354f95b-96ee-4b80-80bc-4377312b55bc" patternsProximity="300" recommendedConfidence="75">
  <Version minEngineVersion="15.01.0131.000">
    <Pattern confidenceLevel="85">
          <IdMatch idRef="Func_jp_bank_account" />
          <Match idRef="Keyword_jp_bank_account" />
          <Any minMatches="1">
            <Match idRef="Func_jp_bank_account_branch_code" />
            <Match idRef="Keyword_jp_bank_branch_code" />
          </Any>
      </Pattern>
  </Version>
     <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_jp_bank_account" />
        <Match idRef="Keyword_jp_bank_account" />
    </Pattern>
</Entity>
```

Keywords:

KEYWORD_JP_BANK_ACCOUNT	KEYWORD_JP_BANK_BRANCH_CODE
Checking Account Number Checking Account Checking Account # Checking Acct Number Checking Acct # Checking Acct No. Checking Account No. Bank Account Number Bank Account Bank Account # Bank Acct Number Bank Acct # Bank Acct No. Bank Account No. Savings Account Number Savings Account Savings Account # Savings Acct Number Savings Acct # Savings Acct No. Savings Account No. Debit Account Number Debit Account Debit Account # Debit Acct Number Debit Acct # Debit Acct No. Debit Account No. 口座番号を当座預金口座の確認 ＃アカウントの確認、勘定番号の確認 ＃勘定の確認 勘定番号の確認 口座番号の確認 銀行口座番号 銀行口座 銀行口座＃ 銀行の勘定番号 銀行のacct＃ 銀行の勘定いいえ 銀行口座番号 普通預金口座番号 預金口座 貯蓄口座＃ 貯蓄勘定の数 貯蓄勘定＃ 貯蓄勘定番号 普通預金口座番号 引き落とし口座番号 口座番号 口座番号＃ デビットのacct番号 デビット勘定＃ デビットACCTの番号 デビット口座番号	Otemachi

Japan Driver's License Number

Format: 12 digits

Pattern: 12 consecutive digits

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_jp_drivers_license_number` finds content that matches the pattern.
- A keyword from `Keyword_jp_drivers_license_number` is found.

```
<!-- Japan Driver's License Number -->
<Entity id="c6011143-d087-451c-8313-7f6d4aed2270" patternsProximity="300" recommendedConfidence="75">
    <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_jp_drivers_license_number" />
        <Match idRef ="Keyword_jp_drivers_license_number" />
    </Pattern>
</Entity>
```

Keywords:

KEYWORD_JP_DRIVERS_LICENSE_NUMBER
driver license drivers license driver's license drivers licenses driver's licenses driver licenses dl# dls# lic# lics# 運転免許証 運転免許 免許証 免許 運転免許証番号 運転免許番号 免許証番号 免許番号 運転免許証ナンバー 運転免許ナンバー 免許証ナンバー 運転免許証No. 運転免許No. 免許証No. 免許No. 運転免許証# 運転免許# 免許証# 免許#

Japan Passport Number

Format: Two letters followed by seven digits

Pattern: Two letters (not case sensitive) followed by seven digits

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_jp_passport` finds content that matches the pattern.
- A keyword from `Keyword_jp_passport` is found.

```
<!-- Japan Passport Number -->
<Entity id="75177310-1a09-4613-bf6d-833aae3743f8" patternsProximity="300" recommendedConfidence="75">
    <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_jp_passport" />
        <Match idRef="Keyword_jp_passport" />
    </Pattern>
</Entity>
```

Keywords:

KEYWORD_JP_PASSPORT
パスポート パスポート番号 パスポートのNum パスポート＃

Japan Resident Registration Number

Format: 11 digits

Pattern: 11 consecutive digits

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_jp_resident_registration_number` finds content that matches the pattern.
- A keyword from `Keyword_jp_resident_registration_number` is found.

```
<!-- Japan Resident Registration Number -->
<Entity id="01c1209b-6389-4faf-a5f8-3f7e13899652" patternsProximity="300" recommendedConfidence="75">
    <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_jp_resident_registration_number" />
        <Match idRef ="Keyword_jp_resident_registration_number" />
    </Pattern>
</Entity>
```

Keywords:

KEYWORD_JP_RESIDENT_REGISTRATION_NUMBER
Resident Registration Number Resident Register Number Residents Basic Registry Number Resident Registration No. Resident Register No. Residents Basic Registry No. Basic Resident Register No. 住民登録番号、登録番号をレジデント 住民基本登録番号、登録番号 住民基本レジストリ番号を常駐 登録番号を常駐住民基本台帳登録番号

Japan Social Insurance Number (SIN)

Format: 7-12 digits

Pattern: 7-12 digits:

- Four digits
- A hyphen (optional)
- Six digits

  OR

- 7-12 consecutive digits

Checksum: No

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_jp_sin` finds content that matches the pattern.
- A keyword from `Keyword_jp_sin` is found.

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_jp_sin_pre_1997` finds content that matches the pattern.
- A keyword from `Keyword_jp_sin` is found.

```
<!-- Japan Social Insurance Number -->
<Entity id="c840e719-0896-45bb-84fd-1ed5c95e45ff" patternsProximity="300" recommendedConfidence="75">
    <Pattern confidenceLevel="85">
        <IdMatch idRef="Func_jp_sin" />
        <Match idRef="Keyword_jp_sin" />
    </Pattern>
    <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_jp_sin_pre_1997" />
        <Match idRef="Keyword_jp_sin" />
    </Pattern>
</Entity>
```

Keywords:

KEYWORD_JP_SIN
Social Insurance No. Social Insurance Num Social Insurance Number 社会保険のテンキー 社会保険番号

Malaysia ID Card Number

Format: 12 digits containing optional hyphens

Pattern: 12 digits:

- Six digits in the format YYMMDD which are the date of birth
- A dash (optional)
- Two-letter place-of-birth code
- A dash (optional)
- Three random digits
- One-digit gender code

Checksum: No

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_malaysia_id_card_number` finds content that matches the pattern.
- A keyword from `Keyword_malaysia_id_card_number` is found.

```
<!-- Malaysia ID Card Number -->
</Entity>
      <Entity id="7f0e921c-9677-435b-aba2-bb8f1013c749" patternsProximity="300" recommendedConfidence="85">
        <Pattern confidenceLevel="85">
            <IdMatch idRef="Regex_malaysia_id_card_number" />
            <Match idRef="Keyword_malaysia_id_card_number" />
        </Pattern>
</Entity>
```

Keywords:

KEYWORD_MALAYSIA_ID_CARD_NUMBER
MyKad Identity Card ID Card Identification Card Digital Application Card Kad Akuan Diri Kad Aplikasi Digital

Netherlands Citizen's Service (BSN) Number

Format: 8-9 digits containing optional spaces

Pattern: 8-9 digits:

- Three digits
- A space (optional)
- Three digits
- A space (optional)
- 2-3 digits

Checksum: Yes

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_netherlands_bsn` finds content that matches the pattern.
- A keyword from `Keyword_netherlands_bsn` is found.
- The function `Func_eu_date` finds a date in the right date format.
- The checksum passes.

A DLP policy is 65% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_netherlands_bsn` finds content that matches the pattern.
- The checksum passes.

```
<!-- Netherlands Citizen's Service (BSN) Number -->
<Entity id="c5f54253-ef7e-44f6-a578-440ed67e946d" recommendedConfidence="85" patternsProximity="300">
  <Pattern confidenceLevel="85">
     <IdMatch idRef="Func_netherlands_bsn"/>
     <Match idRef="Keyword_netherlands_bsn"/>
     <Match idRef="Func_eu_date"/>
  </Pattern>
  <Pattern confidenceLevel="65">
     <IdMatch idRef="Func_netherlands_bsn"/>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_NETHERLANDS_BSN
Citizen service number BSN Burgerservicenummer Sofinummer Persoonsgebonden nummer Persoonsnummer

New Zealand Ministry of Health Number

Format: Three letters, a space (optional), and four digits

Pattern: Three letters (not case sensitive) a space (optional) four digits

Checksum: Yes

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_new_zealand_ministry_of_health_number` finds content that matches the pattern.
- A keyword from `Keyword_nz_terms` is found.
- The checksum passes.

```
<!-- New Zealand Health Number -->
<Entity id="2b71c1c8-d14e-4430-82dc-fd1ed6bf05c7" patternsProximity="300" recommendedConfidence="85">
    <Pattern confidenceLevel="85">
        <IdMatch idRef="Func_new_zealand_ministry_of_health_number" />
        <Any minMatches="1">
          <Match idRef="Keyword_nz_terms" />
        </Any>
    </Pattern>
</Entity>
```

Keywords:

KEYWORD_NZ_TERMS
NHI New Zealand Health treatment

Norway Identification Number

Format: 11 digits

Pattern: 11 digits:

- Six digits in the format DDMMYY which are the date of birth
- Three-digit individual number
- Two check digits

Checksum: Yes

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_norway_id_number` finds content that matches the pattern.
- A keyword from `Keyword_norway_id_number` is found.
- The checksum passes.

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_norway_id_numbe` finds content that matches the pattern.
- The checksum passes.

```
<!-- Norway Identification Number -->
<Entity id="d4c8a798-e9f2-4bd3-9652-500d24080fc3" recommendedConfidence="85" patternsProximity="300">
  <Pattern confidenceLevel="85">
     <IdMatch idRef="Func_norway_id_number"/>
     <Match idRef="Keyword_norway_id_number"/>
  </Pattern>
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Func_norway_id_number"/>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_NORWAY_ID_NUMBER
Personal identification number Norwegian ID Number ID Number Identification Personnummer Fødselsnummer

Philippines Unified Multi-Purpose ID Number

Format: 12 digits separated by hyphens

Pattern: 12 digits:

- Four digits
- A hyphen
- Seven digits
- A hyphen
- One digit

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_philippines_unified_id` finds content that matches the pattern.
- A keyword from `Keyword_philippines_id` is found.

```
<!-- Philippines Unified Multi-Purpose ID number -->
<Entity id="019b39dd-8c25-4765-91a3-d9c6baf3c3b3" recommendedConfidence="75" patternsProximity="300">
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Regex_philippines_unified_id"/>
     <Match idRef="Keyword_philippines_id"/>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_PHILIPPINES_ID
Unified Multi-Purpose ID UMID Identity Card Pinag-isang Multi-Layunin ID

Poland Identity Card

Format: Three letters and six digits

Pattern: Three letters (not case sensitive) followed by six digits

Checksum: Yes

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_polish_national_id` finds content that matches the pattern.
- A keyword from `Keyword_polish_national_id_passport_number` is found.
- The checksum passes.

```
<!-- Poland Identity Card-->
<Entity id="25E64989-ED5D-40CA-A939-6C14183BB7BF" patternsProximity="300" recommendedConfidence="85">
      <Pattern confidenceLevel="85">
          <IdMatch idRef="Func_polish_national_id" />
          <Match idRef="Keyword_polish_national_id_passport_number" />
      </Pattern>
</Entity>
```

Keywords:

KEYWORD_POLISH_NATIONAL_ID_PASSPORT_NUMBER
Nazwa i nr dowodu tożsamości Dowód Tożsamości dow. os.

Poland National ID (PESEL)

Format: 11 digits

Pattern: 11 consecutive digits

Checksum: Yes

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_pesel_identification_number` finds content that matches the pattern.
- A keyword from `Keyword_pesel_identification_number` is found.
- The checksum passes.

```
<!-- Poland National ID (PESEL) -->
<Entity id="E3AAF206-4297-412F-9E06-BA8487E22456" patternsProximity="300" recommendedConfidence="85">
      <Pattern confidenceLevel="85">
          <IdMatch idRef="Func_pesel_identification_number" />
          <Match idRef="Keyword_pesel_identification_number" />
      </Pattern>
</Entity>
```

Keywords:

KEYWORD_PESEL_IDENTIFICATION_NUMBER
Nr PESEL PESEL

Poland Passport

Format: Two letters and seven digits

Pattern: Two letters (not case sensitive) followed by seven digits

Checksum: Yes

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_polish_passport_number` finds content that matches the pattern.
- A keyword from `Keyword_polish_national_id_passport_number` is found.
- The checksum passes.

```
<!-- Poland Passport Number -->
<Entity id="03937FB5-D2B6-4487-B61F-0F8BFF7C3517" patternsProximity="300" recommendedConfidence="85">
      <Pattern confidenceLevel="85">
          <IdMatch idRef="Func_polish_passport_number" />
          <Match idRef="Keyword_polish_national_id_passport_number" />
      </Pattern>
</Entity>
</Version>
```

Keywords:

KEYWORD_POLISH_NATIONAL_ID_PASSPORT_NUMBER
Nazwa i nr dowodu tożsamości Dowód Tożsamości dow. os.

Portugal Citizen Card Number

Format: Eight digits

Pattern: Eight digits

Checksum: No

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_portugal_citizen_card` finds content that matches the pattern.
- A keyword from `Keyword_portugal_citizen_card` is found.

```
<!-- Portugal Citizen Card Number -->
<Entity id="91a7ece2-add4-4986-9a15-c84544d81ecd" recommendedConfidence="85" patternsProximity="300">
  <Pattern confidenceLevel="85">
     <IdMatch idRef="Regex_portugal_citizen_card"/>
     <Match idRef="Keyword_portugal_citizen_card"/>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_PORTUGAL_CITIZEN_CARD
Citizen Card National ID Card CC Cartão de Cidadão Bilhete de Identidade

Saudi Arabia National ID

Format: 10 digits

Pattern: 10 consecutive digits

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_saudi_arabia_national_id` finds content that matches the pattern.
- A keyword from `Keyword_saudi_arabia_national_id` is found.

```
<!-- Saudi Arabia National ID -->
<Entity id="8c5a0ba8-404a-41a3-8871-746aa21ee6c0" patternsProximity="300" recommendedConfidence="75">
    <Pattern confidenceLevel="75">
        <IdMatch idRef="Regex_saudi_arabia_national_id" />
        <Any minMatches="1">
          <Match idRef="Keyword_saudi_arabia_national_id" />
        </Any>
    </Pattern>
</Entity>
```

Keywords:

KEYWORD_SAUDI_ARABIA_NATIONAL_ID
Identification Card I card number ID number الوطنية الهوية بطاقة رقم

Singapore National Registration Identity Card (NRIC) Number

Format: Nine letters and digits

Pattern: Nine letters and digits:

- The letter "F", "G", "S", or "T" (not case sensitive)
- Seven digits
- An alphabetic check digit

Checksum: Yes

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_singapore_nric` finds content that matches the pattern.
- A keyword from `Keyword_singapore_nric` is found.
- The checksum passes.

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_singapore_nric` finds content that matches the pattern.
- The checksum passes.

```
<!-- Singapore National Registration Identity Card (NRIC) Number -->
<Entity id="cead390a-dd83-4856-9751-fb6dc98c34da" recommendedConfidence="75" patternsProximity="300">
  <Pattern confidenceLevel="85">
     <IdMatch idRef="Regex_singapore_nric"/>
     <Match idRef="Keyword_singapore_nric"/>
  </Pattern>
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Regex_singapore_nric"/>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_SINGAPORE_NRIC
National Registration Identity Card Identity Card Number NRIC IC Foreign Identification Number FIN 身份证 身份證

South Africa Identification Number

Format: 13 digits that may contain spaces

Pattern: 13 digits:

- Six digits in the format YYMMDD which are the date of birth
- Four digits
- A single-digit citizenship indicator
- The digit "8" or "9"
- One digit which is a checksum digit

Checksum: Yes

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_south_africa_identification_number` finds content that matches the pattern.
- A keyword from `Keyword_south_africa_identification_number` is found.
- The checksum passes.

```
<!-- South Africa Identification Number -->
<Entity id="e2adf7cb-8ea6-4048-a2ed-d89eb65f2780" recommendedConfidence="85" patternsProximity="300">
  <Pattern confidenceLevel="85">
     <IdMatch idRef="Func_south_africa_identification_number"/>
     <Match idRef="Keyword_south_africa_identification_number"/>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_SOUTH_AFRICA_IDENTIFICATION_NUMBER
Identity card ID Identification

South Korea Resident Registration Number

Format: 13 digits containing a hyphen

Pattern: 13 digits:

- Six digits in the format YYMMDD which are the date of birth
- A hyphen
- One digit determined by the century and gender
- Four-digit region-of-birth code
- One digit used to differentiate people for whom the preceding numbers are identical
- A check digit.

Checksum: Yes

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_south_korea_resident_number` finds content that matches the pattern.
- A keyword from `Keyword_south_korea_resident_number` is found.
- The checksum passes.

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_south_korea_resident_number` finds content that matches the pattern.
- The checksum passes.

```
<!-- South Korea Resident Registration Number -->
<Entity id="5b802e18-ba80-44c4-bc83-bf2ad36ae36a" recommendedConfidence="85" patternsProximity="300">
  <Pattern confidenceLevel="85">
     <IdMatch idRef="Func_south_korea_resident_number"/>
     <Match idRef="Keyword_south_korea_resident_number"/>
  </Pattern>
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Func_south_korea_resident_number"/>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_SOUTH_KOREA_RESIDENT_NUMBER
National ID card Citizen's Registration Number Jumin deungnok beonho RRN 주민등록번호

Spain Social Security Number (SSN)

Format: 11-12 digits

Pattern: 11-12 digits:

- Two digits
- A forward slash (optional)
- 7-8 digits
- A forward slash (optional)
- Two digits

Checksum: Yes

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_spanish_social_security_number` finds content that matches the pattern.
- The checksum passes.

```
<!-- Spain SSN -->
<Entity id="5df987c0-8eae-4bce-ace7-b316347f3070" patternsProximity="300" recommendedConfidence="85">
    <Pattern confidenceLevel="85">
        <IdMatch idRef="Func_spanish_social_security_number" />
    </Pattern>
</Entity>
```

Keywords: None

Sweden National ID

Format: 10 or 12 digits and an optional delimiter

Pattern: 10 or 12 digits and an optional delimiter:

- 2-4 digits (optional)
- Six digits in date format YYMMDD
- Delimiter of "-" or "+" (optional), plus
- Four digits

Checksum: Yes

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_swedish_national_identifier` finds content that matches the pattern.
- The checksum passes.

```
<!-- Sweden National ID -->
<Entity id="f69aaf40-79be-4fac-8f05-fd1910d272c8" patternsProximity="300" recommendedConfidence="85">
    <Pattern confidenceLevel="85">
        <IdMatch idRef="Func_swedish_national_identifier" />
    </Pattern>
</Entity>
```

Keywords: None

Sweden Passport Number

Format: Eight digits

Pattern: Eight consecutive digits

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_sweden_passport_number` finds content that matches the pattern.
- One of the following is true:
  - A keyword from `Keyword_passport` is found.
  - A keyword from `Keyword_sweden_passport` is found.

```
<!-- Sweden Passport Number -->
<Entity id="ba4e7456-55a9-4d89-9140-c33673553526" patternsProximity="300" recommendedConfidence="75">
    <Pattern confidenceLevel="75">
        <IdMatch idRef="Regex_sweden_passport_number" />
        <Any minMatches="1">
          <Match idRef="Keyword_passport" />
          <Match idRef="Keyword_sweden_passport" />
        </Any>
    </Pattern>
</Entity>
```

Keywords:

KEYWORD_SWEDEN_PASSPORT	KEYWORD_PASSPORT
visa requirements Alien Registration Card Schengen visas Schengen visa Visa Processing Visa Type Single Entry Multiple Entry G3 Processing Fees	Passport Number Passport No Passport # Passport# PassportID Passportno passportnumber パスポート パスポート番号 パスポートのNum パスポート＃ Numéro de passeport Passeport n ° Passeport Non Passeport # Passeport# PasseportNon Passeportn °

SWIFT Code

Format: Four letters followed by 5-31 letters or digits

Pattern: Four letters followed by 5-31 letters or digits:

- Four-letter bank code (not case sensitive)
- An optional space
- 4-28 letters or digits (the Basic Bank Account Number (BBAN))
- An optional space
- 1-3 letters or digits (remainder of the BBAN)

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_swift` finds content that matches the pattern.
- A keyword from `Keyword_swift` is found.

```
<Entity id="cb2ab58c-9cb8-4c81-baf8-a4e106791df4" patternsProximity="300" recommendedConfidence="75">
<Pattern confidenceLevel="75">
        <IdMatch idRef="Regex_swift" />
        <Match idRef="Keyword_swift" />
    </Pattern>
</Entity>
```

Keywords:

KEYWORD_SWIFT
international organization for standardization 9362 iso 9362 iso9362 swift# swiftcode swiftnumber swiftroutingnumber swift code swift number # swift routing number bic number bic code bic # bic# bank identifier code 標準化9362 迅速＃ SWIFTコード SWIFT番号 迅速なルーティング番号 BIC番号 BICコード 銀行識別コードのための国際組織 Organisation internationale de normalisation 9362 rapide # code SWIFT le numéro de swift swift numéro d'acheminement le numéro BIC # BIC code identificateur de banque

Taiwan National ID

Format: One letter (in English) followed by nine digits

Pattern: One letter (in English) followed by nine digits:

- One letter (in English, not case sensitive)
- The digit "1" or "2"
- Eight digits

Checksum: Yes

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_taiwanese_national_id` finds content that matches the pattern.
- A keyword from `Keyword_taiwanese_national_id` is found.
- The checksum passes.

```
<!-- Taiwanese National ID -->
<Entity id="4C7BFC34-8DD1-421D-8FB7-6C6182C2AF03" patternsProximity="300" recommendedConfidence="85">
      <Pattern confidenceLevel="85">
          <IdMatch idRef="Func_taiwanese_national_id" />
          <Match idRef="Keyword_taiwanese_national_id" />
      </Pattern>
</Entity>
```

Keywords:

KEYWORD_TAIWANESE_NATIONAL_ID
身份證字號 身份證 身份證號碼 身份證號 身分證字號 身分證 身分證號碼 身份證號 身分證統一編號 國民身分證統一編號 簽名 蓋章 簽名或蓋章 簽章

Taiwan Passport Number

Format:

- Biometric passport number: Nine digits
- Non-biometric passport number: Nine digits

Pattern:

- Biometric passport number
  - The digit "3"
  - Eight digits
- Non-biometric passport number: Nine digits

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_taiwan_passport` finds content that matches the pattern.
- A keyword from `Keyword_taiwan_passport` is found.

```
<!-- Taiwan Passport Number -->
<Entity id="e7251cb4-4c2c-41df-963e-924eb3dae04a" recommendedConfidence="75" patternsProximity="300">
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Regex_taiwan_passport"/>
     <Match idRef="Keyword_taiwan_passport"/>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_TAIWAN_PASSPORT
ROC passport number Passport number Passport no Passport Num Passport # 护照 中華民國護照 Zhōnghuá Mínguó hùzhào

Taiwan Resident Certificate (ARC/TARC) Number

Format: 10 letters and digits

Pattern: 10 letters and digits:

- Two letters (not case sensitive)
- Eight digits

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_taiwan_resident_certificate` finds content that matches the pattern.
- A keyword from `Keyword_taiwan_resident_certificate` is found.

```
<!-- Taiwan Resident Certificate (ARC/TARC) -->
<Entity id="48269fec-05ea-46ea-b326-f5623a58c6e9" recommendedConfidence="75" patternsProximity="300">
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Regex_taiwan_resident_certificate"/>
     <Match idRef="Keyword_taiwan_resident_certificate"/>
  </Pattern>
</Entity>
```

Keywords:

KEYWORD_TAIWAN_RESIDENT_CERTIFICATE
Resident Certificate Resident Cert Resident Cert. Identification card Alien Resident Certificate ARC Taiwan Area Resident Certificate TARC 居留證 外僑居留證 台灣地區居留證

U.K. Driver's License Number

Format: Combination of 18 letters and digits in the specified format

Pattern: 18 letters and digits:

- Five letters (not case sensitive) or the digit "9" in place of a letter
- One digit
- Five digits in the date format DDMMY for date of birth
- Two letters (not case sensitive) or the digit "9" in place of a letter
- Five digits

Checksum: Yes

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_uk_drivers_license` finds content that matches the pattern.
- A keyword from `Keyword_uk_drivers_license` is found.
- The checksum passes.

```
<!-- U.K. Driver's License Number -->
<Entity id="f93de4be-d94c-40df-a8be-461738047551" patternsProximity="300" recommendedConfidence="75">
    <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_uk_drivers_license" />
        <Match idRef="Keyword_uk_drivers_license" />
    </Pattern>
</Entity>
```

Keywords:

KEYWORD_UK_DRIVERS_LICENSE

DVLA
light vans
quadbikes
motor cars
125cc
sidecar
tricycles
motorcycles
photocard licence
learner drivers
licence holder
licence holders
driving licences
driving licence
dual control car

U.K. Electoral Roll Number

Format: Two letters followed by 1-4 digits

Pattern: Two letters (not case sensitive) followed by 1-4 numbers

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_uk_electoral` finds content that matches the pattern.
- A keyword from `Keyword_uk_electoral` is found.

```
<!-- U.K. Electoral Number -->
<Entity id="a3eea206-dc0c-4f06-9e22-aa1be3059963" patternsProximity="300" recommendedConfidence="75">
    <Pattern confidenceLevel="75">
        <IdMatch idRef="Regex_uk_electoral" />
        <Any minMatches="1">
          <Match idRef="Keyword_uk_electoral" />
        </Any>
    </Pattern>
</Entity>
```

Keywords:

KEYWORD_UK_ELECTORAL

council nomination
nomination form
electoral register
electoral roll

U.K. National Health Service Number

Format: 10-17 digits separated by spaces

Pattern: 10-17 digits:

- Either 3 or 10 digits
- A space
- Three digits
- A space
- Four digits

Checksum: Yes

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_uk_nhs_number` finds content that matches the pattern.
- One of the following is true:
  - A keyword from `Keyword_uk_nhs_number` is found.
  - A keyword from `Keyword_uk_nhs_number1` is found.
  - A keyword from `Keyword_uk_nhs_number_dob` is found.
- The checksum passes.

```
<!-- U.K. NHS Number -->
<Entity id="3192014e-2a16-44e9-aa69-4b20375c9a78" patternsProximity="300" recommendedConfidence="85">
    <Pattern confidenceLevel="85">
        <IdMatch idRef="Func_uk_nhs_number" />
        <Any minMatches="1">
          <Match idRef="Keyword_uk_nhs_number" />
          <Match idRef="Keyword_uk_nhs_number1" />
          <Match idRef="Keyword_uk_nhs_number_dob" />
        </Any>
    </Pattern>
</Entity>
```

Keywords:

KEYWORD_UK_NHS_NUMBER	KEYWORD_UK_NHS_NUMBER1	KEYWORD_UK_NHS_NUMBER_DOB
national health service nhs health services authority health authority	patient id patient identification patient no patient number	GP DOB D.O.B Date of Birth Birth Date

U.K. National Insurance Number (NINO)

Format: Nine letters and digits, with each pair of letters and digits optionally separated by spaces or dashes

Pattern: Nine letters and digits, with each pair of letters and digits optionally separated by spaces or dashes:

- Two letters (not case sensitive), neither of which can be D, F, I, Q, U, or V. Additionally, the second letter can't be O. The following combinations are also not allowed: BG, GB, KN, NK, NT, TN, and ZZ.
- Six digits

- A space or dash (optional)
- Two digits
- A space or dash (optional)
- Two digits
- A space or dash (optional)
- Two digits
- One letter that can be A, B, C, D; or one space.

Checksum: No

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_uk_nino` finds content that matches the pattern.
- A keyword from `Keyword_uk_nino` is found.

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_uk_nino` finds content that matches the pattern.
- No keyword from `Keyword_uk_nino` is found.

```
<!-- U.K. NINO -->
<Entity id="16c07343-c26f-49d2-a987-3daf717e94cc" patternsProximity="300" recommendedConfidence="75">
    <Pattern confidenceLevel="85">
        <IdMatch idRef="Func_uk_nino" />
        <Any minMatches="1">
          <Match idRef="Keyword_uk_nino" />
        </Any>
    </Pattern>
     <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_uk_nino" />
        <Any minMatches="0" maxMatches="0">
          <Match idRef="Keyword_uk_nino" />
        </Any>
    </Pattern>
</Entity>
```

Keywords:

KEYWORD_UK_NINO
national insurance number national insurance contributions protection act insurance social security number insurance application medical application social insurance medical attention social security great britain insurance

U.S. / U.K. Passport Number

Format: Nine digits

Pattern: Nine consecutive digits

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_usa_uk_passport` finds content that matches the pattern.
- A keyword from `Keyword_passport` is found.

```
<Entity id="178ec42a-18b4-47cc-85c7-d62c92fd67f8" patternsProximity="300" recommendedConfidence="75">
    <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_usa_uk_passport" />
        <Match idRef="Keyword_passport" />
    </Pattern>
</Entity>
```

Keywords:

KEYWORD_PASSPORT
Passport Number Passport No Passport # Passport# PassportID Passportno passportnumber パスポート パスポート番号 パスポートのNum パスポート＃ Numéro de passeport Passeport n ° Passeport Non Passeport # Passeport# PasseportNon Passeportn °

U.S. Bank Account Number

Format: 4-17 digits

Pattern: 4-17 consecutive digits

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_usa_bank_account_number` finds content that matches the pattern.
- A keyword from `Keyword_usa_Bank_Account` is found.

```
<!-- U.S. Bank Account Number -->
<Entity id="a2ce32a8-f935-4bb6-8e96-2a5157672e2c" patternsProximity="300" recommendedConfidence="75">
    <Pattern confidenceLevel="75">
        <IdMatch idRef="Regex_usa_bank_account_number" />
        <Match idRef="Keyword_usa_Bank_Account" />
    </Pattern>
</Entity>
```

Keywords:

KEYWORD_USA_BANK_ACCOUNT
Checking Account Number Checking Account Checking Account # Checking Acct Number Checking Acct # Checking Acct No. Checking Account No. Bank Account Number Bank Account # Bank Acct Number Bank Acct # Bank Acct No. Bank Account No. Savings Account Number Savings Account. Savings Account # Savings Acct Number Savings Acct # Savings Acct No. Savings Account No. Debit Account Number Debit Account Debit Account # Debit Acct Number Debit Acct # Debit Acct No. Debit Account No.

U.S. Driver's License Number

Format: Depends on the state

Pattern: Depends on the state -- for example, New York:

- Nine digits formatted like ddd ddd ddd will match
- Nine digits like ddddddddd will not match.

Checksum: No

Definition:

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_new_york_drivers_license_number` finds content that matches the pattern.
- A keyword from `Keyword_[state_name]_drivers_license_name` is found.
- A keyword from `Keyword_us_drivers_license` is found.

A DLP policy is 65% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_new_york_drivers_license_number` finds content that matches the pattern.
- A keyword from `Keyword_[state_name]_drivers_license_name` is found.
- A keyword from `Keyword_us_drivers_license_abbreviations` is found.

- No keyword from `Keyword_us_drivers_license` is found.

```
<Pattern confidenceLevel="75">
        <IdMatch idRef="Func_new_york_drivers_license_number" />
        <Match idRef="Keyword_new_york_drivers_license_name" />
        <Match idRef="Keyword_us_drivers_license" />
    </Pattern>
    <Pattern confidenceLevel="65">
        <IdMatch idRef="Func_new_york_drivers_license_number" />
        <Match idRef="Keyword_new_york_drivers_license_name" />
        <Match idRef="Keyword_us_drivers_license_abbreviations" />
        <Any minMatches="0" maxMatches="0">
          <Match idRef="Keyword_us_drivers_license" />
        </Any>
    </Pattern>
```

Keywords:

KEYWORD_US_DRIVERS_LICENSE_ABBREVIATIONS	KEYWORD_US_DRIVERS_LICENSE	KEYWORD_[STATE_NAME]_DRIVERS_LICENSE_NAME
DL DLS CDL CDLS ID IDs DL# DLS# CDL# CDLS# ID# IDs# ID number ID numbers LIC LIC#	DriverLic DriverLics DriverLicense DriverLicenses Driver Lic Driver Lics Driver License Driver Licenses DriversLic DriversLics DriversLicense DriversLicenses Drivers Lic Drivers Lics Drivers License Drivers Licenses Driver'Lic Driver'Lics Driver'License Driver'Licenses Driver' Lic Driver' Lics Driver' License Driver' Licenses Driver'sLic Driver'sLics Driver'sLicense Driver'sLicenses Driver's Lic Driver's Lics Driver's License Driver's Licenses identification number identification numbers identification # id card id cards identification card identification cards DriverLic# DriverLics# DriverLicense# DriverLicenses#	State abbreviation (for example, "NY") State name (for example, "New York")

KEYWORD_US_DRIVERS_LICENSE_ABBREVIATIONS	KEYWORD_US_DRIVERS_LICENSE	KEYWORD_[STATE_NAME]_DRIVERS_LICENSE_NAME
	DriverLicenses# Driver Lic# Driver Lics# Driver License# Driver Licenses# DriversLic# DriversLics# DriversLicense# DriversLicenses# Drivers Lic# Drivers Lics# Drivers License# Drivers Licenses# Driver'Lic# Driver'Lics# Driver'License# Driver'Licenses# Driver' Lic# Driver' Lics# Driver' License# Driver' Licenses# Driver'sLic# Driver'sLics# Driver'sLicense# Driver'sLicenses# Driver's Lic# Driver's Lics# Driver's License# Driver's Licenses# id card# id cards# identification card# identification cards#	

U.S. Individual Taxpayer Identification Number (ITIN)

Format: Nine digits that start with a "9" and contain a "7" or "8" as the fourth digit, optionally formatted with spaces or dashes

Pattern:

Formatted:

- The digit "9"
- Two digits
- A space or dash
- A "7" or "8"
- A digit
- A space, or dash
- Four digits

Unformatted:

- The digit "9"
- Two digits
- A "7" or "8"

- Five digits

Checksum: No

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_formatted_itin` finds content that matches the pattern.
- At least one of the following is true:
    - A keyword from `Keyword_itin` is found.
    - The function `Func_us_address` finds an address in the right date format.
    - The function `Func_us_date` finds a date in the right date format.
    - A keyword from `Keyword_itin_collaborative` is found.

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_unformatted_itin` finds content that matches the pattern.
- At least one of the following is true:
    - A keyword from `Keyword_itin_collaborative` is found.
    - The function `Func_us_address` finds an address in the right date format.
    - The function `Func_us_date` finds a date in the right date format.

```
<!-- U.S. Individual Taxpayer Identification Number (ITIN) -->
<Entity id="e55e2a32-f92d-4985-a35d-a0b269eb687b" patternsProximity="300" recommendedConfidence="75">
    <Pattern confidenceLevel="85">
        <IdMatch idRef="Func_formatted_itin" />
        <Any minMatches="1">
          <Match idRef="Keyword_itin" />
          <Match idRef="Func_us_address" />
          <Match idRef="Func_us_date" />
          <Match idRef="Keyword_itin_collaborative" />
        </Any>
    </Pattern>
    <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_unformatted_itin" />
        <Match idRef="Keyword_itin" />
        <Any minMatches="1">
          <Match idRef="Keyword_itin_collaborative" />
          <Match idRef="Func_us_address" />
          <Match idRef="Func_us_date" />
        </Any>
    </Pattern>
</Entity>
```

Keywords:

KEYWORD_ITIN	KEYWORD_ITIN_COLLABORATIVE
taxpayer tax id tax identification itin ssn tin social security tax payer itins taxid individual taxpayer	License DL DOB Birthdate Birthday Date of Birth

U.S. Social Security Number (SSN)

Format: 9 digits, which may be in a formatted or unformatted pattern

> **NOTE**
>
> If issued before mid-2011, an SSN has strong formatting where certain parts of the number must fall within certain ranges to be valid (but there's no checksum).

Pattern: Four functions look for SSNs in four different patterns:

- `Func_ssn` finds SSNs with pre-2011 strong formatting that are formatted with dashes or spaces (ddd-dd-dddd OR ddd dd dddd)
- `Func_unformatted_ssn` finds SSNs with pre-2011 strong formatting that are unformatted as nine consecutive digits (ddddddddd)
- `Func_randomized_formatted_ssn` finds post-2011 SSNs that are formatted with dashes or spaces (ddd-dd-dddd OR ddd dd dddd)
- `Func_randomized_unformatted_ssn` finds post-2011 SSNs that are unformatted as nine consecutive digits (ddddddddd)

Checksum: No

Definition:

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_ssn` finds content that matches the pattern.
- At least one of the following is true:
  - A keyword from `Keyword_ssn` is found.
  - The function `Func_us_date` finds a date in the right date format.
  - The function `Func_us_address` finds an address in the right date format.

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_unformatted_ssn` finds content that matches the pattern.

- A keyword from `Keyword_ssn` is found.
- At least one of the following is true:
  - The function `Func_us_date` finds a date in the right date format.
  - The function `Func_us_address` finds an address in the right date format.

A DLP policy is 65% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_randomized_formatted_ssn` finds content that matches the pattern.
- The function `Func_ssn` does not find content that matches the pattern.
- At least one of the following is true:
  - A keyword from `Keyword_ssn` is found.
  - The function `Func_us_date` finds a date in the right date format.
  - The function `Func_us_address` finds an address in the right date format.

A DLP policy is 55% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_randomized_unformatted_ssn` finds content that matches the pattern.
- A keyword from `Keyword_ssn` is found.
- The function `Func_unformatted_ssn` does not find content that matches the pattern.
- At least one of the following is true:
  - The function `Func_us_date` finds a date in the right date format.
  - The function `Func_us_address` finds an address in the right date format.

```
<!-- U.S. Social Security Number (SSN) -->
<Entity id="a44669fe-0d48-453d-a9b1-2cc83f2cba77" patternsProximity="300" recommendedConfidence="75">
    <Pattern confidenceLevel="85">
        <IdMatch idRef="Func_ssn" />
        <Any minMatches="1">
          <Match idRef="Keyword_ssn" />
          <Match idRef="Func_us_date" />
          <Match idRef="Func_us_address" />
        </Any>
    </Pattern>
    <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_unformatted_ssn" />
        <Match idRef="Keyword_ssn" />
        <Any minMatches="1">
          <Match idRef="Func_us_date" />
          <Match idRef="Func_us_address" />
        </Any>
    </Pattern>
    <Pattern confidenceLevel="65">
        <IdMatch idRef="Func_randomized_formatted_ssn" />
        <Any minMatches="0" maxMatches="0">
          <Match idRef="Func_ssn" />
        </Any>
        <Any minMatches="1">
          <Match idRef="Keyword_ssn" />
          <Match idRef="Func_us_date" />
          <Match idRef="Func_us_address" />
        </Any>
    </Pattern>
    <Pattern confidenceLevel="55">
        <IdMatch idRef="Func_randomized_unformatted_ssn" />
        <Match idRef="Keyword_ssn" />
        <Any minMatches="0" maxMatches="0">
          <Match idRef="Func_unformatted_ssn" />
        </Any>
        <Any minMatches="1">
          <Match idRef="Func_us_date" />
          <Match idRef="Func_us_address" />
        </Any>
    </Pattern>
</Entity>
```

Keywords:

KEYWORD_SSN
Social Security Social Security# Soc Sec SSN SSNS SSN# SS# SSID

Information Rights Management in Exchange Server

8/3/2020 • 12 minutes to read • Edit Online

Every day, people use email to exchange sensitive information, such as confidential information or reports. Because email is accessible from just about anywhere, mailboxes have transformed into repositories that contain large amounts of potentially sensitive information. As a result, information leakage can be a serious threat to organizations. To help prevent information leakage, Exchange Server includes Information Rights Management (IRM) features, which provide persistent online and offline protection for email messages and attachments. These IRM features are basically unchanged from Exchange 2013.

What is information leakage?

Information leakage is the disclosure of sensitive information to unauthorized users. Information leakage can be costly for an organization, and can have a wide-ranging impact on the organization's business, employees, customers, and partners. To avoid violating any applicable regulations, organizations need to protect themselves against accidental or intentional information leakage.

These are some consequences that can result from information leakage:

- **Financial damage**: The organization might incur a loss of business, fines, or adverse media coverage.
- **Damage to image and credibility**: Leaked email messages can potentially be a source of embarrassment for the sender and the organization.
- **Loss of competitive advantage**: This is one of the most serious consequences. The disclosure of strategic business or merger and acquisition plans can lead to losses in revenue or market capitalization for the organization. Other threats to competitive advantage include the loss of research information, analytical data, and other intellectual property.

Traditional solutions to information leakage

Although traditional solutions to information leakage may protect the initial access to data, they often don't provide constant protection. This table describes some traditional solutions to information leakage.

SOLUTION	DESCRIPTION	LIMITATIONS

SOLUTION	DESCRIPTION	LIMITATIONS
Transport Layer Security (TLS)	TLS is an Internet standard protocol that's used to encrypt network communications. In a messaging environment, TLS is used to encrypt server/server and client/server communications. By default, Exchange uses TLS for all internal message transfers. Opportunistic TLS is also enabled by default for SMTP sessions with external hosts (TLS encryption is tried first, but if it isn't available, unencrypted communication is allowed). You can also configure domain security to enforce mutual TLS with external organizations.	TLS only protects the SMTP session between two SMTP hosts. In other words, TLS protects information in motion, and it doesn't provide protection at the message-level or for information at rest. Unless the messages are encrypted using another method, messages in sender and recipient mailboxes remain unprotected. For email sent outside the organization, you can require TLS only for the first hop. After a remote SMTP host receives the message, it can relay it to another SMTP host over an unencrypted session. Because TLS is a transport layer technology that's used in mail flow, it can't provide control over what the recipient does with the message.
Message encryption	Users can use technologies such as S/MIME to encrypt messages.	Users decide whether a message gets encrypted. There are additional costs of a public key infrastructure (PKI) deployment, with the accompanying overhead of certificate management for users and protection of private keys. After a message is decrypted, there's no control over what the recipient can do with the information. Decrypted information can be copied, printed, or forwarded. By default, saved attachments aren't protected. Messaging servers can't open and inspect messages that are encrypted by S/MIME. Therefore, the messaging servers can't enforce messaging policies, scan messages for viruses, or take other actions that require access to the content in messages.

Finally, traditional solutions often lack enforcement tools that apply uniform messaging policies to prevent information leakage. For example, a user marks a message with **Company Confidential** and **Do Not Forward**. After the message is delivered to the recipient, the sender or the organization no longer has control over the message. The recipient can willfully or accidentally forward the message (using features such as automatic forwarding rules) to external email accounts, which subjects your organization to substantial information leakage risks.

IRM in Exchange

IRM in Exchange helps prevent information leakage by offering these features:

- Prevent an authorized recipient of IRM-protected content from forwarding, modifying, printing, faxing, saving, or cutting and pasting the content.
- Protect supported attachment file formats with the same level of protection as the message.
- Support expiration of IRM-protected messages and attachments so they can no longer be viewed after the specified period.

- Prevent IRM-protected content from being copied using the Snipping Tool inWindows.

However, IRM in Exchange can't prevent the disclosure of information by using these methods:

- Third-party screen capture programs.
- Photographing IRM-protected content that's displayed on the screen.
- Users remembering or manually transcribing the information.

IRM uses Active Directory Rights Management Services (AD RMS), an information protection technology in Windows Server that uses extensible rights markup language (XrML)-based certificates and licenses to certify computers and users, and to protect content. When a document or message is protected using AD RMS, an XrML license containing the rights that authorized users have to the content is attached. To access IRM-protected content, AD RMS-enabled applications must procure a use license for the authorized user from the AD RMS server. Office applications, such as Word, Excel, PowerPoint and Outlook are RMS-enabled and can be used to create and consume protected content.

> **NOTE**
>
> The Exchange Prelicense Agent attaches a use license to messages that are protected by the AD RMS server in your organization. For more information, see the Prelicensing section later in this topic.

To learn more about Active Directory Rights Management Services, see Active Directory Rights Management Services.

Active Directory Rights Management Services rights policy templates

AD RMS servers provide a Web service that's used to enumerate and acquire the XrML-based rights policy templates that you use to apply IRM protection to messages. By applying the appropriate rights policy template, you can control whether a recipient is allowed to reply to, reply to all, forward, extract information from, save, or print the message.

By default, Exchange ships with the **Do Not Forward** template. When this template is applied to a message, only the recipients addressed in the message can decrypt the message. The recipients can't forward the message, copy content from the message, or print the message. You can create additional RMS templates on the AD RMS servers in your organization to meet your requirements.

For more information about rights policy templates, see AD RMS Policy Template Considerations.

For more information about creating AD RMS rights policy templates, see AD RMS Rights Policy Templates Deployment Step-by-Step Guide.

Apply IRM protection to messages

By default, an Exchange organization is enabled for IRM, but to apply IRM protection to messages, you need to use one or more of these methods:

- **Manually by users in Outlook**: Users can IRM-protect messages in Outlook by using the AD RMS rights policy templates that are available to them. This process uses the IRM functionality in Outlook, not Exchange. For more information about using IRM in Outlook, see Introduction to using IRM for email messages.
- **Manually by users in Outlook on the web**: When an administrator enables IRM in Outlook on the web (formerly known as Outlook Web App), users can IRM-protect messages that they send, and view IRM-protected messages that they receive. For more information about IRM in Outlook on the web, see Understanding IRM in Outlook Web App.
- **Manually by users in Exchange ActiveSync**: When an administrator enables IRM in Exchange

ActiveSync users can view, reply to, forward, and create IRM-protected messages on ActiveSync devices. For more information, see Understanding Information Rights Management in Exchange ActiveSync.

- **Automatically in Outlook**: Administrators can create Outlook protection rules to automatically IRM-protect messages. Outlook protection rules are automatically deployed to Outlook clients, and IRM-protection is applied by Outlook when the user is composing a message. For more information, see Outlook Protection Rules.

- **Automatically on Mailbox servers**: Administrators can create mail flow rules (also known as transport rules) to automatically IRM-protect messages that match specified conditions. For more information, see Understanding Transport Protection Rules.

> **NOTE**
>
> IRM protection isn't applied again to messages that are already IRM-protected. For example, if a user IRM-protects a message in Outlook or Outlook on the web, a transport protection rule won't apply IRM protection to the same message.

Scenarios for IRM protection

This table describes the scenarios for sending messages, and whether IRM protection is available.

SCENARIO	IS SENDING IRM-PROTECTED MESSAGES SUPPORTED?	REQUIREMENTS
Sending messages within the same on-premises Exchange organization	Yes	For the requirements, see the IRM requirements section later in this topic.
Sending messages between different Active Directory forests in an on-premises organization.	Yes	For the requirements, see Configuring AD RMS to Integrate with Exchange Server 2010 Across Multiple Forests.
Sending messages between an on-premises Exchange organization and a Microsoft 365 or Office 365 organization in a hybrid deployment.	Yes	For more information, see IRM in Exchange hybrid deployments.
Sending messages to external recipients	No	Exchange doesn't include a solution for sending IRM-protected messages to external recipients in non-federated organizations. To create a federated trust between two Active Directory forests by using Active Directory Federation Services (AD FS), see Understanding AD RMS Trust Policies.

Decrypt IRM-protected messages to enforce messaging policies

To enforce messaging policies and for regulatory compliance, Exchange needs access to the content of encrypted messages. To meet eDiscovery requirements due to litigation, regulatory audits, or internal investigations, a designated auditor must also be able to search encrypted messages. To help with these tasks, Exchange includes the following decryption features:

- **Transport decryption**: Allows access to message content by the transport agents that are installed on Exchange servers. For more information, see Understanding Transport Decryption.

- **Journal report decryption**: Allows standard or premium journaling to save a clear-text copy of IRM-protected messages in journal reports. For more information, see Enable journal report decryption.

- **IRM decryption for Exchange Search**: Allows Exchange Search to index content in IRM-protected messages. When a discovery manager performs an In-Place eDiscovery search, IRM-protected messages that have been indexed are returned in the search results. For more information, see Configure IRM for Exchange Search and In-Place eDiscovery.

To enable these decryption features, you need to add the Federation mailbox (a system mailbox that's created by Exchange), to the Super Users group on the AD RMS server. For instructions, see Add the Federation Mailbox to the AD RMS Super Users Group.

Prelicensing

To allow authorized users to view IRM-protected messages and attachments, Exchange automatically attaches a prelicense to protected messages. This prevents the client from making repeated trips to the AD RMS server to retrieve a use license, and enables offline viewing of IRM-protected messages. Prelicensing also allows users to view IRM-protected messages in Outlook on the web. When you enable IRM features, prelicensing is enabled by default.

IRM agents

IRM features use the built-in transport agents that exist in the Transport service on Mailbox servers. Most of the built-in transport agents are invisible and unmanageable by the transport agent management cmdlets in the Exchange Management Shell (***-TransportAgent**).

The built-in transport agents that are associated with IRM are described in this table:

AGENT NAME	MANAGEABLE?	SMTP OR CATEGORIZER EVENT	DESCRIPTION
Journal Report Decryption Agent	No	**OnCategorizedMessage**	Provides a clear-text copy of the IRM-protected messages that are attached to journal reports.
Prelicense Agent	No	**OnRoutedMessage**	Attaches a prelicense to IRM-protected messages.
RMS Decryption Agent	No	**OnSubmittedMessage**,	Decrypts IRM-protected messages to allow access to the message content by transport agents.
RMS Encryption Agent	No	**OnRoutedMessage**	Applies IRM protection to messages flagged by the transport agent and re-encrypts transport decrypted messages.
RMS Protocol Decryption Agent	No	**OnEndOfData**	Decrypts IRM-protected messages to allow access to the message content by transport agents.

AGENT NAME	MANAGEABLE?	SMTP OR CATEGORIZER EVENT	DESCRIPTION
Transport Rule Agent	Yes	**OnRoutedMessage**	Flags messages that match the conditions in a transport protection rule to be IRM-protected by the RMS Encryption agent.

For more information about transport agents, see Transport Agents.

IRM requirements

By default, an Exchange organization is enabled for IRM. To actually implement IRM in your Exchange Server organization, your deployment must meet the requirements that are described in this table.

SERVER	REQUIREMENTS
AD RMS cluster	*AD RMS cluster* is the term that's used for any AD RMS deployment, including a single AD RMS server. AD RMS is a Web service, so you don't need to set up a Windows Server failover cluster. For high availability and load-balancing, you can deploy multiple AD RMS servers in the cluster and use network load balancing (NLB). **Service connection point**: AD RMS-aware applications like Exchange use the service connection point that's registered in Active Directory to discover an AD RMS cluster and URLs. There's only one service connection point for AD RMS in an Active Directory forest. You can register the service connection point during AD RMS Setup, or after setup is complete. **Permissions**: Read and Execute permissions to the AD RMS server certification pipeline (the `ServerCertification.asmx` file at `\inetpub\wwwroot\_wmcs\certification\`) must be assigned to these security principals: • The Exchange Servers group or individual Exchange servers. • The AD RMS Service group on AD RMS servers. For details, see Set Permissions on the AD RMS Server Certification Pipeline. **AD RMS super users**: To enable transport decryption, journal report decryption, IRM in Outlook on the web, and IRM decryption for Exchange Search, you need to add the Federation mailbox to the Super Users group on the AD RMS server. For details, see Add the Federation Mailbox to the AD RMS Super Users Group.
Exchange	Exchange 2010 or later is required. In a production environment, installing AD RMS and Exchange on the same server isn't supported.
Outlook	AD RMS templates for protecting messages are available in Outlook 2007 or later. Outlook protection rules in Exchange require Outlook 2010 or later.

SERVER	REQUIREMENTS
Exchange ActiveSync	IRM is available on mobile applications and devices that support Exchange ActiveSync protocol version 14.1 or later, and the included **RightsManagementInformation** tag (both introduced in Exchange 2010 Service Pack 1). Users with supported devices can use ActiveSync to view, reply to, forward, and create IRM-protected messages without connecting to a computer to activate the device for IRM. For more information, see Understanding Information Rights Management in Exchange ActiveSync.

Exchange IRM features support Office file formats. You can extend IRM protection to other file formats by deploying custom protectors. For more information about custom protectors, search for Information Protection and Control Partners on the Microsoft solution providers page.

Configure and test IRM

You use the Exchange Management Shell to configure IRM features in Exchange. For procedures, see Managing Rights Protection.

After you install and configure a Mailbox server, you can use the **Test-IRMConfiguration** cmdlet to perform end-to-end tests of your IRM deployment. The cmdlet performs these tests:

- Inspects IRM configuration for your Exchange organization.
- Checks the AD RMS server for version and hotfix information.
- Verifies whether an Exchange server can be activated for RMS by retrieving a Rights Account Certificate (RAC) and client licensor certificate.
- Acquires AD RMS rights policy templates from the AD RMS server.
- Verifies that the specified sender can send IRM-protected messages.
- Retrieves a Super User use license for the specified recipient.
- Acquires a prelicense for the specified recipient.

For more information, see Test-IRMConfiguration.

Extend Rights Management with the Rights Management connector

The Azure Rights Management connector (RMS connector) is an optional application that enhances data protection for your Exchange server by employing the cloud-based Azure Rights Management (Azure RMS) service. Once you install the RMS connector, it provides continuous data protection during the lifetime of the information. And, because these services are customizable, you can define the level of protection that you need. For example, you can limit email message access to specific users, or set view-only rights for certain messages.

To learn more about the RMS connector and how to install it, see Deploying the Azure Rights Management connector.

Journaling in Exchange Server

8/3/2020 • 9 minutes to read • Edit Online

Journaling in Exchange Server can help your organization respond to legal, regulatory, and organizational compliance requirements by recording all or targeted email messages. Journaling in Exchange Server is basically unchanged from Exchange Server 2010.

Exchange provides the following journaling options:

- **Standard journaling**: Journal all messages that are sent to and received by mailboxes on a specific mailbox database. To journal all messages in your organization, you need to configure journaling on all mailbox databases on all Exchange servers.

- **Premium journaling**: Use *journal rules* to journal messages based on recipients (all recipients or specified recipients), and scope (internal messages, external messages, or all messages). Premium journaling requires Exchange Enterprise client access licenses (CALs). For more information about CALs, see Exchange licensing FAQs.

To configure journaling, see Journaling procedures in Exchange Server.

When you plan for messaging retention and compliance, it's important to understand journaling, and how journaling fits in your organization's compliance policies.

Why journaling is important

First, it's important to understand the difference between journaling and archiving when it comes to email messages:

- *Journaling* refers to recording email communications as part of the organization's email retention strategy.

- *Archiving* refers to removing email messages from their native location (for example, a user's mailbox), and storing them elsewhere.

Many organizations need to maintain records of the email communication that occurs as employees perform their daily business tasks. You can use Exchange journaling as a tool in your email retention or archival strategy.

Although a regulation may not specifically require journaling, Exchange journaling can help your organization achieve compliance with the regulation. For example, corporate officers in some financial sectors can be held liable for claims that are made by their employees to customers. Designated compliance managers can use journaling to collect and regularly review the email messages that are sent by employees to customers as part of their greater employee-to-customer communications review. The compliance managers can report their approval to the corporate officer, and the corporate officer can then report compliance to the regulating body.

The following list shows some of the more well-known U.S. and international regulations where Exchange journaling may help form part of your compliance strategies:

- Sarbanes-Oxley Act of 2002 (SOX)

- Security Exchange Commission Rule 17a-4 (SEC Rule 17 A-4)

- National Association of Securities Dealers 3010 & 3110 (NASD 3010 & 3110)

- Gramm-Leach-Bliley Act (Financial Modernization Act)

- Financial Institution Privacy Protection Act of 2001

- Financial Institution Privacy Protection Act of 2003
- Health Insurance Portability and Accountability Act of 1996 (HIPAA)
- Uniting and Strengthening America by Providing Appropriate Tools Required to Intercept and Obstruct Terrorism Act of 2001 (Patriot Act)
- European Union Data Protection Directive (EUDPD)
- Japan's Personal Information Protection Act

Journaling agent

The *Journaling agent* is the built-in Exchange transport agent that processes messages as they flow through the Transport service on Mailbox servers. The journaling configuration settings are stored in Active Directory, and are read by the Journaling agent. The Journaling agent is registered on the **OnSubmittedMessage** and **OnRoutedMessage** categorizer events in the transport pipeline. For more information about the transport pipeline, see Mail flow and the transport pipeline.

Note that built-in transport agents like the Journaling agent are invisible and unmanageable by the transport agent management cmdlets (***-TransportAgent**).

Journal reports

A *journal report* is the message that's recorded by journaling. The journal report contains the original message as an unaltered file attachment. The body of the journal report contains summary information from the original message (for example, the sender's email address, message subject, **Message-ID**, and recipient email addresses). This type of journaling is known as *envelope journaling*, and is the only journaling method that's supported by Exchange.

Journal reports and IRM-protected messages

You need to consider the effects of IRM-protected messages on journal reports. Third-party archiving systems that don't have built-in RMS support can't decrypt the IRM-protected messages in journal reports, which negatively affects the search and discovery of content in journaled messages. In Exchange, you can configure journal report decryption to save a clear-text copy of the message in the journal report. For more information, see Enable journal report decryption.

Journal rules

The basic components of a journal rule are:

- **Journal recipient**: Who you want to journal.
- **Journal rule scope**: What you want to journal.
- **Journaling mailbox**: Where you want to store the journaled messages.

Journal recipient

The *journal recipient* specifies who you want to journal. Messages that are sent to or received by the journal recipient are journaled (the direction doesn't matter). You can configure a journal rule to journal messages for all senders and recipients in the Exchange organization, or you can limit a journal rule to an Exchange mailbox, group, mail user, or mail contact. If you specify a distribution group, you enable journaling for the *members* of the distribution group (not for the group itself).

By targeting specific recipients or groups of recipients, you can configure a journaling environment that helps you meet your organization's regulatory and legal requirements, while minimizing the storage and other costs that are associated with retaining large amounts of data.

Journal recipients that are enabled for Unified Messaging in Exchange 2016

By default, if your Exchange 2016 organization uses Unified Messaging (UM) to consolidate the email, voice mail, and fax infrastructure, Exchange is configured to journal voice mail notification and missed call notification messages. You can disable journaling for these types of messages, but messages that contain UM-generated faxes are always journaled.

To disable journaling for voice mail and missed call notifications, see Enable or disable journaling for voice mail and missed call notifications.

> **NOTE**
>
> Unified Messaging is not available in Exchange 2019.

Journal rule scope

After you define who you want to journal, you need to define the scope of the messages to journal. The available scopes are:

- **Internal messages only**: The source or destination of the message is inside your Exchange organization.
- **External messages only**: The source or destination of the message is outside your Exchange organization.
- **All messages**: The source or destination of the message doesn't matter. Note that a journal rule with this scope could potentially journal messages that were already journaled by other rules with internal only or external only scopes.

Journaling mailbox

The journaling mailbox is where the journaled messages are delivered. How you configure the journaling mailbox depends on your organization's policies, regulatory requirements, and legal requirements. For example, you may be able to configure one journaling mailbox for all journal rules in your organization, or you may be required to use different journaling mailboxes for different journal rules.

Notes:

- Journaling mailboxes contain sensitive information, so you need to secure access to them. Messages in the journaling mailbox may be part of legal proceedings or subject to regulatory requirements. We recommend that you create and enforce clearly-defined policies that indicate who has access to a journaling mailbox. Speak with your legal representatives to verify that your journaling solution complies with all the laws and regulations that apply to your organization.
- A Microsoft 365 or Office 365 mailbox can't be used as a journaling mailbox. If you're running a hybrid deployment between on-premises Exchange and Microsoft 365 or Office 365, you can designate on-premises journaling mailboxes for your Microsoft 365 or Office 365 and on-premises organizations. You can also deliver journaled messages to an on-premises email archiving system or a third-party email archiving service.
- Journaling mailboxes need to accept messages that are at least as large as the maximum message size that's available in your organization. Be sure to account for any custom maximum message sizes that you've configured on individual mailboxes. For more information, see Configure message size limits for a mailbox.
- We recommend that you configure the journaling mailbox to only accept messages from the Microsoft Exchange recipient (the only sender of journal reports). Note that you can only do this in the Exchange Management Shell. For more information, see Configure message delivery restrictions for a mailbox.
- We recommend that you disable the storage quota limits for the journaling mailbox. For more information, see Configure storage quotas for a mailbox.

Alternate journaling mailbox

Like other messages, undeliverable journal reports are queued, and delivery is periodically retried until the message expires (the default value is two days, and is configured by the *MessageExpirationTimeout* parameter on the **Set-TransportService** cmdlet). Unlike other messages, expired journal reports can't be returned to the sender in a non-delivery report (also known as an NDR or bounce message), because the sender is the Microsoft Exchange recipient. Expired journal reports can't be recovered.

If you don't want undeliverable journal reports to queue and eventually expire, you can specify an *alternate journaling mailbox* that accepts the NDRs for *all* undeliverable journal reports when *any* journaling mailbox is unavailable (one alternate journaling mailbox for all journaling mailboxes in your organization). The original journal report is an attachment in the NDR. When the journaling mailbox becomes available again, you can use the **Resend this message** feature in Outlook on the NDRs in the alternate journaling mailbox to send the unaltered delivery reports to the journaling mailbox.

Before you configure an alternate journaling mailbox, contact your legal representatives. Laws or regulations that apply to your organization may prohibit all journaled messages from being stored in the same mailbox.

When you configure an alternate journaling mailbox, you should use the same criteria that you used when you configured the journaling mailbox.

Notes:

- If the alternate journaling mailbox also becomes unavailable and rejects the NDRs for undeliverable journal reports, the original journal reports are lost and can't be recovered.
- You should treat the alternate journaling mailbox as a special dedicated mailbox. Journal rules, Inbox rules, and mail flow rules (also known as transport rules) that involve the alternate journaling mailbox are ignored.

Journal rule replication

Because journal rules are stored in Active Directory, they're read and applied by the Transport service on all Mailbox servers in the organization. When you create, modify, or remove a journal rule, the change is replicated between the domain controllers in your organization. This allows Exchange to provide a consistent set of journal rules across the organization.

Notes:

- Replication between domain controllers depends on factors that aren't controlled by Exchange (for example, the number of Active Directory sites, and the speed of network links). Therefore, you need to consider replication delays when you implement journal rules in your organization. For more information about Active Directory replication, see Introduction to Active Directory Replication and Topology Management Using Windows PowerShell.
- Each Mailbox server caches expanded distribution groups to avoid repeated Active Directory queries to determine a group's membership. By default, entries in the expanded groups cache expire every four hours. Therefore, changes to the group's membership can't be applied to journal rules until the expanded groups cache is updated. To force an immediate update of the cache on a Mailbox server, restart the Microsoft Exchange Transport service. You need to restart the service on each Mailbox server where you want to forcibly update the cache.

Troubleshooting

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server. If you're having trouble with the alternate journaling mailbox, see KB2829319.

Journaling procedures in Exchange Server

8/3/2020 • 11 minutes to read • Edit Online

Journaling in Exchange Server records inbound and outbound email messages. For more information, see Journaling in Exchange Server.

This topic shows you how to configure standard journaling (journal messages for all mailboxes on a mailbox database) and premium journaling (use journal rules to specify the recipients that are journaled). Some configuration settings are available in the Exchange admin center (EAC), while others are only available in the Exchange Management Shell.

What do you need to know before you begin?

- Estimated time to complete each procedure: 5 minutes.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Journaling" entry in the Messaging policy and compliance permissions in Exchange Server topic.
- To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection. If you're having trouble with the **JournalingReportDNRTo** mailbox, see Transport and Mailbox Rules in Exchange Online don't work as expected.

Procedures for standard journaling

Standard journaling records all messages that are sent to and received by all mailboxes on the specified mailbox database. You enable journaling by specifying the journaling mailbox for the database (the mailbox that stores the journaled messages). To disable journaling for the database, clear the value for the journaling mailbox on the mailbox database. For more information about the journaling mailbox, see Journaling mailbox.

Caution

Disabling journaling on a mailbox database may result in your organization being out of compliance with any applicable messaging retention policies.

Use the EAC enable or disable journaling on mailbox databases

1. In the EAC, go to **Servers** > **Databases**.
2. Select the mailbox database, and then click **Edit** (✎).
3. In the mailbox database properties window that opens, click the **Maintenance** tab, and then perform one of the following procedures:
   - **Enable journaling**: Click **Browse** next to the **Journal recipient** field. In the resulting dialog box, select the mailbox where you want to store the journaled messages, and then click **OK**.

- **Disable journaling**: Click **Remove X** next to the value in **Journal recipient** field.

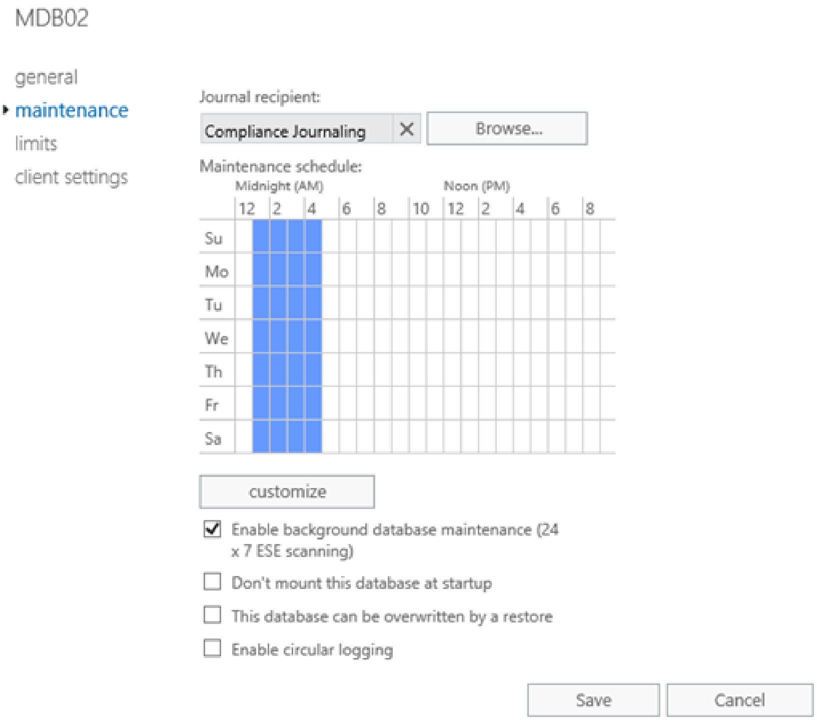

Journal recipient: Compliance Journaling ✕

Maintenance schedule:

☑ Enable background database maintenance (24 x 7 ESE scanning)

□ Don't mount this database at startup

□ This database can be overwritten by a restore

□ Enable circular logging

When you're finished, click **Save**.

Use the Exchange Management Shell to enable or disable journaling on mailbox databases

To enable journaling on a mailbox database, use the following syntax:

```
Set-MailboxDatabase -Identity <MailboxDatabaseIdentity> -JournalRecipient <JournalMailboxIdentity>
```

This example enables journaling on the mailbox database named Sales Database, and configures the mailbox named Sales Database Journal Mailbox as the journaling mailbox that stores the journaled messages.

```
Set-MailboxDatabase -Identity "Sales Database" -JournalRecipient "Sales Database Journal Mailbox"
```

To disable journaling on a mailbox database, use the following syntax:

```
Set-MailboxDatabase -Identity <MailboxDatabaseIdentity> -JournalRecipient $null
```

This example disables journaling on the mailbox database named Sales Database.

```
Set-MailboxDatabase -Identity "Sales Database" -JournalRecipient $null
```

This example disables journaling on all mailbox databases in the Exchange organization.

```
Get-MailboxDatabase | Set-MailboxDatabase -JournalRecipient $null
```

How do you know this worked?

To verify that you've successfully enabled or disabled journaling on a mailbox database, use any of the following procedures:

- In the EAC, go to **Servers** > **Databases** > select the database > **Edit** (✎) > **Maintenance**, and verify the

Journal recipient field is populated (journaling is enabled), or empty (journaling is disabled).

- In the Exchange Management Shell, run the following command to verify the value of the **JournalRecipient** property on all mailbox databases in your organization:

```
Get-MailboxDatabase | Format-Table -Auto Name,JournalRecipient
```

- Send a message to a mailbox on the database, open the journaling mailbox in Outlook or Outlook Web App (formerly known as Outlook on the web), and verify that the journaled message (journal report) is or isn't delivered to the journaling mailbox.

Procedures for premium journaling

Premium journaling uses *journal rules* to record messages based on recipients (all recipients or specified recipients) and scope (internal messages, external messages, or all messages). Premium journaling requires Exchange Enterprise client access licenses (CALs). For more information about CALs, see Exchange licensing FAQs.

Create journal rules

The basic components of a journal rule are:

- **Journal recipient**: Who you want to journal. You can specify all messages, or messages received by or sent by specific recipients (including members of distribution groups).

- **Journal rule scope**: What you want to journal: internal messages only, external messages only, or internal and external messages.

- **Journaling mailbox**: Where you want to store the journaled messages.

Use the EAC to create journal rules

1. In the EAC, go to **Compliance management** > **Journal rules**, and then click **Add** (+).

2. In **New journal rule** window that opens, configure the following settings:

   - **Send journal reports to**: Type the alias or email address of the journaling mailbox where the journaled messages (journal reports) will be delivered.

   - **Name**: Type a unique, descriptive name for the journal rule.

   - **If the message is sent to or received from**: Specify the journal recipient (who you want to journal). Click the drop down arrow and select either of the following values:

   - **A specific user or group**: In the dialog box that opens, select one recipient, and then click **OK** when you're finished.

   - **[Apply to all messages]**

   - **Journal the following messages**: Specify the scope of the journal rule. Click the drop down arrow and select one of the available values:

   - **All messages**

   - **Internal messages only**

   - **External messages only**

   When you're finished, click **Save**.

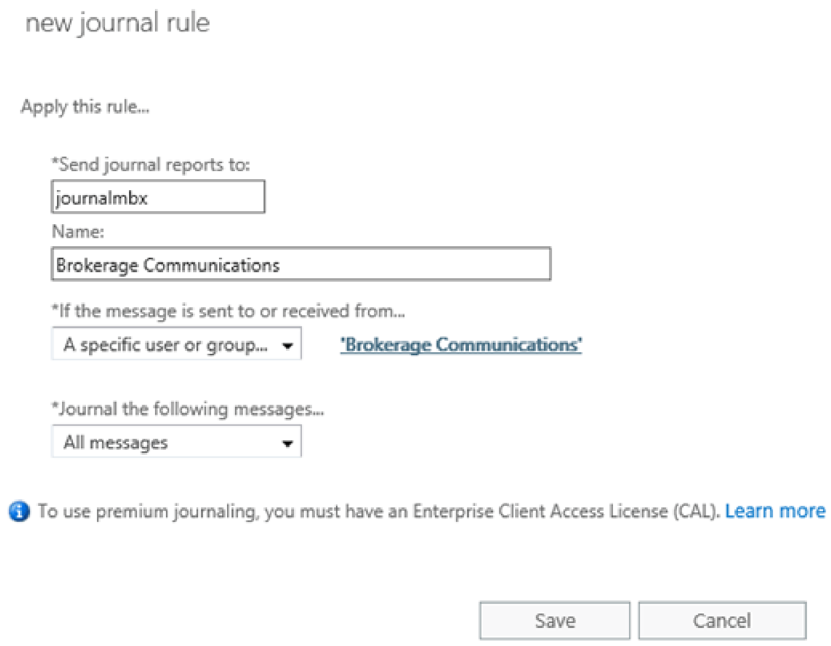

new journal rule

Apply this rule...

*Send journal reports to:

journalmbx

Name:

Brokerage Communications

*If the message is sent to or received from...

A specific user or group... 'Brokerage Communications'

*Journal the following messages...

All messages

To use premium journaling, you must have an Enterprise Client Access License (CAL). Learn more

Use the Exchange Management Shell to create journal rules

To create journal rules in the Exchange Management Shell, use the following syntax:

```
New-JournalRule -Name <RuleName> -JournalEmailAddress <JournalMailboxIdentity> [-Recipient
<JournalRecipientEmailAddress>] [-Scope <Global | Internal | External>] [-Enabled <$true | $false>]
```

This example creates the journal rule named Regulation 123 with the following settings:

- **Journal recipient**: The user Connie Mayr, whose email address is cmayr@contoso.com.
- **Journal rule scope**: Internal and external messages (we didn't use the *Scope* parameter, and the default value is `Global`).
- **Journaling mailbox**: The mailbox named Journal Mailbox.
- The journal rule is enabled (we didn't use the *Enabled* parameter, and the default value is `$true`).

```
New-JournalRule -Name "Regulation 123" -JournalEmailAddress "Journal Mailbox" -Recipient cmayr@contoso.com
```

Note: To create a journal rule that applies to all recipients, don't use the *Recipient* parameter.

For detailed syntax and parameter information, see New-JournalRule.

How do you know this worked?

To verify that you've successfully created a journal rule, use any of the following procedures:

- In the EAC, go to **Compliance management** > **Journal rules** and verify that the new journal rule you created is listed.
- In the Exchange Management Shell, run the following command to verify that the new journal rule is listed:

  ```
  Get-JournalRule | Format-Table -Auto Name,Recipient,JournalEmailAddress,Scope,Enabled
  ```

- Send a message to a recipient that's in the scope of the journal rule, open the journaling mailbox in Outlook or Outlook Web App, and verify that the journaled message (journal report) is delivered to the journaling mailbox.

Enable or disable journal rules

By default, when you create a journal rule in the EAC or the Exchange Management Shell, the rule is enabled. You can only use the Exchange Management Shell to create a journal rule that's disabled (the *Enabled* parameter value is `$false` in the **New-JournalRule** command).

After you create a journal rule, you can use the EAC or the Exchange Management Shell to disable or enable the rule.

> **IMPORTANT**
>
> When a journal rule is disabled, any messages that would have normally been journaled by the rule aren't journaled. Verify that you don't compromise the regulatory or compliance requirements of your organization by disabling a journaling rule.

Use the EAC to enable or disable journal rules

1. In the EAC, go to **Compliance management** > **Journal rules**.

2. In the list view, select the journal rule and in the **On** column, clear the check box to disable the rule, and select the check box to enable the rule.

Use the Exchange Management Shell to enable or disable journal rules

To enable or disable journal rules in the Exchange Management Shell, use the following syntax:

```
<Disable-JournalRule | Enable-JournalRule> -Identity <JournalRuleIdentity>
```

This example disables the journal rule named Contoso Legal.

```
Disable-JournalRule -Identity "Contoso Legal"
```

This example enables the journal rule named Contoso Legal.

```
Enable-JournalRule -Identity "Contoso Legal"
```

How do you know this worked?

To verify that you've successfully enabled or disabled a journal rule, use any of the following procedures:

- In the EAC, go to **Compliance management** > **Journal rules**, and verify the status of the check box in the **On** column for the rule.

- In the Exchange Management Shell, run the following command to verify the value of the **Enabled** property on all journal rules:

  ```
  Get-JournalRule | Format-Table -Auto Name,Enabled
  ```

- Send a message to a recipient that's in the scope of the journal rule, open the journaling mailbox in Outlook or Outlook Web App, and verify that the journaled message (journal report) is or isn't delivered to the journaling mailbox.

Modify journal rules

No additional settings are available when you modify a journal rule. They're the same settings that were available when you created the rule:

- **EAC**: Go to **Compliance management** > **Journal rules**, and then click **Edit** (✎). The available settings are the same as when you created the rule. For more information, see the Use the EAC to create journal rules section.

- **Exchange Management Shell**: The syntax to modify a journal rule is:

  ```
  Set-JournalRule -Identity <JournalRuleIdentity> [-Name <RuleName>] [-JournalEmailAddress
  <JournalMailboxIdentity>] [-Recipient <JournalRecipientEmailAddress | $null>] [-Scope <Global | Internal
  | External>]
  ```

  You can't use the **Set-Journal** cmdlet to enable or disable the rule (there's no *Enabled* parameter). To enable or disable the rule, you use the **Enable-JournalRule** and **Disable-JournalRule** cmdlets as described in the Enable or disable journal rules section.

  For detailed syntax and parameter information, see Set-JournalRule.

Remove journal rules

Use the EAC to remove journal rules

1. In the EAC, go to **Compliance management** > **Journal rules**.

2. In the list view, select the rule or rules that you want to remove, and then click **Delete** (🗑).

Use the Exchange Management Shell to remove journal rules

To remove journal rules in the Exchange Management Shell, use the following syntax:

```
Remove-JournalRule -Identity <JournalRuleIdentity>
```

This example removes the journal rule named Brokerage Journal Rule.

```
Remove-JournalRule "Brokerage Journal Rule"
```

For detailed syntax and parameter information, see Remove-JournalRule.

How do you know this worked?

To verify that you've successfully removed a journal rule, use any of the following procedures:

- In the EAC, go to **Compliance management** > **Journal rules** and verify that the rule you removed is no longer listed.

- In the Exchange Management Shell, run the following command to verify that the rule you removed is no longer listed:

  ```
  Get-JournalRule | Format-Table -Auto Name
  ```

- Send a message to a recipient that was in the scope of the deleted journal rule, open the journaling mailbox in Outlook or Outlook Web App, and verify that the journaled message (journal report) isn't delivered to the journaling mailbox.

Enable or disable journaling for voice mail and missed call notifications

By default, premium journaling will journal voice mail notification and missed call notification messages that are generated by Unified Messaging (UM) in Exchange 2016. However, you can disable journaling for these types of messages. Note that even if you disable journaling for UM notification messages, messages containing faxes that were generated by the UM service are always journaled.

> **NOTE**
>
> Unified Messaging is not available in Exchange 2019.

You can only change this setting in the Exchange Management Shell.

To disable journaling for voice mail and missed call notifications, run the following command:

```
Set-TransportConfig -VoicemailJournalingEnabled $false
```

To enable journaling for voice mail and missed call notifications, run the following command:

```
Set-TransportConfig -VoicemailJournalingEnabled $true
```

How do you know this worked?

To verify that you've successfully enabled or disabled journaling for voice mail and missed call notifications, run the following command to verify the value of the **VoicemailJournalingEnabled** property:

```
Get-TransportConfig | Format-List VoicemailJournalingEnabled
```

Specify the alternate journaling mailbox

For premium journaling, you can specify an *alternate journaling mailbox* that accepts non-delivery reports (also known as NDRs or bounce messages) for *all* undeliverable journal reports when *any* journaling mailbox is unavailable (one alternate journaling mailbox for all journaling mailboxes in your organization). For more information, see Alternate journaling mailbox.

Caution

If the alternate journaling mailbox also becomes unavailable and rejects the NDRs for undeliverable journal reports, the original journal reports are lost and can't be retrieved.

Use the EAC to specify the alternate journaling mailbox

1. In the EAC, go to **Compliance management** > **Journal rules**.
2. Click **Select address** next to **Send undeliverable journal reports to**.
3. In the **NDRs for undeliverable journal reports** window that opens, click **Browse**, select the mailbox in the dialog box that appears, click **OK**, and then click **Save**.

Note: To remove the functionality of the alternate journaling mailbox, click on the email address next to **Send undeliverable journal reports to**. In the In the **NDRs for undeliverable journal reports** window that opens, click **Remove X** next to the email address, and then click **Save**.

Use the Exchange Management Shell to specify an alternate journaling mailbox

To specify the alternate journaling mailbox in the Exchange Management Shell, use the following syntax:

```
Set-TransportConfig -JournalingReportNdrTo <MailboxEmailAddress | $null>
```

This example specifies the mailbox that has the email address altjournalingmbx@contoso.com as the alternate journaling mailbox.

```
Set-TransportConfig -JournalingReportNdrTo altjournalingmbx@contoso.com
```

This example removes the functionality of the alternate journaling mailbox.

```
Set-TransportConfig -JournalingReportNdrTo $null
```

How do you know this worked?

To verify that you've successfully specified an alternate journaling mailbox, use any of the following procedures:

- In the EAC, go to **Compliance management** > **Journal rules** and verify the value of **Send undeliverable journal reports to**.

- In the Exchange Management Shell, run the following command to verify the value of the **JournalingReportNdrTo** property:

  ```
  Get-TransportConfig | Format-List JournalingReportNdrTo
  ```

Enable journal report decryption

Journal report decryption allows **premium journaling** to save a clear-text copy of IRM-protected messages in journal reports (along with the original IRM-protected message). If the message contains any attachments that were protected by the Active Directory Rights Management Services (AD RMS) cluster in your organization, the attachments are also decrypted.

To enable journal report decryption, perform the following steps:

1. Configure the AD RMS super users group. For instructions, see Add the Federation Mailbox to the AD RMS Super Users Group.

2. Run the following command in the Exchange Management Shell:

   ```
   Set-IRMConfiguration -JournalReportDecryptionEnabled $true
   ```

For more information, see Enable or Disable Journal Report Decryption.

Mail flow rules in Exchange Server

8/3/2020 • 12 minutes to read • Edit Online

You can use mail flow rules (also known as transport rules) to identify and take action on messages that flow through the transport pipeline in your Exchange 2016 and Exchange 2019 organization. Mail flow rules are similar to the Inbox rules that are available in Outlook and Outlook on the web (formerly known as Outlook Web App). The main difference is mail flow rules take action on messages while they're in transit, and not after the message is delivered to the mailbox. Mail flow rules contain a richer set of conditions, exceptions, and actions, which provides you with the flexibility to implement many types of messaging policies.

This article explains the components of mail flow rules, and how they work.

You can use the Exchange admin center (EAC) or the Exchange Management Shell to manage mail flow rules. For instructions on how to manage mail flow rules, see Procedures for mail flow rules in Exchange Server.

For each rule, you have the option of enforcing it, testing it, or testing it and notifying the sender. To learn more about the testing options, see Test a mail flow rule and Policy Tips.

For steps to implement specific messaging policies, see the following topics:

- Organization-wide disclaimers, signatures, footers, or headers in Exchange Server
- Common message approval scenarios
- Using mail flow rules to inspect message attachments

Mail flow rule components

A rule is made of conditions, exceptions, actions, and properties:

- **Conditions**: Identify the messages that you want to apply the actions to. Some conditions examine message header fields (for example, the To, From, or Cc fields). Other conditions examine message properties (for example, the message subject, body, attachments, message size, or message classification). Most conditions require you to specify a comparison operator (for example, equals, doesn't equal, or contains) and a value to match. If there are no conditions or exceptions, the rule is applied to all messages.

  For a complete list of mail flow rule conditions, see Mail flow rule conditions and exceptions (predicates) in Exchange Server.

- **Exceptions**: Optionally identify the messages that the actions shouldn't apply to. The same message identifiers that are available in conditions are also available in exceptions. Exceptions override conditions and prevent the rule actions from being applied to a message, even if the message matches all of the configured conditions.

- **Actions**: Specify what to do to messages that match the conditions in the rule, and don't match any of the exceptions. There are many actions available, such as rejecting, deleting, or redirecting messages, adding additional recipients, adding prefixes in the message subject, or inserting disclaimers in the message body.

  For a complete list of mail flow rule actions available, see Mail flow rule actions in Exchange Server.

- **Properties**: Specify other rules settings that aren't conditions, exceptions or actions. For example, when the rule should be applied, whether to enforce or test the rule, and the time period when the rule is active. For more information, see the Mail flow rule properties section in this topic.

Multiple conditions, exceptions, and actions

The following table shows how multiple conditions, condition values, exceptions, and actions are handled in a rule.

COMPONENT	LOGIC	COMMENTS
Multiple conditions	AND	A message must match all the conditions in the rule. If you need to match one condition or another, use separate rules for each condition. For example, if you want to add the same disclaimer to messages with attachments and messages that contain specific text, create one rule for each condition. In the EAC, you can easily copy a rule.
One condition with multiple values	OR	Some conditions allow you to specify more than one value. The message must match any one (not all) of the specified values. For example, if an email message has the subject Stock price information, and the **The subject includes any of these words** condition is configured to match the words Contoso or stock, the condition is satisfied because the subject contains at least one of the specified values.
Multiple exceptions	OR	If a message matches any one of the exceptions, the actions are not applied to the message. The message doesn't have to match all the exceptions.
Multiple actions	AND	Messages that match a rule's conditions get all the actions that are specified in the rule. For example, if the actions **Prepend the subject of the message with** and **Add recipients to the Bcc box** are selected, both actions are applied to the message. Keep in mind that some actions, such as the **Delete the message without notifying anyone** action, prevent subsequent rules from being applied to a message. Other actions such as **Forward the message** do not allow additional actions. You can also set an action on a rule so that when that rule is applied, subsequent rules are not applied to the message.

Mail flow rule properties

The following table describes the rule properties that are available in mail flow rules.

PROPERTY NAME IN THE EAC	PARAMETER NAME IN THE EXCHANGE MANAGEMENT SHELL	DESCRIPTION

PROPERTY NAME IN THE EAC	PARAMETER NAME IN THE EXCHANGE MANAGEMENT SHELL	DESCRIPTION
Priority	*Priority*	Indicates the order that the rules are applied to messages. The default priority is based on when the rule is created (older rules have a higher priority than newer rules), and higher priority rules are processed before lower priority rules. You change the rule priority in the EAC by moving the rule up or down in the list of rules. In the Exchange Management Shell, you set the priority number (0 is the highest priority). For example, if you have one rule to reject messages that include a credit card number, and another one requiring approval, you'll want the reject rule to happen first, and stop applying other rules. For more information, see Set the priority of mail flow rules.
Mode	*Mode*	You can specify whether you want the rule to start processing messages immediately, or whether you want to test rules without affecting the delivery of the message (with or without Data Loss Prevention or DLP Policy Tips). Policy Tips are similar to MailTips, and can be configured to present a brief note in Outlook or Outlook on the web that provides information about possible policy violations to the person that's creating the message. For more information, see Policy Tips.For more information about the modes, see Test a mail flow rule.
Activate this rule on the following date **Deactivate this rule on the following date**	*ActivationDate* *ExpiryDate*	Specifies the date range when the rule is active.
On check box selected or not selected	New rules: *Enabled* parameter on the **New-TransportRule** cmdlet. Existing rules: Use the **Enable-TransportRule** or **Disable-TransportRule** cmdlets. The value is displayed in the **State** property of the rule.	You can create a disabled rule, and enable it when you're ready to test it. Or, you can disable a rule without deleting it to preserve the settings. For instructions, see Enable or disable mail flow rules.
Defer the message if rule processing doesn't complete	*RuleErrorAction*	You can specify how the message should be handled if the rule processing can't be completed. By default, the rule will be ignored, but you can choose to resubmit the message for processing.

PROPERTY NAME IN THE EAC	PARAMETER NAME IN THE EXCHANGE MANAGEMENT SHELL	DESCRIPTION
Match sender address in message	*SenderAddressLocation*	If the rule uses conditions or exceptions that examine the sender's email address, you can look for the value in the message header, the message envelope, or both. For more information, see Senders.
Stop processing more rules	*SenderAddressLocation*	This is an action for the rule, but it looks like a property in the EAC. You can choose to stop applying additional rules to a message after a rule processes a message.
Comments	*Comments*	**Comments** You can enter descriptive comments about the rule.

How mail flow rules are applied

Mail flow rules are applied by a transport agent on Mailbox servers and Edge Transport servers. On Mailbox servers, rules are applied by the Transport Rule agent. On Edge Transport servers, rules are applied by Edge Rule agent. Although similar in functionality, the agents have some differences. The important differences are summarized in the following table:

TRANSPORT AGENT	SMTP OR CATEGORIZER EVENT WHERE RULES ARE INVOKED	WHERE RULES ARE STORED
Transport Rule agent on Mailbox servers	The **OnResolvedMessage** categorizer event. In Exchange 2010, the Transport Rule agent was invoked on the **OnRoutedMessage** categorizer event. The change to **OnResolvedMessage** allowed new rule actions that can change how a message is routed (for example, require TLS).	In Active Directory. Rules are available to all Mailbox servers in the Active Directory forest.
Edge Rule agent on Edge Transport servers	The **OnEndOfData** SMTP event	In the local instance of Active Directory Lightweight Directory Services (AD LDS) on the server. Rules are only applied to messages that flow through the local server.

For more information about transport agents, see Transport Agents.

Differences in processing based on message type

There are several types of messages that flow through an organization. The following table shows which messages types can be processed by mail flow rules.

TYPE OF MESSAGE	CAN A RULE BE APPLIED?
Regular messages Messages that contain a single rich text format (RTF), HTML, or plain text message body or a multipart or alternative set of message bodies.	Yes

TYPE OF MESSAGE	CAN A RULE BE APPLIED?
S/MIME encrypted messages	Rules can only access envelope headers and process messages based on conditions that inspect those headers. Rules with conditions that require inspection of the message's content, or actions that modify the message's content can't be processed.
RMS Protected messages: Messages that are protected by applying an Active Directory Rights Management Services (AD RMS) rights policy template.	Rules can always access envelope headers and process messages based on conditions that inspect those headers.For a rule to inspect or modify a protected message's content, your need to: • Have transport decryption set to **Mandatory** or **Optional**. By default, Transport decryption is set to **Optional**. • Have the encryption key.
Clear-signed messages: Messages that have been signed but not encrypted.	Yes
UM messages: Messages that are created or processed by the Unified Messaging service in Exchange 2016, such as voice mail, fax, missed call notifications, and messages created or forwarded by using Microsoft Outlook Voice Access. (**Note**: Unified Messaging is not available in Exchange 2019.)	Yes
Anonymous messages: Messages that were sent by anonymous senders.	Yes
Read reports: Reports that are generated in response to read receipt requests by senders. Read reports have a message class of `IPM.Note*.MdnRead` or `IPM.Note*.MdnNotRead`.	Yes

Rule storage and replication

Mail flow rules that you create and configure on Mailbox servers are stored in Active Directory, and they're read and applied by the Transport service on all Mailbox servers in the organization. When you create, modify, or remove a mail flow rule, the change is replicated between the domain controllers in your organization. This allows Exchange to provide a consistent set of mail flow rules across the organization.

Notes:

- Replication between domain controllers depends on factors that aren't controlled by Exchange (for example, the number of Active Directory sites, and the speed of network links). Therefore, you need to consider replication delays when you implement mail flow rules in your organization. For more information about Active Directory replication, see Introduction to Active Directory Replication and Topology Management Using Windows PowerShell.

- Each Mailbox server caches expanded distribution groups to avoid repeated Active Directory queries to determine a group's membership. By default, entries in the expanded groups cache expire every four hours. Therefore, changes to the group's membership aren't detected by mail flow rules until the expanded groups cache is updated. To force an immediate update of the cache on a Mailbox server, restart the Microsoft Exchange Transport service. You need to restart the service on each Mailbox server where you want to forcibly update the cache.

Mail flow rules that you create and configure on Edge Transport servers are stored in the local instance of AD LDS on the server. No automated replication of mail flow rules occurs on Edge Transport servers. Rules on the Edge

Transport server apply only to messages that flow through the local server. If you need to apply the same set of mail flow rules on multiple Edge Transport servers, you can clone the Edge Transport server configuration, or export and import the mail flow rules. For more information, see Edge Transport Server Cloned Configuration and Import or export mail flow rule collections.

Whenever the Transport service on a Mailbox server or Edge Transport server detects a modified mail flow rule, an event is logged in the Application log in the Event Viewer (Event ID 4002 on Mailbox servers, and Event ID 16028 on Edge Transport servers).

Rule replication and storage in mixed environments

There are two mixed environment scenarios that are common:

- **Hybrid deployments where part of your organization resides in Microsoft 365 or Office 365**

  In a hybrid environment, there's no replication of rules between your on-premises Exchange organization and Microsoft 365 or Office 365. Therefore, when you create a rule in Exchange, you need to create a matching rule in Microsoft 365 or Office 365. Rules you create in Microsoft 365 or Office 365 are stored in the cloud, whereas the rules you create in your on-premises organization are stored locally in Active Directory. When you manage rules in a hybrid environment, you need to keep the two sets of rules synchronized by making the change in both places, or making the change in one environment and then exporting the rules and importing them in the other environment.

  Important: Even though there is a substantial overlap between the conditions and actions that are available in Microsoft 365 or Office 365 and Exchange Server, there are differences. If you plan on creating the same rule in both locations, make sure that all conditions and actions you plan to use are available. To see the list of available conditions and actions that are available in Microsoft 365 or Office 365, see the following topics:

  Mail flow rule conditions and exceptions (predicates) in Exchange Online

  Mail flow rule actions in Exchange Online

- **Coexistence with Exchange 2010**

  NOTE

  This section applies to Exchange 2016 only.

  When you coexist with Exchange 2010, all mail flow rules are stored in Active Directory and replicated across your organization regardless of the Exchange Server version you used to create the rules. However, all mail flow rules are associated with the Exchange server version that was used to create them and are stored in a version-specific container in Active Directory. When you first deploy Exchange 2016 in your organization, any existing rules are imported to Exchange 2016 as part of the setup process. However, any changes afterwards would need to be made with both versions. For example, if you change an existing rule in Exchange 2016 (Exchange Management Shell or the EAC), you need to make the same change in Exchange 2010 (Exchange Management Shell or the Exchange Management Console).

  Exchange 2010 can't process rules that have the **Version** or **RuleVersion** value 15.*n.n.n*. To be sure all your rules can be processed, only use rules that have the value 14.*n.n.n*.

Mail flow rule conditions and exceptions (predicates) in Exchange Server

8/3/2020 • 29 minutes to read • Edit Online

Conditions and exceptions in mail flow rules (also known as transport rules) identify the messages that the rule is applied to or not applied to. For example, if the rule adds a disclaimer to messages, you can configure the rule to only apply to messages that contain specific words, messages sent by specific users, or to all messages except those sent by the members of a specific group. Collectively, the conditions and exceptions in mail flow rules are also known as *predicates*, because for every condition, there's a corresponding exception that uses the exact same settings and syntax. The only difference is conditions specify messages to include, while exceptions specify messages to exclude.

Most conditions and exceptions have one property that requires one or more values. For example, the **The sender is** condition requires the sender of the message. Some conditions have two properties. For example, the **A message header includes any of these words** condition requires one property to specify the message header field, and a second property to specify the text to look for in the header field. Some conditions or exceptions don't have any properties. For example, the **Any attachment has executable content** condition simply looks for attachments in messages that have executable content.

For more information about mail flow rules in Exchange Server, including how multiple conditions/exceptions or multi-valued conditions/exceptions are handled, see Mail flow rules in Exchange Server.

For more information about conditions and exceptions in mail flow rules in Exchange Online Protection or Exchange Online, see Mail flow rule conditions and exceptions (predicates) in Exchange Online.

Conditions and exceptions for mail flow rules on Mailbox servers

The tables in the following sections describe the conditions and exceptions that are available in mail flow rules on Mailbox servers. The properties types are described in the Property types section.

Senders

Recipients

Message subject or body

Attachments

Any recipients

Message sensitive information types, To and Cc values, size, and character sets

Sender and recipient

Message properties

Message headers

Notes:

- After you select a condition or exception in the Exchange admin center (EAC), the value that's ultimately shown in the **Apply this rule if** or **Except if** field is often different (shorter) than the click path value you selected. Also, when you create new rules based on a template (a filtered list of scenarios), you can often select a short condition name instead of following the complete click path. The short names and full click path values are shown in the EAC column in the tables.

- If you select **[Apply to all messages]** in the EAC, you can't specify any other conditions. The equivalent in the Exchange Management Shell is to create a rule without specifying any condition parameters.

- The settings and properties are the same in conditions and exceptions, so the output of the **Get-TransportRulePredicate** cmdlet doesn't list exceptions separately. Also, the names of some of the predicates that are returned by this cmdlet are different than the corresponding parameter names, and a predicate might require multiple parameters.

Senders

For conditions and exceptions that examine the sender's address, you can specify where rule looks for the sender's address.

In the EAC, in the **Properties of this rule** section, click **Match sender address in message**. Note that you might need to click **More options** to see this setting. In the Exchange Management Shell, the parameter is *SenderAddressLocation*. The available values are:

- **Header**: Only examine senders in the message headers (for example, the **From**, **Sender**, or **Reply-To** fields). This is the default value, and is the way mail flow rules worked before Exchange 2013 Cumulative Update 1 (CU1).
- **Envelope**: Only examine senders from the message envelope (the **MAIL FROM** value that was used in the SMTP transmission, which is typically stored in the **Return-Path** field). Note that message envelope searching is only available for the following conditions (and the corresponding exceptions):
  - **The sender is** (*From*)
  - **The sender is a member of** (*FromMemberOf*)
  - **The sender address includes** (*FromAddressContainsWords*)
  - **The sender address matches** (*FromAddressMatchesPatterns*)
  - **The sender's domain is** (*SenderDomainIs*)
- **Header or envelope** (`HeaderOrEnvelope`): Examine senders in the message header and the message envelope.

CONDITION OR EXCEPTION IN THE EAC	CONDITION AND EXCEPTION PARAMETERS IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY TYPE	DESCRIPTION	AVAILABLE IN
The sender is **The sender > is this person**	*From* *ExceptIfFrom*	`Addresses`	Messages that are sent by the specified mailboxes, mail users, or mail contacts in the Exchange organization.	Exchange 2010 or later
The sender is located **The sender > is external/internal**	*FromScope* *ExceptIfFromScope*	`UserScopeFrom`	Messages that are sent by either internal senders or external senders.	Exchange 2010 or later
The sender is a member of **The sender > is a member of this group**	*FromMemberOf* *ExceptIfFromMemberOf*	`Addresses`	Messages that are sent by a member of the specified group.	Exchange 2010 or later
The sender address includes **The sender > address includes any of these words**	*FromAddressContainsWords* *ExceptIfFromAddressContainsWords*	`Words`	Messages that contain the specified words in the sender's email address.	Exchange 2010 or later
The sender address matches **The sender > address matches any of these text patterns**	*FromAddressMatchesPatterns* *ExceptIfFromAddressMatchesPatterns*	`Patterns`	Messages where the sender's email address contains text patterns that match the specified regular expressions.	Exchange 2010 or later
The sender's specified properties include any of these words **The sender > has specific properties including any of these words**	*SenderADAttributeContainsWords* *ExceptIfSenderADAttributeContainsWords*	First property: `ADAttribute` Second property: `Words`	Messages where the specified Active Directory attribute of the sender contains any of the specified words. Note that the **Country** attribute requires the two-letter country code value (for example, DE for Germany).	Exchange 2010 or later

CONDITION OR EXCEPTION IN THE EAC	CONDITION AND EXCEPTION PARAMETERS IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY TYPE	DESCRIPTION	AVAILABLE IN
The sender's specified properties match these text patterns The sender > has specific properties matching these text patterns	*SenderADAttributeMatchesPatterns* *ExceptIfSenderADAttributeMatchesPatterns*	First property: `ADAttribute` Second property: `Patterns`	Messages where the specified Active Directory attribute of the sender contains text patterns that match the specified regular expressions.	Exchange 2010 or later
The sender has overridden the Policy Tip The sender > has overridden the Policy Tip	*HasSenderOverride* *ExceptIfHasSenderOverride*	n/a	Messages where the sender has chosen to override a data loss prevention (DLP) policy. For more information about DLP policies, see Data loss prevention in Exchange Server.	Exchange 2013 or later
Sender's IP address is in the range The sender > IP address is in any of these ranges or exactly matches	*SenderIPRanges* *ExceptIfSenderIPRanges*	`IPAddressRanges`	Messages where the sender's IP address matches the specified IP address, or falls within the specified IP address range.	Exchange 2013 or later
The sender's domain is The sender > domain is	*SenderDomainIs* *ExceptIfSenderDomainIs*	`DomainName`	Messages where the domain of the sender's email address matches the specified value. If you need to find sender domains that *contain* the specified domain (for example, any subdomain of a domain), use **The sender address matches** (*FromAddressMatchesPatterns*) condition and specify the domain by using the syntax: `'\.domain\.com$'`.	Exchange 2013 or later

Recipients

CONDITION OR EXCEPTION IN THE EAC	CONDITION AND EXCEPTION PARAMETERS IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY TYPE	DESCRIPTION	AVAILABLE IN
The recipient is The recipient > is this person	*SentTo* *ExceptIfSentTo*	`Addresses`	Messages where one of the recipients is the specified mailbox, mail user, or mail contact in the Exchange organization. The recipients can be in the **To**, **Cc**, or **Bcc** fields of the message. **Note**: You can't specify distribution groups or mail-enabled security groups. If you need to take action on messages that are sent to a group, use the **To box contains** (*AnyOfToHeader*) condition instead.	Exchange 2010 or later

CONDITION OR EXCEPTION IN THE EAC	CONDITION AND EXCEPTION PARAMETERS IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY TYPE	DESCRIPTION	AVAILABLE IN
The recipient is located **The recipient > is external/external**	*SentToScope* *ExceptIfSentToScope*	`UserScopeTo`	Messages that are sent to internal recipients, external recipients, external recipients in partner organizations, or external recipients in non-partner organizations.	Exchange 2010 or later
The recipient is a member of **The recipient > is a member of this group**	*SentToMemberOf* *ExceptIfSentToMemberOf*	`Addresses`	Messages that contain recipients who are members of the specified group. The group can be in the **To**, **Cc**, or **Bcc** fields of the message.	Exchange 2010 or later
The recipient address includes **The recipient > address includes any of these words**	*RecipientAddressContainsWords* *ExceptIfRecipientAddressContainsWords*	`Words`	Messages that contain the specified words in the recipient's email address. **Note**: This condition or exception doesn't consider messages that are sent to recipient proxy addresses. It only matches messages that are sent to the recipient's primary email address.	Exchange 2010 or later
The recipient address matches **The recipient > address matches any of these text patterns**	*RecipientAddressMatchesPatterns* *ExceptIfRecipientAddressMatchesPatterns*	`Patterns`	Messages where a recipient's email address contains text patterns that match the specified regular expressions. **Note**: This condition or exception doesn't consider messages that are sent to recipient proxy addresses. It only matches messages that are sent to the recipient's primary email address.	Exchange 2010 or later
The recipient's specified properties include any of these words **The recipient > has specific properties including any of these words**	*RecipientADAttributeContainsWords* *ExceptIfRecipientADAttributeContainsWords*	First property: `ADAttribute` Second property: `Words`	Messages where the specified Active Directory attribute of a recipient contains any of the specified words. Note that the **Country** attribute requires the two-letter country code value (for example, DE for Germany).	Exchange 2010 or later
The recipient's specified properties match these text patterns **The recipient > has specific properties matching these text patterns**	*RecipientADAttributeMatchesPatterns* *ExceptIfRecipientADAttributeMatchesPatterns*	First property: `ADAttribute` Second property: `Patterns`	Messages where the specified Active Directory attribute of a recipient contains text patterns that match the specified regular expressions.	Exchange 2010 or later

CONDITION OR EXCEPTION IN THE EAC	CONDITION AND EXCEPTION PARAMETERS IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY TYPE	DESCRIPTION	AVAILABLE IN
A recipient's domain is **The recipient > domain is**	*RecipientDomainIs* *ExceptIfRecipientDomainIs*	`DomainName`	Messages where the domain of a recipient's email address matches the specified value. If you need to find recipient domains that *contain* the specified domain (for example, any subdomain of a domain), use **The recipient address matches** (*RecipientAddressMatchesPatterns*) condition, and specify the domain by using the syntax `'\.domain\.com$'`.	Exchange 2013 or later

Message subject or body

> **NOTE**
>
> The search for words or text patterns in the subject or other header fields in the message occurs *after* the message has been decoded from the MIME content transfer encoding method that was used to transmit the binary message between SMTP servers in ASCII text. You can't use conditions or exceptions to search for the raw (typically, Base64) encoded values of the subject or other header fields in messages.

CONDITION OR EXCEPTION IN THE EAC	CONDITION AND EXCEPTION PARAMETERS IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY TYPE	DESCRIPTION	AVAILABLE IN
The subject or body includes **The subject or body > subject or body includes any of these words**	*SubjectOrBodyContainsWords* *ExceptIfSubjectOrBodyContainsWords*	`Words`	Messages that have the specified words in the **Subject** field or message body.	Exchange 2010 or later
The subject or body matches **The subject or body > subject or body matches these text patterns**	*SubjectOrBodyMatchesPatterns* *ExceptIfSubjectOrBodyMatchesPatterns*	`Patterns`	Messages where the **Subject** field or message body contain text patterns that match the specified regular expressions.	Exchange 2010 or later
The subject includes **The subject or body > subject includes any of these words**	*SubjectContainsWords* *ExceptIfSubjectContainsWords*	`Words`	Messages that have the specified words in the **Subject** field.	Exchange 2010 or later
The subject matches **The subject or body > subject matches these text patterns**	*SubjectMatchesPatterns* *ExceptIfSubjectMatchesPatterns*	`Patterns`	Messages where the **Subject** field contains text patterns that match the specified regular expressions.	Exchange 2010 or later

Attachments

For more information about how mail flow rules inspect message attachments, see Using mail flow rules to inspect message attachments.

CONDITION OR EXCEPTION IN THE EAC	CONDITION AND EXCEPTION PARAMETERS IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY TYPE	DESCRIPTION	AVAILABLE IN
Any attachment's content includes Any attachment > content includes any of these words	*AttachmentContainsWords* *ExceptIfAttachmentContainsWords*	`Words`	Messages where an attachment contains the specified words.	Exchange 2010 or later
Any attachments content matches Any attachment > content matches these text patterns	*AttachmentMatchesPatterns* *ExceptIfAttachmentMatchesPatterns*	`Patterns`	Messages where an attachment contains text patterns that match the specified regular expressions. **Note**: Only the first 150 kilobytes (KB) of the attachments are scanned.	Exchange 2010 or later
Any attachment's content can't be inspected Any attachment > content can't be inspected	*AttachmentIsUnsupported* *ExceptIfAttachmentIsUnsupported*	n/a	Messages where an attachment isn't natively recognized by Exchange, and the required IFilter isn't installed on the Mailbox server. For more information, see Register Filter Pack IFilters with Exchange Server.	Exchange 2010 or later
Any attachment's file name matches Any attachment > file name matches these text patterns	*AttachmentNameMatchesPatterns* *ExceptIfAttachmentNameMatchesPatterns*	`Patterns`	Messages where an attachment's file name contains text patterns that match the specified regular expressions.	Exchange 2010 or later
Any attachment's file extension matches Any attachment > file extension includes these words	*AttachmentExtensionMatchesWords* *ExceptIfAttachmentExtensionMatchesWords*	`Words`	Messages where an attachment's file extension matches any of the specified words.	Exchange 2013 or later
Any attachment is greater than or equal to Any attachment > size is greater than or equal to	*AttachmentSizeOver* *ExceptIfAttachmentSizeOver*	`Size`	Messages where any attachment is greater than or equal to the specified value. In the EAC, you can only specify the size in kilobytes (KB).	Exchange 2010 or later
The message didn't complete scanning Any attachment > didn't complete scanning	*AttachmentProcessingLimitExceeded* *ExceptIfAttachmentProcessingLimitExceeded*	n/a	Messages where the rules engine couldn't complete the scanning of the attachments. You can use this condition to create rules that work together to identify and process messages where the content couldn't be fully scanned.	Exchange 2013 or later
Any attachment has executable content Any attachment > has executable content	*AttachmentHasExecutableContent* *ExceptIfAttachmentHasExecutableContent*	n/a	Messages where an attachment is an executable file. The system inspects the file's properties rather than relying on the file's extension.	Exchange 2013 or later

CONDITION OR EXCEPTION IN THE EAC	CONDITION AND EXCEPTION PARAMETERS IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY TYPE	DESCRIPTION	AVAILABLE IN
Any attachment is password protected Any attachment > is password protected	*AttachmentIsPasswordProtected* *ExceptIfAttachmentIsPasswordProtected*	n/a	Messages where an attachment is password protected (and therefore can't be scanned). Password detection only works for Office documents and .zip files.	Exchange 2013 or later
has these properties, including any of these words Any attachment > has these properties, including any of these words	*AttachmentPropertyContainsWords* *ExceptIfAttachmentPropertyContainsWords*	First property: `DocumentProperties` Second property: `Words`	Messages where the specified property of an attached Office document contains the specified words. This condition helps you integrate mail flow rules with SharePoint, File Classification Infrastructure (FCI) in Windows Server 2012 R2 or later, or a third-party classification system. You can select from a list of built-in properties, or specify a custom property.	Exchange 2016 or later

Any recipients

The conditions and exceptions in this section provide a unique capability that affects *all* recipients when the message contains at least one of the specified recipients. For example, let's say you have a rule that rejects messages. If you use a recipient condition from the Recipients section, the message is only rejected for those specified recipients. For example, if the rule finds the specified recipient in a message, but the message contains five other recipients. The message is rejected for that one recipient, and is delivered to the five other recipients.

If you add a recipient condition from this section, that same message is rejected for the detected recipient and the five other recipients.

Conversely, a recipient exception from this section *prevents* the rule action from being applied to *all* recipients of the message, not just for the detected recipients.

Note: This condition or exception doesn't consider messages that are sent to recipient proxy addresses. It only matches messages that are sent to the recipient's primary email address.

CONDITION OR EXCEPTION IN THE EAC	CONDITION AND EXCEPTION PARAMETERS IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY TYPE	DESCRIPTION	AVAILABLE IN
Any recipient address includes Any recipient > address includes any of these words	*AnyOfRecipientAddressContainsWords* *ExceptIfAnyOfRecipientAddressContainsWords*	`Words`	Messages that contain the specified words in the **To**, **Cc**, or **Bcc** fields of the message.	Exchange 2013 or later
Any recipient address matches Any recipient > address matches any of these text patterns	*AnyOfRecipientAddressMatchesPatterns* *ExceptIfAnyOfRecipientAddressMatchesPatterns*	`Patterns`	Messages where the **To**, **Cc**, or **Bcc** fields contain text patterns that match the specified regular expressions.	Exchange 2013 or later

Message sensitive information types, To and Cc values, size, and character sets

The conditions in this section that look for values in the **To** and **Cc** fields behave like the conditions in the Any recipients section (*all* recipients of the message are affected by the rule, not just the detected recipients).

Note: The recipient conditions in this section do not consider messages that are sent to recipient proxy addresses. They only match messages that are sent to the recipient's primary email address.

CONDITION OR EXCEPTION IN THE EAC	CONDITION AND EXCEPTION PARAMETERS IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY TYPE	DESCRIPTION	AVAILABLE IN
The message contains sensitive information **The message > contains any of these types of sensitive information**	*MessageContainsDataClassifications* *ExceptIfMessageContainsDataClassifications*	`SensitiveInformationTypes`	Messages that contain sensitive information as defined by data loss prevention (DLP) policies. This condition is required for rules that use the **Notify the sender with a Policy Tip** (*NotifySender*) action.	Exchange 2013 or later
The To box contains **The message > To box contains this person**	*AnyOfToHeader* *ExceptIfAnyOfToHeader*	`Addresses`	Messages where the **To** field includes any of the specified recipients.	Exchange 2010 or later
The To box contains a member of **The message > To box contains a member of this group**	*AnyOfToHeaderMemberOf* *ExceptIfAnyOfToHeaderMemberOf*	`Addresses`	Messages where the **To** field contains a recipient who is a member of the specified group.	Exchange 2010 or later
The Cc box contains **The message > Cc box contains this person**	*AnyOfCcHeader* *ExceptIfAnyOfCcHeader*	`Addresses`	Messages where the **Cc** field includes any of the specified recipients.	Exchange 2010 or later
The Cc box contains a member of **The message > contains a member of this group**	*AnyOfCcHeaderMemberOf* *ExceptIfAnyOfCcHeaderMemberOf*	`Addresses`	Messages where the **Cc** field contains a recipient who is a member of the specified group.	Exchange 2010 or later
The To or Cc box contains **The message > To or Cc box contains this person**	*AnyOfToCcHeader* *ExceptIfAnyOfToCcHeader*	`Addresses`	Messages where the **To** or **Cc** fields contain any of the specified recipients.	Exchange 2010 or later
The To or Cc box contains a member of **The message > To or Cc box contains a member of this group**	*AnyOfToCcHeaderMemberOf* *ExceptIfAnyOfToCcHeaderMemberOf*	`Addresses`	Messages where the **To** or **Cc** fields contain a recipient who is a member of the specified group.	Exchange 2010 or later
The message size is greater than or equal to **The message > size is greater than or equal to**	*MessageSizeOver* *ExceptIfMessageSizeOver*	`Size`	Messages where the total size (message plus attachments) is greater than or equal to the specified value. In the EAC, you can only specify the size in kilobytes (KB). **Note**: Message size limits on mailboxes are evaluated before mail flow rules. A message that's too large for a mailbox will be rejected before a rule with this condition is able to act on the message.	Exchange 2013 or later

CONDITION OR EXCEPTION IN THE EAC	CONDITION AND EXCEPTION PARAMETERS IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY TYPE	DESCRIPTION	AVAILABLE IN
The message character set name includes any of these words The message > character set name includes any of these words	*ContentCharacterSetContainsWords* *ExceptIfContentCharacterSetContainsWords*	`CharacterSets`	Messages that have any of the specified character set names.	Exchange 2013 or later

Sender and recipient

CONDITION OR EXCEPTION IN THE EAC	CONDITION AND EXCEPTION PARAMETERS IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY TYPE	DESCRIPTION	AVAILABLE IN
The sender is one of the recipient's The sender and the recipient > the sender's relationship to a recipient is	*SenderManagementRelationship* *ExceptIfSenderManagementRelationship*	`ManagementRelationship`	Messages where the either sender is the manager of a recipient, or the sender is managed by a recipient.	Exchange 2010 or later
The message is between members of these groups The sender and the recipient > the message is between members of these groups	*BetweenMemberOf1* and *BetweenMemberOf2* *ExceptIfBetweenMemberOf1* and *ExceptIfBetweenMemberOf2*	`Addresses`	Messages that are sent between members of the specified groups.	Exchange 2010 or later
The manager of the sender or recipient is The sender and the recipient > the manager of the sender or recipient is this person	*ManagerForEvaluatedUser* and *ManagerAddress* *ExceptIfManagerForEvaluatedUser* and *ExceptIfManagerAddress*	First property: `EvaluatedUser` Second property: `Addresses`	Messages where either a specified user is the manager of the sender, or a specified user is the manager of a recipient.	Exchange 2010 or later
The sender's and any recipient's property compares as The sender and the recipient > the sender and recipient property compares as	*ADAttributeComparisonAttribute* and *ADComparisonOperator* *ExceptIfADAttributeComparisonAttribute* and *ExceptIfADComparisonOperator*	First property: `ADAttribute` Second property: `Evaluation`	Messages where the specified Active Directory attribute for the sender and recipient either match or don't match.	Exchange 2010 or later

Message properties

CONDITION OR EXCEPTION IN THE EAC	CONDITION AND EXCEPTION PARAMETERS IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY TYPE	DESCRIPTION	AVAILABLE IN

CONDITION OR EXCEPTION IN THE EAC	CONDITION AND EXCEPTION PARAMETERS IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY TYPE	DESCRIPTION	AVAILABLE IN
The message type is The message properties > include the message type	*MessageTypeMatches* *ExceptIfMessageTypeMatches*	`MessageType`	Messages of the specified type. **Note**: When Outlook or Outlook on the web is configured to forward a message, the **ForwardingSmtpAddress** property is added to the message. The message type isn't changed to `AutoForward`.	Exchange 2010 or later
The message is classified as The message properties > include this classification	*HasClassification* *ExceptIfHasClassification*	`MessageClassification`	Messages that have the specified message classification. This is a custom message classification that you can create in your organization by using the **New-MessageClassification** cmdlet.	Exchange 2010 or later
The message isn't marked with any classifications The message properties > don't include any classification	*HasNoClassification* *ExceptIfHasNoClassification*	n/a	Messages that don't have a message classification.	Exchange 2010 or later
The message has an SCL greater than or equal to The message properties > include an SCL greater than or equal to	*SCLOver* *ExceptIfSCLOver*	`SCLValue`	Messages that are assigned a spam confidence level (SCL) that's greater than or equal to the specified value.	Exchange 2010 or later
The message importance is set to The message properties > include the importance level	*WithImportance* *ExceptIfWithImportance*	`Importance`	Messages that are marked with the specified Importance level.	Exchange 2010 or later

Message headers

NOTE

The search for words or text patterns in the subject or other header fields in the message occurs *after* the message has been decoded from the MIME content transfer encoding method that was used to transmit the binary message between SMTP servers in ASCII text. You can't use conditions or exceptions to search for the raw (typically, Base64) encoded values of the subject or other header fields in messages.

CONDITION OR EXCEPTION IN THE EAC	CONDITION AND EXCEPTION PARAMETERS IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY TYPE	DESCRIPTION	AVAILABLE IN

CONDITION OR EXCEPTION IN THE EAC	CONDITION AND EXCEPTION PARAMETERS IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY TYPE	DESCRIPTION	AVAILABLE IN
A message header includes **A message header > includes any of these words**	*HeaderContainsMessageHeader* and *HeaderContainsWords* *ExceptIfHeaderContainsMessageHeader* and *ExceptIfHeaderContainsWords*	First property: `MessageHeaderField` Second property: `Words`	Messages that contain the specified header field, and the value of that header field contains the specified words. The name of the header field and the value of the header field are always used together.	Exchange 2010 or later
A message header matches **A message header > matches these text patterns**	*HeaderMatchesMessageHeader* and *HeaderMatchesPatterns* *ExceptIfHeaderMatchesMessageHeader* and *ExceptIfHeaderMatchesPatterns*	First property: `MessageHeaderField` Second property: `Patterns`	Messages that contain the specified header field, and the value of that header field contains the specified regular expressions. The name of the header field and the value of the header field are always used together.	Exchange 2010 or later

Conditions and exceptions for mail flow rules on Edge Transport servers

The conditions and exceptions that are available in mail flow rules on Edge Transport servers are a small subset of what's available on Mailbox servers. There's no EAC on Edge Transport servers, so you can only manage mail flow rules in the Exchange Management Shell on the local Edge Transport server. The conditions and exceptions are described in the following table. The properties types are described in the Property types section.

CONDITION AND EXCEPTION PARAMETERS IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY TYPE	DESCRIPTION	AVAILABLE IN
AnyOfRecipientAddressContainsWords *ExceptIfAnyOfRecipientAddressContainsWords*	`Words`	Messages that contain the specified words in the **To**, **Cc**, or **Bcc** fields. When a message contains the specified recipient, the rule action is applied (or not applied) to *all* recipients of the message. For example, the message is rejected for all recipients of the message, not just for the specified recipient.	Exchange 2013 or later
AnyOfRecipientAddressMatchesPatterns *ExceptIfAnyOfRecipientAddressMatchesPatterns*	`Patterns`	Messages where the **To**, **Cc**, or **Bcc** fields contain text patterns that match the specified regular expressions. When a message contains the specified recipient, the rule action is applied (or not applied) to *all* recipients of the message. For example, the message is rejected for all recipients of the message, not just for the specified recipient.	Exchange 2013 or later
AttachmentSizeOver *ExceptIfAttachmentSizeOver*	`Size`	Messages with attachments where any attachment is greater than or equal to the specified value.	Exchange 2010 or later
FromAddressContainsWords *ExceptIfFromAddressContainsWords*	`Words`	Messages that contain the specified words in the sender's email address.	Exchange 2010 or later

CONDITION AND EXCEPTION PARAMETERS IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY TYPE	DESCRIPTION	AVAILABLE IN
FromAddressMatchesPatterns *ExceptIfFromAddressMatchesPatterns*	`Patterns`	Messages where the sender's email address contains text patterns that match the specified regular expressions.	Exchange 2010 or later
FromScope *ExceptIfFromScope*	`UserScopeFrom`	Messages that are sent by either internal senders or external senders.	Exchange 2010 or later
HeaderContainsMessageHeader and *HeaderContainsWords* *ExceptIfHeaderContainsMessageHeader* and *ExceptIfHeaderContainsWords*	First property: `MessageHeaderField` Second property: `Words`	Messages that contain the specified header field, and the value of that header field contains the specified words. The name of the header field and the value of the header field are always used together.	Exchange 2010 or later
HeaderMatchesMessageHeader and *HeaderMatchesPatterns* *ExceptIfHeaderMatchesMessageHeader* and *ExceptIfHeaderMatchesPatterns*	First property: `MessageHeaderField` Second property: `Patterns`	Messages that contain the specified header field, and the value of that header field contains the specified regular expressions. The name of the header field and the value of the header field are always used together.	Exchange 2010 or later
MessageSizeOver *ExceptIfMessageSizeOver*	`Size`	Messages where the total size (message plus attachments) is greater than or equal to the specified value.	Exchange 2013 or later
SCLOver *ExceptIfSCLOver*	`SCLValue`	Messages that are assigned an SCL that's greater than or equal to the specified value.	Exchange 2010 or later
SubjectContainsWords *ExceptIfSubjectContainsWords*	`Words`	Messages that contain the specified words in the **Subject** field.	Exchange 2010 or later
SubjectMatchesPatterns *ExceptIfSubjectMatchesPatterns*	`Patterns`	Messages where the **Subject** field contains text patterns that match the specified regular expressions.	Exchange 2010 or later
SubjectOrBodyContainsWords *ExceptIfSubjectOrBodyContainsWords*	`Words`	Messages that contain the specified words in the **Subject** field or message body.	Exchange 2010 or later
SubjectOrBodyMatchesPatterns *ExceptIfSubjectOrBodyMatchesPatterns*	`Patterns`	Messages where the **Subject** field or message body contain text patterns that match the specified regular expressions.	Exchange 2010 or later

Property types

The property types that are used in conditions and exceptions are described in the following table.

> **NOTE**
>
> If the property is a string, trailing spaces are not allowed.

PROPERTY TYPE	VALID VALUES	DESCRIPTION

PROPERTY TYPE	VALID VALUES	DESCRIPTION
`ADAttribute`	Select from a predefined list of Active Directory attributes	You can check against any of the following Active Directory attributes: **City** **Company** **Country** **CustomAttribute1 - CustomAttribute15** **Department** **DisplayName** **Email** **FaxNumber** **FirstName** **HomePhoneNumber** **Initials** **LastName** **Manager** **MobileNumber** **Notes** **Office** **OtherFaxNumber** **OtherHomePhoneNumber** **OtherPhoneNumber** **PagerNumber** **PhoneNumber** **POBox** **State** **Street** **Title** **UserLogonName** **ZipCode** In the EAC, to specify multiple words or text patterns for the same attribute, separate the values with commas. For example, the value `San Francisco,Palo Alto` for the **City** attribute looks for "City equals San Francisco" or City equals Palo Alto". In the Exchange Management Shell, use the syntax `"AttributeName1:Value1,Value 2 with spaces,Value3...","AttributeName2:Word4,Value 5 with spaces,Value6..."` , where `Value` is the word or text pattern that you want to match. For example, `"City:San Francisco,Palo Alto"` or `"City:San Francisco,Palo Alto"` , `"Department:Sales,Finance"` . When you specify multiple attributes, or multiple values for the same attribute, the **or** operator is used. Don't use values with leading or trailing spaces. Note that the **Country** attribute requires the ISO 3166-1 two-letter country code value (for example, DE for Germany). For more information, see Country Codes - ISO 3166.

PROPERTY TYPE	VALID VALUES	DESCRIPTION
`Addresses`	Exchange recipients	Depending on the nature of the condition or exception, you might be able to specify any mail-enabled object in the organization (for example, recipient-related conditions), or you might be limited to a specific object type (for example, groups for group membership conditions). And, the condition or exception might require one value, or allow multiple values. In the Exchange Management Shell, separate multiple values by commas. **Note**: This condition or exception doesn't consider messages that are sent to recipient proxy addresses. It only matches messages that are sent to the recipient's primary email address.

`CharacterSets`	Array of character set names	One or more content character sets that exist in a message. For example:
		`Arabic/iso-8859-6`
		`Chinese/big5`
		`Chinese/euc-cn`
		`Chinese/euc-tw`
		`Chinese/gb2312`
		`Chinese/iso-2022-cn`
		`Cyrillic/iso-8859-5`
		`Cyrillic/koi8-r`
		`Cyrillic/windows-1251`
		`Greek/iso-8859-7`
		`Hebrew/iso-8859-8`
		`Japanese/euc-jp`
		`Japanese/iso-022-jp`
		`Japanese/shift-jis`
		`Korean/euc-kr`
		`Korean/johab`
		`Korean/ks_c_5601-1987`
		`Turkish/windows-1254`
		`Turkish/iso-8859-9`
		`Vietnamese/tcvn`

PROPERTY TYPE	VALID VALUES	DESCRIPTION
`DocumentProperties`	Array of custom or predefined document properties	Specifies a built-in or custom document property. The built-in document properties are: **Business Impact** **Compliancy** **Confidentiality** **Department** **Impact** **Intellectual Property** **Personally Identifiable Information** **Personal Information** **Personal Use** **Required Clearance** **PHI** **PII** **Project** **Protected Health Information** Each property contains a single value. When you specify multiple properties, the **or** operator is used. Exchange Management Shell uses the syntax: `"<PropertyName1>:<PropertyValue1>","<PropertyName2>:<PropertyValue2>"` , where `<PropertyValue>` is the word that you want to match. The syntax for this parameter is `"PropertyName:Word"` . To specify multiple properties, or multiple words for the same property, use the following syntax: `"PropertyName1:Word1,Phrase with spaces,word2...","PropertyName2:Word3,Phrase with spaces,word4...` . Don't use leading or trailing spaces. When you specify multiple properties, or multiple values for the same property, the **or** operator is used.
`DomainName`	Array of SMTP domains	For example, `contoso.com` or `eu.contoso.com` . In the Exchange Management Shell, you can specify multiple domains separated by commas.
`EvaluatedUser`	Single value of **Sender** or **Recipient**	Specifies whether the rule is looking for the manager of the sender or the manager of the recipient.
`Evaluation`	Single value of **Equal** or **Not equal** (`NotEqual`)	When comparing the Active Directory attribute of the sender and recipients, this specifies whether the values should match, or not match.
`Importance`	Single value of **Low**, **Normal**, or **High**	The Importance level that was assigned to the message by the sender in Outlook or Outlook on the web.
`IPAddressRanges`	Array of IP addresses or address ranges	You enter the IPv4 addresses using the following syntax: • **Single IP address**: For example, `192.168.1.1` . • **IP address range**: For example, `192.168.0.1-192.168.0.254` . • **Classless InterDomain Routing (CIDR) IP address range**: For example, `192.168.0.1/25` . In the Exchange Management Shell, you can specify multiple IP addresses or ranges separated by commas.

PROPERTY TYPE	VALID VALUES	DESCRIPTION
`ManagementRelationship`	Single value of **Manager** or **Direct report**(`DirectReport`)	Specifies the relationship between the sender and any of the recipients. The rule checks the **Manager** attribute in Active Directory to see if the sender is the manager of a recipient, or if the sender is managed by a recipient.
`MessageClassification`	Single message classification	In the EAC, you select from the list of message classifications that you've created. In the Exchange Management Shell, you use the **Get-MessageClassification** cmdlet to identify the message classification. For example, use the following command to search for messages with the `Company Internal` classification and prepend the message subject with the value `CompanyInternal` : `New-TransportRule "Rule Name" -HasClassification @(Get-MessageClassification "Company Internal").Identity -PrependSubject "CompanyInternal"`
`MessageHeaderField`	Single string	Specifies the name of the header field. The name of the header field is always paired with the value in the header field (word or text pattern match). The *message header* is a collection of required and optional header fields in the message. Examples of header fields are **To**, **From**, **Received**, and **Content-Type**. Official header fields are defined in RFC 5322. Unofficial header fields start with **X-** and are known as *X-headers*.
`MessageType`	Single message type value	Specifies one of the following message types: • **Automatic reply** (`OOF`) • **Auto-forward** (`AutoForward`) • **Encrypted** • **Calendaring** • **Permission controlled** (`PermissionControlled`) • **Voicemail** • **Signed** • **Approval request** (`ApprovalRequest`) • **Read receipt** (`ReadReceipt`) **Note**: When Outlook or Outlook on the web is configured to forward a message, the **ForwardingSmtpAddress** property is added to the message. The message type isn't changed to `AutoForward` .
`Patterns`	Array of regular expressions	Specifies one or more regular expressions that are used to identify text patterns in values. For more information, see Regular Expression Syntax. In the Exchange Management Shell, you specify multiple regular expressions separated by commas, and you enclose each regular expression in quotation marks (").
`SCLValue`	One of the following values: • **Bypass spam filtering** (`-1`) • Integers 0 through 9	Specifies the spam confidence level (SCL) that's assigned to a message. A higher SCL value indicates that a message is more likely to be spam.

PROPERTY TYPE	VALID VALUES	DESCRIPTION
`SensitiveInformationTypes`	Array of sensitive information types	Specifies one or more sensitive information types that are defined in your organization. For a list of built-in sensitive information types, see Sensitive information types in Exchange Server. In the Exchange Management Shell, use the syntax `@{<SensitiveInformationType1>},@{<SensitiveInforma` . For example, to look for content that contains at least two credit card numbers, and at least one ABA routing number, use the value `@{Name="Credit Card Number"; minCount="2"},@{Name="ABA Routing Number"; minCount="1"}` .
`Size`	Single size value	Specifies the size of an attachment or the whole message. In the EAC, you can only specify the size in kilobytes (KB). In the Exchange Management Shell, when you enter a value, qualify the value with one of the following units: • `B` (bytes) • `KB` (kilobytes) • `MB` (megabytes) • `GB` (gigabytes) For example, `20MB`. Unqualified values are typically treated as bytes, but small values may be rounded up to the nearest kilobyte.
`UserScopeFrom`	Single value of **Inside the organization** (`InOrganization`) or **Outside the organization** (`NotInOrganization`)	A sender is considered to be inside the organization if either of the following conditions is true: • The sender is a mailbox, mail user, group, or mail-enabled public folder that exists in the organization's Active Directory. • The sender's email address is in an accepted domain that's configured as an authoritative domain or an internal relay domain **and** the message was sent or received over an authenticated connection. For more information about accepted domains, see Accepted domains in Exchange Server. A sender is considered to be outside the organization if either of the following conditions is true: • The sender's email address isn't in an accepted domain. • The sender's email address is in an accepted domain that's configured as an external relay domain. **Note**: To determine whether mail contacts are considered to be inside or outside the organization, the sender's address is compared with the organization's accepted domains.

PROPERTY TYPE	VALID VALUES	DESCRIPTION
`UserScopeTo`	One of the following values: • **Inside the organization** (`InOrganization`) • **Outside the organization** (`NotInOrganization`) • **In an external partner organization** (`ExternalPartner`) • **In an external non-partner organization** (`ExternalNonPartner`)	A recipient is considered to be inside the organization if either of the following conditions is true: • The recipient is a mailbox, mail user, group, or mail-enabled public folder that exists in the organization's Active Directory. • The recipient's email address is in an accepted domain that's not configured as an external relay domain **and** the message was sent or received over an authenticated connection. A recipient is considered to be outside the organization if either of the following conditions is true: • The recipient's email address isn't in an accepted domain. • The recipient's email address is in an accepted domain that's configured as an external relay domain. External partner organizations are external domains where you've configured Domain Security (mutual TLS authentication) to send mail. External non-partner organizations are all other external domains that aren't considered partner domains.
`Words`	Array of strings	Specifies one or more words to look for. The words aren't case-sensitive, and can be surrounded by spaces and punctuation marks. Wildcards and partial matches aren't supported. For example, "contoso" matches " Contoso.". However, if the text is surrounded by other characters, it isn't considered a match. For example, "contoso" doesn't match the following values: • Acontoso • Contosoa • Acontosob The asterisk (*) is treated as a literal character, and isn't used as a wildcard character.

For more information

Mail flow rule actions in Exchange Server

Mail flow rule conditions and exceptions (predicates) in Exchange Online

Mail flow rule actions in Exchange Server

8/3/2020 • 19 minutes to read • Edit Online

Actions in mail flow rules (also known as transport rules) specify what you want to do to messages that match conditions of the rule. For example, you can create a rule that forwards message from specific senders to a moderator, or adds a disclaimer or personalized signature to all outbound messages.

Actions typically require additional properties. For example, when the rule redirects a message, you need to specify where to redirect the message. Some actions have multiple properties that are available or required. For example, when the rule adds a header field to the message header, you need to specify both the name and value of the header. When the rule adds a disclaimer to messages, you need to specify the disclaimer text, but you can also specify where to insert the text, or what to do if the disclaimer can't be added to the message. Typically, you can configure multiple actions in a rule, but some actions are exclusive. For example, one rule can't reject and redirect the same message.

For more information about mail flow rules in Exchange Server, including how multiple actions are handled, see Mail flow rules in Exchange Server.

For more information about conditions and exceptions in mail flow rules, see Mail flow rule conditions and exceptions (predicates) in Exchange Server.

Actions for mail flow rules on Mailbox servers

The actions that are available in mail flow rules on Mailbox servers are described in the following table. Valid values for each property are described in Property values[Property values] section.

Notes:

- After you select an action in the Exchange admin center (EAC), the value that's ultimately shown in the **Do the following** field is often different from the click path you selected. Also, when you create new rules, you can sometimes (depending on the selections you make) select a short action name from a template (a filtered list of actions) instead of following the complete click path. The short names and full click path values are shown in the EAC column in the table.
- The names of some of the actions that are returned by the **Get-TransportRuleAction** cmdlet are different than the corresponding parameter names, and multiple parameters might be required for an action.

ACTION IN THE EAC	ACTION PARAMETER IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY	DESCRIPTION	AVAILABLE IN
Forward the message for approval to these people **Forward the message for approval > to these people**	*ModerateMessageByUser*	`Addresses`	Forwards the message to the specified moderators as an attachment wrapped in an approval request. For more information, see Common message approval scenarios. You can't use a distribution group as a moderator.	Exchange 2010 or later

ACTION IN THE EAC	ACTION PARAMETER IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY	DESCRIPTION	AVAILABLE IN
Forward the message for approval to the sender's manager **Forward the message for approval > to the sender's manager**	*ModerateMessageByManager*	n/a	Forwards the message to the sender's manager for approval. This action only works if the sender's **Manager** attribute is defined in Active Directory. Otherwise, the message is delivered to the recipients without moderation.	Exchange 2010 or later
Redirect the message to these recipients **Redirect the message to > these recipients**	*RedirectMessageTo*	`Addresses`	Redirects the message to the specified recipients. The message isn't delivered to the original recipients, and no notification is sent to the sender or the original recipients.	Exchange 2010 or later
Reject the message with the explanation **Block the message > reject the message and include an explanation**	*RejectMessageReasonText*	`String`	Returns the message to the sender in a non-delivery report (also known as an NDR or bounce message) with the specified text as the rejection reason. The recipient doesn't receive the original message or notification. The default enhanced status code that's used is `5.7.1`. When you create or modify the rule in the Exchange Management Shell, you can specify the DSN code by using the *RejectMessageEnhancedStatusCode* parameter.	Exchange 2010 or later

ACTION IN THE EAC	ACTION PARAMETER IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY	DESCRIPTION	AVAILABLE IN
Reject the message with the enhanced status code **Block the message > reject the message with the enhanced status code of**	*RejectMessageEnhancedStatusCode*	`DSNEnhancedStatusCode`	Returns the message to the sender in an NDR with the specified enhanced delivery status notification (DSN) code. The recipient doesn't receive the original message or notification. Valid DSN codes are `5.7.1` or `5.7.900` through `5.7.999`. The default reason text that's used is `Delivery not authorized, message refused`. When you create or modify the rule in the Exchange Management Shell, you can specify the rejection reason text by using the *RejectMessageReasonText* parameter.	Exchange 2010 or later
Delete the message without notifying anyone **Block the message > delete the message without notifying anyone**	*DeleteMessage*	n/a	Silently drops the message without sending a notification to the recipient or the sender.	Exchange 2010 or later
Add recipients to the Bcc box **Add recipients > to the Bcc box**	*BlindCopyTo*	`Addresses`	Adds one or more recipients to the **Bcc** field of the message. The original recipients aren't notified, and they can't see the additional addresses.	Exchange 2010 or later
Add recipients to the To box **Add recipients > to the To box**	*AddToRecipients*	`Addresses`	Adds one or more recipients to the **To** field of the message. The original recipients can see the additional addresses.	Exchange 2010 or later

ACTION IN THE EAC	ACTION PARAMETER IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY	DESCRIPTION	AVAILABLE IN
Add recipients to the Cc box **Add recipients > to the Cc box**	*CopyTo*	`Addresses`	Adds one or more recipients to the **Cc** field of the message. The original recipients can see the additional address.	Exchange 2010 or later
Add the sender's manager as a recipient **Add recipients > add the sender's manager as a recipient**	*AddManagerAsRecipientType*	`AddedManagerAction`	Adds the sender's manager to the message as the specified recipient type (**To**, **Cc**, **Bcc**), or redirects the message to the sender's manager without notifying the sender or the recipient. This action only works if the sender's **Manager** attribute is defined in Active Directory.	Exchange 2010 or later
Append the disclaimer **Apply a disclaimer to the message > append a disclaimer**	*ApplyHtmlDisclaimerText* *ApplyHtmlDisclaimerFallbackAction* *ApplyHtmlDisclaimerTextLocation*	First property: `DisclaimerText` Second property: `DisclaimerFallbackAction` Third property (Exchange Management Shell only): `DisclaimerTextLocation`	Applies the specified HTML disclaimer to the end of the message. When you create or modify the rule in the Exchange Management Shell, use the *ApplyHtmlDisclaimerTextLocation* parameter with the value `Append`.	Exchange 2010 or later
Prepend the disclaimer **Apply a disclaimer to the message > prepend a disclaimer**	*ApplyHtmlDisclaimerText* *ApplyHtmlDisclaimerFallbackAction* *ApplyHtmlDisclaimerTextLocation*	First property: `DisclaimerText` Second property: `DisclaimerFallbackAction` Third property (Exchange Management Shell only): `DisclaimerTextLocation`	Applies the specified HTML disclaimer to the beginning of the message. When you create or modify the rule in the Exchange Management Shell, use the *ApplyHtmlDisclaimerTextLocation* parameter with the value `Prepend`.	Exchange 2010 or later

ACTION IN THE EAC	ACTION PARAMETER IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY	DESCRIPTION	AVAILABLE IN
Remove this header Modify the message properties > remove a message header	*RemoveHeader*	`MessageHeaderField`	Removes the specified header field from the message header.	Exchange 2010 or later
Set the message header to this value Modify the message properties > set a message header	*SetHeaderName* *SetHeaderValue*	First property: `MessageHeaderField` Second property: `String`	Adds or modifies the specified header field in the message header, and sets the header field to the specified value.	Exchange 2010 or later
Apply a message classification Modify the message properties > apply a message classification	*ApplyClassification*	`MessageClassification`	Applies the specified message classification to the message.	Exchange 2010 or later
Set the spam confidence level (SCL) to Modify the message properties > set the spam confidence level (SCL)	*SetSCL*	`SCLValue`	Sets the spam confidence level (SCL) of the message to the specified value.	Exchange 2010 or later
Apply rights protection to the message with Modify the message security > apply rights protection	*ApplyRightsProtectionTemplate*	`RMSTemplate`	Applies the specified Rights Management Services (RMS) template to the message. RMS requires Exchange Enterprise client access licenses (CALs) for each mailbox. For more information about CALs, see Exchange licensing FAQs.	Exchange 2010 or later

ACTION IN THE EAC	ACTION PARAMETER IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY	DESCRIPTION	AVAILABLE IN
Require TLS encryption **Modify the message security > require TLS encryption**	*RouteMessageOutboundRequireTls*	`n/a`	Forces the outbound messages to be routed over a TLS encrypted connection.	Exchange 2013 or later
Prepend the subject of the message with	*PrependSubject*	`String`	Adds the specified text to the beginning of the **Subject** field of the message. Consider using a space or a colon (:) as the last character of the specified text to differentiate it from the original subject text. To prevent the same string from being added to messages that already contain the text in the subject (for example, replies), add the **The subject includes** (*ExceptIfSubjectContainsWords*) exception to the rule.	Exchange 2010 or later

ACTION IN THE EAC	ACTION PARAMETER IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY	DESCRIPTION	AVAILABLE IN
Notify the sender with a Policy Tip	*NotifySender* *RejectMessageReasonText* *RejectMessageEnhancedStatusCode* (Exchange Management Shell only)	First property: `NotifySenderType` Second property: `String` Third property (Exchange Management Shell only): `DSNEnhancedStatusCode`	Notifies the sender or blocks the message when the message matches a DLP policy. When you use this action, you need to use the **The message contains sensitive information** (*MessageContainsDataClassification* condition. When you create or modify the rule in the Exchange Management Shell, the *RejectMessageReasonText* parameter is optional. If you don't use this parameter, the default text `Delivery not authorized, message refused` is used. In the Exchange Management Shell, you can also use the *RejectMessageEnhancedStatusCode* parameter to specify the enhanced status code. If you don't use this parameter, the default enhanced status code `5.7.1` is used. This action limits the other conditions, exceptions, and actions that you can configure in the rule.	Exchange 2013 or later

ACTION IN THE EAC	ACTION PARAMETER IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY	DESCRIPTION	AVAILABLE IN
Generate incident report and send it to	*GenerateIncidentReport* *IncidentReportContent*	First property: `Addresses` Second property: `IncidentReportContent`	Sends an incident report that contains the specified content to the specified recipients. An incident report is generated for messages that match data loss prevention (DLP) policies in your organization.	Exchange 2013 or later
Notify the recipient with a message	*GenerateNotification*	`NotificationMessageText`	Specifies the text, HTML tags, and message keywords to include in the notification message that's sent to the message's recipients. For example, you can notify recipients that the message was rejected by the rule, or marked as spam and delivered to their Junk Email folder.	Exchange 2016 or later
Properties of this rule section > **Audit this rule with severity level**	*SetAuditSeverity*	`AuditSeverityLevel`	Specifies whether to: • Prevent the generation of an incident report and the corresponding entry in the message tracking log. • Generate an incident report and the corresponding entry in the message tracking log with the specified severity level (low, medium, or high).	Exchange 2013 or later
Properties of this rule section > **Stop processing more rules** **More options** > **Properties of this rule** section > **Stop processing more rules**	*StopRuleProcessing*	n/a	Specifies that after the message is affected by the rule, the message is exempt from processing by other rules.	Exchange 2013 or later

Actions for mail flow rules on Edge Transport servers

A small subset of actions that are available on Mailbox servers are also available on Edge Transport servers, but there are also some actions that are only available on Edge Transport servers. There's no EAC on Edge Transport servers, so you can only manage mail flow rules in the Exchange Management Shell on the local Edge Transport server. The actions are described in the following table. The properties types are described in the Property values section.

ACTION PARAMETER IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY	DESCRIPTION	AVAILABLE ON	AVAILABLE IN
AddToRecipients	`Addresses`	Adds one or more recipients to the **To** field of the message. The original recipients can see the additional addresses.	Mailbox servers and Edge Transport servers	Exchange 2010 or later
BlindCopyTo	`Addresses`	Adds one or more recipients to the **Bcc** field of the message. The original recipients aren't notified, and they can't see the additional addresses.	Mailbox servers and Edge Transport servers	Exchange 2010 or later
CopyTo	`Addresses`	Adds one or more recipients to the **Cc** field of the message. The original recipients can see the additional address.	Mailbox servers and Edge Transport servers	Exchange 2010 or later
DeleteMessage	n/a	Silently drops the message without sending a notification to the recipient or the sender.	Mailbox servers and Edge Transport servers	Exchange 2010 or later
Disconnect	n/a	Ends the SMTP connection between the sending server and the Edge Transport server without generating an NDR.	Edge Transport servers only	Exchange 2010 or later

ACTION PARAMETER IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY	DESCRIPTION	AVAILABLE ON	AVAILABLE IN
LogEventText	`String`	Generates an event with the specified text in the Application log of the local Edge Transport server. The entry contains the following information: **Level**: `Information` **Source**: `MSExchange Messaging Policies` **Event ID**: `4000` **Task Category**: `Rules` **EventData**: `The following message is logged by an action in the rules: <text you specify>.`	Edge Transport servers only	Exchange 2010 or later
PrependSubject	`String`	Adds the specified text to the beginning of the **Subject** field of the message. Consider using a space or a colon (:) as the last character of the specified text to differentiate it from the original subject.	Mailbox servers and Edge Transport servers	Exchange 2010 or later
Quarantine	n/a	Delivers the message to the quarantine mailbox that's defined in the content filtering configuration on the Edge Transport server. For more information, see Configure a spam quarantine mailbox. If the quarantine mailbox isn't configured, the message is returned to the sender in an NDR.	Edge Transport servers only	Exchange 2010 or later

ACTION PARAMETER IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY	DESCRIPTION	AVAILABLE ON	AVAILABLE IN
RedirectMessageTo	`Addresses`	Redirects the message to the specified recipients. The message isn't delivered to the original recipients, and no notification is sent to the sender or the original recipients.	Mailbox servers and Edge Transport servers	Exchange 2010 or later
RemoveHeader	`MessageHeaderField`	Removes the specified header field from the message header.	Mailbox servers and Edge Transport servers	Exchange 2010 or later
SetHeaderName *SetHeaderValue*	First property: `MessageHeaderField` Second property: `String`	Adds or modifies the specified header field in the message header, and sets the header field to the specified value.	Mailbox servers and Edge Transport servers	Exchange 2010 or later
SetSCL	`SCLValue`	Sets the SCL of the message to the specified value.	Mailbox servers and Edge Transport servers	Exchange 2010 or later

ACTION PARAMETER IN THE EXCHANGE MANAGEMENT SHELL	PROPERTY	DESCRIPTION	AVAILABLE ON	AVAILABLE IN
SmtpRejectMessageRejectText *SmtpRejectMessageRejectStatusCode*	First property: `String` Second property: `SMTPStatusCode`	Ends the SMTP connection between the sending server and the Edge Transport server with the specified SMTP status code and the specified rejection text. The recipient doesn't receive the original message or notification. Valid values for the SMTP status code are integers from `400` through `500` as defined in RFC 3463. If you specify the rejection text without specifying the SMTP status code, the default code `550` is used. If you specify the SMTP status code without specifying the rejection text, the text that's used is `Delivery not authorized, message refused` .	Edge Transport servers only	Exchange 2010 or later
StopRuleProcessing	n/a	Specifies that after the message is affected by the rule, the message is exempt from processing by other rules.	Mailbox servers and Edge Transport servers	Exchange 2013 or later

Property values

The property values that are used for actions in mail flow rules are described in the following table.

PROPERTY	VALID VALUES	DESCRIPTION

PROPERTY	VALID VALUES	DESCRIPTION
`AddedManagerAction`	One of the following values: • **To** • **Cc** • **Bcc** • **Redirect**	Specifies how to include the sender's manager in messages. If you select **To**, **Cc**, or **Bcc**, the sender's manager is added as a recipient in the specified field. If you select **Redirect**, the message is only delivered to the sender's manager without notifying the sender or the recipient. This action only works if the sender's **Manager** attribute is defined in Active Directory.
`Addresses`	Exchange recipients	Depending on the action, you might be able to specify any mail-enabled object in the organization, or you might be limited to a specific object type. Typically, you can select multiple recipients, but you can only send an incident report to one recipient.
`AuditSeverityLevel`	One of the following values: • Uncheck **Audit this rule with severity level**, or select **Audit this rule with severity level** with the value **Not specified** (`DoNotAudit`) • **Low** • **Medium** • **High**	The values **Low**, **Medium**, or **High** specify the severity level that's assigned to the incident report and to the corresponding entry in the message tracking log. The other value prevents an incident report from being generated, and prevents the corresponding entry from being written to the message tracking log.

PROPERTY	VALID VALUES	DESCRIPTION
`DisclaimerFallbackAction`	One of the following values: • **Wrap** • **Ignore** • **Reject**	Specifies what to do if the disclaimer can't be applied to a message. There are situations where the contents of a message can't be altered (for example, the message is encrypted). The available fallback actions are: **Wrap**: The original message is wrapped in a new message envelope, and the disclaimer text is inserted into the new message. This is the default value. • Subsequent mail flow rules are applied to the new message envelope, not to the original message. Therefore, configure these rules with a lower priority than other rules. • If the original message can't be wrapped in a new message envelope, the original message isn't delivered. The message is returned to the sender in an NDR. **Ignore**: The rule is ignored and the message is delivered without the disclaimer **Reject**: The message is returned to the sender in an NDR.
`DisclaimerText`	HTML string	Specifies the disclaimer text, which can include HTML tags, inline cascading style sheet (CSS) tags, and images by using the IMG tag. The maximum length is 5000 characters, including tags.
`DisclaimerTextLocation`	Single value: `Append` or `Prepend`	In the Exchange Management Shell, you use the *ApplyHtmlDisclaimerTextLocation* to specify the location of the disclaimer text in the message. `Append`: Add the disclaimer to the end of the message body. This is the default value. `Prepend`: Add the disclaimer to the beginning of the message body.

PROPERTY	VALID VALUES	DESCRIPTION
`DSNEnhancedStatusCode`	Single DSN code value: • `5.7.1` • `5.7.900` through `5.7.999`	Specifies the DSN code that's used. You can create custom DSNs by using the **New-SystemMessage** cmdlet. If you don't specify the rejection reason text along with the DSN code, the default reason text that's used is `Delivery not authorized, message refused` . When you create or modify the rule in the Exchange Management Shell, you can specify the rejection reason text by using the *RejectMessageReasonText* parameter.
`IncidentReportContent`	One or more of the following values: • **Sender** • **Recipients** • **Subject** • **Cc'd recipients** (`Cc`) • **Bcc'd recipients** (`Bcc`) • **Severity** • **Sender override information** (`Override`) • **Matching rules** (`RuleDetections`) • **False positive reports** (`FalsePositive`) • **Detected data classifications** (`DataClassifications`) • **Matching content** (`IdMatch`) • **Original mail** (`AttachOriginalMail`)	Specifies the original message properties to include in the incident report. You can choose to include any combination of these properties. In addition to the properties you specify, the message ID is always included. The available properties are: **Sender**: The sender of the original message. **Recipients**, **Cc'd recipients**, and **Bcc'd recipients**: All recipients of the message, or only the recipients in the **Cc** or **Bcc** fields. For each property, only the first 10 recipients are included in the incident report. **Subject**: The **Subject** field of the original message. **Severity**: The audit severity of the rule that was triggered. Message tracking logs include all the audit severity levels, and can be filtered by audit severity. In the EAC, if you clear the **Audit this rule with severity level** check box (in the Exchange Management Shell, the *SetAuditSeverity* parameter value `DoNotAudit`), rule matches won't appear in the rule reports. If a message is processed by more than one rule, the highest severity is included in any incident reports. **Sender override information**: The override if the sender chose to override a Policy Tip. If the sender provided a justification, the first 100 characters of the justification are also included. **Matching rules**: The list of rules that the message triggered. **False positive reports**: The false

PROPERTY	VALID VALUES	DESCRIPTION
		positive if the sender marked the message as a false positive for a Policy Tip. **Detected data classifications**: The list of sensitive information types detected in the message. **Matching content**: The sensitive information type detected, the exact matched content from the message, and the 150 characters before and after the matched sensitive information. **Original mail**: The entire message that triggered the rule is attached to the incident report. In the Exchange Management Shell, you specify multiple values separated by commas.
`MessageClassification`	Single message classification object	In the EAC, you select from the list of available message classifications. In the Exchange Management Shell, use the **Get-MessageClassification** cmdlet to see the message classification objects that are available.
`MessageHeaderField`	Single string	Specifies the SMTP message header field to add, remove, or modify. The *message header* is a collection of required and optional header fields in the message. Examples of header fields are **To**, **From**, **Received**, and **Content-Type**. Official header fields are defined in RFC 5322. Unofficial header fields start with **X-** and are known as *X-headers*.
`NotificationMessageText`	Any combination of plain text, HTML tags, and keywords	Specified the text to use in a recipient notification message. In addition to plain text and HTML tags, you can specify the following keywords that use values from the original message: • `%%From%%` • `%%To%%` • `%%Cc%%` • `%%Subject%%` • `%%Headers%%` • `%%MessageDate%%`

PROPERTY	VALID VALUES	DESCRIPTION
`NotifySenderType`	One of the following values: • **Notify the sender, but allow them to send** (`NotifyOnly`) • **Block the message** (`RejectMessage`) • **Block the message unless it's a false positive** (`RejectUnlessFalsePositiveOverride`) • **Block the message, but allow the sender to override and send** (`RejectUnlessSilentOverride`) • **Block the message, but allow the sender to override with a business justification and send** (`RejectUnlessExplicitOverride`)	Specifies the type of Policy Tip that the sender receives if the message violates a DLP policy. The settings are described in the following list: **Notify the sender, but allow them to send** The sender is notified, but the message is delivered normally. **Block the message** The message is rejected, and the sender is notified. **Block the message unless it's a false positive** The message is rejected unless it's marked as a false positive by the sender. **Block the message, but allow the sender to override and send** The message is rejected unless the sender has chosen to override the policy restriction. **Block the message, but allow the sender to override with a business justification and send** This is similar to **Block the message, but allow the sender to override and send** type, but the sender also provides a justification for overriding the policy restriction. When you use this action, you need to use the **The message contains sensitive information** (*MessageContainsDataClassification*) condition.
`RMSTemplate`	Single RMS template object	Specifies the Rights Management Services (RMS) template that's applied to the message. In the EAC, you select the RMS template from a list. In the Exchange Management Shell, use the **Get-RMSTemplate** cmdlet to see the RMS templates that are available. RMS requires Exchange Enterprise client access licenses (CALs) for each mailbox. For more information about CALs, see Exchange licensing FAQs.
`SCLValue`	One of the following values: • **Bypass spam filtering** (`-1`) • Integers 0 through 9	Specifies the spam confidence level (SCL) that's assigned to the message. A higher SCL value indicates that a message is more likely to be spam.

PROPERTY	VALID VALUES	DESCRIPTION
`String`	Single string	Specifies the text that's applied to the specified message header field, NDR, or event log entry. In the Exchange Management Shell, if the value contains spaces, enclose the value in quotation marks (").

For more information

Mail flow rule conditions and exceptions (predicates) in Exchange Server

Organization-wide disclaimers, signatures, footers, or headers in Exchange Server

8/3/2020 • 6 minutes to read • Edit Online

You can add an email disclaimer, legal disclaimer, disclosure statement, signature, or other information to the top or bottom of email messages that enter or leave your organization. You might be required to do this for legal, business, or regulatory requirements, to identify potentially unsafe email messages, or for other reasons that are unique to your organization.

To create a disclaimer, you create a mail flow rule (also known as transport rule) with an action that adds the specified text to email messages. You can configure the rule to apply the disclaimer to all messages (no conditions), or you can define conditions that determine when the disclaimer is added (for example, when the sender is a member of a specific group, when the message includes specific words or text patterns, or outgoing messages only). You can also define exceptions that prevent the disclaimer from being added to messages (for example, messages from specific senders, messages sent to specific recipients, or messages that already contain the disclaimer). To apply multiple disclaimers to the same message, you need to use multiple rules. For more information about mail flow rules, see Mail flow rules in Exchange Server.

Looking for procedures? See Procedures for mail flow rules in Exchange Server.

Examples

Note: The examples in this topic are not intended for use as-is. Modify them for your needs.

TYPE	SAMPLE TEXT ADDED
Legal - outgoing messages	This email and any files transmitted with it are confidential and intended solely for the use of the individual or entity to whom they are addressed. If you have received this email in error, please notify the system manager.
Legal - incoming messages	Employees are expressly required not to make defamatory statements and not to infringe or authorize any infringement of copyright or any other legal right by email communications. Employees who receive such an email must notify their supervisor immediately.
Notice that message was sent to an alias	This message was sent to the Sales discussion group.
Signature - uses unique data for each employee	Kathleen Mayer Sales Department Contoso www.contoso.com kathleen@contoso.com cell: 111-222-1234
Advertisement	Click here for March specials

Location for your disclaimer

You can choose whether to insert the disclaimer at the beginning of the message (prepend), or at the end of the message (append).

In the EAC, you select the action **Append the disclaimer** or **Apply a disclaimer to the message** > **prepend a disclaimer**.

In the Exchange Management Shell, you use the *ApplyHtmlDisclaimerTextLocation* parameter with the value `Append` (default) or `Prepend`.

Format your disclaimer

Here's the formatting that you can use in your disclaimer text.

TYPE OF INFORMATION	DESCRIPTION
Plain text	The maximum length is 5,000 characters, including any HTML tags and inline Cascading Style Sheets (CSS).
HTML and inline CSS	You can use HTML and inline CSS styles to format the text. For example, use the `<HR>` tag to add a line before the disclaimer. HTML is ignored if the disclaimer is added to a plain text message.
Images	Use the `<IMG>` tag to point to an image available on the Internet. For example, `<IMG src="http://contoso.com/images/companylogo.gif" alt="Contoso logo">` . By default, Outlook and Outlook on the web (formerly known as Outlook Web App) block external web content, including images. Users need to acknowledge and download the blocked external content. We recommend that you test disclaimers that have `IMG` tags to verify they display the way you want.
User information for personalized signatures	You can use tokens to add unique attributes from each user's Active Directory account, such as `DisplayName`, `FirstName`, `LastName`, `PhoneNumber`, `Email`, `FaxNumber`, and `Department`. The syntax is to enclose the attribute name in two percent signs (for example, `%%DisplayName%%`). For a complete list of attributes that can be used in disclaimers and personalized signatures, see the description for the `ADAttribute` property in Mail flow rule conditions and exceptions (predicates) in Exchange Server.

Here's an example of an HTML disclaimer that includes a signature, an `IMG` tag, and embedded CSS.

```
<div style="font-size:9pt;  font-family: 'Calibri',sans-serif;">
%%displayname%%<br/>
%%title%%<br/>
%%company%%<br/>
%%street%%<br/>
%%city%%, %%state%% %%zipcode%%</div>
&nbsp;<br/>
<div style="background-color:#D5EAFF; border:1px dotted #003333; padding:.8em; ">
<div><img alt="Fabrikam"  src="http://fabrikam.com/images/fabrikamlogo.png"></div>
<span style="font-size:12pt;  font-family: 'Cambria','times new roman','garamond',serif; color:#ff0000;">HTML
Disclaimer Title</span><br/>
<p style="font-size:8pt; line-height:10pt; font-family: 'Cambria','times roman',serif;">This message contains
confidential information and is intended only for the individual(s) addressed in the message. If you aren't
the named addressee, you should not disseminate, distribute, or copy this e-mail. If you aren't the intended
recipient, you aren'tified that disclosing, distributing, or copying this e-mail is strictly prohibited.  </p>
<span style="padding-top:10px; font-weight:bold; color:#CC0000; font-size:10pt; font-family:
'Calibri',Arial,sans-serif; "><a href="http://www.fabrikam.com">Fabrikam, Inc. </a></span><br/><br/>
</div>
```

Fallback options for disclaimer rules

Exchange can't modify the content of some messages (for example, encrypted messages). For rules that add disclaimers to messages, you need to specify what to do if the disclaimer can't be added. This is known as the *fallback option* for the disclaimer rule. The available fallback options are:

- **Wrap**: The original message is wrapped in a new message envelope, and the disclaimer text is inserted into the new message. This is the default value.
  - Subsequent mail flow rules are applied to the new message envelope, not to the original message. Therefore, configure these rules with a lower priority than other rules.
  - If the original message can't be wrapped in a new message envelope, the original message isn't delivered. The message is returned to the sender in an non-delivery report (also known as an NDR or bounce message).
- **Ignore**: The rule is ignored and the message is delivered without the disclaimer
- **Reject**: The message is returned to the sender in an NDR.

In the EAC, you select the fallback option in the rule action. In the Exchange Management Shell, you use the *ApplyHtmlDisclaimerFallbackAction* parameter.

Scope your disclaimer

As you work on your disclaimers, consider which messages they should apply to. For example, you might want different disclaimers for internal and external messages, or for messages sent by users in specific departments. To make sure only the first message in a conversation gets a disclaimer, add an exception that prevents the disclaimer text from being applied to the same messages over and over again.

Here are some examples of the conditions and exceptions you can use.

DESCRIPTION	CONDITIONS AND EXCEPTIONS IN EAC	CONDITIONS AND EXCEPTIONS IN THE EXCHANGE MANAGEMENT SHELL FOR THE NEW-TRANSPORTRULE OR SET-TRANSPORTRULE CMDLETS
The recipient is located outside your Exchange organization. An exception is configured so messages that already contain the disclaimer text "CONTOSO LEGAL NOTICE" don't have the disclaimer applied again.	Condition: **The recipient is located** > **Outside the organization** Exception: **The subject or body** > **Subject or body matches these text patterns** > CONTOSO LEGAL NOTICE	`-FromScope NotInOrganization -ExceptIf -SubjectOrBodyMatches "CONTOSO LEGAL NOTICE"`
Incoming messages with executable attachments	Condition 1: **The sender is located** > **Outside the organization** Condition 2: **Any attachment** > **has executable content**	`-FromScope NotInOrganization -AttachmentHasExecutableContent`
Sender is in the marketing department	Condition: **The sender** > **is a member of this group** > *group name*	`-FromMemberOf "Marketing Team"`
Every message that comes from an external sender to the sales discussion group	Condition 1: **The sender is located** > **Outside the organization** Condition 2: **The message** > **To or Cc box contains this person** > *group name*	`-FromScope NotInOrganization -SentTo "Sales Discussion Group"`
Prepend an advertisement to outgoing messages for one month	Condition 1: **The recipient is located** > **Outside the organization** Enter the dates in the **Activate this rule on the following date** and **Deactivate this rule on the following date** fields.	`-ApplyHtmlDisclaimerLocation Prepend -SentToScope NotInOrganization -ActivationDate '03/1/2016' -ExpiryDate '03/31/2016'`

For a complete list of conditions and exceptions that you can use to target the disclaimer, see Mail flow rule conditions and exceptions (predicates) in Exchange Server.

Limitations of organization wide signatures

Exchange Server signatures can't fulfill the following scenarios:

- Insert the signature directly under the latest email reply or forward.
- Display server-side email signatures in users' Sent Items folders.
- Skip lines which contain variables that couldn't be updated (for example, if the value wasn't provided for a user).

To gain these and other capabilities, use a third-party tool. Do an internet search for **email signature software**. A number of these providers are Microsoft Gold Partners and their software provides these capabilities.

For more information

Organization-wide disclaimers, signatures, footers, or headers in Exchange 2013

Procedures for mail flow rules in Exchange Server

8/3/2020 • 16 minutes to read • Edit Online

Mail flow rules (also known as transport rules) identify and take action on messages that flow through your Exchange organization. For more information about mail flow rules, see Mail flow rules in Exchange Server.

On Mailbox servers, you can manage mail flow rules in the Exchange admin center (EAC) and in the Exchange Management Shell. On Edge Transport servers, you can only use the Exchange Management Shell.

> **TIP**
>
> Verify that your rules work the way you expect. Be sure to thoroughly test each rule and the interactions between rules.

What do you need to know before you begin?

- Estimated time to complete each procedure: 5 minutes.
- For more information about the EAC, see Exchange admin center in Exchange Server. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Mail flow rules" entry in Messaging policy and compliance permissions in Exchange Server (Exchange Server), or in Feature Permissions in Exchange Online.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Create mail flow rules

- Creating mail flow rules is mostly about the scenarios that you want to fulfill. For examples, see the following topics:
  - Use mail flow rules to inspect message attachments
  - Organization-wide disclaimers, signatures, footers, or headers in Exchange Server
  - Manage message approval
- Data Loss Prevention (DLP) policies are collections of mail flow rules. To create DLP policies, see Exchange Server DLP Procedures.

Use the EAC to create mail flow rules

The EAC allows you to create mail flow rules by using a template (a filtered list of conditions and actions), by copying an existing rule, or by creating a rule from scratch.

1. In the EAC, go to **Mail flow** > **Rules**, and then select one of the following options:

   - To create a rule from a template, click **Add** (+) and select a template (a value other than **Create new rule**).

   - To copy a rule, select the rule, and then select **Copy** (⧉). Note that the option to copy a rule is only available in the EAC.

   - To create a new rule from scratch, **Add** (+) and then select **Create a new rule**.

2. In the **New rule** page that opens, configure the following settings:

   - **Name**: Enter a unique, descriptive name for the rule.

   - **Apply this rule if**: Select a condition for the rule. If you want the rule to apply to all messages, select **[Apply to all messages]**. For an explanation of the available conditions, see Mail flow rule conditions and exceptions (predicates) in Exchange Server.

   - **Do the following**: Select an action for the rule. The action is applied to messages that match the conditions. For an explanation of the available conditions, see Mail flow rule actions in Exchange Server.

   Optional properties:

   - **Audit this rule with severity level**: For DLP policies, this setting specifies how rule match data is displayed in the DLP policy detection reports. For more information, View DLP policy detection reports. If you clear the check box, or select the value **Not specified**, rule matches won't appear in the rule reports.

   - **Choose a mode for this rule**: You can use one of the two test modes to test the rule without impacting mail flow. In both test modes, when the conditions are met, an entry is added to the message tracking log. Select one of the following values:

   - **Enforce**: This turns on the rule and it starts processing messages immediately. All actions on the rule will be performed. This is the default value.

   - **Test with Policy Tips**: This turns on the rule, and any Policy Tip actions (**Notify the sender with a Policy Tip**) will be sent, but no actions related to message delivery will be performed. DLP is required to use this mode. To learn more, see Policy Tips.

   - **Test without Policy Tips**: For DLP policies, only the **Generate incident report and send it to** action will be enforced. No actions related to message delivery are performed.

3. You can create the rule by clicking **Save**, or you can click **More options** to configure the following additional settings:

   - To add more conditions, click **Add condition**. If you have more than one condition, you can remove a condition by clicking **Remove X**. Note that there are more conditions available after you click **More options**.

   - To add more actions, click **Add action**. If you have more than one action, you can remove an action by clicking **Remove X**. Note that there are more actions available after you click **More options**.

   - To add exceptions for the rule, click **Add exception**, and then select an exception by using the **Except if** drop down. You can remove an exception by clicking **Remove X**.

   - **Activate this rule on the following date**: Specify the start date if you want the rule to take effect after a certain date. Note that the rule will still be enabled prior to that date, but it won't be processed.

   - **Deactivate this rule on the following date**: Specify the end date if you want the rule to stop processing messages on a certain date. Note that the rule will still be enabled after that date, but it won't be processed.

- **Stop processing more rules**: Select this check box to avoid applying additional rules after this rule processes a message.
- **Defer the message if rule processing doesn't complete**: Select this check box to resubmit the message for processing. By default, the rule will be ignored, and delivery of the message will continue as normal.
- **Match sender address in message**: For conditions and exceptions that examine the sender's address, you can specify where the rule looks for the sender's address: in the message header (default), the message envelope, or the header and envelope. For more information, see Senders.
- **Comments**: Specify a descriptive comment for the rule.

When you're finished, click **Save**.

Use the Exchange Management Shell to create mail flow rules

There are two settings that you can configure on new mail flow rules in the Exchange Management Shell that aren't available in the EAC (until after you create the rule):

- Create the new rule as disabled (*Enabled* `$false`)
- Set the priority of the rule (*Priority <Number>*).

To create mail flow rules in the Exchange Management Shell, use the following syntax:

```
New-TransportRule -Name <RuleName> [<Conditions>] [<Exceptions>] <Actions> [<Properties>]
```

This example creates a new rule with the following settings:

- **Name**: Mark messages from the Internet to Sales DG.
- **Conditions**
  - Messages from external senders.

    And

  - Messages sent to the distribution group named Sales Department.
- **Action**: Prepend the message's **Subject** field with the value `"External message to Sales DG: "`. The trailing colon and space help to distinguish the added text from the original value.

```
New-TransportRule -Name "Mark messages from the Internet to Sales DG" -FromScope NotInOrganization -SentTo
"Sales Department" -PrependSubject "External message to Sales DG: "
```

For detailed syntax and parameter information, see New-TransportRule.

Note: The conditions and actions in the example are for illustrative purposes only. Review the available mail flow rule conditions, exceptions, and actions to determine which ones meet your requirements.

How do you know this worked?

To verify that you've successfully created a mail flow rule, use either of the following procedures:

- In the EAC, go to **Mail flow** > **Rules**, and verify that the rule you created is in the list.
- In the Exchange Management Shell, use either of the following procedures:
  - Run the following command to see the new rule in the list of rules:

```
Get-TransportRule
```

- Replace *<RuleName>* with the name of the rule, and run the following command to see the details of the rule:

```
Get-TransportRule -Identity "<RuleName>" | Format-List
```

View mail flow rules

Mail flow rules that you create on a Mailbox server are stored in Active Directory, so when you view the rules on a Mailbox server, you see all rules in your organization. When you use the Exchange Management Shell to view mail flow rules on an Edge Transport server, you see the rules that are stored on the local server.

Use the EAC to view mail flow rules

1. In the EAC, go to **Mail flow** > **Rules**.

2. When you select a rule, information about the rule is displayed in the details pane. To see more information about the rule, click **Edit** (✎).

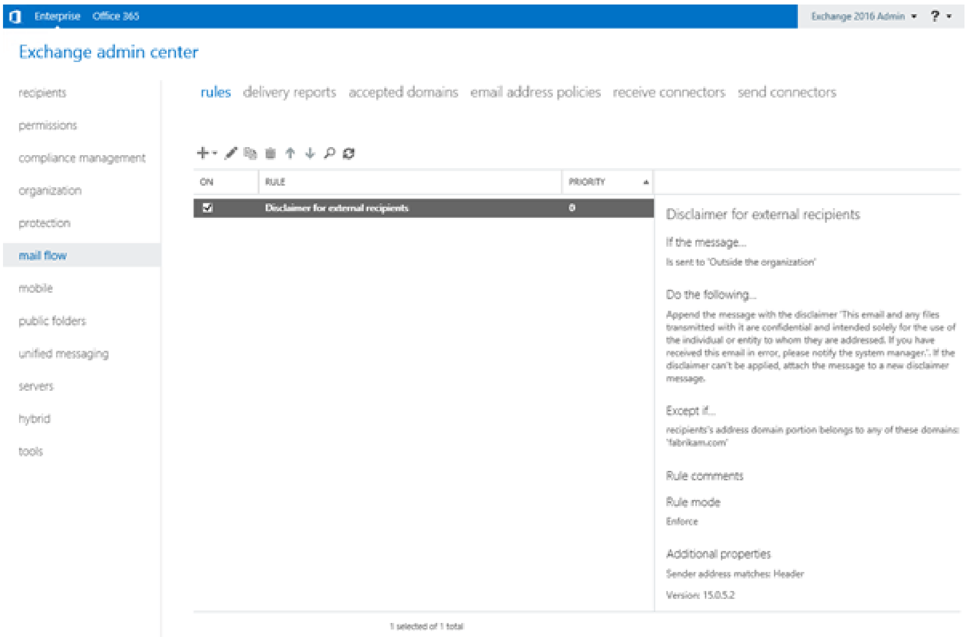

ON	RULE	PRIORITY
☑	Disclaimer for external recipients	0

Disclaimer for external recipients

If the message...

Is sent to 'Outside the organization'

Do the following...

Append the message with the disclaimer 'This email and any files transmitted with it are confidential and intended solely for the use of the individual or entity to whom they are addressed. If you have received this email in error, please notify the system manager.'. If the disclaimer can't be applied, attach the message to a new disclaimer message.

Except if...

recipients's address domain portion belongs to any of these domains: 'fabrikam.com'

Rule comments

Rule mode

Enforce

Additional properties

Sender address matches: Header

Version: 15.0.5.2

1 selected of 1 total

In the EAC, the **Version** property is only visible in the details pane. This property indicates the compatibility of the rule with previous versions of Exchange (14.*n.n.n* is Exchange 2010, 15.0.*n.n* is Exchange 2013).

Use the Exchange Management Shell to view mail flow rules

To return a summary list of all mail flow rules, run the following command:

```
Get-TransportRule
```

To return detailed information about a specific rule, use the following syntax:

```
Get-TransportRule -Identity "<RuleName>" | Format-List [<Specific properties to view>]
```

This example returns all the property values for the rule named "Sender is a member of marketing".

```
Get-TransportRule -Identity "Sender is a member of marketing" | Format-List
```

This example returns only the specified properties for the same rule.

```
Get-TransportRule -Identity "Sender is a member of marketing" | Format-List
Name,State,Mode,Priority,Comments,Conditions,Exceptions,RuleVersion
```

For detailed syntax and parameter information, see Get-TransportRule.

Use the Exchange Management Shell to view the available conditions and exceptions (predicates) for mail flow rules

The conditions and exceptions in mail flow rules are collectively known as *predicates* because for every condition, there's a corresponding exception that uses the exact same settings and syntax. The only difference is: conditions specify messages to include, while exceptions specify messages to exclude. You can only view the list of conditions and exceptions in the Exchange Management Shell.

To view the conditions and exceptions that are available in mail flow rules, run the following command:

```
Get-TransportRulePredicate
```

For detailed syntax and parameter information, see Get-TransportRulePredicate.

Notes:

- Exceptions aren't distinguished from conditions.
- The predicates that are available on Edge Transport servers are a small subset of those available on Mailbox servers. For more information, see Mail flow rule conditions and exceptions (predicates) in Exchange Server.
- Some of the predicate names are different than the corresponding condition and exception parameter names on the **New-TransportRule** and **Set-TransportRule** cmdlets. And, some predicates require multiple parameters.

Use the Exchange Management Shell to view the available actions for mail flow rules

You can only view the list of actions in the Exchange Management Shell.

To view the actions that are available in mail flow rules, run the following command:

```
Get-TransportRuleAction
```

For detailed syntax and parameter information, see Get-TransportRuleAction.

Notes:

- A small subset of actions that are available on Mailbox servers are also available on Edge Transport servers, but some actions are only available on Edge Transport servers. For more information, see Mail flow rule actions in Exchange Server.
- Some of the action names are different than the corresponding action parameter names on the **New-TransportRule** and **Set-TransportRule** cmdlets. And, some actions require multiple parameters.

Modify mail flow rules

Use the EAC to modify mail flow rules

No additional settings are available when you modify a mail flow rule in the EAC. They're the same settings that were available when you created the rule.

1. In the EAC, go to **Mail flow** > **Rules**.

2. Select the rule, and then click **Edit** (✎). Note that the properties of the rule are fully expanded (there's no **More options** link available). For more information about the rule properties, see the Use the EAC to create mail flow rules section in this topic.

Use the Exchange Management Shell to modify mail flow rules

When you modify a mail flow rule in the Exchange Management Shell, you can't disable or enable the rule (there's no *Enabled* parameter on the **Set-TransportRule** cmdlet). Instead, you use the **Disable-TransportRule** and **Enable-TransportRule** cmdlets as describe later in this topic.

To modify a mail flow rule in the Exchange Management Shell, use the following syntax:

```
Set-MailFlowRule -Identity "<RuleName>" [<Conditions>] [<Exceptions>] [<Actions>] [<Properties>]
```

This example adds an exception to the rule named "Sender is a member of marketing" so that it won't apply to messages that are sent by the user named Kelly Rollin.

```
Set-TransportRule -Identity "Sender is a member of marketing" -ExceptIfFrom "Kelly Rollin"
```

For detailed syntax and parameter information, see Set-TransportRule.

How do you know this worked?

To verify that you have successfully modified a mail flow rule, use either of the following procedures:

- In the EAC, go to **Mail flow** > **Rules**, select the rule, and view the information in details pane. To see more settings, click **Edit** (✎).

- In the Exchange Management Shell, replace *<RuleName>* with the name of the rule, and run the following command:

  ```
  Get-TransportRule -Identity "<RuleName>" | Format-List
  ```

Set the priority of mail flow rules

By default, mail flow rules are given a priority that's based on the order they were created in (newer rules are lower priority than older rules). A lower priority number indicates a higher priority for the rule, and rules are processed in priority order (higher priority rules are processed before lower priority rules). No two rules can have the same priority.

Notes:

- You can prevent a message from being acted on by subsequent lower priority rules by including the **Stop processing more rules** (*StopRuleProcessing* `$true`) action in the rule.

- In the EAC, you can only change the priority of the rule after you create it. In the Exchange Management Shell, you can override the default priority when you create the rule (which can affect the priority of existing rules).

Use the EAC to set the priority of mail flow rules

In the EAC, rules are processed in the order that they're displayed (the first rule has the **Priority** value 0). To change the priority of a rule, move the rule up or down in the list (you can also directly modify the **Priority** number by editing the rule in the EAC).

1. In the EAC, go to **Mail flow** > **Rules**.
2. Select a rule, and then click **Move up** (↑) or **Move down** (↓) to move the rule up or down in the list.

Use the Exchange Management Shell to set the priority of mail flow rules

The highest priority value you can set on a rule is 0. The lowest value you can set depends on the number of rules. For example, if you have five rules, you can use the priority values 0 through 4. Changing the priority of an existing rule can have a cascading effect on other rules. For example, if you have five rules (priorities 0 through 4), and you change the priority of a rule to 2, the existing rule with priority 2 is changed to priority 3, and the rule with priority 3 is changed to priority 4.

To set the priority of a rule in the Exchange Management Shell, use the following syntax:

```
Set-TransportRule -Identity "<RuleName>" -Priority <Number>
```

This example sets the priority of the rule named "Sender is a member of marketing" to 2. All existing rules that have a priority less than or equal to 2 are decreased by 1 (their priority numbers are increased by 1).

```
Set-TransportRule -Identity "Sender is a member of marketing" -Priority 2
```

Note: To set the priority of a new rule when you create it, use the *Priority* parameter on the **New-TransportRule** cmdlet.

How do you know this worked?

To verify that you have successfully modified the priority of a mail flow rule, use either of the following procedures:

- In the EAC, go to **Mail flow** > **Rules**, and verify the **Priority** value of the rule in the list.
- In the Exchange Management Shell, use either of the following procedures:
  - Run the following command to see the list of rules and their **Priority** values:

    ```
    Get-TransportRule
    ```

  - Replace *<RuleName>* with the name of the rule, and run the following command:

    ```
    Get-TransportRule -Identity "<RuleName>" | Format-List Name,Priority
    ```

Enable or disable mail flow rules

Disabling a rule prevents the rule from acting on messages, but allows you to preserve the settings of the rule.

By default, mail flow rules are enabled when you create them in the EAC or the Exchange Management Shell, but you can use the Exchange Management Shell to create a disabled rule (use the *Enabled* parameter with the value `$false`).

Use the EAC to enable or disable mail flow rules

1. In the EAC, go to **Mail flow** > **Rules**.
2. Select the rule from the list, and then configure one of the following settings:
   - **Disable the rule**: Clear the check box in the **On** column.
   - **Enable the rule**: Select the check box in the **On** column.

Use the Exchange Management Shell to enable or disable mail flow rules

To enable or disable a mail flow rule in the Exchange Management Shell, use the following syntax:

```
<Enable-TransportRule | Disable-TransportRule> -Identity "<RuleName>"
```

This example disables the mail flow rule named "Sender is a member of marketing".

```
Disable-TransportRule "Sender is a member of marketing"
```

This example enables the mail flow rule named "Sender is a member of marketing".

```
Enable-TransportRule "Sender is a member of marketing"
```

For detailed syntax and parameter information, see Enable-TransportRule and Disable-TransportRule.

How do you know this worked?

To verify that you have successfully enabled or disabled a mail flow rule, use either of the following procedures:

- In the EAC, go to **Mail flow** > **Rules**, and in the list of rules verify the status of the check box in the **On** column.

- In the Exchange Management Shell, use either of the following procedures:

  - Run the following command to see the list of rules and their **State** values:

    ```
    Get-TransportRule
    ```

  - Replace *<RuleName>* with the name of the rule, and run the following command:

    ```
    Get-TransportRule -Identity "<RuleName>" | Format-List Name,State
    ```

Remove mail flow rules

Use the EAC to remove mail flow rules

1. From the EAC, go to **Mail flow** > **Rules**.

2. Select the rule you want to remove from the list, and then click **Delete** (🗑).

Use the Exchange Management Shell to remove mail flow rules

To remove mail flow rules in the Exchange Management Shell, use the following syntax:

```
Remove-TransportRule -Identity "<RuleName>"
```

This example removes the mail flow rule named "Sender is a member of marketing":

```
Remove-TransportRule -Identity "Sender is a member of marketing"
```

For detailed syntax and parameter information, see Remove-TransportRule.

How do you know this worked?

To verify that you have successfully removed a mail flow rule, use either of the following procedures:

- In the EAC, go to **Mail flow** > **Rules**, and verify that the rule you removed is no longer in the list.

- In the Exchange Management Shell, run the following command to verify that the rule you removed is no longer listed:

  ```
  Get-TransportRule
  ```

Import or export mail flow rule collections

You can import a mail flow rule collection that you've previously exported as a backup, or import rules that you've exported from a previous version of Exchange.

Notes:

- You can't import or export mail flow rule collections in the EAC. You can only use the Exchange Management Shell.

- You can't import a mail flow rule collection into Exchange 2010 if that rule collection was exported from Exchange 2013 or later.

Use the Exchange Management Shell to export a mail flow rule collection

1. Run the following command:

   ```
   $File = Export-TransportRuleCollection
   ```

2. Use the following syntax:

   ```
   Set-Content -Path "<OutputFile>" -Value $file.FileData -Encoding Byte
   ```

   For example, to save the exported mail flow rule collection to the file C:\My Documents\Exported Rules.xml, run the following command:

   ```
   Set-Content -Path "C:\My Documents\Exported Rules.xml" -Value $file.FileData -Encoding Byte
   ```

For detailed syntax and parameter information, see Export-TransportRuleCollection.

Use the Exchange Management Shell to import a mail flow rule collection

1. Use the following syntax:

   ```
   [Byte[]]$Data = Get-Content -Path "<OutputFile>" -Encoding Byte -ReadCount 0
   ```

   For example, to import the mail flow rule collection from C:\My Documents\Exported Rules.xml, run the following command:

   ```
   Byte[]]$Data = Get-Content -Path "C:\My Documents\Exported Rules.xml" -Encoding Byte -ReadCount 0
   ```

2. Run the following command:

   ```
   Import-TransportRuleCollection -FileData $Data
   ```

For detailed syntax and parameter information, see Import-TransportRuleCollection.

Need more help?

- Resources for Exchange Server:
  - Mail flow rules in Exchange Server
  - Mail flow rule conditions and exceptions (predicates) in Exchange Server
  - Mail flow rule actions in Exchange Server

Recoverable Items folder in Exchange Server

8/3/2020 • 11 minutes to read • Edit Online

To protect from accidental or malicious deletion and to facilitate discovery efforts commonly undertaken before or during litigation or investigations, Exchange Server and Exchange Online use the Recoverable Items folder. The Recoverable Items folder replaces the feature that was known as *the dumpster* in earlier versions of Exchange. The following Exchange features use the Recoverable Items folder:

- Deleted item retention
- Single item recovery
- In-Place Hold
- Litigation Hold
- Mailbox audit logging
- Calendar logging

Terminology

Knowledge of the following terms will help you understand the content in this topic.

Delete

> Describes when an item is deleted from any folder and placed in the Deleted Items default folder.

Soft delete

> Describes when an item is deleted from the Deleted Items default folder and placed in the Recoverable Items folder. Also describes when an Outlook user deletes an item by pressing Shift+Delete, which bypasses the Deleted Items folder and places the item directly in the Recoverable Items folder.

Hard delete

> Describes when an item is marked to be purged from the mailbox database. This is also known as a *store hard delete*.

Recoverable Items folder

Each user mailbox is divided into two subtrees: the IPM (interpersonal messaging) subtree, which contains the normal, visible folders such as Inbox, Calendar, and Sent Items and the non-IPM subtree, which contains internal data, preferences, and other operational data about the mailbox. The Recoverable Items folder resides in the non-IPM subtree of each mailbox. This subtree isn't visible to users using Outlook, Outlook on the web, or other email clients.

This architectural change provides the following key benefits:

- When a mailbox is moved to another mailbox database, the Recoverable Items folder moves with it.
- The Recoverable Items folder is indexed by Exchange Search and can be discovered by using In-Place eDiscovery.

- The Recoverable Items folder has its own storage quota.
- Exchange can prevent data from being purged from the Recoverable Items folder.
- Exchange can track edits of certain content.

The Recoverable Items folder contains the following subfolders:

- **Deletions**: This subfolder contains all items deleted from the Deleted Items folder. (In Outlook, a user can soft delete an item by pressing Shift+Delete.) This subfolder is exposed to users through the Recover Deleted Items feature in Outlook and Outlook on the web.
- **Versions**: If In-Place Hold or Litigation Hold is enabled, this subfolder contains the original and modified copies of the deleted items. This folder isn't visible to end users.
- **Purges**: If either Litigation Hold or single item recovery is enabled, this subfolder contains all items that are hard deleted. This folder isn't visible to end users.
- **Audits**: If mailbox audit logging is enabled for a mailbox, this subfolder contains the audit log entries. To learn more about mailbox audit logging, see Mailbox audit logging in Exchange Server.
- **DiscoveryHolds**: If In-Place Hold is enabled, this subfolder contains all items that meet the hold query parameters and are hard deleted.
- **Calendar Logging**: This subfolder contains calendar changes that occur within a mailbox. This folder isn't available to users.

The following illustration shows the subfolders in the Recoverable Items folders. It also shows the deleted item retention, single item recovery, and hold workflow processes that are described in the following sections.

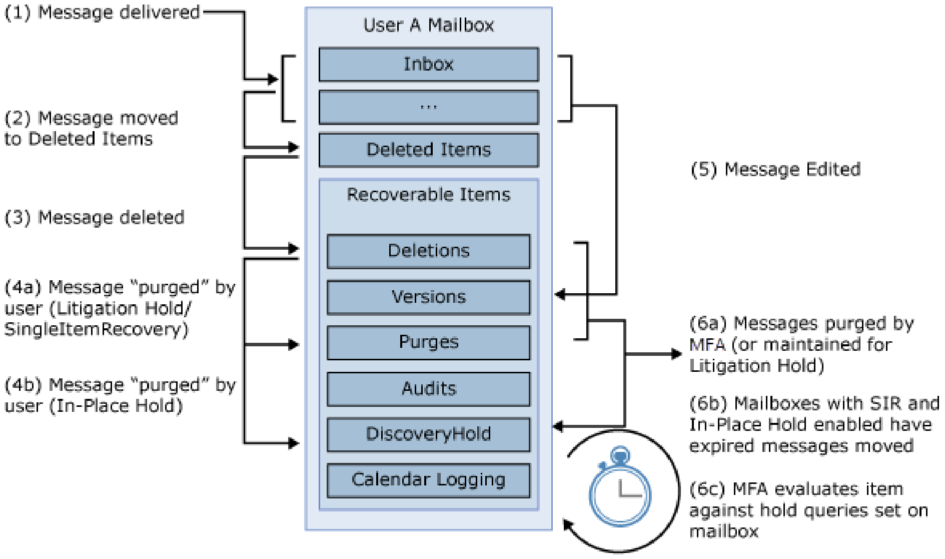
(1) Message delivered
(2) Message moved to Deleted Items
(3) Message deleted
(4a) Message “purged” by user (Litigation Hold/ SingleItemRecovery)
(4b) Message “purged” by user (In-Place Hold)
User A Mailbox
Inbox
...
Deleted Items
Recoverable Items
Deletions
Versions
Purges
Audits
DiscoveryHold
Calendar Logging
(5) Message Edited
(6a) Messages purged by MFA (or maintained for Litigation Hold)
(6b) Mailboxes with SIR and In-Place Hold enabled have expired messages moved
(6c) MFA evaluates item against hold queries set on mailbox

Deleted item retention

An item is considered to be soft deleted in the following cases:

- A user deletes an item or empties all items from the Deleted Items folder.
- A user presses Shift+Delete to delete an item from any other mailbox folder.

Soft-deleted items are moved to the Deletions subfolder of the Recoverable Items folder. This provides an additional layer of protection so users can recover deleted items without requiring Help desk intervention. Users can use the Recover Deleted Items feature in Outlook or Outlook on the web to recover a deleted item. Users can also use this feature to permanently delete an item. For more information, see:

- Recover deleted items in Outlook for Windows

- Recover deleted items or email messages in Outlook on the web

Items remain in the Deletions subfolder until the deleted item retention period is reached. The default deleted item retention period for a mailbox database is 14 days. You can modify this period for a mailbox database or for a specific mailbox. In addition to a deleted item retention period, the Recoverable Items folder is also subject to quotas. To learn more, see Recoverable Items mailbox quotas later in this topic.

After the deleted item retention period expires, the item is moved to the Purges folder and is no longer visible to the user. When the Managed Folder Assistant processes the mailbox, items in the Purges subfolder are purged from the mailbox database.

Single item recovery

If an item is removed from the Deletions subfolder, either by a user purging the item by using the Recover Deleted Items feature or by an automated process such as the Managed Folder Assistant, the item can't be recovered by the user. In previous versions of Exchange, recovering these items required the administrator to restore the mailbox database or a mailbox from backup copies. This process generally delayed recovery by minutes or hours, depending on the backup mechanism used.

In Exchange Server, you can use *single item recovery* to recover items without using backup media to restore the mailbox databases. This results in considerably shorter recovery periods. When the Managed Folder Assistant processes the Recoverable Items folder for a mailbox that has single item recovery enabled, any item in the Purges subfolder isn't purged if the deleted item retention period hasn't expired for that item.

The following table lists the contents of and actions that can be performed in the Recoverable Items folder if single item recovery is enabled.

Recoverable Items folder and single item recovery

STATE OF SINGLE ITEM RECOVERY	RECOVERABLE ITEMS FOLDER CONTAINS SOFT-DELETED ITEMS	RECOVERABLE ITEMS FOLDER CONTAINS HARD-DELETED ITEMS	USERS CAN PURGE ITEMS FROM THE RECOVERABLE ITEMS FOLDER	MANAGED FOLDER ASSISTANT AUTOMATICALLY PURGES ITEMS FROM THE RECOVERABLE ITEMS FOLDER
Enabled	Yes	Yes	No	Yes. By default, all items are purged after 14 days, with the exception of calendar items, which are purged after 120 days.
Disabled	Yes	No	Yes	Yes. By default, all items are purged after 14 days, with the exception of calendar items, which are purged after 120 days. If the Recoverable Items warning quota is reached before the deleted item retention period elapses, messages are deleted in first in, first out (FIFO) order.

In Exchange Server, single item recovery isn't enabled by default for new mailboxes or mailboxes moved from a previous version of Exchange. You need to use the Exchange Management Shell to enable single item recovery for a mailbox, and then configure or modify the deleted item retention period. For details about how to perform a single item recovery, see Recover deleted messages in a user's mailbox.

In-Place Hold and Litigation Hold

In Exchange Server and Exchange Online, discovery managers can use In-Place eDiscovery with delegated Discovery Management role group permissions to perform eDiscovery searches of mailbox content. In Exchange Server and Exchange Online, you can use In-Place Hold to preserve mailbox items that match query parameters and protect the items from deletion by users or automated processes. You can also use Litigation Hold to preserve all items in user mailboxes and protect the items from deletion by users or automated processes.

Putting a mailbox on In-Place Hold or Litigation Hold stops the Managed Folder Assistant from automatically purging messages from the DiscoveryHolds and Purges subfolders. Additionally, copy-on-write page protection is also enabled for the mailbox. Copy-on-write page protection creates a copy of the original item before any modifications are written to the Exchange store. After the mailbox is removed from hold, the Managed Folder Assistant resumes automated purging.

> **NOTE**
>
> If you put a mailbox on both In-Place Hold and Litigation Hold, Litigation Hold takes preference because this puts the entire mailbox on hold.

The following table lists the contents of and actions that can be performed in the Recoverable Items folder if Litigation Hold is enabled.

Recoverable Items folder and holds

STATE OF HOLD	RECOVERABLE ITEMS FOLDER CONTAINS SOFT-DELETED ITEMS	RECOVERABLE ITEMS FOLDER CONTAINS MODIFIED AND HARD-DELETED ITEMS	USERS CAN PURGE ITEMS FROM THE RECOVERABLE ITEMS FOLDER	MANAGED FOLDER ASSISTANT AUTOMATICALLY PURGES ITEMS FROM THE RECOVERABLE ITEMS FOLDER
Enabled	Yes	Yes	No	No
Disabled	Yes	No	Yes	Yes

To learn more about In-Place eDiscovery, In-Place Hold, and Litigation Hold, see the following topics:

- In-Place eDiscovery in Exchange Server
- In-Place Hold and Litigation Hold in Exchange Server

Copy-on-write page protection and modified items

If a user who is placed on In-Place Hold or Litigation Hold modifies specific properties of a mailbox item, a copy of the original mailbox item is created before the changed item is written. The original copy is saved in the Versions subfolder. This process is known as *copy-on-write page protection*. Copy-on-write page protection applies to items residing in any mailbox folder. The Versions subfolder isn't visible to users.

The following table lists the message properties that trigger copy-on-write page protection.

Properties that trigger copy-on-write page protection

ITEM TYPE	PROPERTIES THAT TRIGGER COPY-ON-WRITE PAGE PROTECTION
Messages (IPM.Note*) Posts (IPM.Post*)	Subject Body Attachments Senders and recipients Sent and received dates
Items other than messages and posts	Any change to a visible property, except the following: • Item location (when an item is moved between folders) • Item status change (read or unread) • Changes to a retention tag applied to an item
Items in the Drafts default folder	None. Items in the Drafts folder are exempt from copy-on-write page protection.

> **IMPORTANT**
>
> Copy-on-write page protection doesn't save a version of the meeting when a meeting organizer receives responses from attendees and the meeting's tracking information is updated. Also, changes to RSS feeds aren't captured by copy-on-write page protection.

When a mailbox is no longer on In-Place Hold or litigation hold, copies of modified items stored in the Versions folder are removed.

Recoverable Items mailbox quotas

When an item is moved to the Recoverable Items folder, its size is deducted from the mailbox quota and added to the size of the Recoverable Items folder. In Exchange Server, mailbox databases have a configurable Recoverable Items warning quota (*soft limit*) of 20 GB and a Recoverable Items quota (*hard limit*) of 30 GB. By default, these limits are inherited by all mailboxes in the database. However, you can configure individual mailboxes with different quotas. To learn more, see Configure Deleted Item retention and Recoverable Items quotas.

In Exchange Online, the default limits for the Recoverable Items quota are the same as Exchange Server: a soft limit of 20 GB and a hard limit of 30 GB. However, the quotas for the Recoverable Items folder are automatically increased to 90 GB and 100 GB, respectively, when you place a mailbox on Litigation Hold or In-Place Hold.

When the Recoverable Items folder for a mailbox reaches the Recoverable Items quota, no more items can be stored in the folder. This impacts mailbox functionality in the following ways:

- Mailbox users can't delete items.
- The Managed Folder Assistant can't delete items based on retention tag or managed folder settings.
- For mailboxes that have single item recovery, In-Place Hold or Litigation Hold enabled, the copy-on-write page protection process can't maintain versions of items edited by the user.
- For mailboxes that have mailbox audit logging enabled, no mailbox audit log entries can be saved in the Audits subfolder.

For mailboxes that aren't placed on In-Place Hold or Litigation Hold, the Managed Folder Assistant automatically purges items from the Recoverable Items folder when the deleted item retention period expires. If the folder reaches the Recoverable Items warning quota, the assistant automatically purges items in first-in-first-out order.

When the Recoverable Items folder reaches the soft and hard limit defaults, you are notified by means of the event log and a Microsoft System Center Operations Manager alert. This alert fires when the Recoverable Items folder first reaches the soft and hard limit defaults, and then once daily afterward.

The following table lists the events logged when the Recoverable Items folder reaches the soft and hard limit defaults.

Recoverable Items quota warnings and errors

EVENT ID	TYPE	SOURCE	MESSAGE
10024	Warning	MSExchangeIS Mailbox Store	The mailbox for *<mailbox user>* (*<GUID>*) has exceeded the Recoverable Items Warning Quota. Please remove items from Recoverable Items or increase the Recoverable Items Warning Quota and Recoverable Items Quota. If the Recoverable Items Quota is exceeded, the user will be unable to delete items from the mailbox.
10023	Error	MSExchangeIS Mailbox Store	The mailbox for *<mailbox user>* (*<GUID>*) has exceeded the maximum Recoverable Items Quota. Items cannot be deleted from this mailbox. The mailbox owner should be notified about the condition of the mailbox as soon as possible. Please remove items from Recoverable Items or increase the Recoverable Items Quota to restore functionality.

EVENT ID	TYPE	SOURCE	MESSAGE
10023	Warning	MSExchangeMailboxAssistants	The mailbox: *<mailbox user>* Recoverable Items size has exceeded the warning quota limit. Items were deleted from Recoverable Items folders to prevent mailbox outage. Recoverable Items Warning Quota: 20 GB (21,474,836,480 bytes) Original Recoverable Items size: 21475005311 Current Recoverable Items size: 21474823820 Folder stats: - Folders processed: RecoverableItemsRoot, RecoverableItemsVersions, RecoverableItemsPurges, RecoverableItemsDeletions - Original folder sizes: 21391661934, 55190914, 1987247, 26157788 (item counts: 276828, 400, 84, 646) - Current folder sizes: 21391480443, 55190914, 1987247, 26157788 (item counts: 276817, 400, 84, 646)

If the mailbox is placed on In-Place Hold or Litigation Hold, copy-on-write page protection can't maintain versions of modified items. To maintain versions of modified items, you need to reduce the size of the Recoverable Items folder. You can use the Search-Mailbox cmdlet to copy messages from the Recoverable Items folder of a mailbox to a discovery mailbox, and then delete the items from the mailbox. Alternatively, you can also raise the Recoverable Items quota for the mailbox. For details, see Clean up or delete items from the Recoverable Items folder.

More information

- Copy-on-write is only enabled when a mailbox is on In-Place Hold or Litigation Hold.
- If users need to recover deleted items from the Recoverable Items folder, point them to the following topics:
  - Restore deleted items in Outlook for Windows
  - Recover deleted items or email in Outlook on the web

Clean up or delete items from the Recoverable Items folder

8/3/2020 • 8 minutes to read • Edit Online

The Recoverable Items folder (known in earlier versions of Exchange as *the dumpster*) exists to protect from accidental or malicious deletions and to facilitate discovery efforts commonly undertaken before or during litigation or investigations.

How you clean up a user's Recoverable Items folder depends on whether the mailbox is placed on In-Place Hold or Litigation Hold, or had single item recovery enabled:

- If a mailbox isn't placed on In-Place Hold or Litigation Hold or doesn't have single item recovery enabled, you can simply delete items from the Recoverable Items folder. After items are deleted, you can't use single item recovery to recover them.

- If the mailbox is placed on In-Place Hold or Litigation Hold or has single item recovery enabled, you'll want to preserve the mailbox data until the hold is removed or single item recovery is disabled. In this case, you need to perform more detailed steps to clean up the Recoverable Items folder.

To learn more about In-Place Hold and Litigation Hold, see In-Place Hold and Litigation Hold in Exchange Server. To learn more about single item recovery, see "Single Item Recovery" in Recoverable Items folder in Exchange Server.

What do you need to know before you begin?

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Delete mailbox content" entry in the Messaging policy and compliance permissions in Exchange Server topic.

- Because incorrectly cleaning up the Recoverable Items folder can result in data loss, it's important that you're familiar with the Recoverable Items folder and the impact of removing its contents. Before performing this procedure, we recommend that you review the information in Recoverable Items folder in Exchange Server.

- You can't use the Exchange admin center (EAC) to perform these procedures. You must use the Exchange Management Shell. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to delete items from the Recoverable Items folder for mailboxes that aren't placed on hold or don't have single item recovery enabled

This example permanently deletes items from the user Gurinder Singh's Recoverable Items folder and also copies the items to the GurinderSingh-RecoverableItems folder in the Discovery Search Mailbox (a discovery mailbox created by Exchange Setup).

```
Search-Mailbox -Identity "Gurinder Singh" -SearchDumpsterOnly -TargetMailbox "Discovery Search Mailbox" -
TargetFolder "GurinderSingh-RecoverableItems" -DeleteContent
```

NOTE

To delete items from the mailbox without copying them to another mailbox, use the preceding command without the *TargetMailbox* and *TargetFolder* parameters.

For detailed syntax and parameter information, see Search-Mailbox.

Use the Exchange Management Shell to clean up the Recoverable Items folder for mailboxes that are placed on hold or have single item recovery enabled

You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Delete mailbox content" entry in the Messaging policy and compliance permissions in Exchange Server topic.

If a mailbox reaches its Recoverable Items quota, we recommend that you raise the quota and not delete items from the folder. You can also monitor events in the Application log related to the Recoverable Items warning quota and take necessary actions such as raising the quota or investigating the growth of the Recoverable Items folder for mailboxes that reach the warning quota.

If storage constraints or similar issues prevent you from raising the Recoverable Items quota, we recommend that you first copy data from the user's Recoverable Items folder to another mailbox before you delete messages. If you're deleting items due to storage constraints on one volume, you can copy items to a mailbox located on a volume that has adequate storage.

This procedure copies items from Gurinder Singh's Recoverable Items folder to the GurinderSingh-RecoverableItems folder in the Discovery Search Mailbox. Before you copy and delete items from the Recoverable Items folder, you should first perform several steps to make sure items aren't deleted from the Recoverable Items folder. After you copy items to a discovery or backup mailbox and clean up the folder, you can revert to the mailbox's previous settings.

1. Retrieve the following quota settings. Be sure to note the values so you can revert to these settings after cleaning up the Recoverable Items folder:

    - *RecoverableItemsQuota*
    - *RecoverableItemsWarningQuota*
    - *ProhibitSendQuota*
    - *ProhibitSendReceiveQuota*
    - *UseDatabaseQuotaDefaults*
    - *RetainDeletedItemsFor*
    - *UseDatabaseRetentionDefaults*

> **NOTE**
>
> If the *UseDatabaseQuotaDefaults* parameter is set to `$true`, the previous quota settings aren't applied. The corresponding quota settings configured on the mailbox database are applied, even if individual mailbox settings are populated.

```
Get-Mailbox "Gurinder Singh" | Format-List *Quota*,RetainDeletedItemsFor,UseDatabaseRetentionDefaults
```

2. Retrieve the mailbox access settings for the mailbox. Be sure to note these settings for later.

```
Get-CASMailbox "Gurinder Singh" | Format-List EwsEnabled, ActiveSyncEnabled, MAPIEnabled, OWAEnabled,
ImapEnabled, PopEnabled
```

3. Retrieve the current size of the Recoverable Items folder. Note the size so you can raise the quotas in Step 6.

```
Get-MailboxFolderStatistics "Gurinder Singh" -FolderScope RecoverableItems | Format-List
Name,FolderAndSubfolderSize
```

4. Disable client access to the mailbox to make sure no changes can be made to mailbox data for the duration of this procedure.

```
Set-CASMailbox "Gurinder Singh" -EwsEnabled $false -ActiveSyncEnabled $false -MAPIEnabled $false -
OWAEnabled $false -ImapEnabled $false -PopEnabled $false
```

5. To make sure no items are deleted from the Recoverable Items folder, increase the Recoverable Items quota, increase the Recoverable Items warning quota, and set the deleted item retention period to a value higher than the current size of the user's Recoverable Items folder. This is particularly important for preserving messages for mailboxes placed on In-Place Hold or Litigation Hold. We recommend raising these settings to twice their current size.

```
Set-Mailbox "Gurinder Singh" -RecoverableItemsQuota 50Gb -RecoverableItemsWarningQuota 50Gb -
RetainDeletedItemsFor 3650 -ProhibitSendQuota 50Gb -ProhibitSendReceiveQuota 50Gb -
UseDatabaseQuotaDefaults $false -UseDatabaseRetentionDefaults $false
```

6. Stop the Microsoft Exchange Mailbox Assistants service and prevent it from starting on the Mailbox server by running the following commands:

```
Stop-Service MSExchangeMailboxAssistants; Set-Service MSExchangeMailboxAssistants -StartupType Disabled
```

The effect of this command is to stop the Managed Folder Assistant on the Mailbox server.

> **IMPORTANT**
>
> If the mailbox resides on a mailbox database in a database availability group (DAG), you must disable the Managed Folder Assistant on each DAG member that hosts a copy of the database. If the database fails over to another server, this prevents the Managed Folder Assistant on that server from deleting mailbox data.

7. Disable single item recovery and remove the mailbox from Litigation Hold.

```
Set-Mailbox "Gurinder Singh" -SingleItemRecoveryEnabled $false -LitigationHoldEnabled $false
```

> **IMPORTANT**
>
> After you run this command, it may take up to one hour to disable single item recovery or Litigation Hold. We recommend that you perform the next step only after this period has elapsed.

8. Copy items from the Recoverable Items folder to a folder in the Discovery Search Mailbox and delete the contents from the source mailbox.

   ```
   Search-Mailbox -Identity "Gurinder Singh" -SearchDumpsterOnly -TargetMailbox "Discovery Search Mailbox"
   -TargetFolder "GurinderSingh-RecoverableItems" -DeleteContent
   ```

   If you need to delete only messages that match specified conditions, use the *SearchQuery* parameter to specify the conditions. This example deletes messages that have the string "Your bank statement" in the `Subject` field.

   ```
   Search-Mailbox -Identity "Gurinder Singh" -SearchQuery "Subject:'Your bank statement'" -
   SearchDumpsterOnly -TargetMailbox "Discovery Search Mailbox" -TargetFolder "GurinderSingh-
   RecoverableItems" -DeleteContent
   ```

   > **NOTE**
   >
   > It isn't required to copy items to the Discovery Search Mailbox. You can copy messages to any mailbox. However, to prevent access to potentially sensitive mailbox data, we recommend copying messages to a mailbox that has access restricted to authorized records managers. By default, access to the default Discovery Search Mailbox is restricted to members of the Discovery Management role group. For details, see In-Place eDiscovery in Exchange Server.

9. If the mailbox was placed on Litigation Hold or had single item recovery enabled earlier, enable these features again.

   ```
   Set-Mailbox "Gurinder Singh" -SingleItemRecoveryEnabled $true -LitigationHoldEnabled $true
   ```

   > **IMPORTANT**
   >
   > After you run this command, it may take up to one hour to enable single item recovery or Litigation Hold. We recommend that you enable the Managed Folder Assistant and allow client access (Steps 11 and 12) only after this period has elapsed.

10. Revert the following quotas to the values noted in Step 1:

    - *RecoverableItemsQuota*
    - *RecoverableItemsWarningQuota*
    - *ProhibitSendQuota*
    - *ProhibitSendReceiveQuota*
    - *UseDatabaseQuotaDefaults*
    - *RetainDeletedItemsFor*

   - *UseDatabaseRetentionDefaults*

   In this example, the mailbox is removed from retention hold, the deleted item retention period is reset to the default value of 14 days, and the Recoverable Items quota is configured to use the same value as the mailbox database. If the values you noted in Step 1 are different, you must use the preceding parameters to specify each value and set the *UseDatabaseQuotaDefaults* parameter to `$false`. If the *RetainDeletedItemsFor and UseDatabaseRetentionDefaults* parameters were previously set to a different value, you must also revert them to the values noted in Step 1.

   ```
   Set-Mailbox "Gurinder Singh" -RetentionHoldEnabled $false -RetainDeletedItemsFor 14 -
   RecoverableItemsQuota unlimited -UseDatabaseQuotaDefaults $true
   ```

11. Configure the Microsoft Exchange Mailbox Assistants service to start automatically and start it on the Mailbox server by running the following commands:

    ```
    Set-Service MSExchangeMailboxAssistants -StartupType Automatic; Start-Service
    MSExchangeMailboxAssistants
    ```

12. Enable client access to the mailbox by running the following command:

    ```
    Set-CASMailbox -ActiveSyncEnabled $true -EwsEnabled $true -MAPIEnabled $true -OWAEnabled $true -
    ImapEnabled $true -PopEnabled $true
    ```

For detailed syntax and parameter information, see the following topics:

- Get-Mailbox
- Get-CASMailbox
- Get-MailboxFolderStatistics
- Set-CASMailbox
- Set-Mailbox
- Search-Mailbox

How do you know this worked?

To verify that you have successfully cleaned up the Recoverable Items folder of a mailbox, use Get-MailboxFolderStatistics cmdlet the check the size of the Recoverable Items folder.

This example retrieves the size of the Recoverable Items folder and its subfolders and an item count in the folder and each subfolder.

```
Get-MailboxFolderStatistics -Identity "Gurinder Singh" -FolderScope RecoverableItems | Format-Table
Name,FolderAndSubfolderSize,ItemsInFolderAndSubfolders -Auto
```

S/MIME for message signing and encryption

8/3/2020 • 3 minutes to read • Edit Online

As an administrator in Exchange Server, you can enable Secure/Multipurpose Internet Mail Extensions (S/MIME) for your organization. S/MIME is a widely accepted method (more precisely, a protocol) for sending digitally signed and encrypted messages. S/MIME allows you to encrypt emails and digitally sign them. When you use S/MIME, it helps the people who receive the message by:

- Ensuring that the message in their inbox is the exact message that started with the sender.
- Ensuring that the message came from the specific sender and not from someone pretending to be the sender.

To do this, S/MIME provides for cryptographic security services such as authentication, message integrity, and non-repudiation of origin (using digital signatures). S/MIME also helps enhance privacy and data security (using encryption) for electronic messaging.

S/MIME requires a certificate and publishing infrastructure that is often used in business-to-business and business-to-consumer situations. The user controls the cryptographic keys in S/MIME and can choose whether to use them for each message they send. Email programs such as Outlook search a trusted root certificate authority location to perform digital signing and verification of the signature.

For a more complete background about the history and architecture of S/MIME in the context of email, see Understanding S/MIME.

Supported scenarios and technical considerations for S/MIME

You can set up S/MIME to work with any of the following end points:

- Outlook 2010 or later
- Outlook on the web (formerly known as Outlook Web App)
- Exchange ActiveSync (EAS)

The steps that you follow to set up S/MIME with each of these endpoints are slightly different. Generally, you need to complete these steps:

1. Install a Windows-based Certification Authority and set up a public key infrastructure to issue S/MIME certificates. Certificates issued by third-party certificate providers are supported. For details, see Server Certificate Deployment Overview.
2. Publish the user certificate in an on-premises Active Directory Domain Services (AD DS) account in the **UserSMIMECertificate** and/or **UserCertificate** attributes. Your AD DS needs to be located on computers at a physical location that you control and not at a remote facility or cloud-based service somewhere on the Internet. For more information about AD DS, see Active Directory Domain Services Overview.
3. Set up a virtual certificate collection in order to validate S/MIME. This information is used by Outlook on the web when validating the signature of an email and ensuring that it was signed by a trusted certificate.
4. Set up the Outlook or EAS end point to use S/MIME.

Set up S/MIME with Outlook on the web

Setting up S/MIME with Outlook on the web involves these key steps:

1. S/MIME settings for Outlook on the web in Exchange Server.
2. Set up Virtual Certificate Collection to Validate S/MIME

For information about how to send an S/MIME encrypted message in Outlook on the web, see Encrypt messages by using S/MIME in Outlook on the web.

Related message encryption technologies

A variety of encryption technologies work together to provide protection for messages at rest and in transit. S/MIME can work simultaneously with the following technologies but isn't dependent on them:

- **Transport Layer Security (TLS)**: Encrypts the tunnel or the route between email servers in order to help prevent snooping and eavesdropping, and encrypts the connection between email clients and servers.

> **NOTE**
>
> Secure Sockets Layer (SSL) is being replaced by Transport Layer Security (TLS) as the protocol that's used to encrypt data sent between computer systems. They're so closely related that the terms "SSL" and "TLS" (without versions) are often used interchangeably. Because of this similarity, references to "SSL" in Exchange topics, the Exchange admin center, and the Exchange Management Shell have often been used to encompass both the SSL and TLS protocols. Typically, "SSL" refers to the actual SSL protocol only when a version is also provided (for example, SSL 3.0). To find out why you should disable the SSL protocol and switch to TLS, check out Protecting you against the SSL 3.0 vulnerability.

- **BitLocker**: Encrypts the data on a hard drive in a datacenter so that if someone gets unauthorized access, they can't read it. For more information, see BitLocker: How to deploy on Windows Server 2012 and later

S/MIME settings for Outlook on the web in Exchange Server

8/3/2020 • 2 minutes to read • Edit Online

As an Exchange administrator, you can set up Outlook on the web (formerly known as Outlook Web App) to allow sending and receiving S/MIME-protected messages. Use the **Get-SmimeConfig** and **Set-SmimeConfig** cmdlets to view and manage this feature in the Exchange Management Shell. To open the Exchange Management Shell, see Open the Exchange Management Shell.

For detailed syntax and parameter information, see Get-SmimeConfig and Set-SmimeConfig.

S/MIME in Outlook on the web in Exchange 2016 is only supported in Internet Explorer.

S/MIME in Outlook on the web in Exchange 2019 is supported in Internet Explorer, Edge (non-Chromium version), and Google Chrome.

Considerations for Chrome

To use S/MIME in Outlook on the web in the Google Chrome web browser, you (or another administrator) must set and configure the Chromium policy named **ExtensionInstallForcelist** to install the Microsoft S/MIME extension in Chrome. The policy value is `maafgiompdekodanheihhgilkjchcakm;https://outlook.office.com/owa/SmimeCrxUpdate.ashx`. And note that applying this policy requires domain-joined computers, so using S/MIME in Chrome effectively requires domain-joined computers.

For details about the **ExtensionInstallForcelist** policy, see ExtensionInstallForcelist.

This step is a prerequisite for using Chrome; it does not replace the S/MIME control that's installed by users. Users are prompted to download and install the S/MIME control in Outlook on the web during their first use of S/MIME. Or, users can proactively go to **S/MIME** in their Outlook on the web settings to get the download link for the control.

High availability and site resilience

8/3/2020 • 12 minutes to read • Edit Online

You can protect your Exchange Server mailbox databases and the data they contain by configuring your Exchange servers and databases for high availability and site resilience. Exchange Server minimizes the cost and complexity of deploying a highly available and resilient messaging solution while providing high levels of service and data availability and support for very large mailboxes.

Exchange Server enables customers of all sizes and in all segments to economically deploy a messaging continuity service in their organization by building on the native replication capabilities and high availability architecture introduced in Exchange 2010. For a list of changes since Exchange 2010, see Changes to high availability and site resilience over previous versions.

Key terminology

The following key terms are important to understand high availability or site resilience:

Active Manager

> An internal Exchange component which runs inside the Microsoft Exchange Replication service that's responsible for failure monitoring and corrective action through failover within a database availability group (DAG).

AutoDatabaseMountDial

> A property setting of a Mailbox server that determines whether a passive database copy will automatically mount as the new active copy, based on the number of log files missing by the copy being mounted.

Continuous replication - block mode

> In block mode, as each update is written to the active database copy's active log buffer, it's also shipped to a log buffer on each of the passive mailbox copies in block mode. When the log buffer is full, each database copy builds, inspects, and creates the next log file in the generation sequence.

Continuous replication - file mode

> In file mode, closed transaction log files are pushed from the active database copy to one or more passive database copies.

Database availability group

> A group of up to 16 Exchange servers that hosts a set of replicated databases.

Database mobility

> The ability of an Exchange Server mailbox database to be replicated to and mounted on other Exchange servers.

Datacenter

> Typically this refers to an Active Directory site; however, it can also refer to a physical site. In the context of this documentation, datacenter equals Active Directory site.

Datacenter Activation Coordination mode

> A property of the DAG setting that, when enabled, forces the Microsoft Exchange Replication service to acquire permission to mount databases at startup.

Disaster recovery

> Any process used to manually recover from a failure. This can be a failure that affects a single item, or it can be a failure that affects an entire physical location.

Exchange third-party replication API

> An Exchange-provided API that enables use of third-party synchronous replication for a DAG instead of continuous replication.

High availability

> A solution that provides service availability, data availability, and automatic recovery from failures that affect the service or data (such as a network, storage, or server failure).

Incremental deployment

> The ability to deploy high availability and site resilience after Exchange Server is installed.

Lagged mailbox database copy

> A passive mailbox database copy that has a log replay lag time greater than zero.

Mailbox database copy

> A mailbox database (.edb file and logs), which is either active or passive.

Mailbox resiliency

> The name of a unified high availability and site resilience solution in Exchange Server.

Managed availability

> A set of internal processes made up of probes, monitors, and responders that incorporate monitoring and high availability across all server roles and all protocols.

**over* (pronounced "star over")

> Short for *switchovers* and *failovers*. A switchover is a manual activation of one or more database copies. A failover is an automatic activation of one or more database copies after a failure.

Safety Net

> Formerly known as transport dumpster, this is a feature of the transport service that stores a copy of all messages for *X* days. The default setting is 2 days.

Shadow redundancy

> A transport server feature that provides redundancy for messages for the entire time they're in transit.

Site resilience

> A configuration that extends the messaging infrastructure to multiple Active Directory sites to provide operational continuity for the messaging system in the event of a failure affecting one of the sites.

Database availability groups

A DAG is the base component of the high availability and site resilience framework built into Exchange Server. A DAG is a group of up to 16 Exchange servers that hosts a set of databases and provides automatic, database-level recovery from failures that affect individual databases, networks, or servers. Any server in a DAG can host a copy of a mailbox database from any other server in the DAG. When a server is added to a DAG, it works with the other servers in the DAG to provide automatic recovery from failures that affect mailbox databases, such as a disk failure or server failure. For more information about DAGs, see Database availability groups.

Mailbox database copies

The high availability and site resilience features used first introduced in Exchange 2010 are used in Exchange Server to create and maintain database copies. Exchange Server also leverages the concept of database mobility, which is Exchange-managed database-level failovers.

Database mobility disconnects databases from servers and adds support for up to 16 copies of a single database. It also provides a native experience for creating copies of a database.

Setting a database copy as the active mailbox database is known as a *switchover*. When a failure affecting a database or access to a database occurs and a new database becomes the active copy, this process is known as a *failover*. This process also refers to a server failure in which one or more servers bring online the databases previously online on the failed server. When either a switchover or failover occurs, other Exchange servers become aware of the switchover almost immediately and redirect client and messaging traffic to the new active database.

For example, if an active database in a DAG fails because of an underlying storage failure, Active Manager will automatically recover by failing over to a database copy on another server in the DAG. In Exchange Server, managed availability provides behaviors to recover from loss of protocol access to a database, including recycling application worker pools, restarting services and servers, and initiating database failovers.

For more information about mailbox database copies, see Mailbox database copies.

Active Manager

Exchange Server leverages Active Manager to manage the database and database copy health, status, continuous replication, and other aspects of high availability. For more information about Active Manager, see Active Manager.

Site resilience

In Exchange 2010, you could deploy a DAG across two datacenters and host the witness in a third datacenter and enable failover for the Mailbox server role for either datacenter. But you didn't get failover for the solution itself because the namespace still needed to be manually changed for the non-Mailbox server roles.

In Exchange 2016 and Exchange 2019, the namespace doesn't need to move with the DAG. Exchange leverages fault tolerance built into the namespace through multiple IP addresses, load balancing (and if need be, the ability to take servers in and out of service). Modern HTTP clients work with this redundancy automatically. The HTTP stack can accept multiple IP addresses for a fully qualified domain name (FQDN), and if the first IP address it tries fails hard (that is, it can't connect), it will try the next IP address in the list. In a soft failure (connection is lost after the session is established, perhaps due to an intermittent failure in the service where, for example, a device is dropping packets and needs to be taken out of service), the user might need to refresh their browser.

This means the namespace is no longer a single point of failure as it was in Exchange 2010. In Exchange 2010, perhaps the biggest single point of failure in the messaging system is the FQDN that you give to users because it tells the user where to go. In the Exchange 2010 paradigm, changing where that FQDN goes isn't easy because you have to change DNS, and then handle DNS latency, which in some parts of the world is challenging. And you have name caches in browsers that are typically about 30 minutes or more that also have to be handled.

In Exchange Server, clients have more than one place to go. Almost all the client access protocols in Exchange Server are HTTP based. Examples include Outlook, EAS, EWS, Outlook on the web, and EAC). All supported HTTP clients have the ability to use multiple IP addresses, thereby providing failover on the client side. You can configure DNS to hand multiple IP addresses to a client during name resolution. The client asks for mail.contoso.com and gets back two IP addresses, or four IP addresses, for example. However many IP addresses the client gets back will be used reliably by the client. This makes the client a lot better off because if one of the IP addresses fails, the client has one or more alternative IP addresses to try to connect to. If a client tries one and it fails, it waits about 20 seconds and then tries the next one in the list. Thus, if you lose the VIP for the Client Access service array, recovery for the clients happens automatically, and in about 21 seconds.

The benefits include the following:

- In Exchange Server, if you lose the load balancer in your primary site, you simply turn it off (or maybe turn off the VIP) and repair or replace it. Clients that aren't already using the VIP in the secondary datacenter will automatically fail over to the secondary VIP without any change of namespace, and without any change in DNS. Not only does that mean you no longer have to perform a switchover, but it also means that all of the time normally associated with a datacenter switchover recovery isn't spent. In Exchange 2010, you had to handle DNS latency (hence, the recommendation to set the Time to Live (TTL) to 5 minutes, and the introduction of the failback URL). In Exchange 2016 and Exchange 2019, you don't need to do that because you get fast failover (20 seconds) of the namespace between VIPs (datacenters).

- Because you can fail over the namespace between datacenters, all that's needed to achieve a datacenter failover is a mechanism for failover of the Mailbox server role across datacenters. To get automatic failover for the DAG, you simply architect a solution where the DAG is evenly split between two datacenters, and then place the witness server in a third location so that it can be arbitrated by DAG members in either datacenter, regardless of the state of the network between the datacenters that contain the DAG members. If you only have two datacenters and a third physical location isn't available, you can place the witness server on a Microsoft Azure virtual machine. See Using a Microsoft Azure VM as a DAG witness server for more information.

- In this scenario, the administrator's efforts are geared toward simply fixing the problem, and not spent restoring service. You simply fix the thing that failed; while service has been running and data integrity has been maintained. The urgency and stress level you feel when fixing a broken device is nothing like the urgency and stress you feel when you're working to restore service. It's better for the end user, and less stressful for the administrator.

You can allow failover to occur without having to perform switchbacks (sometimes mistakenly referred to as failbacks). If you lose servers in your primary datacenter, resulting in a 20 second interruption for clients, you might not even care about failing back. At this point, your primary concern would be fixing the core issue (for example, replacing the failed load balancer). After it's back online and functioning, some clients will start using it, and other clients might remain operational through the second datacenter.

Exchange Server also provides functionality that enables administrators to deal with intermittent failures. An intermittent failure is where, for example, the initial TCP connection can be made, but nothing happens afterward. An intermittent failure requires some sort of extra administrative action to be taken because it might be the result of a replacement device being put into service. While this repair process is occurring, the device might be powered on and accepting some requests, but not really ready to service clients until the necessary configuration steps are performed. In this scenario, the administrator can perform a namespace switchover by simply removing the VIP for the device being replaced from DNS. Then during that service period, no clients will be trying to connect to it. After the replacement process has completed, the administrator can add the VIP back to DNS, and clients will eventually start using it.

For details about planning and deploying site resilience, see Plan for high availability and site resilience and Deploying high availability and site resilience.

Third-party replication API

Exchange Server includes a third-party replication API that enables organizations to use third-party synchronous replication solutions instead of the built-in continuous replication feature. Microsoft supports third-party solutions that use this API, provided that the solution provides the necessary functionality to replace all native continuous replication functionality that's disabled as a result of using the API. Solutions are supported only when the API is used within a DAG to manage and activate mailbox database copies. Use of the API outside of these boundaries isn't supported. In addition, the solution must meet the applicable Windows hardware support requirements. (Test validation isn't required for support.)

When deploying a solution that uses the built-in third-party replication API, be aware that the solution vendor is responsible for primary support of the solution. Microsoft supports Exchange data for both replicated and non-replicated solutions. Solutions that use data replication must adhere to the Microsoft support policy for data replication. In addition, solutions that utilize the Windows Failover Cluster resource model must meet Windows cluster supportability requirements as described in Microsoft Knowledge Base article 943984, The Microsoft Support Policy for Windows Server 2008 or Windows Server 2008 R2 Failover Clusters or The Microsoft Support Policy for Windows Server 2012 Failover Clusters.

Microsoft's backup and restore support policy for deployments that use third-party replication API-based solutions is the same as for native continuous replication deployments.

If you're a partner seeking information about the third-party API, contact your Microsoft representative.

High availability and site resilience documentation

The following table contains links to topics that will help you learn about and manage DAGs, mailbox database copies, and backup and restore for Exchange Server.

TOPIC	DESCRIPTION
Database availability groups	Learn about DAGs, Active Manager, Datacenter Activation Coordination (DAC) mode, and mailbox database copies.
Plan for high availability and site resilience	Learn about the general, hardware, network, software, witness server, and other requirements and best practices for DAGs.
Deploying high availability and site resilience	Explore an example deployment scenario for deploying and configuring DAGs.
Managing high availability and site resilience	Learn about DAG management tasks, switchovers and failovers, and maintenance mode.

TOPIC	DESCRIPTION
Monitor database availability groups	Learn about the built-in cmdlets and scripts for monitoring DAGs and database copies.
Backup, restore, and disaster recovery	Learn about backing up and restoring Exchange databases, recovery databases, and server recovery.

Changes to high availability and site resilience over previous versions of Exchange Server

8/3/2020 • 23 minutes to read • Edit Online

Exchange Server 2013 and later uses DAGs and mailbox database copies (along with other features such as single item recovery, retention policies, and lagged database copies) to provide high availability, site resilience, and Exchange native data protection. The high availability platform, Exchange Information Store and Extensible Storage Engine (ESE) have all been enhanced since Exchange 2010 to provide availability and less complex management, and to reduce costs. These enhancements include:

- **Reduction in IOPS**: This enables you to leverage larger disks in terms of capacity and IOPS as efficiently as possible.

- **Managed availability**: With managed availability, internal monitoring and recovery-oriented features are tightly integrated to help prevent failures, proactively restore services, and initiate server failovers automatically or alert administrators to take action. The focus is on monitoring and managing the end-user experience rather than just server and component uptime to help keep the service continuously available.

- **Managed Store**: The Managed Store is the name of the rewritten Information Store processes in Exchange 2013 or later. The Managed Store is written in C# and tightly integrated with the Microsoft Exchange Replication service (MSExchangeRepl.exe) to provide higher availability through improved resiliency.

- **Support for multiple databases per disk**: Enhancements enable you to support multiple databases (mixtures of active and passive copies) on the same disk, thereby leveraging larger disks in terms of capacity and IOPS as efficiently as possible.

- **AutoReseed**: Automatic reseeding capability enables you to quickly restore database redundancy after disk failure. If a disk fails, the database copy stored on that disk is copied from the active database copy to a spare disk on the same server. If multiple database copies were stored on the failed disk, they can all be automatically reseeded on a spare disk. This enables faster reseeds, as the active databases are likely to be on multiple servers and the data is copied in parallel.

- **Automatic recovery from storage failures**: This feature continues the innovation that was introduced in Exchange 2010 to allow the system to recover from failures that affect resiliency or redundancy. Exchange now includes additional recovery behaviors for long I/O times, excessive memory consumption by MSExchangeRepl.exe, and severe cases where the system is in such a bad state that threads can't be scheduled.

- **Lagged copy enhancements**: Lagged copies can now care for themselves to a certain extent using automatic log play down. Lagged copies will automatically play down log files in a variety of situations, such as page patching and low disk space scenarios. If the system detects that page patching is required for a lagged copy, the logs will be automatically replayed into the lagged copy to perform page patching. Lagged copies will also invoke this auto replay feature when a low disk space threshold has been reached, and when the lagged copy has been detected as the only available copy for a specific period of time. In addition, lagged copies can leverage Safety Net, making recovery or activation much easier.

- **Single copy alert enhancements**: The single copy alert introduced in Exchange 2010 is no longer a separate scheduled script. It's now integrated into the managed availability components within the system and is a native function within Exchange.

- **DAG network auto-configuration**: DAG networks can be automatically configured by the system based on configuration settings. In addition to manual configuration options, DAGs can also distinguish between

MAPI and replication networks and configure DAG networks automatically.

Reduction in IOPS

In Exchange 2010, passive database copies have a very low checkpoint depth, which is required for fast failover. In addition, the passive copy performs aggressive pre-reading of data to keep up with a 5-megabyte (MB) checkpoint depth. As a result of using a low checkpoint depth and performing these aggressive pre-read operations, IOPS for a passive database copy was equal to IOPS for an active copy in Exchange 2010.

In Exchange 2013 or later, the system is able to provide fast failover while using a high checkpoint depth on the passive copy (100 MB). Because passive copies have 100-MB checkpoint depth, they've been de-tuned to no longer be so aggressive. As a result of increasing the checkpoint depth and de-tuning the aggressive pre-reads, IOPS for a passive copy is about 50 percent of the active copy IOPS.

Having a higher checkpoint depth on the passive copy also results in other changes. On failover in Exchange 2010, the database cache is flushed as the database is converted from a passive copy to an active copy. Starting in Exchange 2013, ESE logging was rewritten so that the cache is persisted through the transition from passive to active. Because ESE doesn't need to flush the cache, you get fast failover.

One other change was made to the background database maintenance (BDM) process. BDM now processes around 1-2 MB per second per copy.

As a result of these changes, Exchange now provides a significant reduction in IOPS over Exchange 2010.

Managed Availability

Managed Availability is the integration of built-in, active monitoring and the Exchange high availability platform. With Managed Availability, the system can make a determination on when to fail over a database based on service health. Managed Availability is an internal infrastructure that's deployed in the Client Access (frontend) services and backend services on Mailbox servers. Managed Availability includes three main asynchronous components that are constantly doing work:

1. The first component is the probe engine, which is responsible for taking measurements on the server and collecting data. The results of those measurements flow into the second component, the monitor.

2. The monitor contains all of the business logic used by the system based on what is considered healthy on the data collected. Similar to a pattern recognition engine, the monitor looks for the various different patterns on all the collected measurements, and then it decides whether something is considered healthy.

3. Finally, there is the responder engine, which is responsible for recovery actions.

When something is unhealthy, the first action is to attempt to recover that component. This could include multi-stage recovery actions; for example:

1. Restart the application pool.

2. Restart the service.

3. Restart the server.

4. Take the server offline so that it no longer accepts traffic.

If the recovery actions are unsuccessful, the system escalates the issue to a human through event log notifications.

Managed availability is implemented in the form of two services:

- **Exchange Health Manager Service (MSExchangeHMHost.exe)**: This is a controller process that's used to manage worker processes. It's used to build, execute, and start and stop the worker process as needed. It's also used to recover the worker process in case that process crashes, to prevent the worker process from being a single point of failure.

- **Exchange Health Manager Worker process (MSExchangeHMWorker.exe)**: This is the worker process that's responsible for performing the runtime tasks.

Managed availability uses persistent storage to perform its functions:

- XML configuration files are used to initialize the work item definitions during startup of the worker process.
- The registry is used to store runtime data, such as bookmarks.
- The crimson channel event log infrastructure is used to store the work item results.

For more information about managed availability, see Managed availability.

Managed Store

Exchange 2010 and earlier versions support running a single instance of the Information Store process (Store.exe) on the Mailbox server role. This single Store instance hosts all databases on the server: active, passive, lagged, and recovery. In these Exchange architectures, there is little, if any, isolation between the different databases hosted on a Mailbox server. An issue with a single mailbox database has the potential to negatively affect all other databases, and crashes resulting from a mailbox corruption can affect service for all users whose databases are hosted on that server.

Another challenge with a single Store instance is the lack of processor scalability with the Extensible Storage Engine (ESE). ESE scales well to 8-12 processor cores, but beyond that, cross-processor communication and cache synchronization issues lead to negative performance. Given today's servers with 16+ core systems available, this would impose the administrative challenge of managing the affinity of 8-12 cores for ESE and using the other cores for non-Store processes (for example, Assistants, Search Foundation, Managed Availability, etc.). Moreover, the previous architecture restricted scale-up for the Store process.

The Store.exe process has evolved considerably throughout the years as Exchange Server itself evolved, but as a single process, ultimately its scalability is limited, and it represents a single point of failure. Because of these limits, Store.exe was removed in Exchange 2013 and replaced by the Managed Store.

For more information, see Managed Store.

Multiple databases per volume

Although the storage improvements in Exchange are designed primarily for just a bunch of disks (JBOD) configurations, they're available for use by all supported storage configurations. One such feature is the ability to host multiple databases on the same volume. This feature is about Exchange optimizing for large disks. These optimizations result in a much more efficient use of large disks in terms of capacity, IOPS, and reseed times, and they're meant to address the challenges associated with running in a JBOD storage configuration:

- Database sizes must be manageable.
- Reseed operations must be fast and reliable.
- Although storage capacity is increasing, IOPS aren't.
- Disks hosting passive database copies are underutilized in terms of IOPS.
- Lagged copies have asymmetric storage requirements.
- Limited agility exists to recover from low disk space conditions.

The trend of increasing storage capacity continues. For example, the Exchange best practice guideline for maximum database size (2 terabytes) on an 8 terabyte drive means you would waste more than 5 terabytes of disk space.

A ssolution would be to simply grow the databases larger, but that inhibits manageability because it might introduce long reseed times (including operationally unmanagable reseed times) and compromised reliability of copying that amount of data over the network.

In addition, in the Exchange 2010 model, the disk storing a passive copy is underutilized in terms of IOPS. In the case of a lagged passive copy, not only is the disk underutilized in terms of IOPS, but it's also asymmetric in terms of its size, relative to the disks used to store the active and non-lagged passive copies.

Exchange 2013 and later has been optimized to use large disks (8 terabytes) in a JBOD configuration more efficiently. With multiple databases per disk, you can now have the same size disks storing multiple database copies, including lagged copies. The goal is to drive the distribution of users across the number of volumes that exist, providing you with a symmetric design where during normal operations each DAG member hosts a combination of active, passive, and optional lagged copies on the same volumes.

An example of a configuration that uses multiple databases per volume is illustrated below.

Configuration that uses multiple databases per volume

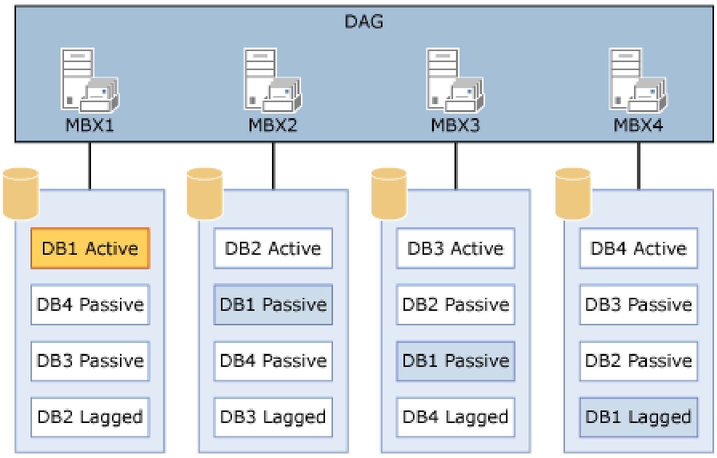

The configuration in the diagram provides a symmetrical design. All four servers have the same four databases all hosted on a single disk per server. The key is that the number of copies of each database that you have should be equal to the number of database copies per disk.

In the configuration in the diagram, there are four copies of each database: one active copy, two passive copies, and one lagged copy. Because there are four copies of each database, the proper configuration is one that has four copies per volume.

In addition, activation preference is configured so that it's balanced across the DAG and across each server. For example:

- The active copy will have an activation preference value of 1.
- The first passive copy will have an activation preference value of 2.
- The second passive copy will have an activation preference value of 3.
- The lagged copy will have an activation preference value of 4.

In addition to having a better distribution of users across the existing volumes, another benefit of using multiple databases per disk is a reduction in the amount of time to restore data protection for failures that require a reseed (for example, disk failure).

As a database gets bigger, reseeding the database takes longer. For example, a 2 terabyte database could take 23 hours to reseed, whereas an 8 terabyte database could take as long as 93 hours (almost 4 days). Both seeds would occur at about 20 MB per second. This generally means that a very large database can't be seeded within an operationally reasonable amount of time.

In the case of a single database copy per disk scenario, the seeding operation is effectively source-bound, because it's always seeding the disk from a single source.

By dividing the volume into multiple database copies, and by having the active copy of the passive databases on a specified volume stored on separate DAG members, the system is no longer source bound in the context of reseeding the disk. When a failed disk is replaced, it can be reseeded from multiple sources. This allows the system to reseed and restore data protection for these databases in a much shorter amount of time.

When you use multiple databases per volume, we recommend that you follow these best practices and requirements:

- A single logical disk partition per physical disk must be used. Don't create multiple partitions on the disk. Each database copy and its companion files (such as transaction logs and content index) should be hosted in a unique directory on the single partition.
- The number of database copies configured per volume should be equal to the number of copies of each database. For example, if you have four copies of your databases, you should use four database copies per volume.
- Database copies should have the same neighbors. (For example, they should all share the same disk on each server.)
- Activation preference across the DAG should be balanced, such that each database copy on a specified disk has a unique activation preference value.

AutoReseed

Automatic reseed (also known as AutoReseed) is the replacement for what is normally an administrator-driven action in response to a disk failure, database corruption event, or other issue that requires a reseed of a database copy. AutoReseed is designed to automatically restore database redundancy after a disk failure by using spare disks that have been provisioned on the system.

For more information, see AutoReseed. For detailed steps to configure AutoReseed, see Configure AutoReseed for a database availability group.

Automatic recovery from storage failures

Automatic recovery from storage failures allows the system to recover from failures that affect resiliency or redundancy. In addition to the bugcheck behaviors introduced in Exchange 2010, Exchange now includes additional recovery behaviors for long I/O times, excessive memory consumption by the Microsoft Exchange Replication service (MSExchangeRepl.exe), and severe cases where threads can't be scheduled.

Even in JBOD environments, storage array controllers can have issues, such as crashing or hanging. The following table lists features that provide hung I/O detection and recovery features that provide enhanced resilience.

NAME	CHECK	ACTION	THRESHOLD
ESE Database Hung IO Detection	ESE checks for outstanding I/Os	Generates a failure item in the crimson channel to restart the server	240 seconds
Failure Item Channel Heartbeat	Ensures failure items can be written to and read from crimson channel	Replication service heartbeats crimson channel and restart server on failures	30 seconds

NAME	CHECK	ACTION	THRESHOLD
System Disk Heartbeat	Verifies server's system disk state	Periodically sends unbuffered I/O to system disk; restarts server on heartbeat time out	120 seconds

Exchange 2013 and later enhances server and storage resilience by including behaviors for other serious conditions. These conditions and behaviors are described in the following table.

NAME	CHECK	ACTION	THRESHOLD
System bad state	No threads, including non-managed threads, can be scheduled	Restart the server	302 seconds
Long I/O times	I/O operation latency measurements	Restart the server	41 seconds
Replication service memory use	Measure the working set of MSExchangeRepl.exe	1: Log event 4395 in the crimson channel with a service termination request 2: Initiate termination of MSExchangeRepl.exe 3: If service termination fails, restart the server	4 gigabyte (GB)
System Event 129 (Bus reset)	Check for Event 129 in System event log	Restart the server	When event occurs
Cluster database hang	Global Update Manager updates are blocked	Restart the server	When event occurs

Lagged copy enhancements

Lagged copy enhancements include integration with Safety Net and automatic play down of log files in certain scenarios. Safety Net was introduced in Exchange 2013 to replace the Exchange 2010 feature known as the transport dumpster. Safety Net is similar to the transport dumpster, in that it's a delivery queue that's associated with the Transport service on a Mailbox server. This queue stores copies of messages that were successfully delivered to the active mailbox database on the Mailbox server. Each active mailbox database on the Mailbox server has its own queue that stores copies of the delivered messages. You can specify how long Safety Net stores copies of the successfully delivered messages before they expire and are automatically deleted.

Safety Net takes some responsibility from shadow redundancy in DAG environments. In DAG environments, shadow redundancy doesn't need to keep another copy of the delivered message in a shadow queue while it waits for the delivered message to replicate to the passive copies of mailbox databases on the other Mailbox servers in the DAG. The copy of the delivered message is already stored in Safety Net, so shadow redundancy can redeliver the message from Safety Net if necessary.

With Safety Net, activating a lagged database copy becomes significantly easier. For example, consider a lagged copy that has a 2-day replay lag. In that case, you would configure Safety Net for a period of 2 days. If you encounter a situation in which you need to use your lagged copy, you can:

1. Suspend replication to it.

2. Copy it twice (to preserve the lagged nature of the database and to create an extra copy in case you need it).

3. Take a copy and discard all the log files, except for those in the required range.

4. Mount the copy, which triggers an automatic request to Safety Net to redeliver the last two days of mail.

With Safety Net, you don't need to hunt for where the point of corruption was introduced. You get the last two days mail, minus the data ordinarily lost on a lossy failover.

Lagged copies can now care for themselves by invoking automatic log replay to play down the log files in certain scenarios:

- When a low disk space threshold is reached
- When the lagged copy has physical corruption and needs to be page patched
- When there are fewer than three available healthy copies (active or passive only; lagged database copies are not counted) for more than 24 hours

In Exchange 2010, page patching wasn't available for lagged copies. In Exchange 2013 or later, page patching is available for lagged copies through this automatic play down feature. If the system detects that page patching is required for a lagged copy, the logs are automatically replayed into the lagged copy to perform page patching. Lagged copies also invoke this auto replay feature when a low disk space threshold has been reached, and when the lagged copy has been detected as the only available copy for a specific period of time.

Lagged copy play down behavior is disabled by default, and can be enabled by running the following command.

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

After being enabled, play down occurs when there are fewer than three copies. You can change the default value of 3, by modifying the following DWORD registry value.

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

To enable play down for low disk space thresholds, you must configure the following registry entry.

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

After configuring either of these registry settings, restart the Microsoft Exchange DAG Management service for the changes to take effect.

As an example, consider an environment where a given database has 4 copies (3 highly available copies and 1 lagged copy), and the default setting is used for *ReplayLagManagerNumAvailableCopies*. If a non-lagged copy is out-of-service for any reason (for example, it is suspended, etc.) then the lagged copy will automatically play down its log files in 24 hours.

Single copy alert enhancements

Ensuring that your servers are operating reliably and that your mailbox database copies are healthy are primary objectives of daily Exchange messaging operations. You must actively monitor the hardware, the Windows operating system, and the Exchange services.

But in an Exchange mailbox resiliency environment, it's important that you monitor the health and status of the DAG and your mailbox database copies. It's especially vital to perform data redundancy risk management and monitor for periods in which a replicated database is down to just a single copy. This is particularly critical in environments that don't use Redundant Array of Independent Disks (RAID) and instead deploy JBOD configurations. In a RAID environment, a single disk failure doesn't affect an active mailbox database copy. However, in a JBOD environment, a single disk failure will trigger a database failover.

The CheckDatabaseRedundancy.ps1 script was introduced in Exchange 2010. As its name implies, the purpose of the script was to monitor the redundancy of replicated mailbox databases by validating that there is at least two configured, healthy, and current copies, and to alert an administrator through event log generation when only a single healthy copy of a replicated database exists. In this case, both active and passive copies are counted when determining redundancy.

Single copy conditions include, but aren't limited to:

- Failure of an active copy to replicate to any passive copy.
- Failure of all passive copies, which includes FailedAndSuspended and Failed states in addition to healthy states where the copy is behind in log copying or replay. Note that lagged copies aren't considered behind if they're within ten minutes in replaying their logs to their lag period.
- Failure of the system to accurately know the current log generation of the active copy.

Because it's a top priority for administrators to know when they're down to a single healthy copy of a database, the CheckDatabaseRedundancy.ps1 script has been replaced with integrated, native functionality that's part of managed availability's DataProtection Health Set.

The native functionality still alerts administrators through event log notifications, and to distinguish Exchange 2013 or later alerts from Exchange 2010, Exchange now uses the following Event IDs:

- Event 4138 (Red Alert)
- Event 4139 (Green Alert)

The native functionality has been enhanced to reduce alert noise that occurs when multiple databases on the same server enter into a single copy condition. In Exchange 2010, single copy alerts were generated on a per-database level. As a result, a server-wide issue that affected multiple databases and multiple database copies could cause alert storms. Because several failures are server-wide (for example, controller or memory problems), there was a good chance that an alert storm would occur for each server incident.

Alerts are now generated on a per-server basis. When an outage affects an entire server and data redundancy becomes at risk for multiple database copies, a single per-server alert is generated.

DAG network auto-configuration

A DAG network is a collection of one or more subnets used for either replication traffic or MAPI traffic. Each DAG contains a maximum of one MAPI network and zero or more replication networks.

In Exchange 2010, the initial DAG networks (for example, DAGNetwork01 and DAGNetwork02) were created by the system based on the subnets that were enumerated by the Cluster service. If you had multiple networks and the interfaces for a specified network (for example, the MAPI network) were on the same subnet, there was little additional configuration required. However, if the interfaces for a specified network were on multiple subnets, you needed to perform a task known as collapsing DAG networks.

In Exchange 2013 or later, collapsing DAG networks is no longer necessary. Exchange still uses the same detection mechanisms to distinguish between the MAPI and replication networks, but it now automatically collapses DAG networks as appropriate.

In addition, by default, DAG networks are now automatically managed by the system. To view DAG network properties using the Exchange admin center (EAC), you must configure the DAG for manual network control by modifying the properties of the DAG using EAC, or by using the **Set-DatabaseAvailabilityGroup** cmdlet to set the *ManualDagNetworkConfiguration* parameter to `$true`.

Changes to best copy selection

Best copy selection (BCS) is an internal algorithm process for finding the best copy of an individual database to activate, given a list of potential copies for activation and their health and status. Active Manager selects the best available (and unblocked) copy to become the new active database copy when the existing active database copy fails or when an administrator performs a targetless switchover. In Exchange 2010, the BCS process evaluated several aspects of each database copy to determine the best copy to activate. These included:

- Copy queue length
- Replay queue length
- Database status
- Content index status

In Exchange 2013 or later, Active Manager performs the same BCS checks and phases to determine replication health, but it now also includes the use of a constraint of the decreasing order of health states. As a result of these changes, BCS is now called best copy and server selection (BCSS).

BCSS includes several new health checks that are now part of the built-in managed availability monitoring components in Exchange. There are four additional checks performed by Active Manager (listed in the order in which they're performed):

1. **All Healthy**: Checks for a server hosting a copy of the affected database that has all monitoring components in a healthy state.
2. **Up to Normal Healthy**: Checks for a server hosting a copy of the affected database that has all monitoring components with Normal priority in a healthy state.
3. **All Better than Source**: Checks for a server hosting a copy of the affected database that has monitoring components in a state that's better than the current server hosting the affected copy.
4. **Same as Source**: Checks for a server hosting a copy of the affected database that has monitoring components in a state that's the same as the current server hosting the affected copy.

If BCSS is invoked as a result of a failover that's triggered by a managed availability monitoring component (for example, via a Failover responder), an additional mandatory constraint is enforced where the target server's component health must be better than the server on which the failover occurred. For example, if a failure of Outlook on the web (formerly known as Outlook Web App) triggers a managed availability failover via a Failover responder, BCSS must select a server hosting a copy of the affected database on which Oulook on the web is healthy.

DAG Management Service

Exchange 2013 CU2 or later includes the Microsoft Exchange DAG Management Service (MSExchangeDAGMgmt). This service contains the internal DAG monitoring functionality that was previously inside the Microsoft Exchange Replication service (MSExchangeRepl).

DAGs without a cluster administrative access point

All DAGs on Exchange servers running Windows Server 2008 R2 or Windows Server 2012 require at least one IP address on every subnet included in the MAPI network. The IP address(es) assigned to the DAG are used by the DAG's cluster with the cluster's administrative access point (also known as the cluster network name) to enable name resolution and connectivity to the cluster (or more precisely, connectivity to the cluster member that currently owns the cluster core resource group) using the cluster name.

Windows Server 2012 R2 or later enables you to create a failover cluster without an administrative access point. Windows failover clusters without administrative access points have the following characteristics:

- No IP address is assigned to the cluster, so there's no IP Address Resource in the cluster core resource group.
- No network name is assigned to the cluster, so there's no Network Name Resource in the cluster core resource group.
- The name of the cluster isn't registered in DNS and the cluster name isn't resolvable on the network.
- A cluster name object (CNO) isn't created in Active Directory.
- You can't manage the Windows failover cluster using the Failover Cluster Management tool. Instead, you need to use Windows PowerShell and you need to run the PowerShell cmdlets against the individual cluster members.

Exchange 2013 SP1 or later running on Exchange on Windows Server 2012 R2 or later enables you to create a DAG without a cluster administrative access point. For more information, see Creating DAGs and Create a database availability group.

Database availability groups

8/3/2020 • 11 minutes to read • Edit Online

A database availability group (DAG) is the base component of the Mailbox server high availability and site resilience framework built into Microsoft Exchange Server. A DAG is a group of up to 16 Mailbox servers that hosts a set of databases and provides automatic database-level recovery from failures that affect individual servers or databases.

> **IMPORTANT**
>
> All servers within a DAG must be running the same version of Exchange. For example, you can't mix Exchange 2013 servers and Exchange 2016 servers in the same DAG.

A DAG is a boundary for mailbox database replication, database and server switchovers and failovers, and an internal component called *Active Manager*. Active Manager, which runs on every Mailbox server, manages switchovers and failovers within DAGs. For more information about Active Manager, see Active Manager.

Any server in a DAG can host a copy of a mailbox database from any other server in the DAG. When a server is added to a DAG, it works with the other servers in the DAG to provide automatic recovery from failures that affect mailbox databases, such as a disk, server, or network failure.

> **NOTE**
>
> For more information about creating DAGs, managing DAG membership, configuring DAG properties, creating and monitoring mailbox database copies, and performing switchovers, see Managing high availability and site resilience.

Database availability group lifecycle

DAGs leverage the concept of *incremental deployment*, which is the ability to deploy service and data availability for all Mailbox servers and databases after Exchange is installed. After you deploy Exchange Server Mailbox servers, you can create a DAG, add Mailbox servers to the DAG, and then replicate mailbox databases between the DAG members.

> **NOTE**
>
> It's supported to create a DAG that contains a combination of physical Mailbox servers and virtualized Mailbox servers, provided that the servers and solution comply with the Exchange Server system requirements and the requirements set forth in Exchange Server virtualization. As with all Exchange high availability configurations, you must ensure that all Mailbox servers in the DAG are sized appropriately to handle the necessary workload during scheduled and unscheduled outages.

A DAG is created by using the New-DatabaseAvailabilityGroup cmdlet. A DAG is initially created as an empty object in Active Directory. This directory object is used to store relevant information about the DAG, such as server membership information and some DAG configuration settings. When you add the first server to a DAG, a failover cluster is automatically created for the DAG. This failover cluster is used exclusively by the DAG, and the cluster must be dedicated to the DAG. Use of the cluster for any other purpose isn't supported.

In addition to a failover cluster being created, the infrastructure that monitors the servers for network or server failures is initiated. The failover cluster heartbeat mechanism and cluster database are then used to track and manage information about the DAG that can change quickly, such as database mount status, replication status, and last mounted location.

During creation, the DAG is given a unique name, and either assigned one or more static IP addresses or configured to use Dynamic Host Configuration Protocol (DHCP), or created without a cluster administrative access point. DAGs without an administrative access point can be created only on servers running Exchange 2019, Exchange 2016, or Exchange 2013 Service Pack 1 or later, with Windows Server 2012 R2 Standard or Datacenter edition. DAGs without cluster administrative access points have the following characteristics:

- There is no IP address assigned to the cluster/DAG, and therefore no IP Address Resource in the cluster core resource group.

- There is no network name assigned to the cluster, and therefore no Network Name Resource in the cluster core resource group

- The name of the cluster/DAG is not registered in DNS, and it is not resolvable on the network.

- A cluster name object (CNO) is not created in Active Directory.

- The cluster cannot be managed using the Failover Cluster Management tool. It must be managed using Windows PowerShell, and the PowerShell cmdlets must be run against individual cluster members.

This example shows you how to use the Exchange Management Shell to create a DAG with a cluster administrative access point that will have three servers. Two servers (EX1 and EX2) are on the same subnet (10.0.0.0), and the third server (EX3) is on a different subnet (192.168.0.0).

```
New-DatabaseAvailabilityGroup -Name DAG1 -WitnessServer EX4 -DatabaseAvailabilityGroupIPAddresses
10.0.0.5,192.168.0.5
Add-DatabaseAvailabilityGroupServer -Identity DAG1 -MailboxServer EX1
Add-DatabaseAvailabilityGroupServer -Identity DAG1 -MailboxServer EX2
Add-DatabaseAvailabilityGroupServer -Identity DAG1 -MailboxServer EX3
```

The commands to create a DAG without a cluster administrative access point are very similar:

```
New-DatabaseAvailabilityGroup -Name DAG1 -WitnessServer EX4 -DatabaseAvailabilityGroupIPAddresses
([System.Net.IPAddress])::None
Add-DatabaseAvailabilityGroupServer -Identity DAG1 -MailboxServer EX1
Add-DatabaseAvailabilityGroupServer -Identity DAG1 -MailboxServer EX2
Add-DatabaseAvailabilityGroupServer -Identity DAG1 -MailboxServer EX3
```

The cluster for DAG1 is created when EX1 is added to the DAG. During cluster creation, the **Add-DatabaseAvailabilityGroupServer** cmdlet retrieves the IP addresses configured for the DAG and ignores the ones that don't match any of the subnets found on EX1. In the first example above, the cluster for DAG1 is created with an IP address of 10.0.0.5, and 192.168.0.5 is ignored. In the second example above, the value of the *DatabaseAvailabilityGroupIPAddresses* parameter instructs the task to create a failover cluster for the DAG that does not have an administrative access point. Thus, the cluster is created with an IP address or network name resource in the core cluster resource group.

Then, EX2 is added, and the **Add-DatabaseAvailabilityGroupServer** cmdlet again retrieves the IP addresses configured for the DAG. There are no changes to the cluster's IP addresses because in EX2 is on the same subnet as EX1.

Then, EX3 is added, and the **Add-DatabaseAvailabilityGroupServer** cmdlet again retrieves the IP addresses configured for the DAG. Because a subnet matching 192.168.0.5 is present on EX3, the 192.168.0.5 address is added as an IP address resource in the cluster group. In addition, an **OR** dependency for the Network Name resource for each IP address resource is automatically configured. The 192.168.0.5 address will be used by the cluster when the cluster core resource group moves to EX3.

For DAGs with cluster administrative access points, Windows failover clustering registers the IP addresses for the cluster in the Domain Name System (DNS) when the Network Name resource is brought online. In addition, when EX1 is added to the cluster, a cluster name object (CNO) is created in Active Directory. The network name, IP address(es), and CNO for the cluster are not used for DAG functions. Administrators and end users don't need to interface with or connect to the cluster/DAG name or IP address for any reason. Some third-party applications connect to the cluster administrative access point to perform management tasks, such as backup or monitoring. If you do not use any third-party applications that require a cluster administrative access point, and your DAG is running Exchange 2016 or Exchange 2019 on Windows Server 2012 R2, then we recommend creating a DAG without an administrative access point. This simplifies DAG configuration, eliminates the need for one or more IP addresses, and reduces the attack surface of a DAG.

DAGs are also configured to use a witness server and a witness directory. The witness server and witness directory are either automatically configured by the system, or they can be manually configured by the administrator. In the examples above, EX4 (a server that is not and will not be a member of the DAG) is being manually configured as the DAG's witness server.

By default, a DAG is designed to use the built-in continuous replication feature to replicate mailbox databases among servers in the DAG. If you're using third-party data replication that supports the Third Party Replication API in Exchange Server, you must create the DAG in third-party replication mode by using the New-DatabaseAvailabilityGroup cmdlet with the *ThirdPartyReplication* parameter. After this mode is enabled, it can't be disabled.

After the DAG is created, Mailbox servers can be added to the DAG. When the first server is added to the DAG, a cluster is formed for use by the DAG. DAGs make use of Windows failover clustering technology, such as the cluster heartbeat, cluster networks, and the cluster database (for storing data that changes, such as database state changes from active to passive or vice versa, or from mounted to dismounted and vice versa). As each subsequent server is added to the DAG, it's joined to the underlying cluster, the cluster's quorum model is automatically adjusted by Exchange, and the server is added to the DAG object in Active Directory.

After Mailbox servers are added to a DAG, you can configure a variety of DAG properties, such as whether to use network encryption or network compression for database replication within the DAG. You can also configure DAG networks and create additional DAG networks.

After you add members to a DAG and configure the DAG, the active mailbox databases on each server can be replicated to the other DAG members. After you create mailbox database copies, you can monitor the health and status of the copies using a variety of built-in monitoring tools. In addition, you can perform database and server switchovers.

Database availability group quorum models

Underneath every DAG is a Windows failover cluster. Failover clusters use the concept of quorum, which uses a consensus of voters to ensure that only one subset of the cluster members (which could mean all members or a majority of members) is functioning at one time. Quorum isn't a new concept for Exchange Server. Highly available Mailbox servers in previous versions of Exchange also use failover clustering and its concept of quorum. Quorum represents a shared view of members and resources, and the term quorum is also used to describe the physical data that represents the configuration within the cluster that's shared between all cluster members. As a result, all DAGs require their underlying failover cluster to have quorum. If the cluster loses quorum, all DAG operations terminate and all mounted databases hosted in the DAG dismount. In this event, administrator intervention is required to correct the quorum problem and restore DAG operations.

Quorum is important to ensure consistency, to act as a tie-breaker to avoid partitioning, and to ensure cluster responsiveness:

- **Ensuring consistency**: A primary requirement for a Windows failover cluster is that each of the members always has a view of the cluster that's consistent with the other members. The cluster hive acts as the definitive repository for all configuration information relating to the cluster. If the cluster hive can't be loaded locally on a DAG member, the Cluster service doesn't start, because it isn't able to guarantee that the member meets the requirement of having a view of the cluster that's consistent with the other members.

- **Acting as a tie-breaker**: A quorum witness resource is used in DAGs with an even number of members to avoid split brain syndrome scenarios and to make sure that only one collection of the members in the DAG is considered official. When the witness server is needed for quorum, any member of the DAG that can communicate with the witness server can place a Server Message Block (SMB) lock on the witness server's witness.log file. The DAG member that locks the witness server (referred to as the *locking node*) retains an additional vote for quorum purposes. The DAG members in contact with the locking node are in the majority and maintain quorum. Any DAG members that can't contact the locking node are in the minority and therefore lose quorum.

- **Ensuring responsiveness**: To ensure responsiveness, the quorum model makes sure that, whenever the cluster is running, enough members of the distributed system are operational and communicative, and at least one replica of the cluster's current state can be guaranteed. No additional time is required to bring members into communication or to determine whether a specific replica is guaranteed.

DAGs with an even number of members use the failover cluster's Node and File Share Majority quorum mode, which employs an external witness server that acts as a tie-breaker. In this quorum mode, each DAG member gets a vote. In addition, the witness server is used to provide one DAG member with a weighted vote (for example, it gets two votes instead of one). The cluster quorum data is stored by default on the system disk of each member of the DAG, and is kept consistent across those disks. However, a copy of the quorum data isn't stored on the witness server. A file on the witness server is used to keep track of which member has the most updated copy of the data, but the witness server doesn't have a copy of the cluster quorum data. In this mode, a majority of the voters (the DAG members plus the witness server) must be operational and able to communicate with each other to maintain quorum. If a majority of the voters can't communicate with each other, the DAG's underlying cluster loses quorum, and the DAG will require administrator intervention to become operational again. For more information, see Datacenter switchovers and Restore-DatabaseAvailabilityGroup.

DAGs with an odd number of members use the failover cluster's Node Majority quorum mode. In this mode, each member gets a vote, and each member's local system disk is used to store the cluster quorum data. If the configuration of the DAG changes, that change is reflected across the different disks. The change is only considered to have been committed and made persistent if that change is made to the disks on half the members (rounding down) plus one. For example, in a five-member DAG, the change must be made on two plus one members, or three members total.

Quorum requires a majority of voters to be able to communicate with each other. Consider a DAG that has four members. Because this DAG has an even number of members, an external witness server is used to provide one of the cluster members with a fifth, tie-breaking vote. To maintain a majority of voters (and therefore quorum), at least three voters must be able to communicate with each other. At any time, a maximum of two voters can be offline without disrupting service and data access. If three or more voters are offline, the DAG loses quorum, and service and data access will be disrupted until you resolve the problem.

Active Manager

8/3/2020 • 9 minutes to read • Edit Online

Microsoft Exchange Server includes a component called *Active Manager* that manages the high availability platform that includes the database availability group (DAG) and mailbox database copies. Active Manager runs inside the Microsoft Exchange Replication service (MSExchangeRepl.exe) on all Mailbox servers. On Mailbox servers that aren't members of a DAG, there is a single Active Manager role: *Standalone Active Manager*.

On servers that are members of a DAG, there are two Active Manager roles: *Primary Active Manager* (PAM) and *Standby Active Manager* (SAM). PAM is the Active Manager role in a DAG that decides which copies will be active and passive. PAM is responsible for getting topology change notifications and reacting to server failures. The DAG member that holds the PAM role is always the member that currently owns the cluster quorum resource (default cluster group). If the server that owns the cluster quorum resource fails, the PAM role automatically moves to a surviving server that takes ownership of the cluster quorum resource. In addition, if you need to take the server that hosts the cluster quorum resource offline for maintenance or an upgrade, you must first move the PAM to another server in the DAG. The PAM controls all movement of the active designations between a database's copies. (Only one copy can be active at any specified time, and that copy may be mounted or dismounted.) The PAM also performs the functions of the SAM role on the local system (detecting local database and local Information Store failures).

The SAM provides information on which server hosts the active copy of a mailbox database to other components of Exchange that are running an Active Manager client component (for example, Client Access or Transport services). The SAM detects failures of local databases and the local Information Store. It reacts to failures by asking the PAM to initiate a failover (if the database is replicated). A SAM doesn't determine the target of failover, nor does it update a database's location state in the PAM. It will access the active database copy location state to answer queries for the active copy of the database that it receives.

> **NOTE**
>
> Exchange Server isn't a clustered application. Instead, it uses the cluster library functions implemented in clusapi.dll for cluster, group, cluster network (heartbeating), node management, cluster registry, and a few control code functions. In addition, Active Manager stores current mailbox database information (for example, active and passive data, and mounted data) in the cluster database (also known as the cluster registry). Although the information is stored directly in the cluster database, it isn't accessed directly by any other components.

In Exchange Server, the Microsoft Exchange Replication service periodically monitors the health of all mounted databases. In addition, it also monitors the Extensible Storage Engine (ESE) for any I/O errors or failures. When the service detects a failure, it notifies Active Manager. Active Manager then determines which database copy should be mounted and what it requires to mount that database. In addition, it tracks the active copy of a mailbox database (based on the last mounted copy of the database) and provides the tracking results information to Client Access services on the Mailbox server to which the client is connected.

Best Copy Selection

When a failure occurs that prevents access to the active copy of a replicated mailbox database, Active Manager selects the best possible passive copy of the affected database to activate. This process was known as best copy selection (BCS) in earlier versions of Exchange, and in Exchange 2016 and Exchange 2019 it's known as best copy and server selection (BCSS). The general process occurs in the following order:

1. Managed availability or Active Manager detects a failure, or an administrator initiates a targetless

switchover.

2. The PAM runs the BCSS internal algorithm.
3. A process called *attempt copy last logs* (ACLL) occurs, which tries to copy any missing log files from the server that hosted the active database copy prior to the failure or switchover.
4. After the ACLL process has completed, the value of the *AutoDatabaseMountDial* for the Mailbox servers hosting copies of the database is compared with the copy queue length of the database being activated. At this point, either:
   - The number of missing log files is equal to or less than the value of *AutoDatabaseMountDial*, in which case Step 5 occurs.
   - The number of missing log files is greater than the value of *AutoDatabaseMountDial*, in which case Active Manager will try to activate next best available copy, if there is one.
5. The PAM issues a mount request to the Microsoft Exchange Information Store via remote procedure call (RPC). At this point, either:
   - The database mounts and is made available to clients.
   - The database doesn't mount, and PAM performs steps 3 and 4 on the next best copy (if one is available).

In earlier versions of Exchange, the BCS process evaluated several aspects of each database copy to determine the best copy to activate. These included:

- Copy queue length
- Replay queue length
- Database status
- Content index status

In Exchange Server, Active Manager runs through all of the same BCS checks and phases, but now it also includes the use of a constraint of the decreasing order of health states. Specifically, BCSS includes several new health checks that are part of the built in managed availability monitoring components in Exchange Server. There are four additional checks performed by Active Manager (listed in the order in which they are performed):

1. **All Healthy**: Checks for a server hosting a copy of the affected database that has all monitoring components in a healthy state.
2. **Up to Normal Healthy**: Checks for a server hosting a copy of the affected database that has all monitoring components with Normal priority in a healthy state.
3. **All Better than Source**: Checks for a server hosting a copy of the affected database that has monitoring components in a state that's better than the current server hosting the affected copy.
4. **Same as Source**: Checks for a server hosting a copy of the affected database that has monitoring components in a state that's the same as the current server hosting the affected copy.

If BCSS is invoked as a result of a failover that's triggered by a monitoring component (for example, via a Failover responder), an additional mandatory constraint is enforced where the target server's component health must be better than the server on which the failover occurred. For example, if a failure of Outlook on the web triggers a failover via a Failover responder, BCSS must select a server hosting a copy of the affected database on which Outlook on the web is healthy.

Best copy selection process

With respect to database failures (not protocol failures), Active Manager begins the best copy selection process by creating a list of database copies that are potential candidates for activation. Any database copies that are unreachable or are administratively blocked from activation are ignored and not used during the selection process. The order of the list depends on the value of the *AutoDatabaseMountDial*:

- If the *AutoDatabaseMountDial* is configured with any value other than `Lossless` on all servers that host a copy of the database, Active Manager sorts the resulting list using the copy queue length as the primary key. The calculation is based on LastLogInspected (from the copy's point of view), so the list of potential copies is sorted by the highest value for LastLogInspected (which will be the copy with the lowest copy queue length). If necessary, Active Manager sorts the list a second time, using the value for activation preference as a secondary key to break any tie conditions where two or more passive copies have the same copy queue length. The copy with the lowest activation preference value has the higher priority on the list.
- If the *AutoDatabaseMountDial* is configured with a value of `Lossless` on any server that hosts a copy of the database, Active Manager sorts the resulting list in ascending order by using the value for activation preference as the primary key. In addition, when an administrator performs a lossless server or database switchover without specifying a target, Active Manager also sorts the resulting list in ascending order by using the value for activation preference as the primary key.

Next, Active Manager attempts to locate a mailbox database copy on the list that has a status of Healthy, DisconnectedAndHealthy, DisconnectedAndResynchronizing, or SeedingSource, and then evaluates the activation potential of each of the copies on the list by using an order set of ten criteria. Active Manager determines if any of the candidates for activation meet the first set of criteria:

- It has a content index with a status of Healthy.
- It has a copy queue length less than 10 log files.
- It has a replay queue length less than 50 log files.

If none of the database copies meets the first set of criteria, Active Manager tries to locate a database copy that meets the second set of criteria:

- It has a content index with a status of Crawling.
- It has a copy queue length less than 10 log files.
- It has a replay queue length less than 50 log files.

If none of the database copies meets the second set of criteria, Active Manager tries to locate a database copy that meets the third set of criteria:

- It has a content index with a status of Healthy.
- It has a replay queue length less than 50 log files.

If none of the database copies meets the third set of criteria, Active Manager tries to locate a database copy that meets the fourth set of criteria:

- It has a content index with a status of Crawling.
- It has a replay queue length less than 50 log files.

If none of the database copies meets the fourth set of criteria, Active Manager tries to locate a database copy that meets the fifth set of criteria:

- It has a replay queue length less than 50 log files.

If none of the database copies meets the fifth set of criteria, Active Manager tries to locate a database copy that meets the sixth set of criteria:

- It has a content index with a status of Healthy.
- It has a copy queue length less than 10 log files.

If none of the database copies meets the sixth criteria, Active Manager tries to locate a database copy that meets the seventh set of criteria:

- It has a content index with a status of Crawling.
- It has a copy queue length less than 10 log files.

If none of the database copies meets the seventh set of criteria, Active Manager tries to locate a database copy that meets the eighth set of criteria:

- It has a content index with a status of Healthy.

If none of the database copies meets all of the eighth set of criteria, Active Manager tries to locate a database copy that meets the ninth set of criteria:

- It has a content index with a status of Crawling.

If none of the database copies meets the ninth set of criteria, Active Manager tries to activate any database copy with a status of Healthy, DisconnectedAndHealthy, DisconnectedAndResynchronizing, or SeedingSource (the tenth set of criteria). If it can't find any database copies that meet the tenth set of criteria, it isn't able to automatically activate a database copy.

After one or more copies are located that meet one or more sets of criteria, the ACLL process copies any log files from the original source to the potential new active copy. After the ACLL process has completed, the PAM issues a mount request and either the database mounts and is made available to clients, or the database doesn't mount and the PAM searches for the next best copy (if one is available).

Datacenter Activation Coordination mode

8/3/2020 • 5 minutes to read • Edit Online

Datacenter Activation Coordination (DAC) mode is a property of a database availability group (DAG). DAC mode is disabled by default but should be enabled for all DAGs with two or more members that use continuous replication. DAC mode shouldn't be enabled for DAGs that use third-party replication mode unless specified by the third-party vendor.

DAC mode is used to control the database mount on startup behavior of a DAG. This control is designed to prevent split brain from occurring at the database level during a datacenter switchback. *Split brain*, also known as split brain syndrome, is a condition that results in a database being mounted as an active copy on two members of the same DAG that are unable to communicate with one another. Split brain is prevented using DAC mode, because DAC mode requires DAG members to obtain permission to mount databases before they can be mounted.

For example, when a primary datacenter contains two DAG members and the witness server, and a second datacenter contains two other DAG members, the DAG is not in DAC mode. The primary datacenter loses power, so you activate the DAG in the second datacenter. Eventually power to the primary datacenter is restored, and the DAG members in the primary datacenter, which had quorum before the power failure, will start up and mount their databases. Because the primary datacenter was restored without network connectivity to the second datacenter, and because the DAG was not in DAC mode, the active databases within the DAG enters a split brain condition.

How DAC mode works

DAC mode includes a protocol called Datacenter Activation Coordination Protocol (DACP). When DAC mode is enabled, DAG members won't automatically mount databases even if they have quorum. Instead DACP is used to determine the current state of the DAG and whether Active Manager should attempt to mount the databases.

You might think of DAC mode as an application level of quorum for mounting databases. To understand the purpose of DACP and how it works, it's important to understand the primary scenario it's intended to handle. Consider the two-datacenter scenario described above. Suppose there is a complete power failure in the primary datacenter. In this event, all of the servers and the WAN are down, so the organization makes the decision to activate the standby datacenter. In almost all such recovery scenarios, when power is restored to the primary datacenter, WAN connectivity is typically not immediately restored. This means that the DAG members in the primary datacenter will power up, but they won't be able to communicate with the DAG members in the activated standby datacenter. The primary datacenter should always contain the majority of the DAG quorum voters, which means that when power is restored, even in the absence of WAN connectivity to the DAG members in the standby datacenter, the DAG members in the primary datacenter have a majority and therefore have quorum. This is a problem because with quorum, these servers may be able to mount their databases, which in turn would cause divergence from the actual active databases that are now mounted in the activated standby datacenter.

DACP was created to address this issue. Active Manager stores a bit in memory (either a 0 or a 1) that tells the DAG whether it's allowed to mount local databases that are assigned as active on the server. When a DAG is running in DAC mode, each time Active Manager starts up the bit is set to 0, meaning it isn't allowed to mount databases. Because it's in DAC mode, the server must try to communicate with all other members of the DAG that it knows to get another DAG member to give it an answer as to whether it can mount local databases that are assigned as active to it. The answer comes in the form of the bit setting for other Active Managers in the DAG. If another server responds that its bit is set to 1, it means servers are allowed to mount databases, so the server starting up sets its bit to 1 and mounts its databases.

But when you recover from a primary datacenter power outage where the servers are recovered but WAN connectivity has not been restored, all of the DAG members in the primary datacenter will have a DACP bit value of 0; and therefore none of the servers starting back up in the recovered primary datacenter will mount databases, because none of them can communicate with a DAG member that has a DACP bit value of 1.

DAC mode for DAGs with two members

DAGs with two members have inherent limitations that prevent the DACP bit alone from fully protecting against application-level split brain syndrome. For DAGs with only two members, DAC mode also uses the boot time of the DAG's witness server to determine whether it can mount databases on startup. The boot time of the witness server is compared to the time when the DACP bit was set to 1.

- If the time the DACP bit was set is earlier than the boot time of the witness server, the system assumes that the DAG member and witness server were rebooted at the same time (perhaps because of power loss in the primary datacenter), and the DAG member isn't permitted to mount databases.
- If the time that the DACP bit was set is more recent than the boot time of the witness server, the system assumes that the DAG member was rebooted for some other reason (perhaps a scheduled outage in which maintenance was performed or perhaps a system crash or power loss isolated to the DAG member), and the DAG member is permitted to mount databases.

> **IMPORTANT**
>
> Because the witness server's boot time is used to determine whether a DAG member can mount its active databases on startup, you should never restart the witness server and the sole DAG member at the same time. Doing so may leave the DAG member in a state where it can't mount databases on startup. If this happens, you must run the Restore-DatabaseAvailabilityGroup cmdlet on the DAG. This resets the DACP bit and permits the DAG member to mount databases.

Other benefits of DAC mode

In addition to preventing split brain syndrome at the application level, DAC mode also enables the use of the built-in site resilience cmdlets used to perform datacenter switchovers. These include the following:

- Stop-DatabaseAvailabilityGroup
- Restore-DatabaseAvailabilityGroup
- Start-DatabaseAvailabilityGroup

Performing a datacenter switchover for DAGs that aren't in DAC mode involves using a combination of Exchange tools and cluster management tools. For more information, see Datacenter switchovers.

Enabling DAC mode

DAC mode can be enabled only by using the Exchange Management Shell. Specifically, you can use the Set-DatabaseAvailabilityGroup cmdlet to enable DAC mode, as illustrated in the following example.

```
Set-DatabaseAvailabilityGroup -Identity DAG2 -DatacenterActivationMode DagOnly
```

In the preceding example, DAG2 is enabled for DAC mode.

For more information about enabling DAC mode, see Configure database availability group properties and Set-DatabaseAvailabilityGroup.

Mailbox database copies

8/3/2020 • 3 minutes to read • Edit Online

Microsoft Exchange Server leverages the concept of database mobility, which is Exchange-managed database-level failovers. Database mobility disconnects databases from servers, adds support for up to 16 copies of a single database, and provides a native experience for adding database copies to a database.

Key characteristics

The key characteristics of mailbox database copies are:

- Up to 16 copies of an Exchange Server mailbox database can be created on multiple Mailbox servers, provided the servers are grouped into a database availability group (DAG), which is a boundary for continuous replication. Exchange Server mailbox databases can be replicated only to the same version Exchange Mailbox servers within a DAG. You can't replicate a database outside of a DAG, nor can you replicate an Exchange 2016 or Exchange 2019 mailbox database to a server running Exchange 2013 or earlier. For detailed information about DAGs, see Database availability groups.
- All Mailbox servers in a DAG must be in the same Active Directory domain.
- Mailbox database copies support the concepts of replay lag time and truncation lag time. Appropriate planning must be performed before enabling these features.
- All database copies can be backed up using an Exchange-aware, Volume Shadow Copy Service (VSS)-based backup application.
- Database copies can be created only on Mailbox servers that don't host the active copy of a database. You can't create two copies of the same database on the same server.
- 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.
- Database copies can be created in the same or different Active Directory sites, and on the same or different network subnets.
- Database copies aren't supported between Mailbox servers with round trip network latency greater than 500 milliseconds (ms).

Mailbox database copies

You can create a mailbox database copy at any time. Mailbox database copies can be distributed across Mailbox servers in a flexible and granular way.

You can create a mailbox database copy using the **Add mailbox database copy** wizard in the Exchange admin center or by using the **Add-MailboxDatabaseCopy** cmdlet in the Exchange Management Shell.

When creating a mailbox database copy, specify the following parameters:

- *Identity*: This parameter specifies the name of the database being copied. Database names must be unique within the Exchange organization.
- *MailboxServer*: This parameter specifies the name of the Mailbox server that will host the database copy. This server must be a member of the same DAG and must not already host a copy of the database.

Optionally, you can also specify:

- *ActivationPreference*: This parameter specifies the activation preference number, which is used as part of Active Manager's best copy selection process. It's also used to redistribute active mailbox databases throughout the DAG when using the RedistributeActiveDatabases.ps1 script. The value for the activation preference is a number equal to or greater than one, where one is at the top of the preference order. The position number cannot be larger than the number of mailbox database copies.

- *ReplayLagTime*: This parameter specifies the amount of time that the Microsoft Exchange Replication service should wait before replaying log files that are copied to the database copy. The format for this parameter is (Days.Hours:Minutes:Seconds). The default setting for this value is 0 seconds. The maximum allowable setting for this value is 14 days. The minimum allowable setting is 0 seconds. Setting the value for replay lag time to 0 turns off log replay delay.

- *TruncationLagTime*: This parameter specifies the amount of time that the Microsoft Exchange Replication service should wait before truncating log files that have replayed into a copy of the database. The time period begins after the log has been successfully replayed into the copy of the database. The format for this parameter is (Days.Hours:Minutes:Seconds). The default setting for this value is 0 seconds. The maximum allowable setting for this value is 14 days. The minimum allowable setting is 0 seconds. Setting the value for truncation lag time to 0 turns off log truncation delay.

- *SeedingPostponed*: This parameter specifies that the task shouldn't automatically seed the database copy on the specified Mailbox server. This option is typically used when you intend to seed a new mailbox database copy by using an existing passive copy of the database (for example, adding a second copy of a specific database to a remote location). When you use this parameter, you must manually seed the database copy using the Update-MailboxDatabaseCopy cmdlet.

For more information about creating, using, and managing mailbox database copies, see Managing mailbox database copies

AutoReseed

8/3/2020 • 5 minutes to read • Edit Online

Automatic Reseed, or AutoReseed, is a feature that replaces standard actions administrators take in response to a disk failure, or a database corruption event, or another issue that necessitates a reseed of a database copy.

Overview of Autoreseed

In an AutoReseed configuration, a standardized storage presentation structure is used, and the administrator picks the starting point. AutoReseed is about restoring redundancy as soon as possible after a drive fails. This involves using mount points to pre-map a set of volumes (including spare volumes) and databases. In the event of a disk failure where the disk is no longer available to the operating system, or is no longer writable, a spare volume is allocated by the system, and the affected database copies are reseeded automatically.

1. The Microsoft Exchange Replication service periodically scans for copies that have a status of FailedAndSuspended. If all database copies on a volume configured for AutoReseed are in a FailedandSuspended state for 15 consecutive minutes, the AutoReseed workflow is initiated.

2. AutoReseed will try to resume the failed and suspended copies up to three times, with a 5-minute sleep in between each attempt. Sometimes, after a FailedandSuspended database copy is resumed, the copy remains in a Failed state. This can happen for a variety of reasons, so this step is designed to handle those cases; AutoReseed will automatically suspend a database copy that has been Failed for 10 consecutive minutes to keep the workflow running. If the suspend and resume actions don't result in a healthy database copy, the workflow continues.

3. When it finds a copy with that status, it performs some prerequisite checks. For example, it will verify that a spare disk is available, that the database and its log files are configured on the same volume, and in the appropriate locations that match the required naming conventions.

4. If the prerequisite checks pass successfully, the Disk Reclaimer function within the Microsoft Exchange Replication service allocates, remaps and formats a spare disk according to the timelines in the table below. AutoReseed will attempt to assign a spare volume up to 5 times, with 1-hour sleeps in between each try.

5. Once a spare has been assigned, AutoReseed will perform an InPlaceSeed operation using the SafeDeleteExistingFiles seeding switch. All databases that were on the affected disk are re-seeded using the active copy of the database as the seeding source.

6. After the seeding operation has been completed, the Microsoft Exchange Replication service verifies that the newly seeded copy is healthy.

Once all retries are exhausted, the workflow stops. If, after 3 days, the database copy is still FailedandSuspended, the workflow state is reset and it starts again from Step 1. This reset/resume behavior is useful (and intentional) since it can take a few days to replace a failed disk, controller, etc.

At this point, if the failure was a disk failure, it would require manual intervention by an operator or administrator to remove and replace the failed disk and reconfigure the replacement disk as a spare.

AutoReseed is configured using three properties of the DAG. Two of the properties refer to the two mount points that are in use. Exchange Server leverages the fact that Windows Server allows multiple mount points per volume. The *AutoDagVolumesRootFolderPath* property refers to the mount point that contains all of the available volumes. This includes volumes that host databases and spare volumes. The *AutoDagDatabasesRootFolderPath* property refers to the mount point that contains the databases. A third DAG property, *AutoDagDatabaseCopiesPerVolume*, is used to configure the number of database copies per volume.

An example AutoReseed configuration is illustrated below.

Example AutoReseed configuration

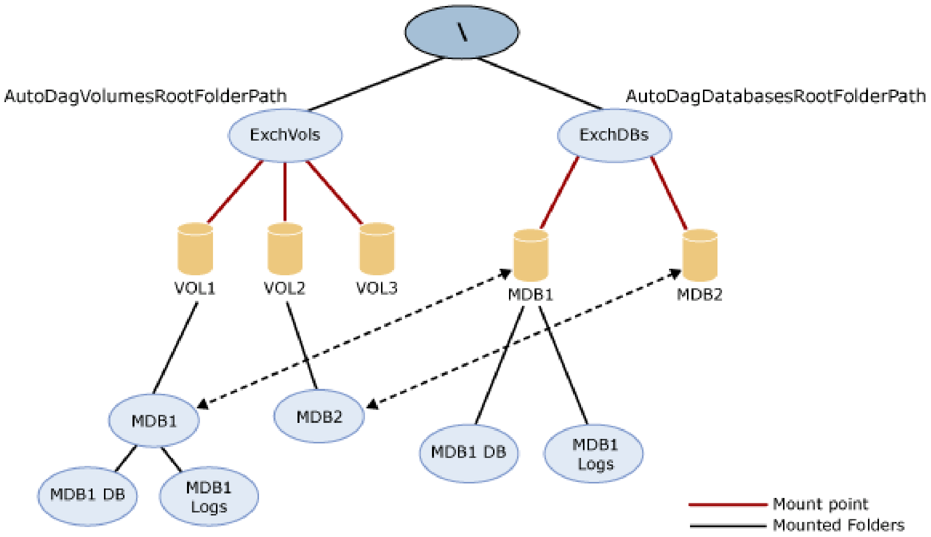

In this example, there are three volumes, two of which will contain databases (VOL1 and VOL2), and one of which is a blank, formatted spare (VOL3).

To configure AutoReseed:

1. All three volumes are mounted under a single mount point. In this example, a mount point of C:\ExchVols is used. This represents the directory used to get storage for Exchange databases.

2. The root directory of the mailbox databases is mounted as another mount point. In this example, a mount point of C:\ExchDBs is used. Next, a directory structure is created so that a parent directory is created for the database, and under the parent directory, two subdirectories are created: one database file and one for the log files.

3. Databases are created. The above example illustrates a simple design using a single database per volume. Thus, on VOL1, there are three directories: the parent directory and two subdirectories (one for MDB1's database file, and one for its logs). Although not depicted in the example image, on VOL2, there would also be three directories: the parent directory, and under that, a directory for MDB2's database file, and one for its log files.

In this configuration, if MDB1 or MDB2 were to experience a failure, a copy of the failed database will be automatically reseeded to VOL3.

Disk Reclaimer

The AutoReseed component that allocates and formats spare disks is called the *Disk Reclaimer*. The Disk Reclaimer component automatically formats spare disks in preparation for automatic reseeding at different intervals, depending on the state of the disk. In order for the Disk Reclaimer to format a disk, certain conditions must met:

- The Disk Reclaimer must be enabled. It is enabled by default, but it can be disabled using Set-DatabaseAvailabilityGroup.

- The volume must have a mount point in the root volumes path (by default, C:\ExchangeVolumes).

- The volume must not have any mount points in the database volumes path (by default, C:\ExchangeDatabases).

- If the volume contains any files, none of the files must have been touched for 24 hours.

In addition to the above conditions, the Disk Reclaimer will only attempt to format a given volume once a day. The following table describes the formatting behavior of the Disk Reclaimer.

STATE OF DISK AND DATABASE COPIES	FORMATTING INTERVAL
Disk is unformatted, or formatted but empty, or formatted but contains files that have not been touched for 24 hours, and there are healthy active database copies in the local Active Directory site that can be used as a seeding source.	1 day
Disk is unformatted, or formatted but empty, or formatted but contains files that have not been touched for 24 hours, but there are no healthy active database copies in the local Active Directory site that can be used as a seeding source.	2 days
Disk is unformatted, or formatted but empty, or formatted but contains files that have not been touched for 24 hours, and there are healthy active database copies in the local Active Directory site that can be used as a seeding source, but there are unknown files outside of the database file (EDB file) and log files.	2 weeks
Disk is unformatted, or formatted but empty, or formatted but contains files that have not been touched for 24 hours, and there are healthy active database copies in the local Active Directory site that can be used as a seeding source, but there are one or more database files (EDB files) for databases that are not present in Active Directory.	2 weeks

MetaCacheDatabase (MCDB) setup

8/3/2020 • 5 minutes to read • Edit Online

The MetaCacheDatabase (MCDB) feature is included in Exchange Server 2019. It allows a database availability group (DAG) to be accelerated by utilizing solid state disks (SSDs). `Manage-MetaCacheDatabase.ps1` is an automation script created for Exchange Server administrators to set up and manage MCDB instances in their Exchange 2019 DAGs.

After installing Exchange Server 2019, you can find `Manage-MetaCacheDatabase.ps1` here: ***drive*:\Program Files\Microsoft\Exchange Server\V15\Scripts**. To make the **Manage-MCDB** CMDLet available in your Exchange Management Shell session, do the following:

```
cd $exscripts
. .\Manage-MetaCacheDatabase.ps1
```

You use this script to configure MCDB prerequisites on a properly configured DAG, to enable or disable MCDB, and to configure and repair MCDB on your servers.

SSD guidance

All SSDs used for MCDB need to be of the same capacity and type. A symmetrical configuration between servers is required, which means there needs to be an identical number of SSDs in each server, and the SSDs all need to be the same size.

> **NOTE**
>
> The **Manage-MCDB** cmdlet will only work with devices exposed as **MediaType SSD** by Windows.

It is recommended to target a 1:3 ratio between SSD and HDD devices per server. Therefore, deploy one SSD for every three HDDs. In order to avoid having to reduce the number of HDDs in the server, consider using M.2 form factor SSDs.

Providing 5% to 6% of SSD capacity relative to total HDD capacity is sufficient for on-premises deployments. For example, if your server contains 100 TB of HDD capacity for mailbox databases, an allocation of 5 TB to 6 TB for SSD capacity is enough.

The SSDs you use should qualify for "mixed use" and support one drive write per day (DWPD) or greater in terms of write endurance.

Prerequisites

The following prerequisites are required for successful configuration and use of MCDB:

1. The DAG is configured for AutoReseed.

   For more information, see the following topics:

   - AutoReseed
   - Configure AutoReseed for a database availability group

2. RAW SSD drives are installed with the same SSD count and size for each server in the DAG. Make sure that

all SSDs are completely empty, unformatted, and not write-protected. To verify this, you can use use DiskPart or Clear-Disk.

3. Exchange Server 2019.

MCDB setup

The process of setting up MCDB can be broken down into four basic steps:

1. Set the correct values for the DAG you want to enable for MCDB.
2. Update Active Directory (AD) settings and wait for propagation (by running `ConfigureMCDBPrerequisite`).
3. Allow MCDB acceleration for each server of the DAG (by running `ServerAllowMCDB`).
4. Create the necessary infrastructure (Volumes, Mount Points) for MCDB on each server (by running `ConfigureMCDBOnServer`).
5. Let databases fail over to pick up the new settings.

After successful execution of all four steps, MCDB acceleration will begin for every database instance with a corresponding MCDB instance.

The following sections describe how to utilize the `Manage-MetaCacheDatabase.ps1` script to achieve the above four steps.

Step 1: Configure proper values on the DAG you want to enable MCDB for

These DAG parameters are used to calculate the proper MCDB size on your SSD drives:

- *AutoDagTotalNumberOfDatabases*: Maximum number of possible active database copies per server that will use MCDB.
- *AutoDagDatabaseCopiesPerDatabase*: The number of active and passive copies each individual database has.
- *AutoDagTotalNumberOfServers*: The amount of servers within your DAG, so between 2 and 16.

For example:

```
Set-DatabaseAvailabilityGroup testdag1 -AutoDagTotalNumberOfDatabases 20 -AutoDagDatabaseCopiesPerDatabase 4 -
AutoDagTotalNumberOfServers 8
```

Step 2: Run Manage-MCDB -ConfigureMCDBPrerequisite

This parameter sets the Active Directory state for the DAG object. Full replication of the Active Directory state is required before MCDB can function properly on all servers.

ParameterSetIdentifier:

- *ConfigureMCDBPrerequisite*

Parameters:

PARAMETER	REQUIRED	DESCRIPTION
DagName	True	Name of the Database availability group.
SSDSizeInBytes	True	The capacity in bytes of each SSD in the server to be used for MCDB.

PARAMETER	REQUIRED	DESCRIPTION
SSDCountPerServer	True	The count of SSD devices to be utilize for MCDB in each server.

Scope:

- DAG: *ConfigureMCDBPrerequisite* operates on a DAG object.

> **NOTE**
>
> MCDB will utilize up to 95% of an SSD's physical capacity. The remaining 5% is kept free to account for file system and partition overhead, as well as for a small amount of additional buffer and over-provisioning.

Example:

```
Manage-MCDB -DagName TestDag1 -ConfigureMCDBPrerequisite -SSDSizeInBytes 5242880000 -SSDCountPerServer 2
```

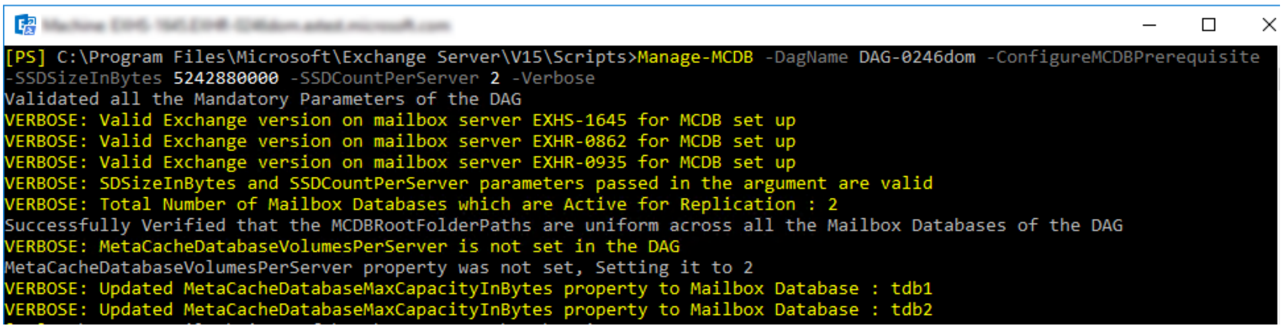

```
[PS] C:\Program Files\Microsoft\Exchange Server\V15\Scripts>Manage-MCDB -DagName DAG-0246dom -ConfigureMCDBPrerequisite
-SSDSizeInBytes 5242880000 -SSDCountPerServer 2 -Verbose
Validated all the Mandatory Parameters of the DAG
VERBOSE: Valid Exchange version on mailbox server EXHS-1645 for MCDB set up
VERBOSE: Valid Exchange version on mailbox server EXHR-0862 for MCDB set up
VERBOSE: Valid Exchange version on mailbox server EXHR-0935 for MCDB set up
VERBOSE: SDSizeInBytes and SSDCountPerServer parameters passed in the argument are valid
VERBOSE: Total Number of Mailbox Databases which are Active for Replication : 2
Successfully Verified that the MCDBRootFolderPaths are uniform across all the Mailbox Databases of the DAG
VERBOSE: MetaCacheDatabaseVolumesPerServer is not set in the DAG
MetaCacheDatabaseVolumesPerServer property was not set, Setting it to 2
VERBOSE: Updated MetaCacheDatabaseMaxCapacityInBytes property to Mailbox Database : tdb1
VERBOSE: Updated MetaCacheDatabaseMaxCapacityInBytes property to Mailbox Database : tdb2
```

Step 3: Run Manage-MCDB -ServerAllowMCDB

This command sets the local state on each DAG member to allow/disallow MCDB population and read acceleration.

ParameterSetIdentifier:

- *ServerAllowMCDB*

Parameters:

PARAMETER	REQUIRED	DESCRIPTION
DagName	True	Name of the Database availability group.
ServerName	True	Specifies the server to enable MetaCacheDatabase on.
ForceFailover	Optional	This Boolean switch can be utilized to cause all databases on a server to fail over. This is required to make all configuration changes take effect and to begin utilizing MCDB after mount points and database instances have been successfully created in Step 4: Run Manage-MCDB -ConfigureMCDBOnServer. It is also needed to disable SSD acceleration.

Scope:

- **Server**: You need to run *ServerAllowMCDB* on each server in the DAG.

Examples:

```
Manage-MCDB -DagName TestDag1 -ServerAllowMCDB $true -ServerName "exhs-5046"
```

```
Manage-MCDB -DagName TestDag1 -ServerAllowMCDB $false -ServerName "exhs-5046" -ForceFailover $true
```

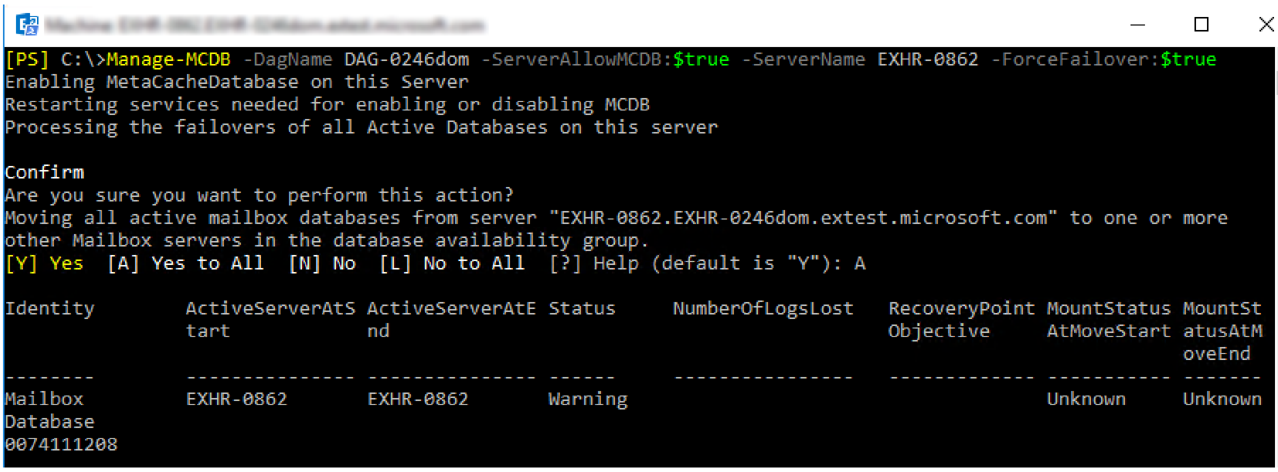

[redacted]

```
[PS] C:\>Manage-MCDB -DagName DAG-0246dom -ServerAllowMCDB:$true -ServerName EXHR-0862 -ForceFailover:$true
Enabling MetaCacheDatabase on this Server
Restarting services needed for enabling or disabling MCDB
Processing the failovers of all Active Databases on this server

Confirm
Are you sure you want to perform this action?
Moving all active mailbox databases from server "EXHR-0862.EXHR-0246dom.extest.microsoft.com" to one or more
other Mailbox servers in the database availability group.
[Y] Yes  [A] Yes to All  [N] No  [L] No to All  [?] Help (default is "Y"): A

Identity        ActiveServerAtS ActiveServerAtE Status     NumberOfLogsLost   RecoveryPoint MountStatus MountSt
                tart            nd                                            Objective     AtMoveStart atusAtM
                                                                                                        oveEnd
--------        --------------- --------------- ------     ----------------   ------------- ----------- -------
Mailbox         EXHR-0862       EXHR-0862       Warning                                     Unknown     Unknown
Database
0074111208
```

Step 4: Run Manage-MCDB -ConfigureMCDBOnServer

This command identifies unformatted SSD devices and formats them, and also creates the necessary mount points on a server for hosting MCDB instances. This parameter set can also be used to re-create mount points on a raw SSD that was added to replace a failed SSD.

ParameterSetIdentifier:

- *ConfigureMCDBOnServer*

Parameters:

PARAMETER	REQUIRED	DESCRIPTION
DagName	True	Name of the Database availability group.
ServerName	True	Specifies the server to identify unformatted SSD devices and create mount points on.
SSDSizeInBytes	True	This is the capacity, in bytes, of each SSD in the server to be used for MCDB.

Scope:

- **Server**: You need to run *ConfigureMCDBOnServer* on each server in the DAG.

Example:

```
Manage-MCDB -DagName TestDag1 -ConfigureMCDBOnServer -ServerName "exhs-4056" -SSDSizeInBytes 5242880000
```

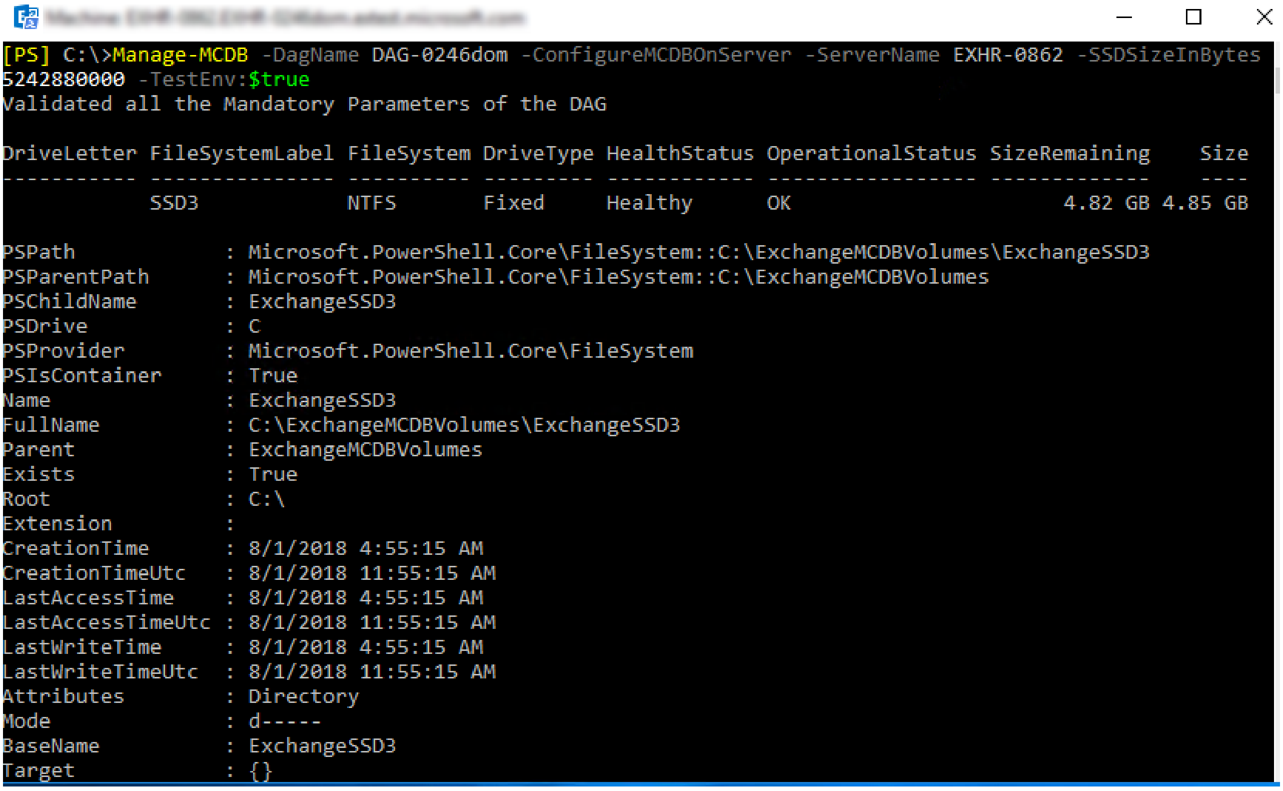

```
[PS] C:\>Manage-MCDB -DagName DAG-0246dom -ConfigureMCDBOnServer -ServerName EXHR-0862 -SSDSizeInBytes
5242880000 -TestEnv:$true
Validated all the Mandatory Parameters of the DAG

DriveLetter FileSystemLabel FileSystem DriveType HealthStatus OperationalStatus SizeRemaining    Size
----------- --------------- ---------- --------- ------------ ----------------- -------------    ----
            SSD3            NTFS       Fixed     Healthy      OK                      4.82 GB 4.85 GB

PSPath            : Microsoft.PowerShell.Core\FileSystem::C:\ExchangeMCDBVolumes\ExchangeSSD3
PSParentPath      : Microsoft.PowerShell.Core\FileSystem::C:\ExchangeMCDBVolumes
PSChildName       : ExchangeSSD3
PSDrive           : C
PSProvider        : Microsoft.PowerShell.Core\FileSystem
PSIsContainer     : True
Name              : ExchangeSSD3
FullName          : C:\ExchangeMCDBVolumes\ExchangeSSD3
Parent            : ExchangeMCDBVolumes
Exists            : True
Root              : C:\
Extension         :
CreationTime      : 8/1/2018 4:55:15 AM
CreationTimeUtc   : 8/1/2018 11:55:15 AM
LastAccessTime    : 8/1/2018 4:55:15 AM
LastAccessTimeUtc : 8/1/2018 11:55:15 AM
LastWriteTime     : 8/1/2018 4:55:15 AM
LastWriteTimeUtc  : 8/1/2018 11:55:15 AM
Attributes        : Directory
Mode              : d-----
BaseName          : ExchangeSSD3
Target            : {}
```

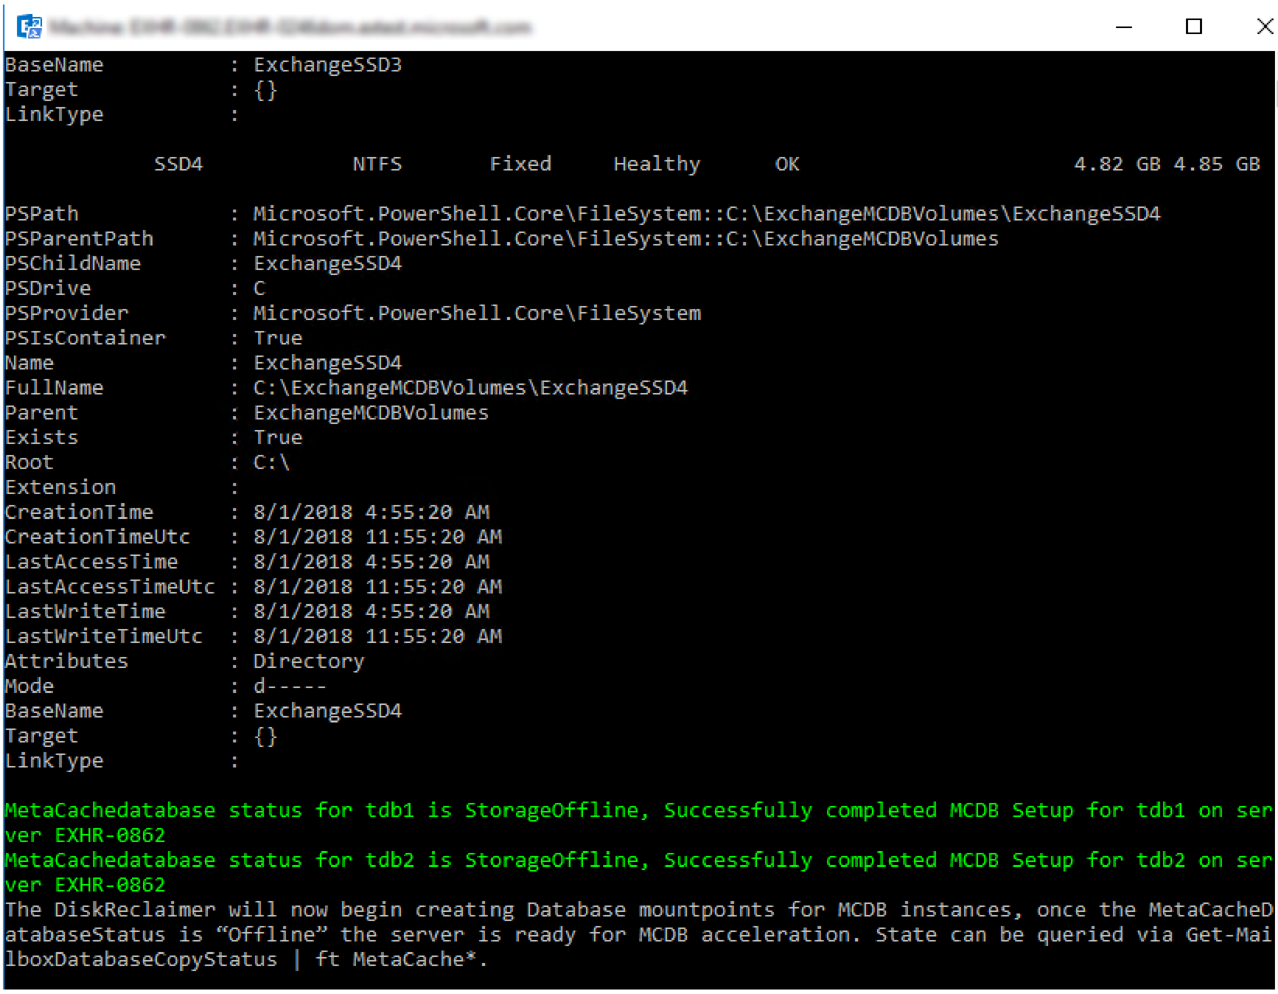

```
BaseName          : ExchangeSSD3
Target            : {}
LinkType          :

            SSD4            NTFS       Fixed     Healthy      OK                      4.82 GB 4.85 GB

PSPath            : Microsoft.PowerShell.Core\FileSystem::C:\ExchangeMCDBVolumes\ExchangeSSD4
PSParentPath      : Microsoft.PowerShell.Core\FileSystem::C:\ExchangeMCDBVolumes
PSChildName       : ExchangeSSD4
PSDrive           : C
PSProvider        : Microsoft.PowerShell.Core\FileSystem
PSIsContainer     : True
Name              : ExchangeSSD4
FullName          : C:\ExchangeMCDBVolumes\ExchangeSSD4
Parent            : ExchangeMCDBVolumes
Exists            : True
Root              : C:\
Extension         :
CreationTime      : 8/1/2018 4:55:20 AM
CreationTimeUtc   : 8/1/2018 11:55:20 AM
LastAccessTime    : 8/1/2018 4:55:20 AM
LastAccessTimeUtc : 8/1/2018 11:55:20 AM
LastWriteTime     : 8/1/2018 4:55:20 AM
LastWriteTimeUtc  : 8/1/2018 11:55:20 AM
Attributes        : Directory
Mode              : d-----
BaseName          : ExchangeSSD4
Target            : {}
LinkType          :

MetaCachedatabase status for tdb1 is StorageOffline, Successfully completed MCDB Setup for tdb1 on ser
ver EXHR-0862
MetaCachedatabase status for tdb2 is StorageOffline, Successfully completed MCDB Setup for tdb2 on ser
ver EXHR-0862
The DiskReclaimer will now begin creating Database mountpoints for MCDB instances, once the MetaCacheD
atabaseStatus is “Offline” the server is ready for MCDB acceleration. State can be queried via Get-Mai
lboxDatabaseCopyStatus | ft MetaCache*.
```

After performing the previous three steps (configuring *ConfigureMCDBPrerequisite*, *ServerAllowMCDB*, and *ConfigureMCDBOnServer*), the MCDB state will display as **Storage Offline**. This means that the environment is prepared and ready for MCDB instances to be created and populated. The next fail over of the database instance will cause the creation of the MCDB instance and enable acceleration. The instances will transition through the health states shown in MCDB health states.

You can use the *ServerAllowMCDB* parameter set to cause fail overs of all DB instances present on a given server. Alternatively, you can use the **Move-ActiveMailboxDatabase** cmdlet to cause individual databases to fail over.

```
Manage-MCDB -DagName TestDag1 -ServerAllowMCDB:$true -ServerName "exhs-5046" -ForceFailover $true
```

MCDB health states

Use **Get-MailboxDatabaseCopyStatus** to query the state of the MCDB instances. There are five states that an MCDB instance can be in, as shown in the following table:

STATE	DESCRIPTION
Disabled	MCDB is turned off.
StorageOffline	Basic infrastructure is missing or inaccessible, such as mount points or file paths. This is the state MCDB is in following an SSD failure.
Offline	Errors at the logical level, for example missing MCDB instances.
Initializing	Transient state, the system is determining what other state it should be in.
Healthy	Ready to serve requests.

Plan for high availability and site resilience

8/3/2020 • 20 minutes to read • Edit Online

During the planning phase, the system architects, administrators, and other key stakeholders should identify the business requirements and the architectural requirements for the deployment; in particular, the requirements about high availability and site resilience.

There are general requirements that must be met for deploying these features, as well as hardware, software, and networking requirements that must also be met.

General requirements

Before deploying a database availability group (DAG) and creating mailbox database copies, make sure that the following system-wide recommendations are met:

- Domain Name System (DNS) must be running. Ideally, the DNS server should accept dynamic updates. If the DNS server doesn't accept dynamic updates, you must create a DNS host (A) record for each Exchange server. Otherwise, Exchange won't function properly.
- Each Mailbox server in a DAG must be a member server in the same domain.
- Adding an Exchange Mailbox server that's also a directory server to a DAG isn't supported.
- The name you assign to the DAG must be a valid, available, and unique computer name of 15 characters or less.

Hardware requirements

Generally, there are no special hardware requirements specific to DAGs or mailbox database copies. The servers used must meet all of the requirements set forth in Exchange Server prerequisites.

Storage requirements

Generally, there are no special storage requirements specific to DAGs or mailbox database copies. DAGs don't require or use cluster-managed shared storage. Cluster-managed shared storage is supported for use in a DAG only when the DAG is configured to use a solution that leverages the Third Party Replication API built into Exchange Server.

Software requirements

Each member of a DAG must be running the same operating system. Exchange Server 2016 is supported on the Windows Server 2012, Windows Server 2012 R2, and Windows Server 2016. Exchange Server 2019 is supported on the Windows Server 2019 operating system. Within a specific DAG, all members must be running the same supported operating system.

In addition to meeting the prerequisites for installing Exchange Server, there are operating system requirements that must be met. DAGs use Windows Failover Clustering technology, and as a result, they require the Standard or Datacenter version of the Windows Server 2012, Windows Server 2012 R2, Windows Server 2016, or Windows Server 2019 operating systems.

Network requirements

There are specific networking requirements that must be met for each DAG and for each DAG member. Each DAG must have a single *MAPI network*, which is used by a DAG member to communicate with other servers (for example, other Exchange servers or directory servers), and zero or more *Replication networks*, which are networks dedicated to log shipping and seeding.

In previous versions of Exchange, we recommended at least two networks (one MAPI network and one Replication network) for DAGs. In Exchange 2016 and Exchange 2019, multiple networks are supported, but our recommendation depends on your physical network topology. If you have multiple physical networks between DAG members that are physically separate from one another, then using a separate MAPI and Replication network provides additional redundancy. If you have multiple networks that are partially physically separate but converge into a single physical network (for example, a single WAN link), then using a single network (preferably 10 gigabit Ethernet) for both MAPI and Replication traffic is recommended. This provides simplicity for the network and the network path.

Consider the following when designing the network infrastructure for your DAG:

- Each member of the DAG must have at least one network adapter that's able to communicate with all other DAG members. If you're using a single network path, we recommend that you use a minimum of 1 gigabit Ethernet, but preferably 10 gigabit Ethernet. In addition, when using a single network adapter in each DAG member, we recommend that you design the overall solution with the single network adapter and path in mind.

- Using two network adapters in each DAG member provides you with one MAPI network and one Replication network, with redundancy for the Replication network and the following recovery behaviors:

    - In the event of a failure affecting the MAPI network, a server failover will occur (assuming there are healthy mailbox database copies that can be activated).

    - In the event of a failure affecting the Replication network, if the MAPI network is unaffected by the failure, log shipping and seeding operations will revert to use the MAPI network, even if the MAPI network has it's *ReplicationEnabled* property set to False. When the failed Replication network is restored to health and ready to resume log shipping and seeding operations, you must manually switch over to the Replication network. To change replication from the MAPI network to a restored Replication network, you can either suspend and resume continuous replication by using the **Suspend-MailboxDatabaseCopy** and **Resume-MailboxDatabaseCopy** cmdlets, or restart the Microsoft Exchange Replication service. We recommend using suspend and resume operations to avoid the brief outage caused by restarting the Microsoft Exchange Replication service.

- Each DAG member must have the same number of networks. For example, if you plan on using a single network adapter in one DAG member, all members of the DAG must also use a single network adapter.

- Each DAG must have no more than one MAPI network. The MAPI network must provide connectivity to other Exchange servers and other services, such as Active Directory and DNS.

- Additional Replication networks can be added, as needed. You can also prevent an individual network adapter from being a single point of failure by using network adapter teaming or similar technology. However, even when using teaming, this doesn't prevent the network itself from being a single point of failure. Moreover, teaming adds unnecessary complexity to the DAG.

- Each network in each DAG member server must be on its own network subnet. Each server in the DAG can be on a different subnet, but the MAPI and Replication networks must be routable and provide connectivity, such that:

    - Each network in each DAG member server is on its own network subnet that's separate from the subnet used by each other network in the server.

    - Each DAG member server's MAPI network can communicate with each other DAG member's MAPI network.

  - Each DAG member server's Replication network can communicate with each other DAG member's Replication network.
  - There is no direct routing that allows heartbeat traffic from the Replication network on one DAG member server to the MAPI network on another DAG member server, or vice versa, or between multiple Replication networks in the DAG.

- Regardless of their geographic location relative to other DAG members, each member of the DAG must have round trip network latency no greater than 500 milliseconds between each other member. As the round trip latency between two Mailbox servers hosting copies of a database increases, the potential for replication not being up to date also increases. Regardless of the latency of the solution, customers should validate that the networks between all DAG members is capable of satisfying the data protection and availability goals of the deployment. Configurations with higher latency values may require special tuning of DAG, replication, and network parameters, such as increasing the number of databases or decreasing the number of mailboxes per database, to achieve the desired goals.

- Round trip latency requirements may not be the most stringent network bandwidth and latency requirement for a multi-datacenter configuration. You must evaluate the total network load, which includes client access, Active Directory, transport, continuous replication, and other application traffic, to determine the necessary network requirements for your environment.

- DAG networks support Internet Protocol version 4 (IPv4) and IPv6. IPv6 is supported only when IPv4 is also used; a pure IPv6 environment isn't supported. Using IPv6 addresses and IP address ranges is supported only when both IPv6 and IPv4 are enabled on that computer, and the network supports both IP address versions. If Exchange Server is deployed in this configuration, all server roles can send data to and receive data from devices, servers, and clients that use IPv6 addresses.

- Automatic Private IP Addressing (APIPA) is a feature of Windows that automatically assigns IP addresses when no Dynamic Host Configuration Protocol (DHCP) server is available on the network. APIPA addresses (including manually assigned addresses from the APIPA address range) aren't supported for use by DAGs or by Exchange Server.

DAG name and IP address requirements

During creation, each DAG is given a unique name, and either assigned one or more static IP addresses, or configured to use DHCP. Regardless of whether you use static or dynamically assigned addresses, any IP address assigned to the DAG must be on the MAPI network.

Each DAG running on Windows Server 2012 requires a minimum of one IP address on the MAPI network. A DAG requires additional IP addresses when the MAPI network is extended across multiple subnets. DAGs running on Windows Server 2012 R2, Windows Server 2016 or Windows Server 2019 that are created without a cluster administrative access point do not require an IP address.

The following figure illustrates a DAG where all nodes in the DAG have the MAPI network on the same subnet.

DAG with MAPI network on same subnet

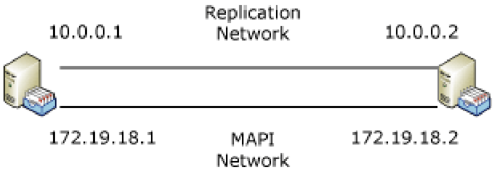

In this example, the MAPI network in each DAG member is on the 172.19.18. *x* subnet. As a result, the DAG requires a single IP address on that subnet.

The next figure illustrates a DAG that has a MAPI network that extends across two subnets: 172.19.18. *x* and

172.19.19. *x*.

DAG with MAPI network on multiple subnets

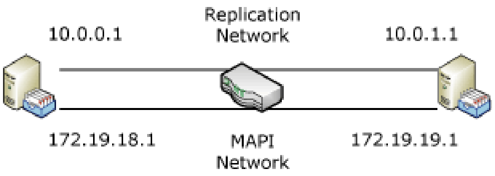

In this example, the MAPI network in each DAG member is on a separate subnet. As a result, the DAG requires two IP addresses, one for each subnet on the MAPI network.

Each time the DAG's MAPI network is extended across an additional subnet, an additional IP address for that subnet must be configured for the DAG. Each IP address that's configured for the DAG is assigned to and used by the DAG's underlying failover cluster. The name of the DAG is also used as the name for the underlying failover cluster.

At any specific time, the cluster for the DAG will use only one of the assigned IP addresses. Windows Failover Clustering registers this IP address in DNS when the cluster IP address and Network Name resources are brought online. In addition to using an IP address and network name, a cluster name object (CNO) is created in Active Directory. The name, IP address, and CNO for the cluster are used internally by the system to secure the DAG and for internal communication purposes. Administrators and end users don't need to interface with or connect to the DAG name or IP address.

NOTE

Although the cluster's IP address and network name are used internally by the system, there is no hard dependency in Exchange Server that these resources be available. Even if the underlying cluster's administrative access point (for example, its IP address and Network Name resources) is offline, internal communication still occurs within the DAG by using the DAG member server names. However, we recommend that you periodically monitor the availability of these resources to ensure that they aren't offline for more than 30 days. If the underlying cluster is offline for more than 30 days, the cluster CNO account may be invalidated by the garbage collection mechanism in Active Directory.

Network adapter configuration for DAGs

Each network adapter must be configured properly based on its intended use. A network adapter that's used for a MAPI network is configured differently from a network adapter that's used for a Replication network. In addition to configuring each network adapter correctly, you must also configure the network connection order in Windows so that the MAPI network is at the top of the connection order. For detailed steps about how to modify the network connection order, see Modify the protocol bindings and network provider order.

MAPI network adapter configuration

A network adapter intended for use by a MAPI network should be configured as described in the following table.

NETWORKING FEATURES	SETTINGS
Client for Microsoft Networks	Enabled
QoS Packet Scheduler	Optionally enabled
File and Printer Sharing for Microsoft Networks	Enabled
Internet Protocol version 6 (TCP/IP v6)	Enabled
Internet Protocol version 4 (TCP/IP v4)	Enabled

NETWORKING FEATURES	SETTINGS
Link-Layer Topology Discovery Mapper I/O Driver	Enabled
Link-Layer Topology Discovery Responder	Enabled

The TCP/IP v4 properties for a MAPI network adapter are configured as follows:

- The IP address for a DAG member's MAPI network can be manually assigned or configured to use DHCP. If DHCP is used, we recommend using persistent reservations for the server's IP address.
- The MAPI network typically uses a default gateway, although one isn't required.
- At least one DNS server address must be configured. Using multiple DNS servers is recommended for redundancy.
- The **Register this connection's addresses in DNS** check box should be selected.

Replication network adapter configuration

A network adapter intended for use by a Replication network should be configured as described in the following table.

NETWORKING FEATURES	SETTINGS
Client for Microsoft Networks	Disabled
QoS Packet Scheduler	Optionally enabled
File and Printer Sharing for Microsoft Networks	Disabled
Internet Protocol version 6 (TCP/IP v6)	Enabled
Internet Protocol version 4 (TCP/IP v4)	Enabled
Link-Layer Topology Discovery Mapper I/O Driver	Enabled
Link-Layer Topology Discovery Responder	Enabled

The TCP/IP v4 properties for a Replication network adapter are configured as follows:

- The IP address for a DAG member's Replication network can be manually assigned or configured to use DHCP. If DHCP is used, we recommend using persistent reservations for the server's IP address.
- Replication networks typically don't have default gateways, and if the MAPI network has a default gateway, no other networks should have default gateways. Routing of network traffic on a Replication network can be configured by using persistent, static routes to the corresponding network on other DAG members using gateway addresses that have the ability to route between the Replication networks. All other traffic not matching this route will be handled by the default gateway that's configured on the adapter for the MAPI network.
- DNS server addresses shouldn't be configured.
- The **Register this connection's addresses in DNS** check box shouldn't be selected.

Witness server requirements

A *witness server* is a server outside a DAG that's used to achieve and maintain quorum when the DAG has an even number of members. DAGs with an odd number of members don't use a witness server. All DAGs with an even number of members must use a witness server. The witness server can be any computer running Windows Server. There is no requirement that the version of the Windows Server operating system of the witness server matches the operating system used by the DAG members.

Quorum is maintained at the cluster level, underneath the DAG. A DAG has quorum when the majority of its members are online and can communicate with the other online members of the DAG. This notion of quorum is one aspect of the concept of quorum in Windows failover clustering. A related and necessary aspect to quorum in failover clusters is the *quorum resource*. The quorum resource is a resource inside a failover cluster that provides a means for arbitration leading to cluster state and membership decisions. The quorum resource also provides persistent storage for storing configuration information. A companion to the quorum resource is the *quorum log*, which is a configuration database for the cluster. The quorum log contains information such as which servers are members of the cluster, what resources are installed in the cluster, and the state of those resources (for example, online or offline).

It's critical that each DAG member have a consistent view of how the DAG's underlying cluster is configured. The quorum acts as the definitive repository for all configuration information relating to the cluster. The quorum is also used as a tie-breaker to avoid *split-brain* syndrome. Split brain syndrome is a condition that occurs when DAG members can't communicate with each other but are running. Split brain syndrome is prevented by always requiring a majority of the DAG members (and in the case of DAGs with an even number of member, the DAG witness server) to be available and interacting for the DAG to be operational.

Planning for site resilience

Every day, more businesses recognize that access to a reliable and available messaging system is fundamental to their success. For many organizations, the messaging system is part of the business continuity plans, and their messaging service deployment is designed with site resilience in mind. Fundamentally, many site resilient solutions involve the deployment of hardware in a second datacenter.

Ultimately, the overall design of a DAG, including the number of DAG members and the number of mailbox database copies, will depend on each organization's recovery service level agreements (SLAs) that cover various failure scenarios. During the planning stage, the solution's architects and administrators identify the requirements for the deployment, including in particular the requirements for site resilience. They identify the locations to be used and the required recovery SLA targets. The SLA will identify two specific elements that should be the basis for the design of a solution that provides high availability and site resilience: the recovery time objective and the recovery point objective. Both of these values are measured in minutes. The recovery time objective is how long it takes to restore service. The recovery point objective refers to how current the data is after the recovery operation has completed. An SLA may also be defined for restoring the primary datacenter to full service after its problems are corrected.

The solution's architects and administrators will also identify which set of users require site resilience protection, and determine if the multiple site solution will be an active/passive or active/active configuration. In an active/passive configuration, no users are normally hosted in the standby datacenter. In an active/active configuration, users are hosted in both locations, and some percentage of the total number of databases within the solution has a preferred active location in a second datacenter. When service for the users of one datacenter fails, those users are activated in the other datacenter.

Constructing the appropriate SLAs often requires answering the following basic questions:

- What level of service is required after the primary datacenter fails?
- Do users need their data or just messaging services?
- How rapidly is data required?

- How many users must be supported?
- How will users access their data?
- What is the standby datacenter activation SLA?
- How is service moved back to the primary datacenter?
- Are the resources dedicated to the site resilience solution?

By answering these questions, you begin to shape a site resilient design for your messaging solution. A core requirement of recovery from site failure is to create a solution that gets the necessary data to the backup datacenter that hosts the backup messaging service.

Certificate planning

There are no unique or special design considerations for certificates when deploying a DAG in a single datacenter. However, when extending a DAG across multiple datacenters in a site resilient configuration, there are some specific considerations with respect to certificates. Generally, your certificate design will depend on the clients in use, as well as the certificate requirements by other applications that use certificates. But there are some specific recommendations and best practices you should follow with respect to the type and number of certificates.

As a best practice, you should minimize the number of certificates you use for your Exchange servers and reverse proxy servers. We recommend using a single certificate for all of these service endpoints in each datacenter. This approach minimizes the number of certificates that are needed, which reduces both cost and complexity for the solution.

For Outlook Anywhere clients, we recommend that you use a single subject alternative name (SAN) certificate for each datacenter, and include multiple host names in the certificate. To ensure Outlook Anywhere connectivity after a database, server, or datacenter switchover, you must use the same Certificate Principal Name on each certificate, and configure the Outlook Provider Configuration object in Active Directory with the same Principal Name in Microsoft-Standard Form (msstd). For example, if you use a Certificate Principal Name of mail.contoso.com, you would configure the attribute as follows.

```
Set-OutlookProvider EXPR -CertPrincipalName "msstd:mail.contoso.com"
```

Some applications that integrate with Exchange have specific certificate requirements that may require using additional certificates. Exchange Server can co-exist with Office Communications Server (OCS). OCS requires certificates with 1024-bit or greater certificates that use the OCS server name for the Certificate Principal Name. Because using an OCS server name for the Certificate Principal Name would prevent Outlook Anywhere from working properly, you would need to use an additional and separate certificate for the OCS environment.

Network planning

In addition to the specific networking requirements that must be met for each DAG, as well as for each server that's a member of a DAG, there are some requirements and recommendations that are specific to site resilience configurations. As with all DAGs, whether the DAG members are deployed in a single site or in multiple sites, the round-trip return network latency between DAG members must be no greater than 500 milliseconds. In addition, there are specific configuration settings that are recommended for DAGs that are extended across multiple sites:

- **MAPI networks should be isolated from Replication networks**: Windows network policies, Windows firewall policies, or router access control lists (ACLs) should be used to block traffic between the MAPI network and the Replication networks. This configuration is necessary to prevent network heartbeat cross talk.
- **Client-facing DNS records should have a Time to Live (TTL) value of 5 minutes**: The amount of downtime that clients experience is dependent not just on how quickly a switchover can occur, but also on how quickly DNS replication occurs and the clients query for updated DNS information. DNS records for all

Exchange client services, including Outlook on the web (formerly known as Outlook Web App), Exchange ActiveSync, Exchange Web Services, Outlook Anywhere, SMTP, POP3, and IMAP4 in both the internal and external DNS servers should be set with a TTL of 5 minutes.

- **Use static routes to configure connectivity across Replication networks**: To provide network connectivity between each of the Replication network adapters, use persistent static routes. This is a quick and one-time configuration that's performed on each DAG member when using static IP addresses. If you're using DHCP to obtain IP addresses for your Replication networks, you can also use it to assign static routes for the replication, thereby simplifying the configuration process.

General site resilience planning

In addition to the requirements listed above for high availability, there are other recommendations for deploying Exchange Server in a site resilient configuration (for example, extending a DAG across multiple datacenters). What you do during the planning phase will directly affect the success of your site resilience solution. For example, poor namespace design can cause difficulties with certificates, and an incorrect certificate configuration can prevent users from accessing services.

To minimize the time it takes to activate a second datacenter, and allow the second datacenter to host the service endpoints of a failed datacenter, the appropriate planning must be completed. The following are examples:

- The SLA goals for the site resilience solution must be well understood and documented.
- The servers in the second datacenter must have sufficient capacity to host the combined user population of both datacenters.
- The second datacenter must have all services enabled that are provided in the primary datacenter (unless the service isn't included as part of the site resilience SLA). This includes Active Directory, networking infrastructure (for example, DNS or TCP/IP), telephony services (if Unified Messaging in Exchange 2016 is in use), and site infrastructure (such as power or cooling).
- For some services to be able to service users from the failed datacenter, they must have the proper server certificates configured. Some services don't allow instancing (for example, POP3 and IMAP4) and only allow the use of a single certificate. In these cases, either the certificate must be a SAN certificate that includes multiple names, or the multiple names must be similar enough so that a wildcard certificate can be used (assuming the security policies of the organization allows the use of wildcard certificates).
- The necessary services must be defined in the second datacenter. For example, if the first datacenter has three different SMTP URLs on different transport servers, the appropriate configuration must be defined in the second datacenter to enable at least one (if not all three) transport server to host the workload.
- The necessary network configuration must be in place to support the datacenter switchover. This might mean making sure that the load balancing configurations are in place, that global DNS is configured, and that the Internet connection is enabled with the appropriate routing configured.
- The strategy for enabling the DNS changes necessary for a datacenter switchover must be understood. The specific DNS changes, including their TTL settings, must be defined and documented to support the SLA in effect.
- A strategy for testing the solution must also be established and factored into the SLA. Periodic validation of the deployment is the only way to guarantee that the quality and viability of the deployment doesn't degrade over time. After the deployment is validated, we recommend that the part of the configuration that directly affects the success of the solution be explicitly documented. In addition, we recommend that you enhance your change management processes around those segments of the deployment.

Deploying high availability and site resilience

8/3/2020 • 11 minutes to read • Edit Online

Microsoft Exchange Server uses the concept known as *incremental deployment* for both high availability and site resilience. You simply install two or more Exchange Mailbox servers as stand-alone servers, and then incrementally configure them and mailbox databases for high availability and site resilience, as needed.

Overview of the deployment process

While the actual steps used by each organization may vary slightly, the overall process for deploying Exchange Server in a highly available or site resilient configuration is generally the same. After performing the necessary planning and design tasks for building and deploying a database availability group (DAG) and creating mailbox database copies, you would:

1. Create a DAG. For detailed steps, see Create a database availability group. It's important to note that all servers within a DAG must be running the same version of Exchange. For example, you can't mix Exchange 2013 and Exchange 2016 servers in the same DAG.

2. If necessary, pre-stage the cluster name object (CNO). Pre-staging the CNO is required when deploying a DAG with Mailbox servers running Windows Server 2012. If you're deploying a DAG without an administrative access point using Mailbox servers running Windows Server 2012 R2, then you do not need to pre-stage a CNO. Pre-staging is also required in environments where computer account creation is restricted or where computer accounts are created in a container other than the default computers container. For detailed steps, see Pre-stage the cluster name object for a database availability group.

3. Add two or more Mailbox servers to the DAG. For detailed steps, see Manage database availability group membership.

4. Configure the DAG properties as needed:

5. Optionally configure DAG encryption and compression, replication port, DAG IP addresses, and other DAG properties. For detailed steps, see Configure database availability group properties.

6. Enable Datacenter Activation Coordination (DAC) mode for the DAG. This protects the DAG from database-level split brain conditions during switchback to the primary datacenter after a datacenter switchover has been performed, and it enables the use of the built-in DAG recovery cmdlets. For more information, see Datacenter Activation Coordination mode.

7. Add mailbox database copies across Mailbox servers in the DAG. For detailed steps, see Add a mailbox database copy.

Example deployment: four-member DAG in two datacenters

This example details how an organization, Contoso, Ltd., is configuring and deploying a four-member DAG that will be extended across two physical locations: Boston and Oklahoma City.

Base infrastructure

Each location contains the infrastructure elements that are necessary to operate a messaging infrastructure based on Exchange Server, namely:

- Directory services (either Active Directory or Active Directory Domain Services (AD DS))

- Domain Name System (DNS) name resolution

- Multiple Exchange servers running Client Access services
- Multiple Exchange Mailbox servers

The following figure illustrates the Contoso configuration.

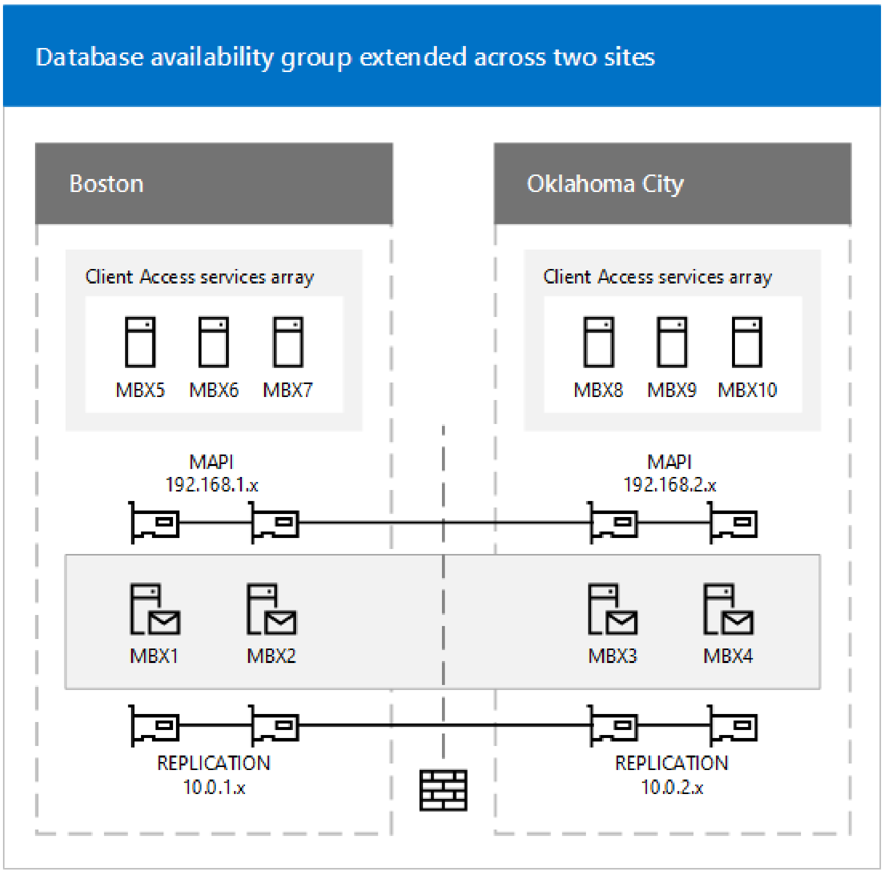

Network configuration

As illustrated in the previous figure, the solution involves the use of multiple subnets and multiple networks. Each Mailbox server in the DAG has two network adapters on separate subnets. In each Mailbox server, one network adapter will be used for the MAPI network (192.168. *x. x*) and one network adapter will be used for the Replication network (10.0. *x. x*). Only the MAPI network provides connectivity to Active Directory, DNS services, other Exchange servers and clients. The adapter used for the Replication network in each member provides connectivity only to the Replication network adapters in the other members of the DAG.

The settings for each network adapter in each node are detailed in the following table.

NAME	IPV4 ADDRESS	SUBNET MASK	DEFAULT GATEWAY
MBX1 (MAPI)	192.168.1.4	255.255.255.0	192.168.1.1
MBX2 (MAPI)	192.168.1.5	255.255.255.0	192.168.1.1
MBX3 (MAPI)	192.168.2.4	255.255.255.0	192.168.2.1
MBX4 (MAPI)	192.168.2.5	255.255.255.0	192.168.2.1
MBX1 (Replication)	10.0.1.4	255.255.255.0	None
MBX2 (Replication)	10.0.1.5	255.255.255.0	None
MBX3 (Replication)	10.0.2.4	255.255.255.0	None

NAME	IPV4 ADDRESS	SUBNET MASK	DEFAULT GATEWAY
MBX4 (Replication)	10.0.2.5	255.255.255.0	None

As shown in the preceding table, adapters used for Replication networks don't use default gateways. To provide network connectivity between each of the Replication network adapters, Contoso uses persistent static routes, which they configure by using the Netsh.exe tool.

To configure routing for the Replication network adapters on MBX1 and MBX2, the following command was run on each server.

```
netsh interface ipv4 add route 10.0.2.0/24 <NetworkName> 10.0.1.254
```

To configure routing for the Replication network adapters on MBX3 and MBX4, the following command was run on each server.

```
netsh interface ipv4 add route 10.0.1.0/24 <NetworkName> 10.0.2.254
```

The following additional network settings have also been configured:

- The **Register this connection's addresses in DNS** check box is selected for each DAG member's MAPI network adapter, and cleared for each Replication network adapter.
- At least one DNS server address is configured for each DAG member's MAPI network adapter, and none are configured for the Replication network adapters. For redundancy, Contoso is using multiple DNS server addresses for their MAPI network adapters.
- Contoso doesn't use the Windows Firewall and have turned it off on their servers.

After the network adapters have been configured, Contoso is ready to create a DAG and add the Mailbox servers to the DAG.

Database availability group creation and configuration

The administrator has decided to create a Windows PowerShell command-line interface script that performs several tasks:

- It uses the New-DatabaseAvailabilityGroup cmdlet to create the DAG. Because BOSTON is considered to be the primary datacenter, Contoso has chosen to use a witness server in the same datacenter, namely, MBX5.
- It uses the Set-DatabaseAvailabilityGroup cmdlet to preconfigure an alternate witness server and alternate witness directory in case a datacenter switchover is ever necessary.
- It uses the Add-DatabaseAvailabilityGroupServer cmdlet to add each of the four Mailbox servers to the DAG.
- It uses the Set-DatabaseAvailabilityGroup cmdlet to configure the DAG for DAC mode. For more information about DAC mode, see Datacenter Activation Coordination mode.

The following are the commands used in the script:

```
New-DatabaseAvailabilityGroup -Name DAG1 -WitnessServer MBX5 -WitnessDirectory C:\DAGWitness\DAG1.contoso.com
-DatabaseAvailabilityGroupIPAddresses 192.168.1.8,192.168.2.8
```

The preceding command creates the DAG DAG1, configures MBX5 to act as the witness server, configures a specific witness directory (C:\DAGWitness\DAG1.contoso.com), and configures two IP addresses for the DAG (one for each subnet on the MAPI network).

```
Set-DatabaseAvailabilityGroup -Identity DAG1 -AlternateWitnessDirectory C:\DAGWitness\DAG1.contoso.com -
AlternateWitnessServer MBX10
```

The preceding command configures DAG1 to use an alternate witness server of MBX10 and an alternate witness directory on MBX10 that uses the same path that was configured on MBX5.

> **NOTE**
>
> Using the same path isn't required; Contoso has chosen to do this to standardize their configuration.

```
Add-DatabaseAvailabilityGroupServer -Identity DAG1 -MailboxServer MBX1
Add-DatabaseAvailabilityGroupServer -Identity DAG1 -MailboxServer MBX3
Add-DatabaseAvailabilityGroupServer -Identity DAG1 -MailboxServer MBX2
Add-DatabaseAvailabilityGroupServer -Identity DAG1 -MailboxServer MBX4
```

The preceding commands add each of the Mailbox servers, one at a time, to the DAG. The commands also install the Windows Failover Clustering component on each Mailbox server (if it isn't already installed), create a failover cluster, and join each Mailbox server to the newly created cluster.

```
Set-DatabaseAvailabilityGroup -Identity DAG1 -DatacenterActivationMode DagOnly
```

The preceding command enables DAC mode for the DAG.

Mailbox databases and mailbox database copies

After creating the DAG and adding the Mailbox servers to the DAG, Contoso prepares to create mailbox databases and mailbox database copies. To meet their criteria for failure resistance, Contoso is planning to configure each mailbox database with three non-lagged database copies, and one lagged database copy. The lagged copy will have a configured log replay delay of three days.

This configuration provides a total of four copies for each database (one active, two non-lagged passives, and a lagged passive). Contoso plans on having four active databases per server. Therefore the Contoso solution contains 16 total database copies.

As shown in the following figure, Contoso is taking a balanced approach to their database layout.

Database copy layout for Contoso, Ltd

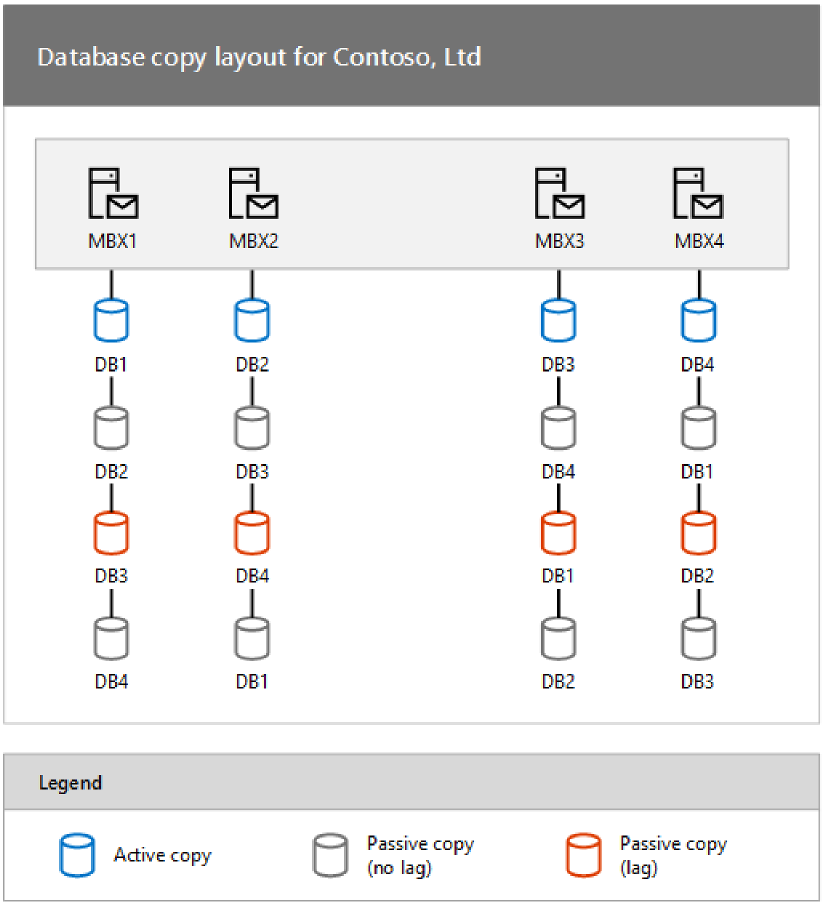

Each Mailbox server hosts an active mailbox database copy, two non-lagged passive database copies, and one lagged passive database copy. The lagged copy of each active mailbox database is hosted on a Mailbox server in the other site.

To create this configuration, the administrator runs several commands.

On MBX1, run the following commands.

```
Add-MailboxDatabaseCopy -Identity DB1 -MailboxServer MBX2
Add-MailboxDatabaseCopy -Identity DB1 -MailboxServer MBX4
Add-MailboxDatabaseCopy -Identity DB1 -MailboxServer MBX3 -ReplayLagTime 3.00:00:00 -SeedingPostponed
Suspend-MailboxDatabaseCopy -Identity DB1\MBX3 -SuspendComment "Seed from MBX4" -Confirm:$False
Update-MailboxDatabaseCopy -Identity DB1\MBX3 -SourceServer MBX4
Suspend-MailboxDatabaseCopy -Identity DB1\MBX3 -ActivationOnly
```

On MBX2, run the following commands.

```
Add-MailboxDatabaseCopy -Identity DB2 -MailboxServer MBX1
Add-MailboxDatabaseCopy -Identity DB2 -MailboxServer MBX3
Add-MailboxDatabaseCopy -Identity DB2 -MailboxServer MBX4 -ReplayLagTime 3.00:00:00 -SeedingPostponed
Suspend-MailboxDatabaseCopy -Identity DB2\MBX4 -SuspendComment "Seed from MBX3" -Confirm:$False
Update-MailboxDatabaseCopy -Identity DB2\MBX4 -SourceServer MBX3
Suspend-MailboxDatabaseCopy -Identity DB2\MBX4 -ActivationOnly
```

On MBX3, run the following commands.

```
Add-MailboxDatabaseCopy -Identity DB3 -MailboxServer MBX4
Add-MailboxDatabaseCopy -Identity DB3 -MailboxServer MBX2
Add-MailboxDatabaseCopy -Identity DB3 -MailboxServer MBX1 -ReplayLagTime 3.00:00:00 -SeedingPostponed
Suspend-MailboxDatabaseCopy -Identity DB3\MBX1 -SuspendComment "Seed from MBX2" -Confirm:$False
Update-MailboxDatabaseCopy -Identity DB3\MBX1 -SourceServer MBX2
Suspend-MailboxDatabaseCopy -Identity DB3\MBX1 -ActivationOnly
```

On MBX4, run the following commands.

```
Add-MailboxDatabaseCopy -Identity DB4 -MailboxServer MBX3
Add-MailboxDatabaseCopy -Identity DB4 -MailboxServer MBX1
Add-MailboxDatabaseCopy -Identity DB4 -MailboxServer MBX2 -ReplayLagTime 3.00:00:00 -SeedingPostponed
Suspend-MailboxDatabaseCopy -Identity DB4\MBX2 -SuspendComment "Seed from MBX1" -Confirm:$False
Update-MailboxDatabaseCopy -Identity DB4\MBX2 -SourceServer MBX1
Suspend-MailboxDatabaseCopy -Identity DB4\MBX2 -ActivationOnly
```

In the preceding examples for the **Add-MailboxDatabaseCopy** cmdlet, the *ActivationPreference* parameter wasn't specified. The task automatically increments the activation preference number with each copy that's added. The original database always has a preference number of 1. The first copy added with the **Add-MailboxDatabaseCopy** cmdlet is automatically assigned a preference number of 2. Assuming no copies are removed, the next copy added is automatically assigned a preference number of 3, and so forth. Thus, in the preceding examples, the passive copy in the same datacenter as the active copy has an activation preference number of 2; the non-lagged passive copy in the remote datacenter has an activation preference number of 3, and the lagged passive copy in the remote datacenter has an activation preference number of 4.

Although there are two copies of each active database across the WAN in the other location, seeding over the WAN was only performed once. This is because Contoso is leveraging the Exchange Server ability to use a passive copy of a database as the source for seeding. Using the Add-MailboxDatabaseCopy cmdlet with the *SeedingPostponed* parameter prevents the task from automatically seeding the new database copy being created. Then, the administrator can suspend the un-seeded copy, and by using the Update-MailboxDatabaseCopy cmdlet with the *SourceServer* parameter, the administrator can specify the local copy of the database as the source of the seeding operation. As a result, seeding of the second database copy added to each location happens locally and not over the WAN.

> **NOTE**
>
> In the preceding example, the non-lagged database copy is seeded over the WAN, and that copy is then used to seed the lagged copy of the database that's in the same datacenter as the non-lagged copy.

Contoso has configured one of the passive copies of each mailbox database as a lagged database copy to provide protection against the extremely rare but catastrophic case of database logical corruption. As a result, the administrator is configuring the lagged copies as blocked for activation by using the Suspend-MailboxDatabaseCopy cmdlet with the *ActivationOnly* parameter. This ensures that the lagged database copies won't be activated if a database or server failover occurs.

Validating the solution

After the solution has been deployed and configured, the administrator performs several tasks that validate the solution's readiness prior to moving production mailboxes to the databases in the DAG. The solution should be tested and inspected using several methods, including failure simulations. To validate the solution, the administrator performs several tasks.

To verify the overall health of the DAG, the administrator runs the Test-ReplicationHealth cmdlet. This cmdlet checks several aspects of the replication and replay status to provide information about each Mailbox server and database copy in the DAG.

To verify replication and replay activity, the administrator runs the Get-MailboxDatabaseCopyStatus cmdlet. This cmdlet can provide real-time status information about a specific mailbox database copy or for all mailbox database copies on a specific server. For more information about monitoring the health and status of replicated databases in a DAG, see Monitor database availability groups.

To verify that switchovers work as expected, the administrator uses the Move-ActiveMailboxDatabase cmdlet to perform a series of database switchovers and server switchovers. When these tasks have completed successfully, the administrator uses the same cmdlet to move the active database copies back to their original locations.

To verify the expected behaviors in various failure scenarios, the administrator performs several tasks that either simulate failures or actually cause failures to occur. For example, the administrator might:

- Unplug the power cord on MBX1, thereby triggering a server failover. The administrator then verifies that DB1 becomes active on another server (preferably MBX2, based on the activation preference values).

- Unplug the network cable for the MAPI network adapter on MBX2, thereby triggering a server failover. The administrator then verifies that DB2 becomes active on another server (preferably MBX1, based on the activation preference values).

- Take the disk used by the active copy of DB3 offline, thereby triggering a database failover. The administrator then verifies that DB3 becomes active on another server (preferably MBX4, based on activation preference values).

There may be other failure scenarios that are tested by an organization, based on the business needs. After simulating a single failure (such as pulling the power plug), and verifying the solution's recovery behavior, the administrator may revert the solution back to its original configuration. In some cases, the solution may be tested for multiple concurrent failures. Ultimately, your solution test plan will dictate whether the solution is reverted back to its original configuration after each failure simulation has been completed.

In addition, an administrator may decide to disconnect the network connection between the two datacenters, thereby simulating a site failure. Performing a datacenter switchover is a much more involved and coordinated process; however, we recommend the process if the solution being deployed is intended to provide site resilience for the messaging services and data.

Transitioning to operations

After the solution has been deployed, it can be extended further using incremental deployment. At this point, management of the solution would also transition to operation processes, in which the following tasks would be performed:

- Monitor the health and status of DAGs and mailbox database copies. For more information, see Monitor database availability groups.

- Perform database switchovers as needed. For detailed steps about how to perform a database switchover, see Activate a mailbox database copy.

For more information about managing the solution, see Managing high availability and site resilience.

Managing high availability and site resilience

8/3/2020 • 7 minutes to read • Edit Online

After you build, validate, and deploy a Microsoft Exchange Server high availability or site resilience solution, the solution transitions from the deployment phase to the operational phase of the overall solution lifecycle. The operational phase consists of several tasks, and all tasks are related to one of the following areas: database availability groups (DAGs), mailbox database copies, performing proactive monitoring, and managing switchovers and failovers.

Database availability group management

The operational management tasks associated with DAGs include:

- **Creating one or more DAGs**: Creating a DAG is typically a one-time procedure performed during the deployment phase of the solution lifecycle. However, there may be reasons for creating DAGs that occur during the operational phase, for example:

  - The DAG is configured for third-party replication mode, and you want to revert to using continuous replication. You can't convert a DAG back to continuous replication; you need to create a DAG.

  - You have servers in multiple domains. All members of the same DAG must also be members of the same domain.

- **Managing DAG membership**: Managing DAG members is an infrequent task typically performed during the deployment phase of the solution lifecycle. However, because of the flexibility provided by incremental deployment, managing DAG membership may also be performed throughout the solution lifecycle.

- **Configuring DAG properties**: Each DAG has various properties that can be configured as needed. These properties include:

  - **Witness server and witness directory**: The witness server is a server outside the DAG that acts as a quorum voter when the DAG contains an even number of members. The witness directory is a directory created and shared on the witness server for use by the system in maintaining a quorum.

  - **IP addresses**: Each DAG will have one or more IPv4 addresses, and optionally, one or more IPv6 addresses. The IP addresses assigned to the DAG are used by the DAG's underlying cluster. The number of IPv4 addresses assigned to the DAG equals the number of subnets that comprise the MAPI network used by the DAG. You can configure the DAG to use static IP addresses or to obtain addresses automatically by using Dynamic Host Configuration Protocol (DHCP).

  - **Datacenter Activation Coordination mode**: Datacenter Activation Coordination mode is a property setting on a DAG that's designed to prevent split-brain conditions at the database level, in a scenario in which you're restoring service to a primary datacenter after a datacenter switchover has been performed. For more information about Datacenter Activation Coordination mode, see Datacenter Activation Coordination mode.

  - **Alternate witness server and alternate witness directory**: The alternate witness server and alternate witness directory are values that you can preconfigure as part of the planning process for a datacenter switchover. These refer to the witness server and witness directory that will be used when a datacenter switchover has been performed.

  - **Replication port**: By default, all DAGs use TCP port 64327 for continuous replication. You can modify the DAG to use a different TCP port for replication by using the *ReplicationPort* parameter of the Set-DatabaseAvailabilityGroup cmdlet.

  - **Network discovery**: You can force the DAG to rediscover networks and network interfaces. This operation is used when you add or remove networks or introduce new subnets. Rediscovery of all DAG networks can be forced by using the *DiscoverNetworks* parameter of the Set-DatabaseAvailabilityGroup cmdlet.

  - **Network compression**: By default, DAGs use compression only between DAG networks on different subnets. You can enable compression for all DAG networks or for seeding operations only, or you can disable compression for all DAG networks.

  - **Network encryption**: By default, DAGs use encryption only between DAG networks on different subnets. You can enable encryption for all DAG networks or for seeding operations only, or you can disable encryption for all DAG networks.

- **Shutting down DAG members**: The Exchange Server high availability solution is integrated with the Windows shutdown process. If an administrator or application initiates a shutdown of a Windows server in a DAG that has a mounted database that's replicated to one or more DAG members, the system will try to activate another copy of the mounted databases prior to allowing the shutdown process to complete. However, this new behavior doesn't guarantee that all of the databases on the server being shut down will experience a lossless activation. As a result, it's a best practice to perform a server switchover prior to shutting down a server that's a member of a DAG.

For detailed steps about how to create a DAG, see Create a database availability group. For detailed steps about how to configure DAGs and DAG properties, see Configure database availability group properties. For more information about each of the preceding management tasks, and about managing DAGs in general, see Manage database availability groups.

Mailbox database copy management

The operational management tasks associated with mailbox database copies include:

- **Adding mailbox database copies**: When you add a copy of a mailbox database, continuous replication is automatically enabled between the existing database and the database copy.

- **Configuring mailbox database copy properties**: You can configure a variety of properties, such as the database activation policy, the amount of time, if any, for replay lag and truncation lag, and the activation preference for the database copy.

- **Suspending or resuming a mailbox database copy**: You can suspend a mailbox database copy in preparation for seeding, or for other forms of maintenance. You can also suspend a mailbox database copy for activation only. This configuration prevents the system from automatically activating the copy as a result of a failure, but it still allows the system to keep the database copy up to date with log shipping and replay.

- **Updating a mailbox database copy**: Updating, also known as *seeding*, is the process in which a copy of a mailbox database is added to another Mailbox server. This becomes the baseline database for the copy. After the initial first seed of the baseline database copy, only in rare circumstances will the database need to be seeded again.

- **Activating a mailbox database copy**: Activating is the process of designating a specific passive copy as the new active copy of a mailbox database. This process is referred to as a *switchover*. For more information, see "Switchovers and Failovers" later in this topic.

- **Removing a mailbox database copy**: Occasionally, it may be necessary to remove a mailbox database copy. For example, you can't remove a Mailbox server from a DAG until all mailbox database copies are removed from the server. In addition, you must remove all copies of a mailbox database before you can change the path for a mailbox database.

For detailed steps about how to add a mailbox database copy, see Add a mailbox database copy. For detailed steps about how to configure mailbox database copies, see Configure mailbox database copy properties. For more information about each of the preceding management tasks, and about managing mailbox database copies in general, see Manage mailbox database copies. For detailed steps about how to remove a mailbox database copy, see Remove a mailbox database copy.

Proactive monitoring

Making sure that your servers are operating reliably and that your database copies are healthy are key objectives for daily messaging operations. Exchange Server includes a number of features that can be used to perform a variety of health monitoring tasks for DAGs and mailbox database copies, including:

- Get-MailboxDatabaseCopyStatus
- Test-ReplicationHealth
- Crimson channel event logging

In addition to monitoring the health and status, it's also critical to monitor for situations that can compromise availability. For example, we recommend that you monitor the redundancy of your replicated databases. It's critical to avoid situations where you're down to a single copy of a database. This scenario should be treated with the highest priority and resolved as soon as possible.

For more detailed information about monitoring the health and status of DAGs and mailbox database copies, see Monitor database availability groups.

Switchovers and failovers

A *switchover* is a manual process in which an administrator manually activates one or more mailbox database copies. Switchovers, which can occur at the database or server level, are typically performed as part of preparation for maintenance activities. Switchover management involves performing database or server switchovers as needed. For example, if you need to perform maintenance on a Mailbox server in a DAG, you would first perform a server switchover so that the server didn't host any active mailbox database copies. For detailed steps about how to perform a database switchover, see Activate a mailbox database copy. Switchovers can also be performed at the datacenter level.

A *failover* is the automatic activation by the system of one or more database copies in reaction to a failure. For example, the loss of a disk drive in a RAID-less environment will trigger a database failover. The loss of the MAPI network or a power failure will trigger a server failover.

Manage database availability groups

8/3/2020 • 34 minutes to read • Edit Online

A database availability group (DAG) is a set of up to 16 Exchange Mailbox servers that provides automatic, database-level recovery from a database, server, or network failure. DAGs use continuous replication and a subset of Windows failover clustering technologies to provide high availability and site resilience. Mailbox servers in a DAG monitor each other for failures. When a Mailbox server is added to a DAG, that server works with the other servers in the DAG to provide automatic, database-level recovery from database failures.

When you create a DAG, it's initially empty. When you add the first server to a DAG, a failover cluster is automatically created for the DAG. In addition, the infrastructure that monitors the servers for network or server failures is initiated. The failover cluster heartbeat mechanism and cluster database are then used to track and manage information about the DAG that can change quickly, such as database mount status, replication status, and last mounted location.

Creating DAGs

A DAG can be created using the New Database Availability Group wizard in the Exchange admin center (EAC), or by running the **New-DatabaseAvailabilityGroup** cmdlet in the Exchange Management Shell. When creating a DAG, you provide a name for the DAG, and optional witness server and witness directory settings. In addition, you can assign one or more IP addresses to the DAG, either by using static IP addresses or by allowing the DAG to be automatically assigned the necessary IP addresses using Dynamic Host Configuration Protocol (DHCP). You can manually assign IP addresses to the DAG by using the *DatabaseAvailabilityGroupIpAddresses* parameter. If you omit this parameter, the DAG attempts to obtain an IP address by using a DHCP server on your network.

If you're creating a DAG that will contain Mailbox servers that are running Windows Server 2012 R2, you also have the option of creating a DAG without a cluster administrative access point. In that case, the cluster will not have a cluster name object (CNO) in Active Directory, and the cluster core resource group will not contain a network name resource or an IP address resource.

For detailed steps about how to create a DAG, see Create a database availability group.

When you create a DAG, an empty object representing the DAG with the name you specified and an object class of **msExchMDBAvailabilityGroup** is created in Active Directory.

DAGs use a subset of Windows failover clustering technologies in Windows Server 2008 R2 or later, such as the cluster heartbeat, cluster networks, and cluster database (for storing data that changes or can change quickly, such as database state changes from active to passive or the reverse, or from mounted to dismounted or the reverse). Therefore you can create DAGs only on Exchange Mailbox servers installed on supported versions of Windows that include Windows failover clustering.

> **NOTE**
>
> The failover cluster created and used by the DAG must be dedicated to the DAG. The cluster can't be used for any other high availability solution or for any other purpose. For example, the failover cluster can't be used to cluster other applications or services. Using a DAG's underlying failover cluster for purposes other than the DAG isn't supported.

DAG witness server and witness directory

When creating a DAG, you need to specify a name for the DAG no longer than 15 characters that's unique within the Active Directory forest. In addition, each DAG is configured with a witness server and witness directory. The witness server and its directory are used only when there's an even number of members in the DAG and then only for quorum purposes. You don't need to create the witness directory in advance. Exchange automatically creates and secures the directory for you on the witness server. The witness directory shouldn't be used for any purpose other than for the DAG witness server.

NOTE

In the database mirroring topology, you can have a third server called the witness. The witness server enables automatic failover from principal to mirror server or vice-versa. Unlike principal and mirror servers, the witness server does not serve the database. The role of the witness is to verify whether a given partner server is up and functioning. Supporting automatic failover is the only function for witness server, and it identifies which server holds the principal copy and which server holds the mirror copy of the database.

The requirements for the witness server are as follows:

- The witness server can't be a member of the DAG.
- The witness server must be in the same Active Directory forest as the DAG.
- The witness server must be running Windows Server 2008 or later.
- A single server can serve as a witness for multiple DAGs. However, each DAG requires its own witness directory.

Regardless of what server is used as the witness server, if the Windows Firewall is enabled on the intended witness server, you must enable the Windows Firewall exception for File and Printer Sharing. The witness server uses SMB port 445.

IMPORTANT

If the witness server you specify isn't an Exchange 2010 or later server, you must add the Exchange Trusted Subsystem universal security group (USG) to the local Administrators group on the witness server prior to creating the DAG. These security permissions are necessary to ensure that Exchange can create a directory and share on the witness server as needed.

Neither the witness server nor the witness directory needs to be fault tolerant or use any form of redundancy or high availability. There's no need to use a clustered file server for the witness server or employ any other form of resiliency for the witness server. There are several reasons for this. With larger DAGs (for example, six members or more), several failures are required before the witness server is needed. Because a six-member DAG can tolerate as many as two voter failures without losing quorum, it would take as many as three voters failing before the witness server would be needed to maintain a quorum. Also, if there's a failure that affects your current witness server (for example, you lose the witness server because of a hardware failure), you can use the Set-DatabaseAvailabilityGroup cmdlet to configure a new witness server and witness directory (provided you have a quorum).

NOTE

You can also use the **Set-DatabaseAvailabilityGroup** cmdlet to configure the witness server and witness directory in the original location if the witness server lost its storage or if someone changed the witness directory or share permissions.

Witness server placement considerations

The placement of a DAG's witness server will depend on your business requirements and the options available to your organization. Exchange now includes support for new DAG configuration options that aren't recommended or aren't possible in Exchange 2010. These options include using a third location, such as a third datacenter, a branch office, or a Microsoft Azure virtual network.

The following table lists general witness server placement recommendations for different deployment scenarios.

DEPLOYMENT SCENARIO	RECOMMENDATIONS
Single DAG deployed in a single datacenter	Locate witness server in the same datacenter as DAG members
Single DAG deployed across two datacenters; no additional locations available	Locate witness server on a Microsoft Azure virtual network to enable automatic datacenter failover, or Locate witness server in primary datacenter
Multiple DAGs deployed in a single datacenter	Locate witness server in the same datacenter as DAG members. Additional options include: • Using the same witness server for multiple DAGs • Using a DAG member to act as a witness server for a different DAG
Multiple DAGs deployed across two datacenters	Locate witness server on a Microsoft Azure virtual network to enable automatic datacenter failover, or Locate witness server in the datacenter that is considered primary for each DAG. Additional options include: • Using the same witness server for multiple DAGs • Using a DAG member to act as a witness server for a different DAG
Single or Multiple DAGs deployed across more than two datacenters	In this configuration, the witness server should be located in the datacenter where you want the majority of quorum votes to exist.

When a DAG has been deployed across two datacenters, you can now use a third location for hosting the witness server. If your organization has a third location with a network infrastructure that is isolated from network failures that affect the two datacenters in which your DAG is deployed, then you can deploy the DAG's witness server in that third location, thereby configuring your DAG with the ability automatically failover databases to the other datacenter in response to a datacenter-level failure event. If your organization only has two physical locations, you can use a Microsoft Azure virtual network as a third location to place your witness server.

Specifying a witness server and witness directory during DAG creation

When creating a DAG, you must provide a name for the DAG. You can optionally also specify a witness server and witness directory.

When creating a DAG, the following combinations of options and behaviors are available:

- You can specify only a name for the DAG, and leave the **Witness server** and **Witness directory** fields blank. In this scenario, the wizard searches the local Active Directory site for a Client Access server that doesn't have the Mailbox server installed, and it automatically creates the default directory (%SystemDrive%:\DAGFileShareWitnesses\< *DAGFQDN*>) and default share (< *DAGFQDN*>) on that server and uses that Client Access server as the witness server. For example, consider the witness server CAS3 on which the operating system has been installed onto drive C. A DAG named DAG1 in the contoso.com domain would use a default witness directory of C:\DAGFileShareWitnesses\DAG1.contoso.com, which would be shared as \CAS3\DAG1.contoso.com.
- You can specify a name for the DAG, the witness server that you want to use, and the directory you want created and shared on the witness server.
- You can specify a name for the DAG and the witness server that you want to use, and leave the **Witness directory** field blank. In this scenario, the wizard creates the default directory on the specified witness server.

- You can specify a name for the DAG, leave the **Witness server** field blank, and specify the directory you want created and shared on the witness server. In this scenario, the wizard searches for a Client Access server that doesn't have the Mailbox server installed, and it automatically creates the specified DAG on that server, shares the directory, and uses that Client Access server as the witness server.

When a DAG is formed, it initially uses the Node Majority quorum model. When the second Mailbox server is added to the DAG, the quorum is automatically changed to a Node and File Share Majority quorum model. When this change occurs, the DAG's cluster begins using the witness server for maintaining quorum. If the witness directory doesn't exist, Exchange automatically creates it, shares it, and provisions the share with full control permissions for the CNO computer account for the DAG.

> **NOTE**
>
> Using a file share that's part of a Distributed File System (DFS) namespace isn't supported.

If Windows Firewall is enabled on the witness server before the DAG is created, it may block the creation of the DAG. Exchange uses Windows Management Instrumentation (WMI) to create the directory and file share on the witness server. If Windows Firewall is enabled on the witness server and there are no firewall exceptions configured for WMI, the **New-DatabaseAvailabilityGroup** cmdlet fails with an error. If you specify a witness server, but not a witness directory, you receive the following error message.

```
The task was unable to create the default witness directory on server <Server Name>. Please manually specify a
witness directory.
```

If you specify a witness server and witness directory, you receive the following warning message.

```
Unable to access file shares on witness server '<ServerName>'. Until this problem is corrected, the database
availability group may be more vulnerable to failures. You can use the Set-DatabaseAvailabilityGroup cmdlet to
try the operation again. Error: The network path was not found.
```

If Windows Firewall is enabled on the witness server after the DAG is created but before servers are added, it may block the addition or removal of DAG members. If Windows Firewall is enabled on the witness server and there are no firewall exceptions configured for WMI, the **Add-DatabaseAvailabilityGroupServer** cmdlet displays the following warning message.

```
Failed to create file share witness directory 'C:\DAGFileShareWitnesses\DAG_FQDN' on witness server
'<ServerName>'. Until this problem is corrected, the database availability group may be more vulnerable to
failures. You can use the Set-DatabaseAvailabilityGroup cmdlet to try the operation again. Error: WMI
exception occurred on server '<ServerName>': The RPC server is unavailable. (Exception from HRESULT:
0x800706BA)
```

To resolve the preceding error and warnings, do one of the following:

- Manually create the witness directory and share on the witness server, and assign the CNO for the DAG full control for the directory and share.

- Enable the WMI exception in Windows Firewall.

- Disable Windows Firewall.

DAG membership

After a DAG has been created, you can add servers to or remove servers from the DAG using the Manage Database Availability Group wizard in the Exchange admin center (EAC), or using the **Add-DatabaseAvailabilityGroupServer** or **Remove-DatabaseAvailabilityGroupServer** cmdlets in the Exchange Management Shell. For detailed steps about how to manage DAG membership, see Manage database availability group membership.

> **NOTE**
>
> Each Mailbox server that's a member of a DAG is also a node in the underlying cluster used by the DAG. As a result, at any one time, a Mailbox server can be a member of only one DAG.

If the Mailbox server being added to a DAG doesn't have the failover clustering component installed, the method used to add the server (for example, the **Add-DatabaseAvailabilityGroupServer** cmdlet or the Manage Database Availability Group wizard) installs the failover clustering feature.

When the first Mailbox server is added to a DAG, the following occurs:

- The Windows failover clustering component is installed, if it isn't already installed.
- A failover cluster is created using the name of the DAG. This failover cluster is used exclusively by the DAG, and the cluster must be dedicated to the DAG. Use of the cluster for any other purpose isn't supported.
- A CNO is created in the default computers container.
- The name and IP address of the DAG is registered as a Host (A) record in Domain Name System (DNS).
- The server is added to the DAG object in Active Directory.
- The cluster database is updated with information on the databases mounted on the added server.

In a large or multiple site environment, especially those in which the DAG is extended to multiple Active Directory sites, you must wait for Active Directory replication of the DAG object containing the first DAG member to complete. If this Active Directory object isn't replicated throughout your environment, adding the second server may cause a new cluster (and new CNO) to be created for the DAG. This is because the DAG object appears empty from the perspective of the second member being added, thereby causing the **Add-DatabaseAvailabilityGroupServer** cmdlet to create a cluster and CNO for the DAG, even though these objects already exist. To verify that the DAG object containing the first DAG server has been replicated, use the **Get-DatabaseAvailabilityGroup** cmdlet on the second server being added to verify that the first server you added is listed as a member of the DAG.

When the second and subsequent servers are added to the DAG, the following occurs:

- The server is joined to the Windows failover cluster for the DAG.
- The quorum model is automatically adjusted:
  - A Node Majority quorum model is used for DAGs with an odd number of members.
  - A Node and File Share Majority quorum model is used for DAGs with an even number of members.
- The witness directory and share are automatically created by Exchange when needed.
- The server is added to the DAG object in Active Directory.
- The cluster database is updated with information about mounted databases.

> **NOTE**
>
> The quorum model change should happen automatically. However, if the quorum model doesn't automatically change to the proper model, you can run the **Set-DatabaseAvailabilityGroup** cmdlet with only the *Identity* parameter to correct the quorum settings for the DAG.

Pre-staging the cluster name object for a DAG

The CNO is a computer account created in Active Directory and associated with the cluster's Name resource. The cluster's Name resource is tied to the CNO, which is a Kerberos-enabled object that acts as the cluster's identity and provides the cluster's security context. The formation of the DAG's underlying cluster and the CNO for that cluster is performed when the first member is added to the DAG. When the first server is added to the DAG, remote PowerShell contacts the Microsoft Exchange Replication service on the Mailbox server being added. The Microsoft Exchange Replication service installs the failover clustering feature (if it isn't already installed) and begins the cluster creation process. The Microsoft Exchange Replication service runs under the LOCAL SYSTEM security context, and it's under this context in which cluster creation is performed.

Caution

If your DAG members are running Windows Server 2012, you must pre-stage the CNO prior to adding the first server to the DAG. If your DAG members are running Windows Server 2012 R2, and you create a DAG without a cluster administrative access point, then a CNO will not be created, and you do not need to create a CNO for the DAG.

In environments where computer account creation is restricted, or where computer accounts are created in a container other than the default computers container, you can pre-stage and provision the CNO. You create and disable a computer account for the CNO, and then either:

- Assign full control of the computer account to the computer account of the first Mailbox server you're adding to the DAG.
- Assign full control of the computer account to the Exchange Trusted Subsystem USG.

Assigning full control of the computer account to the computer account of the first Mailbox server you're adding to the DAG ensures that the LOCAL SYSTEM security context will be able to manage the pre-staged computer account. Assigning full control of the computer account to the Exchange Trusted Subsystem USG can be used instead because the Exchange Trusted Subsystem USG contains the machine accounts of all Exchange servers in the domain.

For detailed steps about how to pre-stage and provision the CNO for a DAG, see Pre-stage the cluster name object for a database availability group.

Removing servers from a DAG

Mailbox servers can be removed from a DAG by using the Manage Database Availability Group wizard in the EAC or the **Remove-DatabaseAvailabilityGroupServer** cmdlet in the Exchange Management Shell. Before a Mailbox server can be removed from a DAG, all replicated mailbox databases must first be removed from the server. If you attempt to remove a Mailbox server with replicated mailbox databases from a DAG, the task fails.

There are scenarios in which you must remove a Mailbox server from a DAG before performing certain operations. These scenarios include:

- **Performing a server recovery operation**: If a Mailbox server that's a member of a DAG is lost, or otherwise fails and is unrecoverable and needs replacement, you can perform a server recovery operation using the **Setup /m:RecoverServer** switch. However, before you can perform the recovery operation, you must first remove the server from the DAG using the Remove-DatabaseAvailabilityGroupServer cmdlet with the *ConfigurationOnly* parameter.
- **Removing the database availability group**: There may be situations in which you need to remove a DAG (for example, when disabling third-party replication mode). If you need to remove a DAG, you must first remove all servers from the DAG. If you attempt to remove a DAG that contains any members, the task fails.

Configuring DAG properties

After servers have been added to the DAG, you can use the EAC or the Exchange Management Shell to configure the properties of a DAG, including the witness server and witness directory used by the DAG, and the IP addresses assigned to the DAG.

Configurable properties include:

- **Witness server**: The name of the server that you want to host the file share for the file share witness. We recommend that you specify a Client Access server as the witness server. This enables the system to automatically configure, secure, and use the share, as needed, and enables the messaging administrator to be aware of the availability of the witness server.
- **Witness directory**: The name of a directory that will be used to store file share witness data. This directory will automatically be created by the system on the specified witness server.
- **Database availability group IP addresses**: One or more IP addresses must be assigned to the DAG, unless the DAG members are running Windows Server 2012 R2 and you're creating a DAG without an IP address. Otherwise, the DAG's IP addresses can be configured using manually assigned static IP addresses, or they can be automatically assigned to the DAG using a DHCP server in your organization.

The Exchange Management Shell enables you to configure DAG properties that aren't available in the EAC, such as DAG IP addresses, network encryption and compression settings, network discovery, the TCP port used for replication, and alternate witness server and witness directory settings, and to enable Datacenter Activation Coordination mode.

For detailed steps about how to configure DAG properties, see Configure database availability group properties.

DAG network encryption

DAGs support the use of encryption by leveraging the encryption capabilities of the Windows Server operating system. DAGs use Kerberos authentication between Exchange servers. Microsoft Kerberos security support provider (SSP) EncryptMessage and DecryptMessage APIs handle encryption of DAG network traffic. Microsoft Kerberos SSP supports multiple encryption algorithms. (For the complete list, see section 3.1.5.2, "Encryption Types" of Kerberos Protocol Extensions). The Kerberos authentication handshake selects the strongest encryption protocol supported in the list: typically Advanced Encryption Standard (AES) 256-bit, potentially with a SHA Hash-based Message Authentication Code (HMAC) to maintain integrity of the data. For details, see HMAC.

Network encryption is a property of the DAG and not a DAG network. You can configure DAG network encryption using the **Set-DatabaseAvailabilityGroup** cmdlet in the Exchange Management Shell. The possible encryption settings for DAG network communications are shown in the following table.

SETTING	DESCRIPTION
Disabled	Network encryption isn't used.
Enabled	Network encryption is used on all DAG networks for replication and seeding.
InterSubnetOnly	Network encryption is used on DAG networks when replicating across different subnets. This is the default setting.
SeedOnly	Network encryption is used on all DAG networks for seeding only.

DAG network compression

DAGs support built-in compression. When compression is enabled, DAG network communication uses XPRESS, which is the Microsoft implementation of the LZ77 algorithm. This is the same type of compression used in many Microsoft protocols, in particular, MAPI RPC compression between Microsoft Outlook and Exchange.

As with network encryption, network compression is also a property of the DAG and not a DAG network. You configure DAG network compression by using the Set-DatabaseAvailabilityGroup cmdlet in the Exchange Management Shell. The possible compression settings for DAG network communications are shown in the following table.

SETTING	DESCRIPTION
Disabled	Network compression isn't used.
Enabled	Network compression is used on all DAG networks for replication and seeding.
InterSubnetOnly	Network compression is used on DAG networks when replicating across different subnets. This is the default setting.
SeedOnly	Network compression is used on all DAG networks for seeding only.

DAG networks

A DAG network is a collection of one or more subnets used for either replication traffic or MAPI traffic. Each DAG contains a maximum of one MAPI network and zero or more replication networks. In a single network adapter configuration, the network is used for both MAPI and replication traffic. Although a single network adapter and path is supported, we recommend that each DAG have a minimum of two DAG networks. In a two-network configuration, one network is typically dedicated for replication traffic, and the other network is used primarily for MAPI traffic. You can also add network adapters to each DAG member and configure additional DAG networks as replication networks.

> **NOTE**
>
> When using multiple replication networks, there's no way to specify an order of precedence for network use. Exchange randomly selects a replication network from the group of replication networks to use for log shipping.

In Exchange 2010, manual configuration of DAG networks was necessary in many scenarios. By default, in later versions of Exchange, DAG networks are automatically configured by the system. Before you can create or modify DAG networks, you must first enable manual DAG network control by running the following command:

```
Set-DatabaseAvailabilityGroup <DAGName> -ManualDagNetworkConfiguration $true
```

After you've enabled manual DAG network configuration, you can use the **New-DatabaseAvailabilityGroupNetwork** cmdlet in the Exchange Management Shell to create a DAG network. For detailed steps about how to create a DAG network, see Create a database availability group network.

You can use the **Set-DatabaseAvailabilityGroupNetwork** cmdlet in the Exchange Management Shell to configure DAG network properties. For detailed steps about how to configure DAG network properties, see Configure database availability group network properties. Each DAG network has required and optional parameters to configure:

- **Network name**: A unique name for the DAG network of up to 128 characters.
- **Network description**: An optional description for the DAG network of up to 256 characters.
- **Network subnets**: One or more subnets entered using a format of *IPAddress/Bitmask* (for example, 192.168.1.0/24 for Internet Protocol version 4 (IPv4) subnets; 2001:DB8:0:C000::/64 for Internet Protocol version 6 (IPv6) subnets).
- **Enable replication**: In the EAC, select the check box to dedicate the DAG network to replication traffic and

block MAPI traffic. Clear the check box to prevent replication from using the DAG network and to enable MAPI traffic. In the Exchange Management Shell, use the *ReplicationEnabled* parameter in the Set-DatabaseAvailabilityGroupNetwork cmdlet to enable and disable replication.

> **NOTE**
>
> Disabling replication for the MAPI network doesn't guarantee that the system won't use the MAPI network for replication. When all configured replication networks are offline, failed, or otherwise unavailable, and only the MAPI network remains (which is configured as disabled for replication), the system uses the MAPI network for replication.

The initial DAG networks (for example, MapiDagNetwork and ReplicationDagNetwork01) created by the system are based on the subnets enumerated by the Cluster service. Each DAG member must have the same number of network adapters, and each network adapter must have an IPv4 address (and optionally, an IPv6 address as well) on a unique subnet. Multiple DAG members can have IPv4 addresses on the same subnet, but each network adapter and IP address pair in a specific DAG member must be on a unique subnet. In addition, only the adapter used for the MAPI network should be configured with a default gateway. Replication networks shouldn't be configured with a default gateway.

For example, consider DAG1, a two-member DAG where each member has two network adapters (one dedicated for the MAPI network and the other for a replication network). Example IP address configuration settings are shown in the following table.

Example network adapter settings

SERVER-NETWORK ADAPTER	IP ADDRESS/SUBNET MASK	DEFAULT GATEWAY
EX1-MAPI	192.168.1.15/24	192.168.1.1
EX1-Replication	10.0.0.15/24	Not applicable
EX2-MAPI	192.168.1.16	192.168.1.1
EX2-Replication	10.0.0.16	Not applicable

In the following configuration, there are two subnets configured in the DAG: 192.168.1.0 and 10.0.0.0. When EX1 and EX2 are added to the DAG, two subnets will be enumerated and two DAG networks will be created: MapiDagNetwork (192.168.1.0) and ReplicationDagNetwork01 (10.0.0.0). These networks will be configured as shown in the following table.

Enumerated DAG network settings for a single-subnet DAG

NAME	SUBNETS	INTERFACES	MAPI ACCESS ENABLED	REPLICATION ENABLED
MapiDagNetwork	192.168.1.0/24	EX1 (192.168.1.15) EX2 (192.168.1.16)	True	True
ReplicationDagNetwork01	10.0.0.0/24	EX1 (10.0.0.15) EX2 (10.0.0.16)	False	True

To complete the configuration of ReplicationDagNetwork01 as the dedicated replication network, disable replication for MapiDagNetwork by running the following command.

```
Set-DatabaseAvailabilityGroupNetwork -Identity DAG1\MapiDagNetwork -ReplicationEnabled:$false
```

After replication is disabled for MapiDagNetwork, the Microsoft Exchange Replication service uses ReplicationDagNetwork01 for continuous replication. If ReplicationDagNetwork01 experiences a failure, the Microsoft Exchange Replication service reverts to using MapiDagNetwork for continuous replication. This is done intentionally by the system to maintain high availability.

DAG networks and multiple subnet deployments

In the preceding example, even though there are two different subnets in use by the DAG (192.168.1.0 and 10.0.0.0), the DAG is considered a single-subnet DAG because each member uses the same subnet to form the MAPI network. When DAG members use different subnets for the MAPI network, the DAG is referred to as a *multi-subnet DAG*. In a multi-subnet DAG, the proper subnets are automatically associated with each DAG network.

For example, consider DAG2, a two-member DAG where each member has two network adapters (one dedicated for the MAPI network and the other for a replication network), and each DAG member is located in a separate Active Directory site, with its MAPI network on a different subnet. Example IP address configuration settings are shown in the following table.

Example network adapter settings for a multi-subnet DAG

SERVER-NETWORK ADAPTER	IP ADDRESS/SUBNET MASK	DEFAULT GATEWAY
EX1-MAPI	192.168.0.15/24	192.168.0.1
EX1-Replication	10.0.0.15/24	Not applicable
EX2-MAPI	192.168.1.15	192.168.1.1
EX2-Replication	10.0.1.15	Not applicable

In the following configuration, there are four subnets configured in the DAG: 192.168.0.0, 192.168.1.0, 10.0.0.0, and 10.0.1.0. When EX1 and EX2 are added to the DAG, four subnets will be enumerated, but only two DAG networks will be created: MapiDagNetwork (192.168.0.0, 192.168.1.0) and ReplicationDagNetwork01 (10.0.0.0, 10.0.1.0). These networks will be configured as shown in the following table.

Enumerated DAG network settings for a multi-subnet DAG

NAME	SUBNETS	INTERFACES	MAPI ACCESS ENABLED	REPLICATION ENABLED
MapiDagNetwork	192.168.0.0/24 192.168.1.0/24	EX1 (192.168.0.15) EX2 (192.168.1.15)	True	True
ReplicationDagNetwork01	10.0.0.0/24 10.0.1.0/24	EX1 (10.0.0.15) EX2 (10.0.1.15)	False	True

DAG networks and iSCSI networks

By default, DAGs perform discovery of all networks detected and configured for use by the underlying cluster. This includes any Internet SCSI (iSCSI) networks in use as a result of using iSCSI storage for one or more DAG members. As a best practice, iSCSI storage should use dedicated networks and network adapters. These networks shouldn't be managed by the DAG or its cluster, or used as DAG networks (MAPI or replication). Instead, these networks should be manually disabled from use by the DAG, so they can be dedicated to iSCSI storage traffic. To disable iSCSI networks from being detected and used as DAG networks, configure the DAG to ignore any currently detected iSCSI networks using the Set-DatabaseAvailabilityGroupNetwork cmdlet, as shown in this example:

```
Set-DatabaseAvailabilityGroupNetwork -Identity DAG2\DAGNetwork02 -ReplicationEnabled:$false -
IgnoreNetwork:$true
```

This command will also disable the network for use by the cluster. Although the iSCSI networks will continue to appear as DAG networks, they won't be used for MAPI or replication traffic after running the above command.

Configuring DAG members

Mailbox servers that are members of a DAG have some properties specific to high availability that should be configured as described in the following sections:

- Automatic database mount dial
- Database copy automatic activation policy
- Maximum active databases

Automatic database mount dial

The *AutoDatabaseMountDial* parameter specifies the automatic database mount behavior after a database failover. You can use the Set-MailboxServer cmdlet to configure the *AutoDatabaseMountDial* parameter with any of the following values:

- `BestAvailability`: If you specify this value, the database automatically mounts immediately after a failover if the copy queue length is less than or equal to 12. The copy queue length is the number of logs recognized by the passive copy that needs to be replicated. If the copy queue length is more than 12, the database doesn't automatically mount. When the copy queue length is less than or equal to 12, Exchange attempts to replicate the remaining logs to the passive copy and mounts the database.

- `GoodAvailability`: If you specify this value, the database automatically mounts immediately after a failover if the copy queue length is less than or equal to six. The copy queue length is the number of logs recognized by the passive copy that needs to be replicated. If the copy queue length is more than six, the database doesn't automatically mount. When the copy queue length is less than or equal to six, Exchange attempts to replicate the remaining logs to the passive copy and mounts the database.

- `Lossless`: If you specify this value, the database doesn't automatically mount until all logs generated on the active copy have been copied to the passive copy. This setting also causes the Active Manager best copy selection algorithm to sort potential candidates for activation based on the database copy's activation preference value and not its copy queue length.

The default value is `GoodAvailability`. If you specify either `BestAvailability` or `GoodAvailability`, and all the logs from the active copy can't be copied to the passive copy being activated, you may lose some mailbox data. However, the Safety Net feature (which is enabled by default) helps protect against most data loss by resubmitting messages that are in the Safety Net queue.

Example: configuring automatic database mount dial

The following example configures a Mailbox server with an *AutoDatabaseMountDial* setting of `GoodAvailability`.

```
Set-MailboxServer -Identity EX1 -AutoDatabaseMountDial GoodAvailability
```

Database copy automatic activation policy

The *DatabaseCopyAutoActivationPolicy* parameter specifies the type of automatic activation available for mailbox database copies on the selected Mailbox servers. You can use the Set-MailboxServer cmdlet to configure the *DatabaseCopyAutoActivationPolicy* parameter with any of the following values:

- `Blocked`: If you specify this value, databases can't be automatically activated on the selected Mailbox servers.

- `IntrasiteOnly`: If you specify this value, the database copy is allowed to be activated on servers in the same Active Directory site. This prevents cross-site failover or activation. This property is for incoming mailbox database copies (for example, a passive copy being made an active copy). Databases can't be activated on this Mailbox server for database copies that are active in another Active Directory site.

- `Unrestricted`: If you specify this value, there are no special restrictions on activating mailbox database copies on the selected Mailbox servers.

Example: configuring database copy automatic activation policy

The following example configures a Mailbox server with a *DatabaseCopyAutoActivationPolicy* setting of `Blocked`.

```
Set-MailboxServer -Identity EX1 -DatabaseCopyAutoActivationPolicy Blocked
```

Maximum active databases

The *MaximumActiveDatabases* parameter (also used with the Set-MailboxServer cmdlet) specifies the number of databases that can be mounted on a Mailbox server. You can configure Mailbox servers to meet your deployment requirements by ensuring that an individual Mailbox server doesn't become overloaded.

The *MaximumActiveDatabases* parameter is configured with a whole number numeric value. When the maximum number is reached, the database copies on the server won't be activated if a failover or switchover occurs. If the copies are already active on a server, the server won't allow databases to be mounted.

Example: configuring maximum active databases

The following example configures a Mailbox server to support a maximum of 20 active databases.

```
Set-MailboxServer -Identity EX1 -MaximumActiveDatabases 20
```

Performing maintenance on DAG members

Before performing any type of software or hardware maintenance on a DAG member, you should first put the DAG member in maintenance mode. The following scripts are provided with Exchange Server to assist with DAG maintenance procedures.

- **StartDagServerMaintenance.ps1**: Assists with moving all active databases off the server. It also moves all critical DAG support functionality, such as the Primary Active Manager (PAM) role, and blocks them from moving back to the server before maintenance is complete.

- **StopDagServerMaintenance.ps1**: Assists with taking the DAG member out of maintenance mode, and making it an active target for all databases and all critical DAG support functionality.

Both of the above scripts accept the *ServerName* parameter (which can be either the host name or the fully qualified domain name (FQDN) of the DAG member) and the *WhatIf* parameter. Both scripts can be run locally or remotely. The server on which the scripts are executed must have the Windows Failover Cluster Management tools installed (RSAT-Clustering).

> **NOTE**
>
> The **RedistributeActiveDatabases.ps1** script is also avaialble, which assists with mounting mailbox databases on specific DAG members as inidicated by the Activation Preference number on each database. However, in Exchange 2016 CU2 or later, the new DAG property named *PreferenceMoveFrequency* automatically balances database copies across a DAG. Therefore, you'll only need to use **RedistributeActiveDatabases.ps1** if you've disabled this functionality or if you want to balance database copies manually.

The StartDagServerMaintenance.ps1 script performs the following tasks:

- Sets the value of the *DatabaseCopyAutoActivationPolicy* parameter on the DAG member to `Blocked`, which prevents any database copies from being activated on the server.
- Pauses the node in the cluster, which prevents the node from being and becoming the PAM.
- Moves all active databases currently hosted on the DAG member to other DAG members.
- If the DAG member currently owns the default cluster group, the script moves the default cluster group (and therefore the PAM role) to another DAG member.

If any of the preceding tasks fails, all operations, except for successful database moves, are undone by the script.

To begin maintenance procedures on a DAG member, including flushing the transport queues and suspending client connectivity, perform the following tasks:

1. To empty the transport queues, run the following command:

   ```
   Set-ServerComponentState <ServerName> -Component HubTransport -State Draining -Requester Maintenance
   ```

2. To initiate the draining of the transport queues, run the following command:

   ```
   Restart-Service MSExchangeTransport
   ```

3. To begin the process of draining all Unified Messaging calls (in Exchange 2016 only), run the following command:

   ```
   Set-ServerComponentState <ServerName> -Component UMCallRouter -State Draining -Requester Maintenance
   ```

4. To access the DAG maintenance scripts, run the following command:

   ```
   CD $ExScripts
   ```

5. To run the StartDagServerMaintenance.ps1 script, run the following command:

   ```
   .\StartDagServerMaintenance.ps1 -ServerName <ServerName> -MoveComment Maintenance -PauseClusterNode
   ```

   For the value of the *MoveComment* parameter, you can make any notation you want. The above example uses "Maintenance."

   > **NOTE**
   >
   > This script can take some time to execute, and during this time you may not see any activity on your screen.

6. To redirect messages pending delivery in the local queues to the Exchange server specified by the Target parameter, run

   ```
   Redirect-Message -Server <ServerName> -Target <Server FQDN>
   ```

7. To place the server into maintenance mode, run:

```
Set-ServerComponentState <ServerName> -Component ServerWideOffline -State Inactive -Requester
Maintenance
```

To verify that a server is ready for maintenance, perform the following tasks:

1. To verify the server has been placed into maintenance mode, confirm that only `Monitoring` and `RecoveryActionsEnabled` are in an Active state when you run the following command:

   ```
   Get-ServerComponentState <ServerName> | Format-Table Component,State -Autosize
   ```

2. To verify the server is not hosting any active database copies, run:

   ```
   Get-MailboxServer <ServerName> | Format-List DatabaseCopyAutoActivationPolicy
   ```

3. To verify that the cluster node is paused, run:

   ```
   Get-ClusterNode <ServerName> | Format-List
   ```

4. To verify that all transport queues have been emptied, run:

   ```
   Get-Queue
   ```

After the maintenance is complete and the DAG member is ready to return to service, the StopDagServerMaintenance.ps1 script helps takes the DAG member out of maintenance mode and put it back into production. The StopDagServerMaintenance.ps1 script performs the following tasks:

- Resumes the node in the cluster, which enables full cluster functionality for the DAG member.
- Sets the value of the *DatabaseCopyAutoActivationPolicy* parameter on the DAG member to `Unrestricted`.
- Runs the Resume-MailboxDatabaseCopy cmdlet for each database copy hosted on the DAG member.

When you're ready to restore the DAG member to full production status, including resuming the transport queues and client connectivity, perform the following tasks:

1. To configure the server as out of maintenance mode and ready to accept client connections, run:

   ```
   Set-ServerComponentState <ServerName> -Component ServerWideOffline -State Active -Requester Maintenance
   ```

2. To allow the server to accept Unified Messaging calls (in Exchange 2016 only), run:

   ```
   Set-ServerComponentState <ServerName> -Component UMCallRouter -State Active -Requester Maintenance
   ```

3. To access the DAG maintenance scripts, run the following command:

   ```
   CD $ExScripts
   ```

4. To execute the StopDagServerMaintenance.ps1 script, run:

   ```
   .\StopDagServerMaintenance.ps1 -serverName <ServerName>
   ```

5. To enable the transport queues to resume accepting and processing messages, run:

   ```
   Set-ServerComponentState <ServerName> -Component HubTransport -State Active -Requester Maintenance
   ```

6. To resume transport activity, run:

   ```
   Restart-Service MSExchangeTransport
   ```

To verify that a server is ready for production use, perform the following tasks:

1. To verify the server is not in maintenance mode, run

   ```
   Get-ServerComponentState <ServerName> | Format-Table Component,State -Autosize
   ```

If you're installing an Exchange update, and the update process fails, it can leave some server components in an inactive state, which will be displayed in the output of the above `Get-ServerComponentState` cmdlet. To resolve this, run the following commands:

1. `Set-ServerComponentState <ServerName> -Component ServerWideOffline -State Active -Requester Functional`
2. `Set-ServerComponentState <ServerName> -Component Monitoring -State Active -Requester Functional`
3. `Set-ServerComponentState <ServerName> -Component RecoveryActionsEnabled -State Active -Requester Functional`

Shutting down DAG members

Exchange high availability solution is integrated with the Windows shutdown process. If an administrator or application initiates a shutdown of a Windows server in a DAG that has a mounted database that's replicated to one or more DAG members, the system attempts to activate another copy of the mounted database prior to allowing the shutdown process to complete.

However, this new behavior doesn't guarantee that all of the databases on the server being shut down will experience a `lossless` activation. As a result, it's a best practice to perform a server switchover prior to shutting down a server that's a member of a DAG.

Installing updates on DAG members

Installing Exchange updates on a server that's a member of a DAG is a relatively straightforward process. When you install an update on a server that's a member of a DAG, several services are stopped during the installation, including all Exchange services and the Cluster service. The general process for applying updates to a DAG member is as follows:

1. Use the steps described above to put the DAG member in maintenance mode.
2. Install the update.
3. Use the steps described above to take the DAG member out of maintenance mode and put it back into production.
4. Optionally, use the RedistributeActiveDatabases.ps1 script to rebalance the active database copies across the DAG.

For more information about the latest Exchange updates, see Exchange Server build numbers and release dates.

Create a database availability group

8/3/2020 • 5 minutes to read • Edit Online

A database availability group (DAG) is a set of up to 16 Microsoft Exchange Server Mailbox servers that provide automatic database-level recovery from a database, server, or network failure. When a Mailbox server is added to a DAG, it works with the other servers in the DAG to provide automatic, database-level recovery from database, server, and network failures.

> **IMPORTANT**
>
> All servers within a DAG must be running the same version of Exchange. You can't mix Exchange 2013 servers and Exchange servers in the same DAG.

Looking for other management tasks related to DAGs? Check out Manage database availability groups.

What do you need to know before you begin?

- Estimated time to complete: 1 minute

- To open the EAC, see Exchange admin center in Exchange Server. To open the Exchange Management Shell, see Open the Exchange Management Shell.

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Database availability groups" entry in the High availability and site resilience permissions topic.

- When creating a DAG with Mailbox servers running Windows Server 2012, you must pre-stage the cluster name object (CNO) before adding members to the DAG. If you're creating a DAG without an administrative access point with Mailbox servers running Windows Server 2012 R2, then you do not need to pre-stage a CNO for the DAG. For detailed steps, see Pre-stage the cluster name object for a database availability group.

- When creating a DAG, you provide a unique name for the DAG of up to 15 characters. In addition to providing a name for the DAG, you must also assign one or more IP addresses (either IPv4 or both IPv4 and IPv6) to the DAG, unless you're creating a Windows Server 2012 R2 DAG without an administrative access point and you aren't assigning any IP addresses to the DAG. Otherwise, the IP addresses you assign must be on each subnet intended for the MAPI network and must be available for use. If you specify one or more IPv4 addresses and your system is configured to use IPv6, the task will also attempt to automatically assign the DAG one or more IPv6 addresses.

- When creating a DAG, you must specify a witness server and witness directory. We recommend that you use an Exchange server with Client Access services. This allows an Exchange administrator to be aware of the availability of the witness, and it ensures that all of the necessary security permissions needed for using the witness server are in place.

  The following combinations of options and behaviors are available:

  - You can specify a name for the DAG, the witness server that you want to use, and the directory you want created and shared on the witness server.

  - You can specify a name for the DAG and the witness server that you want to use, and leave the **Witness directory** field empty. In this scenario, the task will create the default witness directory on the specified witness server.

Note: If the witness server you specify isn't an Exchange server in your organization, you must add the Exchange Trusted Subsystem universal security group to the local Administrators group on the witness server. These security permissions are necessary to ensure that Exchange can create a directory and share on the witness server as needed. If the proper permissions aren't configured, the following error is returned:

```
Error: An error occurred during discovery of the database availability group topology. Error: An
error occurred while attempting a cluster operation. Error: Cluster API "AddClusterNode()
(MaxPercentage=12) failed with 0x80070005. Error: Access is denied."
```

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to create a database availability group

1. In the EAC, go to **Servers** > **Database Availability Groups**.

2. Click + to create a DAG.

3. On the **new database availability group** page, provide the following information for the DAG:

   - **Database availability group name**: Use this field to type a valid and unique name for the DAG of up to 15 characters. The name is equivalent to a computer name, and a corresponding CNO will be created in Active Directory with that name. This name will be both the name of the DAG and the name of the underlying cluster.

   - **Witness server**: Use this field to specify a witness server for the DAG.

     Note: You must use either a host name or a fully qualified domain name (FQDN) for the witness server. Using an IP address or a wildcard name isn't supported. In addition, the witness server can't be a member of the DAG.

   - **Witness directory**: Use this field to type the path to a directory on the witness server that will be used to store witness data. If the directory doesn't exist, the system will create it for you on the witness server. If you leave this field blank, the default directory (%SystemDrive%\DAGFileShareWitnesses\<DAG FQDN>) will be created on the witness server.

   - **Database availability group IP addresses**: Use this field to assign one or more static IPv4 addresses to the DAG. Enter an IPv4 address and click + to add it. Leave this field blank if you want the DAG to use Dynamic Host Configuration Protocol (DHCP) to obtain the necessary IPv4 addresses. Optionally, enter 255.255.255.255 to create a DAG without an IP address or cluster administrative access point, which applies only to DAGs that will contain Mailbox servers running Windows Server 2012 R2.

4. Click **Save** to create the DAG.

Use the Exchange Management Shell to create a database availability group

The following example creates a DAG named DAG1, which is configured to use the witness server FILESRV1 and the local directory C:\DAG1. DAG1 is also configured to use DHCP for the DAG's IP addresses.

```
New-DatabaseAvailabilityGroup -Name DAG1 -WitnessServer FILESRV1 -WitnessDirectory C:\DAG1
```

This example creates the DAG DAG3. DAG3 is configured to use the witness server MBX2 and the local directory C:\DAG3. DAG3 is assigned multiple static IP addresses because its DAG members are on different subnets on the MAPI network.

```
New-DatabaseAvailabilityGroup -Name DAG3 -WitnessServer MBX2 -WitnessDirectory C:\DAG3 -
DatabaseAvailabilityGroupIPAddresses 10.0.0.8,192.168.0.8
```

This example creates the DAG DAG5 that will not have an administrative access point (valid for Windows Server 2012 R2 DAGs only). In addition, MBX4 will be used as the witness server for the DAG, and the default witness directory will be created.

```
New-DatabaseAvailabilityGroup -Name DAG5 -DatabaseAvailabilityGroupIPAddresses ([System.Net.IPAddress]::None)
-WitnessServer MBX4
```

How do you know this worked?

To verify that you've successfully created a DAG, do one of the following:

- In the EAC, navigate to **Servers** > **Database Availability Groups**. The newly created DAG is displayed.
- In the Exchange Management Shell, run the following command to verify the DAG was created and to display DAG property information.

  ```
  Get-DatabaseAvailabilityGroup <DAGName> | Format-List
  ```

For more information

Database availability groups

Configure database availability group properties

Set-DatabaseAvailabilityGroup

New-DatabaseAvailabilityGroup

New-DatabaseAvailabilityGroupNetwork

Add-DatabaseAvailabilityGroupServer

Pre-stage the cluster name object for a database availability group

8/3/2020 • 3 minutes to read • Edit Online

In environments where computer account creation is restricted, or where computer accounts are created in a container other than the default computers container, you can pre-stage the cluster name object (CNO) and then provision the CNO by assigning permissions to it.

Pre-staging the CNO is also required for Windows Server 2012 and Windows Server 2012 R2 DAG members due to permissions changes in Windows for computer objects. When deploying a database availability group (DAG) using Mailbox servers that are running Windows Server 2012 or Windows Server 2012 R2, you must pre-stage and provision the CNO, unless you're deploying a DAG without a cluster administrative access point. DAGs without cluster administrative access points do not use CNOs; therefore pre-staging is not required for those DAGs.

You create and disable a computer account for the CNO, and then either:

- Assign full control of the computer account to the computer account of the first Mailbox server you're adding to the DAG.

  -Or-

- Assign full control of the computer account to the Exchange Trusted Subsystem universal security group (USG).

What do you need to know before you begin?

- Estimated time to complete: 1 minute
- You must use an account that has permissions to create computer objects in Active Directory.
- After completing the following steps, allow time for Active Directory replication to occur. After the object is replicated, you can add the first member to the DAG.

> **TIP**
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Pre-stage the CNO

1. Open Active Directory Users and Computers.
2. Expand the forest node.
3. Right-click the organizational unit (OU) in which you want to create the new account, select **New**, and then select **Computer**.
4. In **New Object - Computer**, type the computer account name for the CNO in the **Computer name** box. This is the name that you'll use for the DAG. Click **OK** to create the account.
5. Right-click the new computer account, and then click **Disable Account**. Click **Yes** to confirm the disable action, and then click **OK**.

Assign permissions to the CNO

1. Open Active Directory Users and Computers.

2. If Advanced Features aren't enabled, turn them on by clicking **View**, and then clicking **Advanced Features**.

3. Right-click the new computer account, and then click **Properties**.

4. In **<Computer Name> Properties**, on the **Security** tab, click **Add** to add either the computer account for the first node to be added to the DAG or to add the Exchange Trusted Subsystem USG:

   - To add the Exchange Trusted Subsystem, type Exchange Trusted Subsystem in the **Enter the object names to select** field. Click **OK** to add the USG. Select the Exchange Trusted Subsystem USG and in the **Permissions for Exchange Trusted Subsystem** field, select **Full Control** in the **Allow** column. Click **OK** to save the permission settings.

   - To add the computer account for the first node to be added to the DAG, click **Object Types**. In the **Object Types** dialog box, clear the **Built-in security principals**, **Groups**, and **Users** check boxes. Select the **Computers** check box and click **OK**. In the **Enter the object names to select** field, type the name of the first Mailbox server to be added to the DAG, and then click **OK**. Select the first node's computer account, and in the **Permissions for <NodeName>** field, select **Full Control** in the **Allow** column. Click **OK** to save the permission settings.

How do you know this worked?

To verify that you've successfully created the CNO, do the following:

1. Open Active Directory Users and Computers.

2. Expand the forest node.

3. Open the OU in which you created the account, and then verify that the account is listed.

Using a Microsoft Azure VM as a DAG witness server

8/3/2020 • 13 minutes to read • Edit Online

This configuration requires three separate physical locations: two datacenters for mailbox servers and a third location to place the witness server for the DAG. Organizations with only two physical locations now can also take advantage of automatic datacenter failover by using a Microsoft Azure file server virtual machine to act as the DAG's witness server. This article focuses on the placement of the DAG witness on Microsoft Azure and assumes that you're familiar with site resilience concepts and already have a fully functional DAG infrastructure spanning two datacenters. If you don't already have your DAG infrastructure configured, we recommend that you first review the following articles:

High availability and site resilience

Database availability groups

Plan for high availability and site resilience

Changes to Microsoft Azure

This configuration requires a multi-site VPN. It has always been possible to connect your organization's network to Microsoft Azure using a site-to-site VPN connection. However, in the past, Azure supported only a single site-to-site VPN. Since configuring a DAG and its witness across three datacenters required multiple site-to-site VPNs, placement of the DAG witness on an Azure VM wasn't initially possible.

In June 2014, Microsoft Azure introduced multi-site VPN support, which enabled organizations to connect multiple datacenters to the same Azure virtual network. This change also made it possible for organizations with two datacenters to leverage Microsoft Azure as a third location to place their DAG witness servers. To learn more about the multi-site VPN feature in Azure, see Configure a Multi-Site VPN.

> **NOTE**
>
> This configuration leverages Azure virtual machines and a multi-site VPN for deploying the witness server and does not use the Azure Cloud Witness.

Microsoft Azure file server witness

The following diagram is an overview of using a Microsoft Azure file server VM as a DAG witness. You need an Azure virtual network, a multi-site VPN that connects your datacenters to your Azure virtual network, and a domain controller and a file server deployed on Azure virtual machines.

> **NOTE**
>
> It is technically possible to use a single Azure VM for this purpose and place the file witness share on the domain controller. However, this will result in an unnecessary elevation of privileges. Therefore, it is not a recommended configuration.

DAG witness server on Microsoft Azure

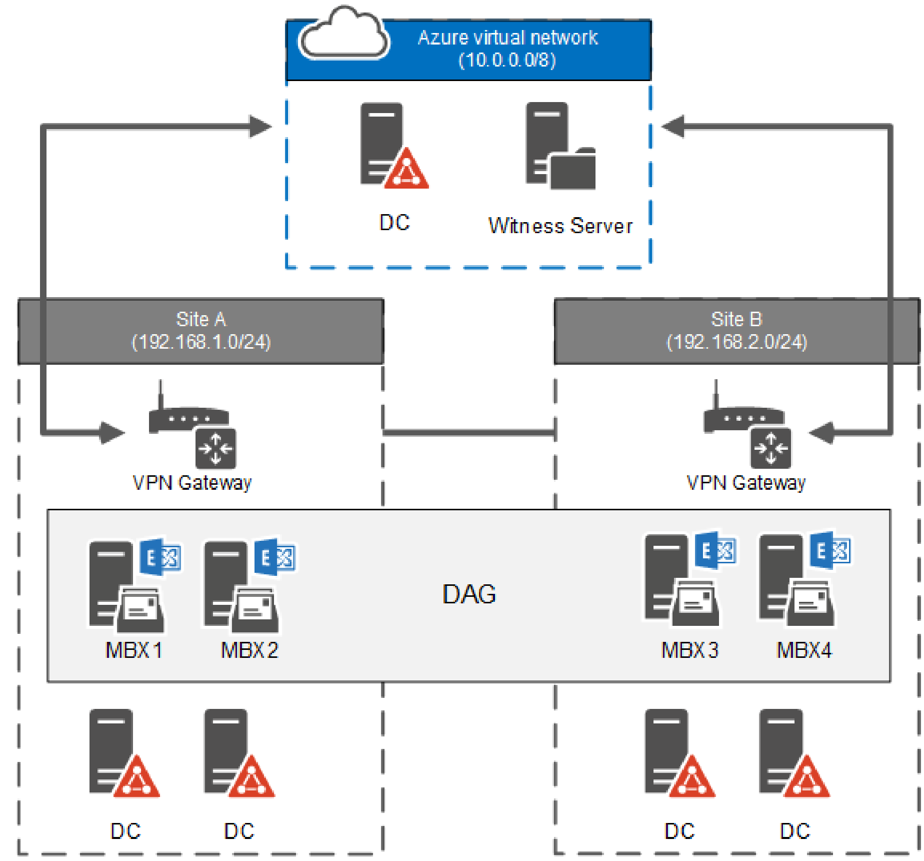

- Because this configuration requires name resolution between the on-premises servers and Azure VMs, you will need to configure Azure to use your own DNS servers.
- All DCs represented are members of the same AD Forest.

The first thing you need to do in order to use a Microsoft Azure VM for your DAG witness is to get a subscription. See How to buy Azure for the best way to acquire an Azure subscription.

After you have your Azure subscription, you need to do the following in order:

1. Prepare the Microsoft Azure virtual network

2. Configure a multi-site VPN

3. Configure virtual machines

4. Configure the DAG witness

> **NOTE**
>
> A significant portion of the guidance in this article involves Microsoft Azure configuration. Therefore, we link to Azure documentation whenever applicable.

Prerequisites

- Two datacenters that are capable of supporting an Exchange high availability and site resilience deployment. See Plan for high availability and site resilience for more information.

- A public IP address that is not behind NAT for the VPN gateways in each site.

- A VPN device in each site that is compatible with Microsoft Azure. See About VPN Devices for Virtual Network for more information about compatible devices.

- Familiarity with DAG concepts and management.

- Familiarity with Windows PowerShell.

Phase 1: Prepare the Microsoft Azure virtual network

Configuring the Microsoft Azure network is the most crucial part of the deployment process. At the end of this phase, you will have a fully functional Azure virtual network that is connected to your two datacenters via a multi-site VPN.

Register DNS servers

Because this configuration requires name resolution between the on-premises servers and Azure VMs, you will need to configure Azure to use your own DNS servers. Name resolution for resources in Azure virtual networks topic provides an overview of name resolution in Azure.

Do the following to register your DNS servers:

1. In the Azure portal, go to **networks**, and then click **NEW**.
2. Click **NETWORK SERVICES** > **VIRTUAL NETWORK** > **REGISTER DNS SERVER**.
3. Type the name and IP address for your DNS server. The name specified here is a logical name used in the management portal and doesn't have to match the actual name of your DNS server.
4. Repeat steps 1 through 3 for any other DNS servers you want to add.

> **NOTE**
>
> The DNS servers you register are not used in a round robin fashion. Azure VMs will use the first DNS server listed and will only use any additional servers if the first one is not available.

5. Repeat steps 1 through 3 to add the IP address you will use for the domain controller you will deploy on Microsoft Azure.

Create local (on-premises) network objects in Azure

Next, do the following to create logical network objects that represent your datacenters in Microsoft Azure:

1. In the Azure portal, and then go to **networks**, and then click **NEW**.
2. Click **NETWORK SERVICES** > **VIRTUAL NETWORK** > **ADD LOCAL NETWORK**.
3. Type the name for your first datacenter site and the IP address of the VPN device on that site. This IP address must be a static public IP address that is not behind NAT.
4. On the next screen, specify the IP subnets for your first site.
5. Repeat steps 1through 4 for your second site.

Create the Azure virtual network

Now, do the following to create an Azure virtual network that will be used by the VMs:

1. In the Azure portal, go to **networks**, and then click **NEW**.
2. Click **NETWORK SERVICES** > **VIRTUAL NETWORK** > **CUSTOM CREATE**.
3. On the **Virtual Network Details** page, specify a name for the virtual network, and select a geographic location for the network.
4. In the **DNS Servers and VPN Connectivity** page, verify that the DNS servers you previously registered are listed as the DNS servers.
5. Select the **Configure a site-to-site VPN** check box under **SITE-TO-SITE CONNECTIVITY**.

IMPORTANT

Do not select **Use ExpressRoute** because this will prevent the necessary configuration changes required to set up a multi-site VPN.

6. Under **LOCAL NETWORK**, select one of the two on-premises networks you configured.

7. In the **Virtual Network Address Spaces** page, specify the IP address range you will use for your Azure virtual network.

Checkpoint: Review the network configuration

At this point, when you go to **networks**, you should see the virtual network you configured under **VIRTUAL NETWORKS**, your local sites under **LOCAL NETWORKS**, and your registered DNS servers under **DNS SERVERS**.

Phase 2: Configure a multi-site VPN

The next step is to establish the VPN gateways to your on-premises sites. To do this, you need to:

1. Establish a VPN gateway to one of your sites by using the Azure portal.

2. Export the virtual network configuration settings.

3. Modify the configuration file for multi-site VPN.

4. Import the updated Azure network configuration.

5. Record the Azure gateway IP address and preshared keys.

6. Configure on-premises VPN devices.

For more information about configuring a multi-site VPN, see Configure a Multi-Site VPN.

Establish a VPN gateway to your first site

When creating your virtual gateway, note that you already specified that it will be connected to your first on-premises site. When you go into the virtual network dashboard, you will see that the gateway has not been created.

To establish the VPN gateway on the Azure side, see VPN Gateway.

IMPORTANT

Only perform the steps in the "Start the virtual network gateway" section of the article, and do not continue to the subsequent sections.

Export virtual network configuration settings

The Azure management portal doesn't currently allow you to configure a multi-site VPN. For this configuration, you need to export the virtual network configuration settings to an XML file and then modify that file. Follow the instructions atExport Virtual Network Settings to a Network Configuration File to export your settings.

Modify the network configuration settings for the multi-site VPN

Open the file you exported in any XML editor. The gateway connections to your on-premises sites are listed in the "ConnectionsToLocalNetwork" section. Search for that term in the XML file to locate the section. This section in the configuration file will look like the following (assuming the site name you created for your local site is "Site A").

```
<ConnectionsToLocalNetwork>
    <LocalNetworkSiteRef name="Site A">
        <Connection type="IPsec" />
</LocalNetworkSiteRef>
```

To configure your second site, add another "LocalNetworkSiteRef" section under the "ConnectionsToLocalNetwork" section. The section in the updated configuration file will look like the following (assuming the site name for your second local site is "Site B").

```
<ConnectionsToLocalNetwork>
    <LocalNetworkSiteRef name="Site A">
        <Connection type="IPsec" />
    <LocalNetworkSiteRef name="Site B">
        <Connection type="IPsec" />
</LocalNetworkSiteRef>
```

Save the updated configuration settings file.

Import virtual network configuration settings

The second site reference you've added to the configuration file will trigger Microsoft Azure to create a new tunnel. Import the updated file using the instructions in Create a virtual network (classic) by using the Azure portal. After you complete the import, the virtual network dashboard will show the gateway connections to both of your local sites.

Record the Azure gateway IP address and pre-shared keys

After the new network configuration settings are imported, the virtual network dashboard will display the IP address for the Azure gateway. This is the IP address that the VPN devices on both of your sites will connect to. Record this IP address for reference.

You also will need to get the pre-shared IPsec/IKE keys for each tunnel that was created. You will use these keys along with the Azure gateway IP address to configure your on-premises VPN devices.

You need to use PowerShell to get the pre-shared keys. If you aren't familiar with using PowerShell to manage Azure, see Azure PowerShell.

Use the Get-AzureVNetGatewayKey cmdlet to extract the pre-shared keys. Run this cmdlet once for each tunnel. The following example shows the commands you need to run to extract the keys for tunnels between the virtual network "Azure Site" and sites "Site A" and "Site B." In this example, the outputs are saved into separate files. Alternatively, you can pipeline these keys to other PowerShell cmdlets or use them in a script.

```
Get-AzureVNETGatewayKey -VNetName "Azure Site" -LocalNetworkSiteName "Site A" | Set-Content -Path
C:\Keys\KeysForTunnelToSiteA.txt
Get-AzureVNETGatewayKey -VNetName "Azure Site" -LocalNetworkSiteName "Site B" | Set-Content -Path
C:\Keys\KeysForTunnelToSiteB.txt
```

Configure on-premises VPN devices

Microsoft Azure provides VPN device configuration scripts for supported VPN devices. Click the **Download VPN Device Script** link on the virtual network dashboard for the appropriate script for your VPN devices.

The script you download will have the configuration setting for the first site that you configured when you set up your virtual network, and can be used as is to configure the VPN device for that site. For example, if you specified Site A as the **LOCAL NETWORK** when you created your virtual network, the VPN device script can be used for Site A. However, you will need to modify it to configure the VPN device for Site B. Specifically, you need to update the pre-shared key to match the key for the second site.

For example, if you're using a Routing and Remote Access Service (RRAS) VPN device for your sites, you will need to:

1. Open the configuration script in any text editor.

2. Find the `#Add S2S VPN interface` section.

3. Find the **Add-VpnS2SInterface** command in this section. Verify that the value for the *SharedSecret*

parameter matches the pre-shared key for the site for which you're configuring the VPN device.

Other devices might require additional verifications. For example, the configuration scripts for Cisco devices set ACL rules by using the local IP address ranges. You need to review and verify all references to the local site in the configuration script before you use it.

Checkpoint: Review the VPN status

At this point, both of your sites are connected to your Azure virtual network through the VPN gateways. You can validate the status of the multi-site VPN by running the following command in PowerShell.

```
Get-AzureVnetConnection -VNetName "Azure Site" | Format-Table LocalNetworkSiteName, ConnectivityState
```

If both tunnels are up and running, the output of this command will look like the following.

```
LocalNetworkSiteName    ConnectivityState
--------------------    -----------------
Site A                  Connected
Site B                  Connected
```

You can also verify connectivity by viewing the virtual network dashboard in the Azure management portal. The **STATUS** column for both sites will show as **Connected**.

> **NOTE**
>
> It can take several minutes after the connection is successfully established for the status change to appear in the Azure management portal.

Phase 3: Configure virtual machines

You need to create a minimum of two virtual machines in Microsoft Azure for this deployment: a domain controller and a file server that will serve as the DAG witness.

1. Create virtual machines for your domain controller and your file server using the instructions in Create a Virtual Machine Running Windows. Make sure that you select the virtual network you created for **REGION/AFFINITY GROUP/VIRTUAL NETWORK** when specifying the settings of your virtual machines.

2. Specify preferred IP addresses for both the domain controller and the file server using Azure PowerShell. When you specify a preferred IP address for a VM, it needs to be updated, which will require restarting the VM. The following example sets the IP addresses for Azure-DC and Azure-FSW to 10.0.0.10 and 10.0.0.11 respectively.

   ```
   Get-AzureVM Azure-DC | Set-AzureStaticVNetIP -IPAddress 10.0.0.10 | Update-AzureVM
   ```

   ```
   Get-AzureVM Azure-FSW | Set-AzureStaticVNetIP -IPAddress 10.0.0.11 | Update-AzureVM
   ```

   > **NOTE**
   >
   > A VM with a preferred IP address will attempt to use that address. However, if that address has been assigned to a different VM, the VM with the preferred IP address configuration will not start. To avoid this situation, make sure that the IP address you use isn't assigned to another VM.

3. Provision the domain controller VM on Azure using the standards used by your organization.

4. Prepare the file server with the prerequisites for an Exchange DAG witness:

    a. Add the File Server role using the Add Roles and Features Wizard or the Install-WindowsFeature cmdlet.

    b. Add the Exchange Trusted Subsystems universal security group to the Local Administrators group.

Checkpoint: Review virtual machine status

At this point, your virtual machines should be up and running and should be able to communicate with servers in both of your on-premises datacenters:

- Verify that your domain controller in Azure is replicating with your on-premises domain controllers.
- Verify that you can reach the file server on Azure by name and establish an SMB connection from your Exchange servers.
- Verify that you can reach your Exchange servers by name from the file server on Azure.

Phase 4: Configure the DAG witness

Finally, you need to configure your DAG to use the new witness server. By default, Exchange uses the C:\DAGFileShareWitnesses as the file share witness path on your witness server. If you're using a custom file path, you should also update the witness directory for the specific share.

1. Connect to Exchange Management Shell.

2. Run the following command to configure the witness server for your DAGs.

```
Set-DatabaseAvailabilityGroup -Identity DAG1 -WitnessServer Azure-FSW
```

See the following topics for more information:

Configure database availability group properties

Set-DatabaseAvailabilityGroup

Checkpoint: Validate the DAG file share witness

At this point, you have configured your DAG to use the file server on Azure as your DAG witness. Do the following to validate your configuration:

1. Validate the DAG configuration by running the following command.

```
Get-DatabaseAvailabilityGroup -Identity DAG1 -Status | Format-List Name, WitnessServer,
WitnessDirectory, WitnessShareInUse
```

   Verify that the *WitnessServer* parameter is set to the file server on Azure, the *WitnessDirectory* parameter is set to the correct path, and the *WitnessShareInUse* parameter shows `Primary`.

2. If the DAG has an even number of nodes, the file share witness will be configured. Validate the file share witness setting in cluster properties by running the following command. The value for the *SharePath* parameter should point to the file server and display the correct path.

```
Get-ClusterResource -Cluster MBX1 | Get-ClusterParameter | Format-List
```

3. Next, verify the status of the "File Share Witness" cluster resource by running the following command. The *State* of the cluster resource should display `Online`.

```
Get-ClusterResource -Cluster MBX1
```

4. Lastly, verify that the share is successfully created on the file server by reviewing the folder in File Explorer and the shares in Server Manager.

Remove a database availability group

8/3/2020 • 2 minutes to read • Edit Online

Removing a DAG is a quick and easy task. You can use the EAC or the Exchange Management Shell to remove a DAG.

Looking for other management tasks related to DAGs? Check out Manage database availability groups.

What do you need to know before you begin?

- Estimated time to complete: 1 minute
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Database availability groups" entry in the High availability and site resilience permissions topic.
- Before you can remove a DAG, the DAG must be empty. If the DAG you want to remove contains any Mailbox servers, you must first remove the servers from the DAG. For detailed steps about how to remove a Mailbox server from a DAG, see Manage database availability group membership.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to remove a database availability group

1. Navigate to **Servers** > **Database availability groups**.
2. Select the DAG you want to remove and click **Delete** 🗑.
3. Click **Yes** to confirm the warning and remove the DAG.

Use the Exchange Management Shell to remove a database availability group

This example removes the DAG DAG1.

```
Remove-DatabaseAvailabilityGroup -Identity DAG1
```

How do you know this worked?

To verify that you've successfully removed the DAG, do one of the following:

- In the EAC, go to **Servers** > **Database Availability Groups**, and see if the DAG is still displayed.
- In the Exchange Management Shell, run the following command to see if the DAG still exists:

```
Get-DatabaseAvailabilityGroup <DAGName>
```

If the DAG was successfully deleted, the preceding command will produce an error message indicating the object could not be found.

Configure AutoReseed for a database availability group

8/3/2020 • 5 minutes to read • Edit Online

Use the steps in this topic to configure AutoReseed for a database availability group (DAG) in Exchange Server.

Caution

The AutoReseed feature doesn't perform any prerequisite configuration tasks for you. Installing disks correctly, adding spare disks to the system, replacing bad disks, and formatting new disks must be done manually by an administrator.

For additional management tasks related to DAGs, see Manage database availability groups.

What do you need to know before you begin?

- Estimated time to complete this task: 10 minutes.
- To open the Exchange Management Shell, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Database availability groups" entry in the High availability and site resilience permissions topic.
- A single logical disk/partition per physical disk must be created.
- The specific database and log folder structure described in the steps below must be used.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Step 1: Configure the root paths for databases and volumes

The first step involves configuring the root directories for the databases (*AutoDagDatabasesRootFolderPath*) and volumes (*AutoDagVolumesRootFolderPath*) used by the DAG. The defaults are C:\ExchangeDatabases, and C:\ExchangeVolumes, respectively. You can omit this step if you're using the default paths.

This example illustrates how to configure the root path for the databases.

```
Set-DatabaseAvailabilityGroup DAG1 -AutoDagDatabasesRootFolderPath "C:\ExchDbs"
```

This example illustrates how to configure the root path for the storage volumes.

```
Set-DatabaseAvailabilityGroup DAG1 -AutoDagVolumesRootFolderPath "C:\ExchVols"
```

How do you know this step worked?

To verify that you've successfully configured the root paths for databases and volumes, run the following command.

```
Get-DatabaseAvailabilityGroup DAG1 | Format-List *auto*
```

The output for *AutoDagDatabasesRootFolderPath* and *AutoDagVolumesRootFolderPath* should reflect the configured paths.

Step 2: Configure the number of databases per volume

Next, configure the number of databases per volume (*AutoDagDatabaseCopiesPerVolume*) for the DAG.

This example illustrates how to configure this AutoReseed setting for a DAG configured with 4 databases per volume.

```
Set-DatabaseAvailabilityGroup DAG1 -AutoDagDatabaseCopiesPerVolume 4
```

How do you know this step worked?

To verify that you've successfully configured the number of databases per volume, run the following command.

```
Get-DatabaseAvailabilityGroup DAG1 | Format-List *auto*
```

The output for *AutoDagDatabaseCopiesPerVolume* should reflect the configured value.

Step 3: Create the root directories for databases and volumes

Next, create the directories that correspond to the root directories you configured in Step 1. This example shows how to create the default directories using the command prompt.

```
md C:\ExchangeDatabases
md C:\ExchangeVolumes
```

How do you know this step worked?

To verify that you've successfully configured the root directories for databases and volumes, run the following command.

```
Dir C:\
```

The created directories should appear in the output list.

Step 4: Mount the volume folders

For every volume that will be used for databases (including spare volumes), use the Windows Disk Management application (diskmgmt.msc) to mount each volume in a mounted folder under C:\ExchangeVolumes. For example, if there are 2 volumes with databases and 1 spare volume, mount the volumes to the following mounted folders:

- C:\ExchangeVolumes\Volume1
- C:\ExchangeVolumes\Volume2
- C:\ExchangeVolumes\Volume3

The names of the mounted folders can be any folder name, as long as the folders are mounted under the root volume's path.

How do you know this step worked?

To verify that you've successfully mounted the volume folders, run the following command.

```
Dir C:\
```

The mounted volumes should appear in the output list.

Step 5: Create the database folders

Next, create the database directories under the root path C:\ExchangeDatabases. This example illustrates how to create directories for a storage configuration with 4 databases on each volume.

```
md c:\ExchangeDatabases\db001
```

```
md c:\ExchangeDatabases\db002
```

```
md c:\ExchangeDatabases\db003
```

```
md c:\ExchangeDatabases\db004
```

How do you know this step worked?

To verify that you've successfully mounted the database folders, run the following command.

```
Dir C:\ExchangeDatabases
```

The created directories should appear in the output list.

Step 6: Create the mount points for the databases

Create the mount points for each database and link the mount point to the correct volume. For example, the mounted folder for db001 should be at C:\ExchangeDatabases\db001. You can use diskmgmt.msc or mountvol.exe to do this. This example illustrates how to mount db001 to C:\ExchangeDatabases\db001 using mountvol.exe.

```
Mountvol.exe c:\ExchangeDatabases\db001 \\?\Volume (GUID)
```

How do you know this step worked?

To verify that you've successfully created the mount points for the database, run the following command.

```
Mountvol.exe C:\ExchangeDatabases\db001 /L
```

The mounted volume should appear in the mount point list.

Step 7: Create the database directory structure

Next, create two directories underneath the folders you created in Step 5, one for each database and one for each of the database's log stream that will be stored on the same volume. You must use the following format for your directory structure:

C:\< *DatabaseFolderName* >\ *DatabaseName* \< *DatabaseName* >.db

C:\< *DatabaseFolderName* >\ *DatabaseName* \< *DatabaseName* >.log

This example illustrates how to create directories for 4 databases that will be stored on Volume 1:

```
md c:\ExchangeDatabases\db001\db001.db
```

```
md c:\ExchangeDatabases\db001\db001.log
```

```
md c:\ExchangeDatabases\db002\db002.db
```

```
md c:\ExchangeDatabases\db002\db002.log
```

```
md c:\ExchangeDatabases\db003\db003.db
```

```
md c:\ExchangeDatabases\db003\db003.log
```

```
md c:\ExchangeDatabases\db004\db004.db
```

```
md c:\ExchangeDatabases\db004\db004.log
```

Repeat the preceding commands for databases on every volume.

How do you know this step worked?

To verify that you've successfully created the database directory structure, run the following command.

```
Dir C:\ExchangeDatabases /s
```

The created directories should appear in the output list.

Step 8: Create databases

Create databases with log and database paths configured with the appropriate folders. This example illustrates how to create a database that's stored in the newly created directory and mount point structure.

```
New-MailboxDatabase -Name db001 -Server MBX1 -LogFolderPath C:\ExchangeDatabases\db001\db001.log -EdbFilePath
C:\ExchangeDatabases\db001\db001.db\db001.edb
```

How do you know this step worked?

To verify that you've successfully created databases in the appropriate folder, run the following command.

```
Get-MailboxDatabase db001 | Format List *path*
```

Database properties that are returned should indicate that the database file and log files are being stored in the above folders.

How do you know this task worked?

To verify that you've configured AutoReseed for a DAG, do the following:

1. Run the following command to verify the DAG is configured correctly.

   ```
   Get-DatabaseAvailabilityGroup DAG1 | Format-List *auto*
   ```

2. Run the following command to verify the directory structure is configured correctly (below are the default paths; if necessary, substitute the paths for the paths you're using).

   ```
   Dir c:\ExchangeDatabases /s
   ```

   ```
   Dir c:\ExchangeVolumes /s
   ```

Configure database availability group network properties

8/3/2020 • 2 minutes to read • Edit Online

Configurable properties include the name of the DAG network, a description field for the DAG network, a list of subnets that are used by the DAG network, and whether the DAG network is enabled for replication.

Looking for other management tasks related to DAGs? Check out Manage database availability groups.

What do you need to know before you begin?

- Estimated time to complete: 1 minute
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Database availability groups" entry in the High availability and site resilience permissions topic.
- You can configure a DAG network only when automatic network configuration has been disabled for a DAG. For detailed steps about how to disable automatic network configuration for a DAG, see Configure database availability group properties.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to configure database availability group network properties

1. In the EAC, go to **Servers** > **Database Availability Groups**.
2. Select the DAG you want to configure, and in the Details pane, under the DAG network you want to configure, choose from the following configuration options.

> **NOTE**
>
> These options will only be visible if you have selected **Configure database availability group networks manually** on the DAG properties page.

- **Disable Replication** or **Enable Replication**: Configures the replication settings for the DAG network.
- **Remove**: Removes a DAG network. Before you can remove a DAG network, you must first remove all associated subnets from the DAG network.
- **View details**: Configures DAG network properties, such as the name, description, and associated subnets for the DAG network. You can also view the network interfaces associated with those subnets, and enable or disable replication for the DAG network.

Use the Exchange Management Shell to configure database availability group network properties

This example adds a subnet of 10.0.0.0 and subnet mask of 255.0.0.0 to the DAG network MapiDagNetwork in the DAG DAG1.

```
Set-DatabaseAvailabilityGroupNetwork -Subnets 10.0.0.0/8 -Identity DAG1\MapiDagNetwork
```

How do you know this worked?

To verify that you've successfully configured the DAG network, do the following:

- In the Exchange Management Shell, run the following command to display DAG network configuration settings and verify the DAG network was configured successfully.

  ```
  Get-DatabaseAvailabilityGroupNetwork <DAGNetworkName> | Format-List
  ```

For more information

Set-DatabaseAvailabilityGroupNetwork

Get-DatabaseAvailabilityGroupNetwork

New-DatabaseAvailabilityGroupNetwork

Remove-DatabaseAvailabilityGroupNetwork

Create a database availability group network

8/3/2020 • 2 minutes to read • Edit Online

You can use the EAC or the Exchange Management Shell to create a DAG network.

Looking for other management tasks related to DAGs? Check out Manage database availability groups.

What do you need to know before you begin?

- Estimated time to complete: 1 minute
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Database availability groups" entry in the High availability and site resilience permissions topic.
- You can create a DAG network only when automatic network configuration has been disabled for a DAG. For detailed steps about how to disable automatic network configuration for a DAG, see Configure database availability group properties.
- When creating a DAG network, you must assign unique subnets that aren't in use by another DAG network. If you use subnets that are assigned to an existing DAG network, they will be removed from that DAG network and added to the newly created DAG network.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to create a database availability group network

1. In the EAC, go to **Servers** > **Database Availability Groups**.
2. Select the DAG you want to configure, and then click .
3. On the **new database availability group network** page, provide the following information:
   - **Database availability group network name**: Use this field to type a name for the network that's unique in the DAG.
   - **Description**: Use this field to provide a text description of the DAG network.
   - **Subnets**: Use this field to associate one or more subnets with the DAG network. Click + to add a subnet, click ✎ to edit a subnet, and click minus (-) to remove a subnet.
4. Click **Save** to create the DAG network.

Use the Exchange Management Shell to create a database availability group network

This example creates the network ReplicationDagNetwork02 with a subnet of 10.0.0.0 and a bitmask of 8 in the

DAG DAG1. Replication is enabled for the network, and an optional description of the network is also being added.

```
New-DatabaseAvailabilityGroupNetwork -DatabaseAvailabilityGroup DAG1 -Name ReplicationDagNetwork02 -
Description "Replication network 2" -Subnets 10.0.0.0/8 -ReplicationEnabled:$True
```

How do you know this worked?

To verify that you've successfully created a DAG network, do one of the following:

- In the EAC, navigate to **Servers** > **Database Availability Groups**. Select the appropriate DAG, and the newly created DAG network is displayed in the details pane.

- In the Exchange Management Shell, run the following command to verify the DAG network was created and to display DAG network configuration information.

  ```
  Get-DatabaseAvailabilityGroupNetwork <DAGNetworkName> | Format-List
  ```

For more information

Set-DatabaseAvailabilityGroupNetwork

Get-DatabaseAvailabilityGroupNetwork

New-DatabaseAvailabilityGroupNetwork

Remove-DatabaseAvailabilityGroupNetwork

Manage database availability group membership

8/3/2020 • 3 minutes to read • Edit Online

When you add a server to a DAG, the server works with the other DAG members to provide automatic database-level recovery from database, server, or network failures. When you remove a server from a DAG, the server is no longer automatically protected from failures.

Looking for other management tasks related to DAGs? Check out Manage database availability groups.

What do you need to know before you begin?

- Estimated time to complete: 5 minutes per server
- To open the EAC, see Exchange admin center in Exchange Server. To open the Exchange Management Shell, see Open the Exchange Management Shell.
- DAGs use Windows Failover Clustering (WFC) technologies. Each Mailbox server that's a member of a DAG is also a node in the underlying cluster used by the DAG. As a result, at any specific time, a Mailbox server can be a member of only one DAG. Because DAGs use WFC technology, all servers added to a DAG must be running the same operating system: either Windows Server 2008 R2 Enterprise or Datacenter Edition, or the Standard or Datacenter Edition of Windows Server 2012 or Windows Server 2012 R2.
- If you're adding Mailbox servers running Windows Server 2012, you must pre-stage the cluster name object (CNO) for the DAG. If you're adding Mailbox servers running Windows Server 2012 R2, and your DAG does not have an administrative access point, then you do not need to pre-stage a CNO, as DAGs without administrative access points do not have a CNO. For detailed steps, see Pre-stage the cluster name object for a database availability group.
- Before you can add members to a DAG, you must first create a DAG. For detailed steps, see Create a database availability group.
- You must remove all replicated database copies from the server before you can remove it from a DAG.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Database availability groups" entry in the High availability and site resilience permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to manage database availability group membership

1. In the EAC, go to **Servers** > **Database Availability Groups**.
2. Select the DAG you want to configure, and then click ⚙.
   - To add one or more Mailbox servers to the DAG, click **+**, select the servers from the list, click **Add**, and then click **OK**.

   - To remove one or more Mailbox servers from the DAG, select the servers, and then click the minus (-) icon.

3. Click **Save** to save the changes.

4. When the task has completed successfully, click **Close**.

Use the Exchange Management Shell to manage database availability group membership

This example adds the Mailbox server MBX1 to the DAG DAG1.

```
Add-DatabaseAvailabilityGroupServer -Identity DAG1 -MailboxServer MBX1
```

This example removes the Mailbox server MBX1 from the DAG DAG1. Before running this command, make sure that no replicated databases exist on the Mailbox server.

```
Remove-DatabaseAvailabilityGroupServer -Identity DAG1 -MailboxServer MBX1
```

This example removes the configuration settings for the Mailbox server MBX4 from the DAG DAG2. MBX4 is expected to be offline for an extended period, so its configuration is being removed from the DAG while it's offline to establish quorum with the remaining online DAG members.

```
Remove-DatabaseAvailabilityGroupServer -Identity DAG2 -MailboxServer MBX4 -ConfigurationOnly
```

How do you know this worked?

To verify that you've successfully managed DAG membership, do one of the following:

- In the EAC, navigate to **Servers** > **Database Availability Groups**. The current DAG membership is displayed in the **Member Servers** column.

- In the Exchange Management Shell, run the following command to display DAG membership information.

  ```
  Get-DatabaseAvailabilityGroup <DAGName> | Format-List Servers
  ```

For more information

Add-DatabaseAvailabilityGroupServer

Remove-DatabaseAvailabilityGroupServer

Configure database availability group properties

8/3/2020 • 3 minutes to read • Edit Online

The Exchange Management Shell enables you to configure DAG properties that aren't available in the EAC, such as alternate witness server and alternate witness directory information, the TCP port used for replication, and datacenter activation coordination (DAC) mode.

What do you need to know before you begin?

- Estimated time to complete: 1 minute
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Database availability groups" entry in the High availability and site resilience permissions topic.
- DAG property values are stored in both Active Directory and the cluster database. However, some properties are stored only in the cluster database. As a result, the underlying cluster for the DAG must be running and have quorum to set the properties for:
  - ReplicationPort
  - NetworkCompression
  - NetworkEncryption
  - DiscoverNetworks
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to configure database availability group properties

1. In the EAC, go to **Servers** > **Database Availability Groups**.
2. Select the DAG you want to configure and click ✏.
3. Use the **General** page to view DAG membership and operational status, and to configure the DAG's witness server, witness directory, and automatic network configuration:
   - **Witness server**: The host name or fully qualified domain name (FQDN) of the witness server for the DAG. Although this is a required property for all DAGs, the witness server is used when there is an even number of DAG members and the quorum model in use by the cluster is Node and File Share Majority.
   - **Witness directory**: The full path of the directory used to store the witness.log file on the witness server. Although this is a required property for all DAGs, the witness directory is used only when the DAG's witness server is in use.
   - **Database availability group members**: A read-only field that displays a list of DAG members

and their current operational status.

- **Configure database availability group networks manually**: A check box that you select when you want to configure all DAG networks manually. When the check box is clear, the system configures DAG networks automatically based on network interface configuration, and the **Set-DatabaseAvailabilityGroupNetwork** and **New-DatabaseAvailabilityGroupNetwork** cmdlets are disabled for the DAG.

4. Use the **IP addresses** page to view and modify the IP addresses assigned to the DAG:

   - Select an existing IP address and click ✎ to modify it.
   - Select an existing IP address and click the minus icon (delete) to remove it.
   - Enter an IP address and click + to add it to the DAG.

5. Click **Save** to save any changes that were made.

Use the Exchange Management Shell to configure database availability group properties

This example sets the witness directory to C:\DAG1DIR for the DAG DAG1.

```
Set-DatabaseAvailabilityGroup -Identity DAG1 -WitnessDirectory C:\DAG1DIR
```

This example preconfigures an alternate witness server of MBX3 and an alternate witness directory of C:\DAGFileShareWitnesses\DAG1.contoso.com for the DAG DAG1.

```
Set-DatabaseAvailabilityGroup -Identity DAG1 -AlternateWitnessDirectory
C:\DAGFileShareWitnesses\DAG1.contoso.com -AlternateWitnessServer MBX3
```

This example configures the DAG DAG1 to use Dynamic Host Configuration Protocol (DHCP) to obtain an IP address.

```
Set-DatabaseAvailabilityGroup -Identity DAG1 -DatabaseAvailabilityGroupIPAddresses 0.0.0.0
```

This example configures the DAG DAG1 to use a static IP address of 10.0.0.8.

```
Set-DatabaseAvailabilityGroup -Identity DAG1 -DatabaseAvailabilityGroupIPAddresses 10.0.0.8
```

This example configures the multi-subnet DAG DAG1 with multiple static IP addresses.

```
Set-DatabaseAvailabilityGroup -Identity DAG1 -DatabaseAvailabilityGroupIPAddresses 10.0.0.8,10.0.1.8
```

This example configures the DAG DAG1 for DAC mode.

```
Set-DatabaseAvailabilityGroup -Identity DAG1 -DatacenterActivationMode DagOnly
```

This example configures the replication port for the DAG DAG1 to be 63132.

```
Set-DatabaseAvailabilityGroup -Identity DAG1 -ReplicationPort 63132
```

NOTE

After changing the default replication port for a DAG, you must manually modify the Windows Firewall exceptions on each member of the DAG to allow communication to occur over the specified port.

How do you know this worked?

To verify that you've successfully configured the DAG, do the following:

- In the Exchange Management Shell, run the following command to display DAG configuration settings and verify the DAG was configured successfully.

  ```
  Get-DatabaseAvailabilityGroup <DAGName> | Format-List
  ```

For more information

Create a database availability group

Remove a database availability group

Create a database availability group network

Manage database availability group membership

Get-DatabaseAvailabilityGroup

Set-DatabaseAvailabilityGroup

Manage mailbox database copies

8/3/2020 • 33 minutes to read • Edit Online

In Exchange Server, you can use the Exchange Management Console (EAC) or the Exchange Management Shell to add mailbox database copies after a database availability group (DAG) has been created, configured, and populated with Mailbox server members.

Managing database copies

After multiple copies of a database are created, you can use the EAC or the Exchange Management Shell to monitor the health and status of each copy and to perform other management tasks associated with database copies. Some of the management tasks you may need to perform include suspending or resuming a database copy, seeding a database copy, monitoring database copies, configuring database copy settings, and removing a database copy.

Suspending and resuming database copies

For a variety of reasons, such as performing planned maintenance, you may need to suspend and resume continuous replication activity for a database copy. In addition, some administrative tasks, such as seeding, require that you first suspend a database copy. We recommend that you suspend all replication activity when the path for the database or its log files is being changed. You can suspend and resume database copy activity by using the EAC, or by running the **Suspend-MailboxDatabaseCopy** and **Resume-MailboxDatabaseCopy** cmdlets in the Exchange Management Shell. For detailed steps about how to suspend or resume continuous replication activity for a database copy, see Suspend or resume a mailbox database copy.

Seeding a database copy

Seeding, also known as *updating*, is the process in which either a blank database or a copy of the production database, is added to the target copy location on another Mailbox server in the same DAG as the active database. This becomes the baseline database for the copy maintained by that server.

Depending on the situation, you can seed a database by using an automatic process or a manual process that you initiate. When a database copy is added, the copy will be automatically seeded, provided that the target server and its storage are properly configured. If you want to manually seed a database copy and don't want automatic seeding to occur when creating the copy, you can use the *SeedingPostponed* parameter when running the Add-MailboxDatabaseCopy cmdlet.

Database copies rarely need to be reseeded after the initial seeding. However, if reseeding is necessary, or if you want to manually seed a database copy instead of having the system automatically seed the copy, you have two options. You can reseed a database by using the Update Mailbox Database Copy wizard in the EAC or by using the Update-MailboxDatabaseCopy cmdlet in the Exchange Management Shell. Before seeding a database copy, you must first suspend the mailbox database copy. For detailed steps about how to seed a database copy, see Update a mailbox database copy.

After a manual seed operation has completed, replication for the seeded mailbox database copy is automatically resumed. If you don't want replication to automatically resume, you can use the *ManualResume* parameter when running the Update-MailboxDatabaseCopy cmdlet.

Choosing what to seed

When you perform a seed operation, you can choose to seed the mailbox database copy, the content index catalog for the mailbox database copy, or both the database copy and the content index catalog copy. The default behavior of the Update Mailbox Database Copy wizard and the Update-MailboxDatabaseCopy cmdlet is to seed both the mailbox database copy and the content index catalog copy. To seed just the mailbox database copy without seeding the content index catalog, use the *DatabaseOnly* parameter when running the Update-MailboxDatabaseCopy cmdlet. To seed just the content index catalog copy, use the *CatalogOnly* parameter when running the Update-MailboxDatabaseCopy cmdlet.

Selecting the seeding source

Any healthy database copy can be used as the seeding source for an additional copy of that database. This is particularly useful when you have a DAG that has been extended across multiple physical locations. For example, consider a four-member DAG deployment, where two members (MBX1 and MBX2) are located in Portland, Oregon and two members (MBX3 and MBX4) are located in New York, New York. A mailbox database named DB1 is active on MBX1 and there are passive copies of DB1 on MBX2 and MBX3. When adding a copy of DB1 to MBX4, you have the option of using the copy on MBX3 as the source for seeding. In doing so, you avoid seeding over the wide area network (WAN) link between Portland and New York.

To use a specific copy as a source for seeding when adding a new database copy, you can do the following:

- Use the *SeedingPostponed* parameter when running the Add-MailboxDatabaseCopy cmdlet to add the database copy. If you don't use the *SeedingPostponed* parameter, the database copy will be explicitly seeded using the active copy of the database as the source.

- You can specify the source server you want to use as part of the Update Mailbox Database Copy wizard in the EAC, or you can use the *SourceServer* parameter when running the Update-MailboxDatabaseCopy cmdlet to specify the desired source server for seeding. In the preceding example, you would specify MBX3 as the source server. If the you don' t use *SourceServer* parameter, the database copy will be explicitly seeded from the active copy of the database.

Seeding and networks

In addition to selecting a specific source server for seeding a mailbox database copy, you can also use the Exchange Management Shell to specify which DAG networks to use. You have the option to override the DAG network's compression and encryption settings during the seed operation.

You can specify the networks you want to use for seeding by using the *Network* parameter when running the Update-MailboxDatabaseCopy cmdlet and specify the DAG networks that you want to use. If you don't use the *Network* parameter, the system uses the following default behavior for selecting a network to use for the seeding operation:

- If the source server and target server are on the same subnet and a replication network has been configured that includes the subnet, the replication network will be used.

- If the source server and target server are on different subnets, even if a replication network that contains those subnets has been configured, the client (MAPI) network will be used for seeding.

- If the source server and target server are in different datacenters, the client (MAPI) network will be used for seeding.

At the DAG level, DAG networks are configured for encryption and compression. The default settings use encryption and compression only for communications on different subnets. If the source and target are on different subnets and the DAG is configured with the default values for *NetworkCompression* and *NetworkEncryption*, you can override these values by using the *NetworkCompressionOverride* and *NetworkEncryptionOverride* parameters, respectively, when running the Update-MailboxDatabaseCopy cmdlet.

Seeding process

When you begin a seeding process by using the Add-MailboxDatabaseCopy or Update-MailboxDatabaseCopy cmdlets, the following tasks are performed:

1. Database properties from Active Directory are read to validate the specified database and servers, and to verify that the source and target servers are running Exchange Server, they are both members of the same

DAG, and that the specified database isn't a recovery database. The database file paths are also read.

2. Preparations occur for reseed checks from the Microsoft Exchange Replication service on the target server.
3. The Microsoft Exchange Replication service on the target server checks for the presence of database and transaction log files in the file directories read by the Active Directory checks in step 1.
4. The Microsoft Exchange Replication service returns the status information from the target server to the administrative interface from where the cmdlet was run.
5. If all preliminary checks have passed, you're prompted to confirm the operation before continuing. If you confirm the operation, the process continues. If an error is encountered during the preliminary checks, the error is reported and the operation fails.
6. The seed operation is started from the Microsoft Exchange Replication service on the target server.
7. The Microsoft Exchange Replication service suspends database replication for the active database copy.
8. The state information for the database is updated by the Microsoft Exchange Replication service to reflect a status of Seeding.
9. If the target server doesn't already have the directories for the target database and log files, they are created.
10. A request to seed the database is passed from the Microsoft Exchange Replication service on the target server to the Microsoft Exchange Replication service on the source server using TCP. This request and the subsequent communications for seeding the database occur on a DAG network that has been configured as a replication network.
11. The Microsoft Exchange Replication service on the source server initiates an Extensible Storage Engine (ESE) streaming backup via the Microsoft Exchange Information Store service interface.
12. The Microsoft Exchange Information Store service streams the database data to the Microsoft Exchange Replication service.
13. The database data is moved from the source server's Microsoft Exchange Replication service to the target server's Microsoft Exchange Replication service.
14. The Microsoft Exchange Replication service on the target server writes the database copy to a temporary directory located in the main database directory called *temp-seeding*.
15. The streaming backup operation on the source server ends when the end of the database is reached.
16. The write operation on the target server completes, and the database is moved from the temp-seeding directory to the final location. The temp-seeding directory is deleted.
17. On the target server, the Microsoft Exchange Replication service proxies a request to the Microsoft Exchange Search service to mount the content index catalog for the database copy, if it exists. If there are existing out-of-date catalog files from a previous instance of the database copy, the mount operation fails, which triggers the need to replicate the catalog from the source server. Likewise, if the catalog doesn't exist on a new instance of the database copy on the target server, a copy of the catalog is required. The Microsoft Exchange Replication service directs the Microsoft Exchange Search service to suspend indexing for the database copy while a new catalog is copied from the source.
18. The Microsoft Exchange Replication service on the target server sends a seed catalog request to the Microsoft Exchange Replication service on the source server.
19. On the source server, the Microsoft Exchange Replication service requests the directory information from the Microsoft Exchange Search service and requests that indexing be suspended.

20. The Microsoft Exchange Search service on the source server returns the search catalog directory information to the Microsoft Exchange Replication service.

21. The Microsoft Exchange Replication service on the source server reads the catalog files from the directory.

22. The Microsoft Exchange Replication service on the source server moves the catalog data to the Microsoft Exchange Replication service on the target server using a connection across the replication network. After the read is complete, the Microsoft Exchange Replication service sends a request to the Microsoft Exchange Search service to resume indexing of the source database.

23. If there are any existing catalog files on the target server in the directory, the Microsoft Exchange Replication service on the target server deletes them.

24. The Microsoft Exchange Replication service on the target server writes the catalog data to a temporary directory called *CiSeed.Temp* until the data is completely transferred.

25. The Microsoft Exchange Replication service moves the complete catalog data to the final location.

26. The Microsoft Exchange Replication service on the target server resumes search indexing on the target database.

27. The Microsoft Exchange Replication service on the target server returns a completion status.

28. The final result of the operation is passed to the administrative interface from which the cmdlet was called.

Configuring database copies

After a database copy is created, you can view and modify its configuration settings when needed. You can view some configuration information by examining the **Properties** page for a database copy in the EAC. You can also use the Get-MailboxDatabase and Set-MailboxDatabaseCopy cmdlets in the Exchange Management Shell to view and configure database copy settings, such as replay lag time, truncation lag time, and activation preference order. For detailed steps about how to view and configure database copy settings, see Configure mailbox database copy properties.

Using replay lag and truncation lag options

Mailbox database copies support the use of a *replay lag time* and a *truncation lag time*, both of which are configured in minutes. Setting a replay lag time enables you to take a database copy back to a specific point in time. Setting a truncation lag time enables you to use the logs on a passive database copy to recover from the loss of log files on the active database copy. Because both of these features result in the temporary buildup of log files, using either of them will affect your storage design.

Replay lag time

Replay lag time is a mailbox database copy property that specifies the amount of time, in minutes, to delay log replay for the database copy. The replay lag timer starts when a log file has been replicated to the passive copy and has successfully passed inspection. By delaying the replay of logs to the database copy, you have the capability to recover the database to a specific point in time in the past. A mailbox database copy configured with a replay lag time greater than 0 is referred to as a *lagged mailbox database copy*, or simply, a *lagged copy*.

A strategy that uses database copies and the litigation hold features in Exchange Server can provide protection against a range of failures that would ordinarily cause data loss. However, these features can't provide protection against data loss in the event of logical corruption, which although rare, can cause data loss. Lagged copies are designed to prevent loss of data in the case of logical corruption. Generally, there are two types of logical corruption:

- **Database logical corruption**: The database pages checksum matches, but the data on the pages is wrong logically. This can occur when ESE attempts to write a database page and even though the operating system returns a success message, the data is either never written to the disk or it's written to the wrong place. This is referred to as a *lost flush*. To prevent lost flushes from losing data, ESE includes a lost flush detection mechanism in the database along with a page patching feature (single page restore).

- **Store logical corruption**: Data is added, deleted, or manipulated in a way that the user doesn't expect. These cases are generally caused by third-party applications. It's generally only corruption in the sense that the user views it as corruption. The Exchange store considers the transaction that produced the logical corruption to be a series of valid MAPI operations. The litigation hold feature in Exchange Server provides protection from store logical corruption (because it prevents content from being permanently deleted by a user or application). However, there may be scenarios where a user mailbox becomes so corrupted that it would be easier to restore the database to a point in time prior to the corruption, and then export the user mailbox to retrieve uncorrupted data.

The combination of database copies, hold policy, and ESE single page restore leaves only the rare but catastrophic store logical corruption case. Your decision on whether to use a database copy with a replay lag (a lagged copy) will depend on which third-party applications you use and your organization's history with store logical corruption.

If you choose to use lagged copies, be aware of the following implications for their use:

- The replay lag time is an administrator-configured value, and by default, it's disabled.

- The replay lag time setting has a default setting of 0 days, and a maximum setting of 14 days.

- Lagged copies aren't considered highly available copies. Instead, they are designed for disaster recovery purposes, to protect against store logical corruption.

- The greater the replay lag time set, the longer the database recovery process. Depending on the number of log files that need to replayed during recovery, and the speed at which your hardware can replay them, it may take several hours or more to recover a database.

- We recommend that you determine whether lagged copies are critical for your overall disaster recovery strategy. If using them is critical to your strategy, we recommend using multiple lagged copies, or using a redundant array of independent disks (RAID) to protect a single lagged copy, if you don't have multiple lagged copies. If you lose a disk or if corruption occurs, you don't lose your lagged point in time.

- Lagged copies cant be patched with the ESE single page restore feature. If a lagged copy encounters database page corruption (for example, a -1018 error), the copy will have to be reseeded. Reseeding will lose the lagged aspect of the copy.

If you want the database to replay all log files and make the database copy current, then activating and recovering a lagged mailbox database copy is an easy process . If you want to replay log files up to a specific point in time, the prosess is more difficult because you have to manually manipulate log files and run Exchange Server Database Utilities (Eseutil.exe).

For detailed steps about how to activate a lagged mailbox database copy, see Activate a lagged mailbox database copy.

Truncation lag time

Truncation lag time is the property of a mailbox database copy that specifies the amount of time, in minutes, to delay log deletion for the database copy after the log file has been replayed into the database copy. The truncation lag timer starts when a log file has been replicated to the passive copy, successfully passed inspection, and has been successfully replayed into the copy of the database. By delaying the truncation of log files from the database copy, you have the capability to recover from failures that affect the log files for the active copy of the database.

Database copies and log truncation

Log truncation works the same in Exchange 2016 and Exchange 2019 as it did in Exchange 2010. Truncation behavior is determined by the replay lag time and truncation lag time settings for the copy.

The following criteria must be met for a database copy's log file to be truncated when lag settings are left at their default values of 0 (disabled):

- The log file must have been successfully backed up, or circular logging must be enabled.
- The log file must be below the checkpoint (the minimum log file required for recovery) for the database.
- All other lagged copies must have inspected the log file.
- All other copies (except lagged copies) must have replayed the log file.

The following criteria must be met for truncation to occur for a lagged database copy:

- The log file must be below the checkpoint for the database.
- The log file must be older than ReplayLagTime + TruncationLagTime.
- The log file must have been truncated on the active copy.

In Exchange Server, log truncation doesn't occur on an active mailbox database copy when one or more passive copies are suspended. If planned maintenance activities are going to take an extended period of time (for example, several days), you may have considerable log file buildup. To prevent the log drive from filling up with transaction logs, you can remove the affected passive database copy instead of suspending it. When the planned maintenance is completed, you can re-add the passive database copy.

Exchange Server now has a feature called *loose truncation* that is disabled by default. During normal operations, each database copy keeps logs that need to be shipped to other database copies until all copies of a database confirm they have replayed (passive copies) or received (lagged copies) the log files. This is default log truncation behavior. If a database copy goes offline for some reason, the log files begin accumulating on the disks used by the other copies of the database. If the affected database copy remains offline for an extended period, this can cause the other database copies to run out of disk space.

Truncation behavior is different when loose truncation and circular logging are enabled. Each database copy tracks its own free disk space and applies loose truncation behavior if free space gets low.

- For the active copy, the oldest straggler (the passive database copy that is farthest behind in log replay) is ignored and truncation respects the oldest remaining passive copies. The active database copy is where global truncation is calculated.
- For a passive copy, if space gets low, it will independently truncate its log files using the configured parameters described later in the Registry Value table.The passive copies will attempt to respect the truncation decision made on the active copy. Despite the implication of the name MinCopiesToProtect, Exchange will only ignore the oldest known straggler at the time truncation is run.

When the offline database is brought back online, it will be missing log files that have been deleted from the other healthy copies, and its database copy status will be FailedAndSuspended. In this event, if Autoreseed is configured, the affected copy will be automatically reseeded. If Autoreseed is not configured, the database copy will need to be manually seeded by an administrator.

If circular logging is disabled, loose truncation respects backups if they have been taken, so if logs have not been backed up they will not be removed by loose Truncation.

truncation is a recommended feature for preferred architecture where backups are not used and circular logging is enabled.

The required number of healthy copies, the free disk space threshold, and the number of logs to keep are all configurable parameters. By default, the free disk space threshold is 204800 MB (200 GB), and the number of logs to keep is 100,000 (100 GB) for passive copies, and 10,000 (10 GB) for active copies.

Enabling loose truncation and configuring loose truncation parameters is performed by editing the Windows registry on each DAG member. There are three registry values that can be configured, that are all stored under HKLM\Software\Microsoft\ExchangeServer\v15\BackupInformation. The BackupInformation key the following DWORD values do not exist by default and must be manually created. The DWORD registry values under BackupInformation are described in the following table:

REGISTRY VALUE	DESCRIPTION	DEFAULT VALUE
LooseTruncation_MinCopiesToProtect	This key is used to enable loose truncation. It represents the number of passive copies to protect from loose truncation on the active copy of a database. Setting the value of this key to 0 disables loose truncation.	0
LooseTruncation_MinDiskFreeSpaceThresholdInMB	Available disk space (in MB) threshold for triggering loose truncation. If free disk space falls below this value, loose truncation is triggered.	If this registry value is not configured, the default value used by loose truncation is 200 GB.
LooseTruncation_MinLogsToProtect	The minimum number of log files to retain on healthy copies whose logs are being truncated. If this registry value is configured, then the configured value applies to both active and passive copies.	If this registry value is not configured, then default values of 100,000 for passive database copies and 10,000 for active database copies is used.

When using the LooseTruncation_MinLogsToProtect registry value, note that the behavior is different for active and passive database copies. On the active database copy, this is the number of extra logs that are retained preceding those that are required by the protected passive copies and the required range of the active copy.On a passive database copy, this is the number of logs maintained from the latest available log. One tenth of this number is also used to maintain logs prior to the required range of this passive copy. The two limits are in place to ensure that lagged database copies don't take up too much space, since their required range is typically very large.

Database activation policy

There are scenarios in which you may want to create a mailbox database copy and prevent the system from automatically activating that copy in the event of a failure, for example:

- If you deploy one or more mailbox database copies to an alternate or standby datacenter.
- If you configure a lagged database copy for recovery purposes.
- If you're performing maintenance or an upgrade of a server.

In each of the preceding scenarios, you have database copies that you don't want the system to activate automatically. To prevent the system from automatically activating a mailbox database copy, you can configure the copy to be blocked (suspended) for activation. This allows the system to maintain the currency of the database through log shipping and replay, but prevents the system from automatically activating and using the copy. Copies blocked for activation must be manually activated by an administrator. You can configure the database activation policy for an entire server by using the Set-MailboxServer cmdlet or an individual database copy by using the Set-MailboxDatabaseCopy cmdlet to set the *DatabaseCopyAutoActivationPolicy* parameter to Blocked.

For more information about configuring database activation policy, see Configure activation policy for a mailbox database copy.

Effect of mailbox moves on continuous replication

On a very busy mailbox database with a high log generation rate, there is a greater chance for data loss if replication to the passive database copies can't keep up with log generation. One scenario that can introduce a high log generation rate is mailbox moves. Exchange Server includes a Data Guarantee API that's used by services such as the Exchange Mailbox Replication service (MRS) to check the health of the database copy architecture based on the value of the *DataMoveReplicationConstraint* parameter that was set by the system or an administrator. Specifically, the Data Guarantee API can be used to:

- **Check replication health**: Confirms that the prerequisite number of database copies is available.
- **Check replication flush**: Confirms that the required log files have been replayed against the prerequisite number of database copies.

When executed, the API returns the following status information to the calling application:

- **Retry**: Signifies that there are transient errors that prevent a condition from being checked against the database.
- **Satisfied**: Signifies that the database meets the required conditions or the database isn't replicated.
- **NotSatisfied**: Signifies that the database doesn't meet the required conditions. In addition, information is provided to the calling application as to why the **NotSatisfied** response was returned.

The value of the *DataMoveReplicationConstraint* parameter for the mailbox database determines how many database copies should be evaluated as part of the request. The *DataMoveReplicationConstraint* parameter has the following possible values:

- `None`: When you create a mailbox database, this value is set by default. When this value is set, the Data Guarantee API conditions are ignored. This setting should be used only for mailbox databases that aren't replicated.
- `SecondCopy`: This is the default value when you add the second copy of a mailbox database. When this value is set, at least one passive database copy must meet the Data Guarantee API conditions.
- `SecondDatacenter`: When this value is set, at least one passive database copy in another Active Directory site must meet the Data Guarantee API conditions.
- `AllDatacenters`: When this value is set, at least one passive database copy in each Active Directory site must meet the Data Guarantee API conditions.
- `AllCopies`: When this value is set, all copies of the mailbox database must meet the Data Guarantee API conditions.

Check Replication Health

When the Data Guarantee API is executed to evaluate the health of the database copy infrastructure, several items are evaluated.

In all scenarios, the passive database copy must meet the following conditions:

- Be healthy.
- Have a replay queue within 10 minutes of the replay lag time.
- Have a copy queue length less than 10 logs.
- Have an average copy queue length less than 10 logs. The average copy queue length is computed based on the number of times the application has queried the database status.

IF THE *DATAMOVEREPLICATIONCONSTRAINT* PARAMETER IS SET TO...	THEN, FOR A GIVEN DATABASE...
`SecondCopy`	At least one passive database copy for a replicated database must meet the previously described conditions.
`SecondDatacenter`	At least one passive database copy in another Active Directory site must meet the previously described conditions.
`AllDatacenters`	The active copy must be mounted, and a passive copy in each Active Directory site must meet the previously described conditions.
`AllCopies`	The active copy must be mounted, and all passive database copies must meet the previously described conditions.

Check Replication Flush

The Data Guarantee API can also be used to validate that a prerequisite number of database copies have replayed the required transaction logs. This is verified by comparing the last log replayed timestamp with that of the calling service's commit timestamp (in most cases, this is the timestamp of the last log file that contains required data) plus an additional five seconds (to deal with system time clock skews or drift). If the replay timestamp is greater than the commit timestamp, the *DataMoveReplicationConstraint* parameter is satisfied. If the replay timestamp is less than the commit timestamp, the *DataMoveReplicationConstraint* isn't satisfied.

Before moving large numbers of mailboxes to or from replication databases within a DAG, we recommend that you configure the *DataMoveReplicationConstraint* parameter on each mailbox database according to the following:

IF YOU'RE DEPLOYING...	SET DATAMOVEREPLICATIONCONSTRAINT TO...
Mailbox databases that don't have any database copies	`None`
A DAG within a single Active Directory site	`SecondCopy`
A DAG in multiple datacenters using a stretched Active Directory site	`SecondCopy`
A DAG that spans twoActive Directory sites, and you will have highly available database copies in each site	`SecondDatacenter`
A DAG that spans two Active Directory sites, and you will have only lagged database copies in the second site	`SecondCopy` This is because the Data Guarantee API won't guarantee data being committed until the log file is replayed into the database copy, and due to the nature of the database copy being lagged, this constraint will fail the move request, unless the lagged database copy ReplayLagTime value is less than 30 minutes.
A DAG that spans three or more Active Directory sites, and each site will contain highly available database copies	`AllDatacenters`

Balancing database copies

Due to the inherent nature of DAGs, as the result of database switchovers and failovers, active mailbox database copies will change hosts several times throughout a DAG's lifetime. As a result, DAGs can become unbalanced in terms of active mailbox database copy distribution. The following table shows an example of a DAG that has four databases with four copies of each database (for a total of 16 databases on each server) with an uneven distribution of active database copies.

DAG with unbalanced active copy distribution

SERVER	NUMBER OF ACTIVE DATABASES	NUMBER OF PASSIVE DATABASES	NUMBER OF MOUNTED DATABASES	NUMBER OF DISMOUNTED DATABASES	PREFERENCE COUNT LIST
EX1	5	11	5	0	4, 4, 3, 5
EX2	1	15	1	0	1, 8, 6, 1
EX3	12	4	12	0	13, 2, 1, 0
EX4	1	15	1	0	1, 1, 5, 9

In the preceding example, there are four copies of each database, and therefore, only four possible values for activation preference (1, 2, 3, or 4). The **Preference count list** column shows the count of the number of databases with each of these values. For example, on EX3, there are 13 database copies with an activation preference of 1, two copies with an activation preference of 2, one copy with an activation preference of 3, and no copies with an activation preference of 4.

As you can see, this DAG isn't balanced in terms of the number of active databases hosted by each DAG member, the number of passive databases hosted by each DAG member, or the activation preference count of the hosted databases.

You can use the RedistributeActiveDatabases.ps1 script to balance the active mailbox databases copies across a DAG. This script moves databases between their copies in an attempt to have an equal number of mounted databases on each server in DAG. If required, the script also attempts to balance active databases across sites.

The script provides two options for balancing active database copies within a DAG:

- **BalanceDbsByActivationPreference**: When this option is specified, the script attempts to move databases to their most preferred copy (based on activation preference) without regard to the Active Directory site.
- **BalanceDbsBySiteAndActivationPreference**: When this option is specified, the script attempts to move active databases to their most preferred copy, while also trying to balance active databases within each Active Directory site.

After running the script with the first option, the preceding unbalanced DAG becomes balanced, as shown in the following table.

DAG with balanced active copy distribution

SERVER	NUMBER OF ACTIVE DATABASES	NUMBER OF PASSIVE DATABASES	NUMBER OF MOUNTED DATABASES	NUMBER OF DISMOUNTED DATABASES	PREFERENCE COUNT LIST
EX1	4	12	4	0	4, 4, 4, 4
EX2	4	12	4	0	4, 4, 4, 4
EX3	4	12	4	0	4, 4, 4, 4
EX4	4	12	4	0	4, 4, 4, 4

As shown in the preceding table, this DAG is now balanced in terms of number of active and passive databases on each server and activation preference across the servers.

The following table lists the available parameters for the RedistributeActiveDatabases.ps1 script.

RedistributeActiveDatabases.ps1 script parameters

PARAMETER	DESCRIPTION
DagName	Specifies the name of the DAG you want to rebalance. If this parameter is omitted, the DAG of which the local server is a member is used.
BalanceDbsByActivationPreference	Specifies that the script should move databases to their most preferred copy without regard to the Active Directory site.
BalanceDbsBySiteAndActivationPreference	Specifies that the script should attempt to move active databases to their most preferred copy, while also trying to balance active databases within each Active Directory site.
ShowFinalDatabaseDistribution	Specifies that a report of current database distribution be displayed after redistribution is complete.
AllowedDeviationFromMeanPercentage	Specifies the allowed variation of active databases across sites, expressed as a percentage. The default is 20%. For example, if there were 99 databases distributed between three sites, the ideal distribution would be 33 databases in each site. If the allowed deviation is 20%, the script attempts to balance the databases so that each site has no more than 10% more or less than this number. 10% of 33 is 3.3, which is rounded up to 4. Therefore, the script attempts to have between 29 and 37 databases in each site.
ShowDatabaseCurrentActives	Specifies that the script produce a report for each database detailing how the database was moved and whether it's now active on its most-preferred copy.
ShowDatabaseDistributionByServer	Specifies that the script produce a report for each server showing its database distribution.
RunOnlyOnPAM	Specifies that the script run only on the DAG member that currently has the PAM role. The script verifies it's being run from the PAM. If it isn't being run from the PAM, the script exits.
LogEvents	Specifies that the script logs an event (MsExchangeRepl event 4115) containing a summary of the actions.
IncludeNonReplicatedDatabases	Specifies that the script should include non-replicated databases (databases without copies) when determining how to redistribute the active databases. Although non-replicated databases can't be moved, they may affect the distribution of the replicated databases.
Confirm	The Confirm switch can be used to suppress the confirmation prompt that appears by default when this script is run. To suppress the confirmation prompt, use the syntax -Confirm:$False. You must include a colon (:) in the syntax.

RedistributeActiveDatabases.ps1 examples

This example shows the current database distribution for a DAG, including preference count list.

```
RedistributeActiveDatabases.ps1 -DagName DAG1 -ShowDatabaseDistributionByServer | Format-Table
```

This example redistributes and balances the active mailbox database copies in a DAG using activation preference without prompting for input.

```
RedistributeActiveDatabases.ps1 -DagName DAG1 -BalanceDbsByActivationPreference -Confirm:$False
```

This example redistributes and balances the active mailbox database copies in a DAG using activation preference, and produces a summary of the distribution.

```
RedistributeActiveDatabases.ps1 -DagName DAG1 -BalanceDbsByActivationPreference -
ShowFinalDatabaseDistribution
```

Monitoring database copies

You can view a variety of information, including copy queue length, replay queue length, status, and content index state information, by examining the details of a database copy in the EAC. You can also use the **Get-MailboxDatabaseCopyStatus** cmdlet in the Exchange Management Shell to view a variety of status information for a database copy.

> **NOTE**
>
> A database copy is your first defense if a failure occurs that affects the active copy of a database. It's therefore critical to monitor the health and status of database copies to ensure that they are available when needed.

For more information about monitoring database copies, see Monitor database availability groups.

Removing a database copy

A database copy can be removed at any time by using the EAC or by using the **Remove-MailboxDatabaseCopy** cmdlet in the Exchange Management Shell. After removing a database copy, you must manually delete any database and transaction log files from the server from which the database copy is being removed. For detailed steps about how to remove a database copy, see Remove a mailbox database copy.

Database switchovers

The Mailbox server that hosts the active copy of a database is referred to as the mailbox database master. The process of activating a passive database copy changes the mailbox database master for the database and turns the passive copy into the new active copy. This process is called a database switchover. In a database switchover, the active copy of a database is dismounted on one Mailbox server and a passive copy of that database is mounted as the new active mailbox database on another Mailbox server. When performing a switchover, you can optionally override the database mount dial setting on the new mailbox database master.

You can quickly identify which Mailbox server is the current mailbox database master by reviewing the right-hand column under the **Database Copies** tab in the EAC. You can perform a switchover by using the **Activate** link in the EAC, or by using the **Move-ActiveMailboxDatabase** cmdlet in the Exchange Management Shell.

There are several internal checks that will be performed before a passive copy is activated. In some cases, the database switchover is blocked or canceled. In other cases, you can use cmdlets to move or skip over some checks.

- The status of the database copy is checked. If the database copy is in a failed state, the switchover is blocked. You can override this behavior and bypass the health check by using the *SkipHealthChecks* parameter of the **Move-ActiveMailboxDatabase** cmdlet. This parameter lets you move the active copy to a database copy in a failed state.

- The active database copy is checked to see if it's currently a seeding source for any passive copies of the database. If the active copy is currently being used as a source for seeding, the switchover is blocked. You can override this behavior and bypass the seeding source check by using the *SkipActiveCopyChecks* parameter of the **Move-ActiveMailboxDatabase** cmdlet. This parameter allows you to move an active copy that's being used as a seeding source. Using this parameter will cause the seeding operation to be cancelled and considered failed.

- The copy queue and replay queue lengths for the database copy are checked to ensure their values are within the configured criteria. Also, the database copy is verified to ensure that it isn't currently in use as a source for seeding. If the values for the queue lengths are outside the configured criteria, or if the database is currently used as a source for seeding, the switchover is blocked. You can override this behavior and bypass these checks by using the *SkipLagChecks* parameter of the **Move-ActiveMailboxDatabase** cmdlet. This parameter allows a copy to be activated that has replay and copy queues outside of the configured criteria.

- The state of the search catalog (content index) for the database copy is checked. If the search catalog isn't up to date, is in an unhealthy state, or is corrupt, the switchover is blocked. You can override this behavior and bypass the search catalog check by using the *SkipClientExperienceChecks* parameter of the **Move-ActiveMailboxDatabase** cmdlet. This parameter causes this search to skip the catalog health check. If the search catalog for the database copy you're activating is in an unhealthy or unusable state and you use this parameter to skip the catalog health check and activate the database copy, you will need to either crawl or seed the search catalog again.

When you perform a database switchover, you also have the option of overriding the mount dial settings configured for the server that hosts the passive database copy being activated. Using the *MountDialOverride* parameter of the **Move-ActiveMailboxDatabase** cmdlet instructs the target server to override its own mount dial settings and use those specified by the *MountDialOverride* parameter.

For detailed steps about how to perform a switchover of a database copy, see Activate a mailbox database copy.

Add a mailbox database copy

8/3/2020 • 3 minutes to read • Edit Online

When you add a copy of a mailbox database, continuous replication is automatically enabled between the existing database and the database copy. Database copies are automatically assigned an identity in the format of < *DatabaseName*>\< *HostMailboxServerName*>. For example, a copy of the database DB1 that's hosted on the server MBX3 would be DB1\MBX3.

Looking for other management tasks related to mailbox database copies? Check out Manage mailbox database copies.

What do you need to know before you begin?

- Estimated time to complete this task: 2 minutes, plus the time to seed the database copy, which depends on a variety of factors, such as the size of the database, the speed, available bandwidth and latency of the network, and storage speeds.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Mailbox database copies" entry in the High availability and site resilience permissions topic.
- The active copy of the database must be mounted.
- The specified Mailbox server must not already host a copy of the database.
- The path for the database copy and its log files must be available on the selected Mailbox server.
- The server hosting the active copy and the server that will host the passive copy must be in the same database availability group (DAG). The DAG must also have quorum and be healthy.
- If you're adding the second copy of a database (for example, creating the first passive copy of the database), circular logging must not be enabled for the specified mailbox database. If circular logging is enabled, you must first disable it. After the mailbox database copy has been added, circular logging can be enabled. After circular logging is enabled for a replicated mailbox database, continuous replication circular logging (CRCL) is used instead of JET circular logging. If you're adding the third or subsequent copy of a database, CRCL can remain enabled.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to add a mailbox database copy

1. In the EAC, go to **Servers** > **Databases**.

2. Select the database that you want to copy, click **More** (the three dots to the right of the Refresh icon), and then click **Add database copy**.

3. On the **add mailbox database copy** page, click **Browse...**, select the Mailbox server that will host the

database copy, and then click **OK**.

4. Optionally, configure the **Activation preference number** for the database copy.

5. Click **More options...** to designate the database copy as a lagged database copy by configuring a replay lag time, or to postpone automatic seeding of the database copy.

6. Click **Save** to save the configuration changes and add the mailbox database copy.

7. Click **OK** to acknowledge any messages that appear.

Use the Exchange Management Shell to add a mailbox database copy

This example adds a copy of mailbox database DB1 to the Mailbox server MBX3. Replay lag time and truncation lag time are left at the default values of zero, and the activation preference is configured with a value of 2.

```
Add-MailboxDatabaseCopy -Identity DB1 -MailboxServer MBX3 -ActivationPreference 2
```

This example adds a copy of mailbox database DB2 to the Mailbox server MBX4. Replay lag time and truncation lag time are left at the default values of zero, and the activation preference is configured with a value of `5`. In addition, seeding is being postponed for this copy so that it can be seeded using a local source server instead of the current active database copy, which is geographically distant from MBX4.

```
Add-MailboxDatabaseCopy -Identity DB2 -MailboxServer MBX4 -ActivationPreference 5 -SeedingPostponed
```

This example adds a copy of mailbox database DB3 to the Mailbox server MBX5. Replay lag time is set to 3 days, truncation lag time is left at the default value of zero, and the activation preference is configured with a value of `4`.

```
Add-MailboxDatabaseCopy -Identity DB3 -MailboxServer MBX5 -ReplayLagTime 3.00:00:00 -ActivationPreference 4
```

How do you know this worked?

To verify that you have successfully created a mailbox database copy, do one of the following:

- In the EAC, navigate to **Servers** > **Databases**. Select the database that was copied. In the Details pane, the status of the database copy and its content index are displayed, along with the current copy queue length.

- In the Exchange Management Shell, run the following command to verify the mailbox database copy was created and is healthy.

  ```
  Get-MailboxDatabaseCopyStatus <DatabaseCopyName>
  ```

  The Status and Content Index State should both be Healthy.

For more information

Mailbox database copies

Manage mailbox database copies

Configure mailbox database copy properties

8/3/2020 • 4 minutes to read • Edit Online

Each mailbox database copy has its own properties, which you can configure. These properties include the amount of time, if any, for replay lag and truncation lag, and the activation preference number. For more information about replay lag, truncation lag and the activation preference number, see Manage mailbox database copies.

What do you need to know before you begin?

- Estimated time to complete this task: 1 minute
- To open the EAC, see Exchange admin center in Exchange Server. To open the Exchange Management Shell, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Mailbox database copies" entry in the High availability and site resilience permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to configure mailbox database copy properties

1. In the EAC, go to **Servers** > **Databases**.
2. Select the database you want to configure.
3. In the Details pane, under **Database Copies**, click **View details** for the desired database copy, and then view or configure the following:
   - **Database**: Displays the name of the selected database.
   - **Mailbox server**: Displays the name of the Mailbox server that hosts the selected database copy.
   - **Content index state**: Displays the current state of the content index for the selected database copy.
   - **Status**: Displays the current status of the selected database copy.
   - **Copy queue length**: Indicates the number of log files waiting to be copied to the selected database copy. This field is relevant only for passive database copies.
   - **Replay queue length**: Indicates the number of log files waiting to be replayed into the selected database copy. This field is relevant only for passive database copies.
   - **Error messages**: Displays any error messages for database copies that have a status of `Failed` or `Failed and Suspended`.
   - **Latest available log time**: Displays the date and time stamp of the most recently generated log file on the active copy of the database. This field is relevant only for passive database copies. On active

database copies (replicated and stand-alone), this field will display never.

- **Last inspected log time**: Displays the date and time stamp of the last log file that was inspected by the LogInspector on the selected database copy. This field is relevant only for passive database copies. On active database copies (replicated and stand-alone), this field will display never.

- **Last copied log time**: Displays the date and time stamp of the last log file that was copied by the LogCopier on the selected database copy. This field is relevant only for passive database copies. On active database copies (replicated and stand-alone), this field will display never.

- **Last replayed log time**: Displays the date and time stamp of the last log file that was replayed by the LogReplayer into the selected database copy. This field is relevant only for passive database copies. On active database copies (replicated and stand-alone), this field will display never.

- **Activation preference number**: Displays the activation preference number. This is used as part of Active Manager's best copy selection process, and it is used to balance the DAG by redistributing active mailbox databases throughout the DAG via the DAG's `PreferenceMoveFrequency` property. This property defines the frequency (measured in time) when the Microsoft Exchange Replication service rebalances database copies by performing a lossless switchover that activates the copy with an activation preference number of 1. The value for activation preference is a number equal to or greater than 1, where 1 is at the top of the preference order. The number can't be larger than the number of copies of the mailbox database.

- **Replay lag time (days)**: Displays the amount of time that the Microsoft Exchange Information Store service should wait before replaying log files that have been copied by the Microsoft Exchange Replication service to the passive database copy. Setting this parameter to a value greater than 0 creates a lagged database copy. The default setting for this value is 0 days. The maximum allowable value for this setting is 14 days. The minimum allowable value is 0 days, and setting this value to 0 disables replay lag.

Use the Exchange Management Shell to configure mailbox database copy properties

This example configures a mailbox database copy with an activation preference number of 3.

```
Set-MailboxDatabaseCopy -Identity DB3\EX3 -ActivationPreference 3
```

This example configures a copy of the database DB1 that's hosted on Server1 with a replay lag time and truncation lag time of 1 day, and an activation preference number of 2.

```
Set-MailboxDatabaseCopy -Identity DB1\Server1 -ReplayLagTime 1.0:0:0 -TruncationLagTime 1.0:0:0 -
ActivationPreference 2
```

How do you know this worked?

To verify that you've successfully configured a mailbox database copy, do one of the following:

- In the EAC, navigate to **Servers** > **Databases**. Select the appropriate database, and in the Details pane, click **View details** to view the database copy properties.

- In the Exchange Management Shell, run the following command to display configuration information for a database copy.

```
Get-MailboxDatabaseCopyStatus <DatabaseCopyName> | Format-List
```

For more information

Set-MailboxDatabaseCopy

Get-MailboxDatabaseCopyStatus

Get-MailboxDatabase

Move the mailbox database path for a mailbox database copy

8/3/2020 • 4 minutes to read • Edit Online

If the mailbox database being moved is replicated to one or more mailbox database copies, you must follow the procedure in this topic to move the mailbox database path. All copies of a mailbox database must be located in the same path on each server that hosts a copy. For example, if database DB1 is located at C:\mountpoints\DB1 on server EX1, copies of DB1 on servers EX2, EX3, and so on, must also be located at C:\mountpoints\DB1.

> **NOTE**
>
> After you create a new mailbox database, you can move it to another volume, folder, location, or path by using either the EAC or the Exchange Management Shell. For step-by-step instructions about how to move the database path for a **non-replicated** mailbox database, see Manage mailbox databases in Exchange Server.

Looking for other management tasks related to mailbox database copies? Check out Managing mailbox database copies.

What do you need to know before you begin?

- Estimated time to complete this task: 2 minutes, plus the time to move the data, which depends on a variety of factors, such as the size of the database, the speed, available bandwidth and latency of the network, and storage speeds.

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Mailbox database copies" entry in the High availability and site resilience permissions topic.

- To perform the move operation, the database must be temporarily dismounted, making it inaccessible to all users. If the database is currently dismounted, it isn't remounted upon completion.

- To perform the move operation, replication for the database must be disabled for all copies. It's not enough to suspend replication; you must disable it by using the Remove-MailboxDatabaseCopy cmdlet to remove the database copies.

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to move a replicated mailbox database to a new path

NOTE

You can't use the Exchange admin center (EAC) to move a replicated mailbox database to a new path.

1. Note any replay lag or truncation lag settings for all copies of the mailbox database being moved. You can obtain this information by using the Get-MailboxDatabase cmdlet, as shown in this example.

   ```
   Get-MailboxDatabase DB1 | Format-List *lag*
   ```

2. If circular logging is enabled for the database, it must be disabled before proceeding. You can disable circular logging for a mailbox database by using the Set-MailboxDatabase cmdlet, as shown in this example.

   ```
   Set-MailboxDatabase DB1 -CircularLoggingEnabled $false
   ```

3. Remove all mailbox database copies for the database being moved. For detailed steps, see Remove a mailbox database copy. After all copies are removed, preserve the database and transaction log files from each server from which the database copy is being removed by moving them to another location. These files are being preserved so the database copies don't require re-seeding after they have been re-added.

4. Move the mailbox database path to the new location. For detailed steps, see Move a mailbox database path.

   IMPORTANT

   During the move operation, the database being moved must be dismounted. Until the move is complete, this process will cause an interruption in service and an outage for all users with mailboxes on the database being moved. After the move operation completes, the database is automatically mounted.

5. Create the necessary folder structure on each Mailbox server that previously contained a passive copy of the moved mailbox database. For example, if you moved the database to C:\mountpoints\DB1, you must create this same path on each Mailbox server that will host a mailbox database copy.

6. After creating the folder structure, move the passive copy of the mailbox database and its log stream to the new location. These are the files that were left from and preserved after Step 3. Repeat this process for each database copy that was removed in Step 3.

7. Add all of the database copies that were removed in Step 3. For detailed steps, see Add a mailbox database copy.

8. On each server that contains a copy of the mailbox database being moved, run the following command to stop and restart the content index services.

   ```
   Restart-Service MSExchangeFastSearch
   ```

9. Optionally, enable circular logging by using the Set-MailboxDatabase cmdlet, as shown in this example.

   ```
   Set-MailboxDatabase DB1 -CircularLoggingEnabled $true
   ```

10. Reconfigure any previously set values for replay lag time and truncation lag time by using the Set-MailboxDatabaseCopy cmdlet, as shown in this example.

```
Set-MailboxDatabaseCopy DB1\MBX2 -ReplayLagTime 00:15:00
```

11. As each copy is added, we recommend that you verify the health and status of the copy prior to adding the next copy. You can verify the health and status by:

    a. Examining the event log for any error or warning events related to the database or the database copy.

    b. Using the Get-MailboxDatabaseCopyStatus cmdlet to check the health and status of continuous replication for the database copy.

    c. Using the Test-ReplicationHealth cmdlet to verify the health and status of the database availability group and continuous replication.

For detailed syntax and parameter information, see the following topics:

- Get-MailboxDatabase
- Set-MailboxDatabase
- Set-MailboxDatabaseCopy
- Get-MailboxDatabaseCopyStatus
- Test-ReplicationHealth

How do you know this worked?

To verify that you've successfully moved the path for a mailbox database copy, do one of the following:

- In the EAC, navigate to **Servers** > **Databases**. Select the database that was copied. In the Details pane, the status of the database copy and its content index are displayed, along with the current copy queue length. Verify that the status is Healthy.

- In the Exchange Management Shell, run the following command to verify the mailbox database copy was created and is healthy.

  ```
  Get-MailboxDatabaseCopyStatus <DatabaseCopyName>
  ```

  The Status and Content Index State should both be Healthy.

Configure activation policy for a mailbox database copy

8/3/2020 • 2 minutes to read • Edit Online

Activation is the process of changing a mailbox database copy from a passive copy to an active copy. Activation can occur automatically (by the system as part of a database or server failover operation) or it can be performed manually (by an administrator as part of a database or server switchover operation). Blocking a database for activation prevents it from becoming the active copy during a database or server failover.

Looking for other management tasks related to mailbox database copies? Check out Manage mailbox database copies.

What do you need to know before you begin?

- Estimated time to complete this task: 1 minute
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Mailbox database copies" entry in the High availability and site resilience permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to configure the activation policy for a mailbox database copy

1. In the EAC, go to **Servers** > **Databases**.
2. Select the database that you want to configure.
3. In the Details pane, under **Database Copies**, locate the database copy you want to configure and click **Suspend**.
4. Optionally, add a comment, and select the check box that says **This copy can only be activated by manual intervention**.
5. Click **Save** to save the configuration changes for the mailbox database copy.

Use the Exchange Management Shell to suspend or resume a database copy for activation

This example blocks the copy of the database DB1 on the server MBX2 for activation.

```
Suspend-MailboxDatabaseCopy -Identity DB1\MBX2 -ActivationOnly
```

This example resumes the copy of the database DB1 on the server MBX2 for activation.

```
Resume-MailboxDatabaseCopy -Identity DB1\MBX2
```

For detailed syntax and parameter information, see Suspend-MailboxDatabaseCopy or Resume-MailboxDatabaseCopy.

Use the Exchange Management Shell to configure the activation policy for a server

This example configures the database copies on server MBX2 as blocked for activation.

```
Set-MailboxServer -Identity MBX2 -DatabaseCopyAutoActivationPolicy Blocked
```

This example configures the database copies on server MBX3 as blocked for out-of-site activation.

```
Set-MailboxServer -Identity MBX3 -DatabaseCopyAutoActivationPolicy IntrasiteOnly
```

This example configures the database copies on server MBX4 as unblocked for activation.

```
Set-MailboxServer -Identity MBX4 -DatabaseCopyAutoActivationPolicy Unrestricted
```

For detailed syntax and parameter information, see Suspend-MailboxDatabaseCopy, Resume-MailboxDatabaseCopy, or Set-MailboxServer.

How do you know this worked?

To verify that you've successfully configured the activation policy, do one of the following:

- In the Exchange Management Shell, run the following command to verify activation settings for a database copy.

  ```
  Get-MailboxDatabaseCopyStatus <DatabaseCopyName> | Format-List ActivationSuspended
  ```

- In the Exchange Management Shell, run the following command to verify activation settings for a DAG member.

  ```
  Get-MailboxServer <ServerName> | Format-List DatabaseCopyAutoActivationPolicy
  ```

Update a mailbox database copy

8/3/2020 • 6 minutes to read • Edit Online

Updating, also known as *seeding*, is the process in which a copy of a mailbox database is added to another Mailbox server in a database availability group (DAG). The newly added copy becomes the baseline database for the passive copy into which log files copied from the active copy are replayed. Seeding is required under the following conditions:

- When a new passive copy of a database is created. Seeding can be postponed for a new mailbox database copy, but eventually each passive database copy must be seeded to function as a redundant database copy.
- After a failover occurs in which data is lost as a result of the passive database copy having become diverged and unrecoverable.
- When the system has detected a corrupted log file that can't be replayed into the passive copy of the database.
- After an offline defragmentation of any copy of the database occurs.
- After the log generation sequence for the database has been reset back to 1.

You can perform seeding by using the following methods:

- **Automatic seeding**: An automatic seed produces a passive copy of the active database on the target Mailbox server. Automatic seeding occurs during the creation of a database.
- **Seeding using the Update-MailboxDatabaseCopy cmdlet**: You can use the Update-MailboxDatabaseCopy cmdlet in the Exchange Management Shell to seed a database copy at any time.
- **Seeding using the Update Mailbox Database Copy wizard**: You can use the Update Mailbox Database Copy wizard in the EAC to seed a database copy at any time.
- **Manually copying the offline database**: You can dismount the active copy of the database and copy the database file to the same location on another Mailbox server in the same DAG. If you use this method, there will be an interruption in service because the process requires you to dismount the database.

Updating a database copy can take a long time, especially if the database being copied is large, or if there is high network latency or low network bandwidth. After the seeding process has started, don't close the EAC or the Exchange Management Shell until the process has completed. If you do, the seeding operation will be terminated.

A database copy can be seeded using either the active copy or an up-to-date passive copy as the source for the seed. When seeding from a passive copy, be aware that the seed operation will terminate with a network communication error under the following circumstances:

- If the status of the seeding source copy changes to Failed or FailedAndSuspended.
- If the database fails over to another copy.

Multiple database copies can be seeded simultaneously. However, when seeding multiple copies simultaneously, you must seed only the database file, and omit the content index catalog. You can do this by using the *DatabaseOnly* parameter with the Update-MailboxDatabaseCopy cmdlet.

NOTE

If you don't use the *DatabaseOnly* parameter when seeding multiple targets from the same source, the task will fail with `SeedInProgressException` error `FE1C6491`.

Looking for other management tasks related to mailbox database copies? Check out Manage mailbox database copies.

What do you need to know before you begin?

- Estimated time to complete this task: 2 minutes, plus the time to seed the database copy, which depends on a variety of factors, such as the size of the database, the speed, available bandwidth and latency of the network, and storage speeds.

- To open the EAC, see Exchange admin center in Exchange Server. To open the Exchange Management Shell, see Open the Exchange Management Shell.

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Mailbox database copies" entry in the High availability and site resilience permissions topic.

- The mailbox database copy must be suspended. For detailed steps, see Suspend or resume a mailbox database copy.

- The Remote Registry service must be running on the server hosting the passive database copy you're updating.

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Update a mailbox database copy

Use the EAC to update a mailbox database copy

1. In the EAC, go to **Servers** > **Databases**.

2. Select the mailbox database whose passive copy you want to update.

3. In the Details pane, under **Database Copies**, click **Suspend** under the passive database copy you want to seed. Provide any optional comments, and click **save**.

4. In the Details pane, under **Database Copies**, click **Update** under the passive database copy you want to seed.

5. By default, the active copy of the database is used as the source database for seeding. If you prefer to use a passive copy of the database for seeding, click **browse...** to select the server containing the passive database copy you want to use for the source.

6. Click **save** to update the passive database copy.

Use the Exchange Management Shell to update a mailbox database copy

This example shows how to seed a copy of the database DB1 on MBX1.

```
Update-MailboxDatabaseCopy -Identity DB1\MBX1
```

This example shows how to seed a copy of the database DB1 on MBX1 using MBX2 as the source Mailbox server for the seed.

```
Update-MailboxDatabaseCopy -Identity DB1\MBX1 -SourceServer MBX2
```

This example shows how to seed a copy of the database DB1 on MBX1 without seeding the content index catalog.

```
Update-MailboxDatabaseCopy -Identity DB1\MBX1 -DatabaseOnly
```

This example shows how to seed the content index catalog for the copy of the database DB1 on MBX1 without seeding the database file.

```
Update-MailboxDatabaseCopy -Identity DB1\MBX1 -CatalogOnly
```

Manually copy an offline database

1. If circular logging is enabled for the database, it must be disabled before proceeding. You can disable circular logging for a mailbox database by using the Set-MailboxDatabase cmdlet, as shown in this example.

   ```
   Set-MailboxDatabase DB1 -CircularLoggingEnabled $false
   ```

2. Dismount the database. You can use the Dismount-Database cmdlet, as shown in this example.

   ```
   Dismount-Database DB1 -Confirm $false
   ```

3. Manually copy the database files (the database file and all log files) to a second location, such as an external disk drive or a network share.

4. Mount the database. You can use the Mount-Database cmdlet, as shown in this example.

   ```
   Mount-Database DB1
   ```

5. On the server that will host the copy, copy the database files from the external drive or network share to the same path as the active database copy. For example, if the active copy database path is D:\DB1\DB1.edb and log file path is D:\DB1, you would copy the database files to D:\DB1 on the server that will host the copy.

6. Add the mailbox database copy by using the Add-MailboxDatabaseCopy cmdlet with the *SeedingPostponed* parameter, as shown in this example.

   ```
   Add-MailboxDatabaseCopy -Identity DB1 -MailboxServer MBX3 -SeedingPostponed
   ```

7. If circular logging is enabled for the database, enable it again by using the Set-MailboxDatabase cmdlet, as shown in this example.

   ```
   Set-MailboxDatabase DB1 -CircularLoggingEnabled $true
   ```

How do you know this worked?

To verify that you've successfully seeded a mailbox database copy, do one of the following:

- In the EAC, navigate to **Servers** > **Databases**. Select the database that was seeded. In the Details pane, the status of the database copy and its content index are displayed, along with the current copy queue length.

- In the Exchange Management Shell, run the following command to verify the mailbox database copy was seeded successfully and is healthy.

  ```
  Get-MailboxDatabaseCopyStatus <DatabaseCopyName>
  ```

  The Status and Content Index State should both be Healthy.

Suspend or resume a mailbox database copy

8/3/2020 • 2 minutes to read • Edit Online

You may need to suspend or resume a database copy for a variety of reasons, such as maintenance on the disk that contains the database copy. Or you may need to suspend an individual database copy from activation for disaster recovery purposes.

Looking for other management tasks related to mailbox database copies? Check out Manage mailbox database copies.

What do you need to know before you begin?

- Estimated time to complete this task: 1 minute
- To open the EAC, see Exchange admin center in Exchange Server. To open the Exchange Management Shell, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Mailbox database copies" entry in the High availability and site resilience permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Suspend a mailbox database copy

Use the EAC to suspend a mailbox database copy

1. In the EAC, go to **Servers** > **Databases**.
2. Select the database whose copy you want to suspend.
3. In the Details pane, under **Database Copies**, click **Suspend** under the database copy you want to suspend.
4. In the **Comments** field, add an optional comment of up to 512 characters specifying the reason for the suspension.
5. To suspend the database copy from automatic activation, select the **This copy can only be activated by manual intervention** check box.
6. Click **save** to suspend the database copy.

Use the Exchange Management Shell to suspend a mailbox database copy

This example suspends continuous replication for a copy of the database DB1 hosted on the server MBX1. An optional comment has also been specified.

```
Suspend-MailboxDatabaseCopy -Identity DB1\MBX1 -SuspendComment "Maintenance on MBX1" -Confirm:$False
```

This example suspends activation for a copy of the database DB2 hosted on the server MBX2.

```
Suspend-MailboxDatabaseCopy -Identity DB2\MBX2 -ActivationOnly -Confirm:$False
```

For detailed syntax and parameter information, see Suspend-MailboxDatabaseCopy.

Resume a mailbox database copy

Use the EAC to resume a mailbox database copy

1. In the EAC, go to **Servers** > **Databases**.
2. Select the database whose copy you want to resume.
3. In the Details pane, under **Database Copies**, click **Resume** under the database copy you want to resume.
4. Click **yes** to resume the database copy.

Use the Exchange Management Shell to resume a mailbox database copy

This example resumes a copy of the database DB1 on the server MBX1.

```
Resume-MailboxDatabaseCopy -Identity DB1\MBX1
```

This example resumes a copy of the database DB2 on the server MBX2 for replication only.

```
Resume-MailboxDatabaseCopy -Identity DB2\MBX2 -ReplicationOnly
```

For detailed syntax and parameter information, see Resume-MailboxDatabaseCopy.

How do you know this worked?

To verify that you have successfully suspended or resumed a mailbox database copy, do one of the following:

- In the EAC, navigate to **Servers** > **Databases**. Select the appropriate database, and in the Details pane, click **View details** to view the database copy properties:
- In the Exchange Management Shell, run the following command to display status information for a database copy:

  ```
  Get-MailboxDatabaseCopyStatus <DatabaseCopyName> | Format-List
  ```

Activate a mailbox database copy

8/3/2020 • 3 minutes to read • Edit Online

Activating a mailbox database copy is the process of designating a specific passive copy as the new active copy of a mailbox database. This process is referred to as a *database switchover*. A database switchover involves dismounting the current active database and mounting the database copy on the specified server as the new active mailbox database copy. The database copy that will become the active mailbox database must be healthy and current.

Looking for other management tasks related to mailbox database copies? Check out Manage mailbox database copies.

What do you need to know before you begin?

- Estimated time to complete this task: 1 minute
- To open the EAC, see Exchange admin center in Exchange Server. To open the Exchange Management Shell, see Open the Exchange Management Shell.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Mailbox database copies" entry in the High availability and site resilience permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to move the active mailbox database

1. In the EAC, go to **Servers** > **Databases**.
2. Select the database whose copy you want to activate.
3. In the Details pane, under **Database Copies**, click **Activate** under the database copy you want to activate.
4. Click **yes** to activate the database copy.

Use the Exchange Management Shell to move the active mailbox database

This example activates and mounts a copy of the database DB4 hosted on MBX3 as the new active mailbox database. This command makes DB4 the new active mailbox database, and it doesn't override the database mount dial settings on MBX3.

```
Move-ActiveMailboxDatabase DB4 -ActivateOnServer MBX3 -MountDialOverride:None
```

This example performs a switchover of the database DB2 to the Mailbox server MBX1. When the command completes, MBX1 hosts the active copy of DB2. Because the *MountDialOverride* parameter is set to `None`, MBX1 mounts the database using its own defined database auto mount dial settings.

```
Move-ActiveMailboxDatabase DB2 -ActivateOnServer MBX1 -MountDialOverride:None
```

This example performs a switchover of the database DB1 to the Mailbox server MBX3. When the command completes, MBX3 hosts the active copy of DB1. Because the *MountDialOverride* parameter is specified with a value of `Good Availability`, MBX3 mounts the database using a database auto mount dial setting of *GoodAvailability*.

```
Move-ActiveMailboxDatabase DB1 -ActivateOnServer MBX3 -MountDialOverride:GoodAvailability
```

This example performs a switchover of the database DB3 to the Mailbox server MBX4. When the command completes, MBX4 hosts the active copy of DB3. Because the *MountDialOverride* parameter isn't specified, MBX4 mounts the database using a database auto mount dial setting of *Lossless*.

```
Move-ActiveMailboxDatabase DB3 -ActivateOnServer MBX4
```

This example performs a server switchover for the Mailbox server MBX1. All active mailbox database copies on MBX1 will be activated on one or more other Mailbox servers with healthy copies of the active databases on MBX1.

```
Move-ActiveMailboxDatabase -Server MBX1
```

This example performs a switchover of the database DB4 to the Mailbox server MBX5. In this example, the database copy on MBX5 has a replay queue greater than 6. As a result, the *SkipLagChecks* parameter must be specified to activate the database copy on MBX5.

```
Move-ActiveMailboxDatabase DB4 MBX5 -SkipLagChecks
```

This example performs a switchover of the database DB5 to the Mailbox server MBX6. In this example, the database copy on MBX6 has a *ContentIndexState* of Failed. As a result, the *SkipClientExperienceChecks* parameter must be specified to activate the database copy on MBX6.

```
Move-ActiveMailboxDatabase DB5 MBX6 -SkipClientExperienceChecks
```

How do you know this worked?

To verify that you've successfully activated a mailbox database copy, do one of the following:

- In the EAC, navigate to **Servers** > **Databases**. Select the appropriate database, and in the Details pane, click **View details** to view the database copy properties.

- In the Exchange Management Shell, run the following command to display status information for a database copy.

  ```
  Get-MailboxDatabaseCopyStatus <DatabaseCopyName> | Format-List
  ```

For more information

Mailbox database copies

Configure mailbox database copy properties

Activate a lagged mailbox database copy

8/3/2020 • 5 minutes to read • Edit Online

A lagged mailbox database copy is a mailbox database copy configured with a replay lag time value greater than 0. If you want the database to replay all log files and make the database copy current, activating and recovering a lagged mailbox database copy is a simple process. However, if you want to replay log files up to a specific point in time, it's a more difficult operation because you have to manually manipulate log files and run Eseutil.

Looking for other information related to lagged mailbox database copies? Check out Manage mailbox database copies

> **NOTE**
> The amount of time it takes to activate a lagged mailbox database copy directly depends on how many log files need to be replayed and how fast the hardware can replay them. At a minimum, you should experience a log replay rate of at least two logs per second per database.

What do you need to know before you begin?

- Estimated time to complete this task: 1 minute, plus the time it takes to duplicate the lagged copy, replay the necessary log files, and extract the data or mount the database for client activity.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Mailbox database copies" entry in the High availability and site resilience permissions topic.
- The mailbox database copy being activated must be configured with a replay lag time greater than 0.
- The mailbox database copy being activated must have all log files to the point in time to which you want to recover it. Keep in mind that database transactions can span multiple log files when determining the point in time to which you want to recover.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to activate a lagged mailbox database copy to a specific point in time

> **NOTE**
> You can't use the EAC to activate a lagged mailbox database copy to a specific point in time. Instead, you perform a series of steps using the Exchange Management Shell and the command line.

1. This example suspends replication for the lagged copy being activated by using the Suspend-

MailboxDatabaseCopy cmdlet.

```
Suspend-MailboxDatabaseCopy DB1\EX3 -SuspendComment "Activate lagged copy of DB1 on Server EX3" -
Confirm:$false
```

2. Optionally (to preserve a lagged copy), make a copy of the database copy and its log files.

   > **NOTE**
   >
   > At this point, continuing to perform this procedure on the existing volume would incur a copy on write performance penalty. As an alternative, you can copy the database and log files to another volume to perform the recovery.

3. Determine which log files are required to replay into the database to meet your point-in-time requirement for this recovery (based on log file date and time, as shown in Windows Explorer). All logs created after this point should be moved to a different directory, until the recovery process is completed, and the logs are no longer needed.

4. Delete the checkpoint (.chk) file for the database.

5. This example uses Eseutil to perform the recovery operation.

   ```
   Eseutil.exe /r eXX /a
   ```

   > **NOTE**
   >
   > • In the preceding example, e *XX* is the log generation prefix for the database (for example, E00, E01, E02, and so on).
   >
   > • This step may take a considerable amount of time, depending on several factors, such as the length of the replay lag time, the number of log files generated during that period, and the speed at which your hardware can replay those logs into the database being recovered.

6. After log replay is finished, the database is in a clean shutdown state and can be copied and used for recovery purposes.

7. After the recovery process is complete, this example resumes replication for the database that was used as part of the recovery process.

   ```
   Resume-MailboxDatabaseCopy DB1\EX3
   ```

For detailed syntax and parameter information, see Suspend-MailboxDatabaseCopy or Resume-MailboxDatabaseCopy.

Use the Exchange Management Shell to activate a lagged mailbox database copy by replaying all uncommitted log files

1. Optionally (to preserve a lagged copy), make a copy of the database copy and its log files.

2. This example suspends replication for the lagged copy being activated by using the Suspend-MailboxDatabaseCopy cmdlet.

```
Suspend-MailboxDatabaseCopy DB1\EX3 -SuspendComment "Activate lagged copy of DB1 on Server EX3" -
Confirm:$false
```

3. Optionally (to preserve a lagged copy), make a copy of the database copy and its log files.

> **NOTE**
>
> At this point, continuing to perform this procedure on the existing volume would incur a copy on write performance penalty. If this isn't desirable, you can copy the database and log files to another volume to perform the recovery.

4. This example activates the lagged mailbox database copy using the Move-ActiveMailboxDatabase cmdlet with the *SkipLagChecks* parameter.

```
Move-ActiveMailboxDatabase DB1 -ActivateOnServer EX3 -SkipLagChecks
```

Use the Exchange Management Shell to activate a lagged mailbox database copy by using SafetyNet recovery

1. Optionally (to preserve a lagged copy), take a file system-based (non-Exchange aware) Volume Shadow Copy Service (VSS) snapshot of the volumes containing the database copy and its log files.

2. This example suspends replication for the lagged copy being activated by using the Suspend-MailboxDatabaseCopy cmdlet.

```
Suspend-MailboxDatabaseCopy DB1\EX3 -SuspendComment "Activate lagged copy of DB1 on Server EX3" -
Confirm:$false
```

3. Optionally (to preserve a lagged copy), make a copy of the database copy and its log files.

> **NOTE**
>
> At this point, continuing to perform this procedure on the existing volume would incur a copy-on-write performance penalty. If this isn't desirable, you can copy the database and log files to another volume to perform the recovery.

4. Determine the required logs for the lagged database copy by looking for the "Log Required:" value in ESEUTIL database header output

```
Eseutil /mh <DBPath> | findstr /c:"Log Required"
```

Take note of the hexadecimal numbers in parentheses. The first number is the lowest generation required (referred to as LowGeneration), and the second number is the highest generation required (referred to as HighGeneration). Move all log generation files that have a generation sequence greater than HighGeneration to a different location so that they are not replayed into the database.

5. On the server hosting the active copy of database, either delete the log files for the lagged copy being activated from the active copy, or stop the Microsoft Exchange Replication service.

6. Perform a database switchover and activate the lagged copy. This example activates the database by using the Move-ActiveMailboxDatabase cmdlet with several parameters.

```
Move-ActiveMailboxDatabase DB1 -ActivateOnServer EX3 -MountDialOverride BestEffort -SkipActiveCopyChecks
-SkipClientExperienceChecks -SkipHealthChecks -SkipLagChecks
```

7. At this point, the database will automatically mount and request redelivery of missing messages from SafetyNet.

How do you know this worked?

To verify that you've successfully activated a lagged mailbox database copy, do one of the following:

- In the EAC, navigate to **Servers** > **Databases**. Select the appropriate database, and in the Details pane, click **View details** to view the database copy properties.

- In the Exchange Management Shell, run the following command to display status information for a database copy.

```
Get-MailboxDatabaseCopyStatus <DatabaseCopyName> | Format-List
```

Remove a mailbox database copy

8/3/2020 • 2 minutes to read • Edit Online

You can use these procedures to remove a copy of a mailbox database, but you can't use them to remove the **last** copy of a mailbox database. For detailed steps about how to remove the last copy of a mailbox database, see Remove a mailbox database or Remove-MailboxDatabase.

Looking for other management tasks related to mailbox database copies? Check out Manage mailbox database copies.

What do you need to know before you begin?

- Estimated time to complete this task: 1 minute

- Mailbox database copies can only be removed from a healthy database availability group (DAG). If the DAG isn't healthy (for example, the DAG's underlying cluster is down because quorum was lost), you won't be able to remove any mailbox database copies.

- If you're removing the last passive copy of the database, continuous replication circular logging (CRCL) must not be enabled for the specified mailbox database. If CRCL is enabled, you must first disable it. After the mailbox database copy has been removed, circular logging can be enabled. Once enabled for a non-replicated mailbox database, JET circular logging is used instead of CRCL. If you aren't removing the last passive copy of a database, CRCL can remain enabled.

- After removing a database copy, you must manually delete any database and transaction log files from the server from which the database copy is being removed.

- To open the EAC, see Exchange admin center in Exchange Server. To open the Exchange Management Shell, see Open the Exchange Management Shell.

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Mailbox database copies" entry in the High availability and site resilience permissions topic.

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to remove a mailbox database copy

1. In the EAC, go to **Servers** > **Databases**.

2. Select the mailbox database whose copy you want to remove.

3. In the Details pane, locate the passive copy you want to remove and click **Remove**.

4. Confirm the removal on the warning dialog box by clicking **yes**.

5. Click **ok** to confirm the removal after reviewing any messages.

6. Manually delete any database and transaction log files from the server from which the database copy is being removed.

Use the Exchange Management Shell to remove a mailbox database copy

This example removes a copy of the mailbox database DB1 from the Mailbox server MBX1.

```
Remove-MailboxDatabaseCopy -Identity DB1\MBX1 -Confirm:$False
```

For detailed syntax and parameter information, see Remove-MailboxDatabaseCopy.

How do you know this worked?

To verify that you've successfully removed a mailbox database copy, do one of the following:

- In the EAC, navigate to **Servers** > **Databases**. Select the appropriate database, and in the Details pane, the removed passive copy is no longer listed.

- In the Exchange Management Shell, run the following command to verify removal of the copy.

  ```
  Get-MailboxDatabase <DatabaseName> | Format-List DatabaseCopies
  ```

  The removed passive copy is no longer listed.

Monitor database availability groups

8/3/2020 • 19 minutes to read • Edit Online

You can use the details in this topic for monitoring mailbox database copies for database availability groups (DAGs), for gathering diagnostic information, and for configuring the low disk space monitoring threshold.

Get-MailboxDatabaseCopyStatus cmdlet

Use the Get-MailboxDatabaseCopyStatus cmdlet to view status information about mailbox database copies. This cmdlet enables you to view information about all copies of a particular database, information about a specific copy of a database on a specific server, or information about all database copies on a server. The following table describes possible values for the copy status of a mailbox database copy.

Database copy status

DATABASE COPY STATUS	DESCRIPTION
Failed	The mailbox database copy is in a Failed state because it isn't suspended, and it isn't able to copy or replay log files. While in a Failed state and not suspended, the system will periodically check whether the problem that caused the copy status to change to Failed has been resolved. After the system has detected that the problem is resolved, and barring no other issues, the copy status will automatically change to Healthy.
Seeding	The mailbox database copy is being seeded, the content index for the mailbox database copy is being seeded, or both are being seeded. Upon successful completion of seeding, the copy status should change to Initializing.
SeedingSource	The mailbox database copy is being used as a source for a database copy seeding operation.
Suspended	The mailbox database copy is in a Suspended state as a result of an administrator manually suspending the database copy by running the **Suspend-MailboxDatabaseCopy** cmdlet.
Healthy	The mailbox database copy is successfully copying and replaying log files, or it has successfully copied and replayed all available log files.
ServiceDown	The Microsoft Exchange Replication service isn't available or running on the server that hosts the mailbox database copy.
Initializing	The mailbox database copy is in an Initializing state when a database copy has been created, when the Microsoft Exchange Replication service is starting or has just been started, and during transitions from Suspended, ServiceDown, Failed, Seeding, or SinglePageRestore to another state. While in this state, the system is verifying that the database and log stream are in a consistent state. In most cases, the copy status will remain in the Initializing state for about 15 seconds, but in all cases, it should generally not be in this state for longer than 30 seconds.

DATABASE COPY STATUS	DESCRIPTION
Resynchronizing	The mailbox database copy and its log files are being compared with the active copy of the database to check for any divergence between the two copies. The copy status will remain in this state until any divergence is detected and resolved.
Mounted	The active copy is online and accepting client connections. Only the active copy of the mailbox database copy can have a copy status of Mounted.
Dismounted	The active copy is offline and not accepting client connections. Only the active copy of the mailbox database copy can have a copy status of Dismounted.
Mounting	The active copy is coming online and not yet accepting client connections. Only the active copy of the mailbox database copy can have a copy status of Mounting.
Dismounting	The active copy is going offline and terminating client connections. Only the active copy of the mailbox database copy can have a copy status of Dismounting.
DisconnectedAndHealthy	The mailbox database copy is no longer connected to the active database copy, and it was in the Healthy state when the loss of connection occurred. This state represents the database copy with respect to connectivity to its source database copy. It may be reported during DAG network failures between the source copy and the target database copy.
DisconnectedAndResynchronizing	The mailbox database copy is no longer connected to the active database copy, and it was in the Resynchronizing state when the loss of connection occurred. This state represents the database copy with respect to connectivity to its source database copy. It may be reported during DAG network failures between the source copy and the target database copy.
FailedAndSuspended	The Failed and Suspended states have been set simultaneously by the system because a failure was detected, and because resolution of the failure explicitly requires administrator intervention. An example is if the system detects unrecoverable divergence between the active mailbox database and a database copy. Unlike the Failed state, the system won't periodically check whether the problem has been resolved, and automatically recover. Instead, an administrator must intervene to resolve the underlying cause of the failure before the database copy can be transitioned to a healthy state.
SinglePageRestore	This state indicates that a single page restore operation is occurring on the mailbox database copy.

The **Get-MailboxDatabaseCopyStatus** cmdlet also returns details about the in-use replication networks, including *IncomingLogCopyingNetwork*, which is returned for passive database copies, and *OutgoingConnections*, which is returned for active databases that have more than one copy, as well as any database copy being used as a source for a database seeding operation. Outgoing connection information is provided for database copies that are in file mode replication. Outgoing connection information is not provided for database copies that are in block mode replication.

Get-MailboxDatabaseCopyStatus examples

The following examples use the **Get-MailboxDatabaseCopyStatus** cmdlet. Each example pipes the results to the **Format-List** cmdlet to display the output in list format.

This example returns status information for all copies of the database DB2.

```
Get-MailboxDatabaseCopyStatus -Identity DB2 | Format-List
```

This example returns the status for all database copies on the Mailbox server MBX2.

```
Get-MailboxDatabaseCopyStatus -Server MBX2 | Format-List
```

This example returns the status for all database copies on the local Mailbox server.

```
Get-MailboxDatabaseCopyStatus -Local | Format-List
```

For more information about using the **Get-MailboxDatabaseCopyStatus** cmdlet, see Get-MailboxDatabaseCopyStatus.

Test-ReplicationHealth cmdlet

You can use the Test-ReplicationHealth cmdlet to view continuous replication status information about mailbox database copies. This cmdlet can be used to check all aspects of the replication and replay status to provide a complete overview of a specific Mailbox server in a DAG.

The **Test-ReplicationHealth** cmdlet is designed for the proactive monitoring of continuous replication and the continuous replication pipeline, the availability of Active Manager, and the health and status of the underlying cluster service, quorum, and network components. It can be run locally on or remotely against any Mailbox server in a DAG. The **Test-ReplicationHealth** cmdlet performs the tests listed in the following table.

Test-ReplicationHealth cmdlet tests

TEST NAME	DESCRIPTION
ClusterService	Verifies that the Cluster service is running and reachable on the specified DAG member, or if no DAG member is specified, on the local server.
ReplayService	Verifies that the Microsoft Exchange Replication service is running and reachable on the specified DAG member, or if no DAG member is specified, on the local server.
ActiveManager	Verifies that the instance of Active Manager running on the specified DAG member, or if no DAG member is specified, the local server, is in a valid role (primary, secondary, or stand-alone).
TasksRpcListener	Verifies that the tasks remote procedure call (RPC) server is running and reachable on the specified DAG member, or if no DAG member is specified, on the local server.

TEST NAME	DESCRIPTION
TcpListener	Verifies that the TCP log copy listener is running and reachable on the specified DAG member, or if no DAG member is specified, on the local server.
ServerLocatorService	Verifies the Active Manager client/server processes on DAG members and on the Client Access Server that perform lookups in Active Directory and Active Manager to determine where a user's mailbox database is active.
DagMembersUp	Verifies that all DAG members are available, running, and reachable.
ClusterNetwork	Verifies that all cluster-managed networks on the specified DAG member, or if no DAG member is specified, the local server, are available.
QuorumGroup	Verifies that the default cluster group (quorum group) is in a healthy and online state.
FileShareQuorum	Verifies that the witness server and witness directory and share configured for the DAG are reachable.
DatabaseRedundancy	Verifies that there is at least one healthy copy available of the databases on the specified DAG member, or if no DAG member is specified, on the local server.
DatabaseAvailability	Verifies that the databases have sufficient availability on the specified DAG member, or if no DAG member is specified, on the local server.
DBCopySuspended	Checks whether any mailbox database copies are in a state of Suspended on the specified DAG member, or if no DAG member is specified, on the local server.
DBCopyFailed	Checks whether any mailbox database copies are in a state of Failed on the specified DAG member, or if no DAG member is specified, on the local server.
DBInitializing	Checks whether any mailbox database copies are in a state of Initializing on the specified DAG member, or if no DAG member is specified, on the local server.
DBDisconnected	Checks whether any mailbox database copies are in a state of Disconnected on the specified DAG member, or if no DAG member is specified, on the local server.
DBLogCopyKeepingUp	Verifies that log copying and inspection by the passive copies of databases on the specified DAG member, or if no DAG member is specified, on the local server, are able to keep up with log generation activity on the active copy.
DBLogReplayKeepingUp	Verifies that replay activity for the passive copies of databases on the specified DAG member, or if no DAG member is specified, on the local server, is able to keep up with log copying and inspection activity.

Test-ReplicationHealth example

This example uses the **Test-ReplicationHealth** cmdlet to test the health of replication for the Mailbox server MBX1.

```
Test-ReplicationHealth -Identity MBX1
```

Crimson channel event logging

Windows includes two categories of event logs: Windows logs, and Applications and Services logs. The Windows logs category includes the event logs available in previous versions of Windows: Application, Security, and System event logs. It also includes two new logs: the Setup log and the ForwardedEvents log. Windows logs are intended to store events from legacy applications and events that apply to the entire system.

Applications and Services logs are a new category of event logs. These logs store events from a single application or component rather than events that might have system-wide impact. This new category of event logs is referred to as an application's crimson channel.

The Applications and Services logs category includes four subtypes: Admin, Operational, Analytic, and Debug logs. Events in Admin logs are of particular interest if you use event log records to troubleshoot problems. Events in the Admin log should provide you with guidance about how to respond to the events. Events in the Operational log are also useful, but may require more interpretation. Admin and Debug logs aren't as user friendly. Analytic logs (which by default are hidden and disabled) store events that trace an issue, and often a high volume of events are logged. Debug logs are used by developers when debugging applications.

Exchange Server logs events to crimson channels in the Applications and Services logs area. You can view these channels by performing these steps:

1. Open Event Viewer.

2. In the console tree, navigate to **Applications and Services Logs** > **Microsoft** > **Exchange**.

3. Under **Exchange**, select a crimson channel, such as **HighAvailability** or **MailboxDatabaseFailureItems** to see DAG and database copy-related events, or **ActiveMontoring** or **ManagedAvailability** to see events related to Managed Availability.

The HighAvailability channel contains events related to startup and shutdown of the Microsoft Exchange Replication service, and the various components that run within the Microsoft Exchange Replication service, such as Active Manager, the third-party synchronous replication API, the tasks RPC server, TCP listener, and Volume Shadow Copy Service (VSS) writer. The HighAvailability channel is also used by Active Manager to log events related to Active Manager role monitoring and database action events, such as a database mount operation and log truncation, and to record events related to the DAG's underlying cluster.

The MailboxDatabaseFailureItems channel is used to log events associated with any failures that affect a replicated mailbox database.

The ActiveMonitoring channel contains definition and result events for Managed Availability probes, monitors and responders.

The ManagedAvailability channel contains recovery action logs and results and related events.

Low Disk Space Monitor

Exchange Server Managed Availability monitors hundreds of system metrics and components every minute, including the amount of free disk space on volumes used by the Mailbox server role. Prior to Exchange 2013 Service Pack 1 (SP1), Exchange monitored available space on all local volumes, including volumes that don't contain any databases or log files. In Exchange 2016 and Exchange 2019, only volumes that contain Exchange databases and log files are monitored. The default threshold for the low volume space monitor is 180 GB. You can configure the threshold by adding the following DWORD registry value (in MB) on each Mailbox server that you want to customize:

Path: **HKEY_LOCAL_MACHINE\Software\Microsoft\ExchangeServer\v15\Replay\Parameters**

Value: *SpaceMonitorLowSpaceThresholdInMB*

For example to configure the threshold to 100 GB, you would configure the following registry value:

REG_DWORD 186a0 (100000)

After configuring or modifying the above registry value, you must restart the Microsoft Exchange DAG Management service for the change to take effect.

CollectOverMetrics.ps1 script

Exchange Server includes a script called CollectOverMetrics.ps1, which can be found in the Scripts folder. CollectOverMetrics.ps1 reads DAG member event logs to gather information about database operations (such as database mounts, moves, and failovers) over a specific time period. For each operation, the script records the following information:

- Identity of the database
- Time at which the operation began and ended
- Servers on which the database was mounted at the start and finish of the operation
- Reason for the operation
- Whether the operation was successful, and if the operation failed, the error details

The script writes this information to .csv files with one operation per row. It writes a separate .csv file for each DAG.

The script supports parameters that allow you to customize the script's behavior and output. For example, the results can be restricted to a specified subset by using the *Database* or *ReportFilter* parameters. Only the operations that match these filters will be included in the summary HTML report. The available parameters are listed in the following table.

CollectOverMetrics.ps1 script parameters

PARAMETER	DESCRIPTION
DatabaseAvailabilityGroup	Specifies the name of the DAG from which you want to collect metrics. If this parameter is omitted, the DAG of which the local server is a member will be used. Wildcard characters can be used to collect information from and report on multiple DAGs.
Database	Provides a list of databases for which the report needs to be generated. Wildcard characters are supported, for example, `-Database:"DB1","DB2"` or `-Database:"DB*"`.

PARAMETER	DESCRIPTION
StartTime	Specifies the duration of the time period to report on. The script gathers only the events logged during this period. As a result, the script may capture partial operation records (for example, only the end of an operation at the start of the period or vice-versa). If neither *StartTime* nor *EndTime* is specified, the script defaults to the past 24 hours. If only one parameter is specified, the period will be 24 hours, either beginning or ending at the specified time.
EndTime	Specifies the duration of the time period to report on. The script gathers only the events logged during this period. As a result, the script may capture partial operation records (for example, only the end of an operation at the start of the period or vice-versa). If neither *StartTime* nor *EndTime* is specified, the script defaults to the past 24 hours If only one parameter is specified, the period will be 24 hours, either beginning or ending at the specified time.
ReportPath	Specifies the folder used to store the results of event processing. If this parameter is omitted, the Scripts folder will be used. When specified, the script takes a list of .csv files generated by the script and uses them as the source data to generate a summary HTML report. The report is the same one that's generated with the -GenerateHtmlReport option. The files can be generated across multiple DAGs at many different times, or even with overlapping times, and the script will merge all of their data together.
GenerateHtmlReport	Specifies that the script gather all the information it has recorded, group the data by the operation type, and then generate an HTML file that includes statistics for each of these groups. The report includes the total number of operations in each group, the number of operations that failed, and statistics for the time taken within each group. The report also contains a breakdown of the types of errors that resulted in failed operations.
ShowHtmlReport	Specifies that the HTML-generated report should be displayed in a Web browser after it's generated.
SummariseCsvFiles	Specifies that the script read the data from existing .csv files that were previously generated by the script. This data is then used to generate a summary report similar to the report generated by the *GenerateHtmlReport* parameter.
ActionType	Specifies the type of operational actions the script should collect. The values for this parameter are `Move`, `Mount`, `ismount`, and `Remount`. The `Move` value refers to any time that the database changes its active server, whether by controlled moves or by failovers. The `Mount`, `Dismount`, and `Remount` values refer to times that the database changes its mounted status without moving to another computer.

PARAMETER	DESCRIPTION
ActionTrigger	Specifies which administrative operations should be collected by the script. The values for this parameter are `Admin` or `Automatic`. Automatic actions are those performed automatically by the system (for example, a failover when a server goes offline). Admin actions are any actions that were performed by an administrator using either the Exchange Management Shell or the Exchange admin center.
RawOutput	Specifies that the script writes the results that would have been written to .csv files directly to the output stream, as would happen with write-output. This information can then be piped to other commands.
IncludedExtendedEvents	Specifies that the script collects the events that provide diagnostic details of times spent mounting databases. This can be a time-consuming stage if the Application event log on the servers is large.
MergeCSVFiles	Specifies that the script takes all the .csv files containing data about each operation and merges them into a single .csv file.
ReportFilter	Specifies that a filter should be applied to the operations using the fields as they appear in the .csv files. This parameter uses the same format as a `Where` operation, with each element set to `$_` and returning a Boolean value. For example: `{$_DatabaseName -notlike "Mailbox Database*"}` can be used to exclude the default databases from the report.

CollectOverMetrics.ps1 examples

The following example collects metrics for all databases that match DB* (which includes a wildcard character) in the DAG DAG1. After the metrics are collected, an HTML report is generated and displayed.

```
CollectOverMetrics.ps1 -DatabaseAvailabilityGroup DAG1 -Database:"DB*" -GenerateHTMLReport -ShowHTMLReport
```

The following examples demonstrate ways that the summary HTML report may be filtered. The first uses the *Database* parameter, which takes a list of database names. The summary report then contains data only about those databases. The next two examples use the *ReportFilter* option. The last example filters out all the default databases.

```
CollectOverMetrics.ps1 -SummariseCsvFiles (dir *.csv) -Database MailboxDatabase123,MailboxDatabase456
```

```
CollectOverMetrics.ps1 -SummariseCsvFiles (dir *.csv) -ReportFilter {$_.DatabaseName -notlike "Mailbox
Database*"}
```

```
CollectOverMetrics.ps1 -SummariseCsvFiles (dir *.csv) -ReportFilter {($_.ActiveOnStart -like "ServerXYZ*") -
and ($_.ActiveOnEnd -notlike "ServerXYZ*")}
```

CollectReplicationMetrics.ps1 script

CollectReplicationMetrics.ps1 is another health metric script included in Exchange Server. This script provides an active form of monitoring because it collects metrics in real time, while the script is running. CollectReplicationMetrics.ps1 collects data from performance counters related to database replication. The script gathers counter data from multiple Mailbox servers, writes each server's data to a .csv file, and then reports various statistics across all of this data (for example, the amount of time each copy was failed or suspended, the average copy or replay queue length, or the amount of time that copies were outside of their failover criteria).

You can either specify the servers individually, or you can specify entire DAGs. You can either run the script to first collect the data and then generate the report, or you can run it to just gather the data or to only report on data that's already been collected. You can specify the frequency at which data should be sampled and the total duration to gather data.

The data collected from each server is written to a file named **CounterData.<ServerName>.<TimeStamp>.csv**. The summary report will be written to a file named **HaReplPerfReport.<DAGName>.<TimeStamp>.csv**, or **HaReplPerfReport.<TimeStamp>.csv** if you didn't run the script with the *DagName* parameter.

The script starts Windows PowerShell jobs to collect the data from each server. These jobs run for the full period in which data is being collected. If you specify a large number of servers, this process can use a considerable amount of memory. The final stage of the process, when data is processed into a summary report, can also be quite time consuming for large amounts of data. It's possible to run the collection stage on one computer, and then copy the data elsewhere for processing.

The CollectReplicationMetrics.ps1 script supports parameters that allow you to customize the script's behavior and output. The available parameters are listed in the following table.

CollectReplicationMetrics.ps1 script parameters

PARAMETER	DESCRIPTION
DagName	Specifies the name of the DAG from which you want to collect metrics. If this parameter is omitted, the DAG of which the local server is a member will be used.
DatabaseNames	Provides a list of databases for which the report needs to be generated. Wildcard characters are supported for use, for example, `-DatabaseNames:"DB1","DB2"` or `-DatabaseNames:"DB*"`.
ReportPath	Specifies the folder used to store the results of event processing. If this parameter is omitted, the Scripts folder will be used.
Duration	Specifies the amount of time the collection process should run. Typical values would be one to three hours. Longer durations should be used only with long intervals between each sample or as a series of shorter jobs run by scheduled tasks.
Frequency	Specifies the frequency at which data metrics are collected. Typical values would be 30 seconds, one minute, or five minutes. Under normal circumstances, intervals that are shorter than these won't show significant changes between each sample.
Servers	Specifies the identity of the servers from which to collect statistics. You can specify any value, including wildcard characters or GUIDs.

PARAMETER	DESCRIPTION
SummariseFiles	Specifies a list of .csv files to generate a summary report. These files are the files named **CounterData.<CounterData>*** and are generated by the CollectReplicationMetrics.ps1 script.
Mode	Specifies the processing stages that the script executes. You can use the following values: `CollectAndReport`: This is the default value. This value signifies that the script should both collect the data from the servers and then process them to produce the summary report. `CollectOnly`: This value signifies that the script should just collect the data and not produce the report. `ProcessOnly`: This value signifies that the script should import data from a set of .csv files and process them to produce the summary report. The *SummariseFiles* parameter is used to provide the script with the list of files to process.
MoveFilestoArchive	Specifies that the script should move the files to a compressed folder after processing.
LoadExchangeSnapin	Specifies that the script should load the Exchange Management Shell commands. This parameter is useful when the script needs to run from outside the Exchange Management Shell, such as in a scheduled task.

CollectReplicationMetrics.ps1 example

The following example gathers one hour's worth of data from all the servers in the DAG DAG1, sampled at one minute intervals, and then generates a summary report. In addition, the *ReportPath* parameter is used, which causes the script to place all the files in the current directory.

```
CollectReplicationMetrics.ps1 -DagName DAG1 -Duration "01:00:00" -Frequency "00:01:00" -ReportPath
```

The following example reads the data from all the files matching CounterData* and then generates a summary report.

```
CollectReplicationMetrics.ps1 -SummariseFiles (dir CounterData*) -Mode ProcessOnly -ReportPath
```

Switchovers and failovers

8/3/2020 • 19 minutes to read • Edit Online

Switchovers and failovers are the two forms of outages in Microsoft Exchange Server:

- A *switchover* is a scheduled outage of a database or server that's explicitly initiated by a cmdlet or by the managed availability system in Exchange Server. Switchovers are typically done to prepare for performing a maintenance operation. Switchovers involve moving the active mailbox database copy to another server in the database availability group (DAG). If no healthy target is found during a switchover, administrators will receive an error and the mailbox database will remain up, or mounted.

- A *failover* refers to unexpected events that result in the unavailability of services, data, or both. A failover involves the system automatically recovering from the failure by activating a passive mailbox database copy to make it the active mailbox database copy. If no healthy target is found during a failover, the mailbox database will be dismounted.

Exchange Server is specifically designed to handle both switchovers and failovers.

Looking for management tasks related to high availability and site resilience? See Managing high availability and site resilience.

Switchovers

There are three types of switchovers in Exchange Server:

- Database switchovers

- Server switchovers

- Datacenter switchovers

Database Switchovers

A *database switchover* is the process by which an individual active database is switched over to another database copy (a passive copy), and that database copy is made the new active database copy. Database switchovers can happen both within and across datacenters. A database switchover can be performed by using the Exchange admin center (EAC) or the Exchange Management Shell. Regardless of which interface is used, the switchover process is as follows:

1. The administrator initiates a database switchover to move the current active mailbox database copy to another server.

2. The client used for the task makes an RPC call to the Microsoft Exchange Replication service on a DAG member.

3. If the DAG member doesn't hold the Primary Active Manager (PAM) role, the DAG member refers the task to the server that holds the PAM role.

4. The task makes an RPC call to the Microsoft Exchange Replication service on the server that holds the PAM role.

5. The PAM reads and updates the database location information that's stored in the cluster database for the DAG.

6. The PAM contacts the Microsoft Exchange Replication service on the DAG member whose passive copy is

being activated as the new active mailbox database copy.

7. The Microsoft Exchange Replication service on the target server queries the Microsoft Exchange Replication services on all other DAG members to determine the best log source for the database copy.
8. The database is dismounted from the current server and the Microsoft Exchange Replication service on the target server copies the remaining logs to the target server.
9. The Microsoft Exchange Replication service on the target server requests a database mount.
10. The Microsoft Exchange Information Store service on the target server replays the log files and mounts the database.
11. Any error codes are returned to the Microsoft Exchange Replication service on the target server.
12. The PAM updates the database copy state information in the cluster database for the DAG.
13. Any error codes are returned by the Microsoft Exchange Replication service on the target server to the Microsoft Exchange Replication service on the PAM.
14. The Microsoft Exchange Replication service on the PAM returns any errors to the administrative interface where the task was called.
15. Remote PowerShell returns the results of the operation to the calling administrative interface.

For detailed steps about how to perform a database switchover, see Activate a mailbox database copy.

Server Switchovers

A server switchover is the process by which all active databases on a DAG member are activated on one or more other DAG members. Like database switchovers, a server switchover can occur both within a datacenter and across datacenters, and it can be initiated by using both the EAC and the Exchange Management Shell. Regardless of which interface is used, the server switchover process is as follows:

1. The administrator initiates a server switchover to move all current active mailbox database copies to one or more other servers.
2. The task performs the same steps described earlier in this topic for database switchovers (Steps 2 through 4) for each of the active databases on the current server.
3. The PAM reads and updates the database location information that's stored in the cluster database for the DAG.
4. The PAM contacts the Microsoft Exchange Replication service on each DAG member that has a passive copy being activated.
5. The Microsoft Exchange Replication service on the target servers query the Microsoft Exchange Replication services on all other DAG members to determine the best log source for the database copy.
6. The database is dismounted from the current server and the Microsoft Exchange Replication service on each target server copies the remaining logs.
7. The Microsoft Exchange Replication service on each target server requests a database mount.
8. The Microsoft Exchange Information Store service on each target server replays the log files and mounts the database.
9. Any error codes are returned to the Microsoft Exchange Replication service on the target server.
10. The PAM updates the database copy state information in the cluster database for the DAG.
11. Any error codes are returned by the Microsoft Exchange Replication service on the target server to the

Microsoft Exchange Replication service on the PAM.

12. The Microsoft Exchange Replication service on the PAM returns any errors to the administrative interface where the task was called.
13. Remote PowerShell returns the results of the operation to the calling administrative interface.

For detailed steps about how to perform a server switchover, see Perform a server switchover.

Datacenter Switchovers

In a site resilient configuration, automatic recovery in response to a site-level failure can occur within a DAG, allowing the messaging system to remain in a functional state. This configuration requires at least three locations, as it requires deploying DAG members in two locations and the DAG's witness server in a third location.

If you don't have three locations, or even if you do have three locations but you want to control datacenter-level recovery actions, you can configure a DAG for manual recovery in the event of a site-level failure. In that event, you would perform a process called a *datacenter switchover*. As with many disaster recovery scenarios, prior planning and preparation for a datacenter switchover can simplify your recovery process and reduce the duration of your outage. For detailed steps to performing a datacenter switchover, see Datacenter switchovers

Failovers

A failover is an automatic activation process that can occur at the database, server, or datacenter level. Failovers occur in response to a failure that affects an individual database (for example, an isolated storage loss) an entire server (for example, a motherboard failure or a loss of power), or an entire site (for example, the loss of all DAG members in a site).

DAGs and mailbox database copies provide full redundancy and rapid recovery of both the data and the services that provide access to the data. The following table lists the expected recovery actions for a variety of failures. Some failures require the administrator to initiate the recovery, and other failures are automatically handled by the system.

DESCRIPTION	AUTOMATIC ACTIVATION	AUTOMATIC REPAIR ACTION	STATE DURING REPAIR: ACTIVE	STATE DURING REPAIR: PASSIVE	REPAIR ACTIONS	COMMENTS
Extensible Storage Engine (ESE) soft database failure: The drives storing the database are returning errors on some reads (for example, a -1018 error).	Possible short outage. Possible automatic failover.	Automatic patching of bad page.	Manual switchover, automatic failover, or online repair.	Failed	RAID rebuild, database and database copy repair, restore and run recovery then page patching, or page patching from copy.	There may be other soft database failure codes. Doesn't include NTFS file system block failures. If failover or switchover is performed, host server is updated.

DESCRIPTION	AUTOMATIC ACTIVATION	AUTOMATIC REPAIR ACTION	STATE DURING REPAIR: ACTIVE	STATE DURING REPAIR: PASSIVE	REPAIR ACTIONS	COMMENTS
ESE " *semi-soft"* database failure: The drives storing the database are returning errors on some writes.	Short outage during automatic failover.	Automatic volume/disk rebuilt after possible drive replacement.	Dismounted if can't be recovered.	Failed	RAID rebuild may solve the problem. Copy and repair, restore and run recovery, or volume/disk rebuilt after possible replacement.	An ESE semi-soft write error means some writes are successful. Doesn't include an NTFS block failure.
ESE "semi-soft" log failure: The drives storing the log data are returning non-recovered errors on some reads or writes.	Short outage during automatic failover.	Automatic volume/disk rebuilt after possible drive replacement.	Dismounted if can't be recovered.	Failed	RAID rebuild may solve the problem. Copy and repair, restore and run recovery, or volume/disk rebuilt after possible replacement.	An ESE semi-soft read/write error means some reads/writes are successful. If the database fails, automated recovery will occur before log data recovery processing starts.
ESE software error or resource exhaustion: An error where ESE terminates instance (for example, Event ID 1022, checkpoint depth too deep).	Short outage during automatic failover.	None.	Dismounted if can't be recovered.	Failed	Fix underlying resource issue.	This failure could be the surfaced error of other cases.

DESCRIPTION	AUTOMATIC ACTIVATION	AUTOMATIC REPAIR ACTION	STATE DURING REPAIR: ACTIVE	STATE DURING REPAIR: PASSIVE	REPAIR ACTIONS	COMMENTS
NTFS block failures: The drives storing the database or logs experiences a read or write error to an NTFS control structure.	Short outage during automatic failover.	Volume completely rebuilt after possible drive replacement.	Dismounted if can't be recovered.	Failed	RAID rebuild may solve the problem. NTFS utilities may solve the NTFS problems. Exchange recovery may be required.	This is more likely to occur when RAID isn't in use. If this impacts the active log volume, some recent log files will be lost. Doesn't include errors automatically corrected by NTFS or its underlying software or hardware stack.
Database or log drive failure: A drive storing the database or logs has completely failed and is inaccessible.	Short outage during automatic failover.	Drive reformatted or replaced, followed by complete volume rebuild.	Dismounted if can't be recovered.	Failed	Drive replacement followed by possible RAID rebuild. Drive replacement followed by complete volume rebuild. Complete volume rebuild.	Not applicable.
Database or log volume failure: The volume fails due to NTFS or lower level volume issues.	Short outage during automatic failover.	Drive reformatted or replaced.	Dismounted if can't be recovered.	Failed	Drive replacement followed by possible RAID rebuild. Drive replacement followed by complete volume rebuild. Complete volume rebuild.	Not applicable.

DESCRIPTION	AUTOMATIC ACTIVATION	AUTOMATIC REPAIR ACTION	STATE DURING REPAIR: ACTIVE	STATE DURING REPAIR: PASSIVE	REPAIR ACTIONS	COMMENTS
Database or log volume out of space: The NTFS file system with the database or log files is out of space.	Automatic failover if other copy isn't in similar state.	None.	Dismounted.	Failed	Run full or incremental backups, manually delete logs, let time pass, resume database copy, or repair failed database copy.	Not applicable.
Administrator dismounts the wrong database.	If automatic failover isn't blocked by the administrator, there will be a short outage. If automatic failover is prevented, there will be an outage until the database is mounted.	None.	Dismounted.	Not applicable	Administrator corrects the error.	Not applicable.
Administrator suspends the wrong database copy.	Depending on configuration and impacted copy, auto recovery may be prevented.	None.	Not applicable.	Suspended	Administrator corrects the error.	Not applicable.
Administrator dismounts a database for storage, NTFS, or volume maintenance.	If automatic failover isn't blocked by the administrator, there will be a short outage. If automatic failover is blocked, there will be an outage until the administrator completes the task.	None.	Dismounted.	Not applicable	Administrator completes the task.	Not applicable.

DESCRIPTION	AUTOMATIC ACTIVATION	AUTOMATIC REPAIR ACTION	STATE DURING REPAIR: ACTIVE	STATE DURING REPAIR: PASSIVE	REPAIR ACTIONS	COMMENTS
Administrator suspends a database copy for storage, NTFS, or volume maintenance.	Depending on configuration and impacted copy, auto recovery may be prevented.	None.	Not applicable.	Suspended	Administrator completes the actions.	Not applicable.
Administrator dismounts a database for offline database maintenance.	Outage until repaired.	None.	Dismounted.	Suspended	Administrator completes the actions.	Active and passive database copies are diverged. Administrator must suspend copies.
Storage area network (SAN), disk, or storage controller failure.	Short outage during automatic failover.	None.	Dismounted.	Any	Repair hardware.	A passive database copy will be in the state that existed at the time when the system failed.
Server hardware maintenance.	Short outage during automatic failover (unless blocked by an administrator).	None.	Dismounted.	Any	Complete actions.	A passive database copy will be in the state that existed at the time when the system was shut down.
Server software maintenance.	Short outage during automatic failover (unless blocked by an administrator).	None.	Dismounted.	Any	Complete actions.	A passive database copy will be in the state that existed at the time when the system was shut down.
Microsoft Exchange Information Store service is stopped or paused by an administrator.	Short outage during automatic failover.	None.	Dismounted.	Any	Restart the Microsoft Exchange Information Store service.	Not applicable.

DESCRIPTION	AUTOMATIC ACTIVATION	AUTOMATIC REPAIR ACTION	STATE DURING REPAIR: ACTIVE	STATE DURING REPAIR: PASSIVE	REPAIR ACTIONS	COMMENTS
Microsoft Exchange Information Store service fails; operating system is still running.	Short outage during automatic failover.	Service Control Manager restarts the Microsoft Exchange Information Store service.	Dismounted.	Any	Manually or automatically restart the Microsoft Exchange Information Store service.	A passive database copy will be in the state that existed when the Microsoft Exchange Information Store service failed.
Partial Microsoft Exchange Information Store service failure; some part of the Exchange store stops functioning, but it's not identified as completely failed.	Possible short outage during automatic failover.	None.	Mounted and partially functional.	Any, but may be only partially functional	Restart server, operating system, or Microsoft Exchange Information Store service.	Not applicable.
Server failure: The server fails for one of the following reasons: Complete power failure Unrecovered failure of the processor chip, motherboard, or backplane Operating system stop error Operating system stops responding Complete communication failure	Short outage during automatic failover.	Restart computer.	Dismounted.	Any	Restore power, change operating system settings, change hardware settings, replace hardware, restart operating system, service operating system, service hardware, or repair communication problems.	Not applicable.
DAG experiences a quorum failure.	Outage until repaired.	None.	Dismounted.	Any	Repair failed quorum, assign new quorum, or restore the network that's causing quorum failure.	A passive database copy will be in the state that existed at the time when the system failed.

DESCRIPTION	AUTOMATIC ACTIVATION	AUTOMATIC REPAIR ACTION	STATE DURING REPAIR: ACTIVE	STATE DURING REPAIR: PASSIVE	REPAIR ACTIONS	COMMENTS
MAPI network communication failure: The server is no longer available on the MAPI network.	Short outage during automatic failover; must be lossless.	None. Communication continues to be attempted.	Dismounted.	Any	Fix communication problem by correcting hardware or software issues.	Not applicable.
Replication network communication failure: The server can't receive heartbeats, log copies, or seed through the failed replication network.	Possible short copying or seeding outage while the workload is switched to other network.	None. Communication continues to be attempted.	None.	Any	Fix communication problem by correcting hardware or software issues.	Resiliency impacted by failure.
Multiple network communication failure: The server can't receive heartbeats, log copies, or seed through multiple networks.	Short outage during automatic failover; must be lossless.	None. Communication continues to be attempted.	Dismounted.	Any	Fix communication problem by correcting hardware or software issues.	At least one network is still functional.
Partial failure of one or more networks: Networks experience high error rates.	Failure not detected; no action.	None.	Mounted, but possible performance issues.	Any	Fix communication problem by correcting hardware or software issues.	Network experiences higher than normal error rates.

DESCRIPTION	AUTOMATIC ACTIVATION	AUTOMATIC REPAIR ACTION	STATE DURING REPAIR: ACTIVE	STATE DURING REPAIR: PASSIVE	REPAIR ACTIONS	COMMENTS
Undetected operating system hang: Operating system stops responding but it's not detected by monitoring or clustering.	None.	None.	Any.	Any	Restart or terminate the resources that aren't responding.	Hang isn't detected so no action is taken. Some functionality may be operational.
Operating system drive experiences a failure.	Short outage during automatic failover.	None.	Dismounted.	Any	Replace drive and rebuild server or rebuild volume by using RAID.	Not applicable.
Operating system drive out of space.	Short outage during automatic failover.	None.	Dismounted.	Any	Manually free space on the volume.	Not applicable.
Drive containing Exchange binaries experiences a volume or drive failure.	Short outage during automatic failover.	None.	Dismounted.	Any	Replace drive and reinstall application or rebuild volume by using RAID.	Not applicable.
Drive containing the Exchange binaries is out of space.	Short outage during automatic failover.	None.	Dismounted.	Any	Manually free space on the volume.	Not applicable.
Invalid new log detected: The log sequence is disrupted by an existing file.	Short outage during automatic failover; assume other copies don't have the same problem.	None.	Dismounted.	Failed	Remove disruptive logs after determining source.	The disruptive logs shouldn't replicate.
Continuous replication detects invalid log: Replay detects an inappropriate log during copy or replay.	Not applicable.	Discard log.	Not applicable.	Failed	Discard invalid log; move impacting log stream.	Not applicable.

Database Failovers

A database failover occurs when a database copy that was active is no longer able to remain active. The following occurs as part of a database failover:

1. The database failure is detected by the Microsoft Exchange Information Store service.
2. The Microsoft Exchange Information Store service writes failure events to the crimson channel event log.
3. The Active Manager on the server that contains the failed database detects the failure events.
4. The Active Manager requests the database copy status from the other servers that hold a copy of the database.
5. The other servers return the requested database copy status to the requesting Active Manager.
6. The PAM initiates a move of the active database to another server in the DAG using a best copy selection algorithm.
7. The PAM updates the database mount location in the cluster database to refer to the selected server.
8. The PAM sends a request to the Active Manager on the selected server to become the database master.
9. The Active Manager on the selected server requests that the Microsoft Exchange Replication service attempt to copy the last logs from the previous server and set the mountable flag for the database.
10. The Microsoft Exchange Replication service copies the logs from the server that previously had the active copy of the database.
11. The Active Manager reads the maximum log generation number from the cluster database.
12. The Microsoft Exchange Information Store service mounts the new active database copy.

Server Failovers

A server failover occurs when the DAG member is no longer able to service the MAPI network, or when the Cluster service on a DAG member is no longer able to contact the remaining DAG members. The following occurs as part of a server failover:

1. The Cluster service on the PAM sends a notification to the PAM for one of two conditions:
2. **Node Down**: The server is reachable but is unable to participate in DAG operations.
3. **MAPI Network Down**: The server can't be contacted over the MAPI network and therefore can't participate in DAG operations.
4. If the server is reachable, the PAM contacts the Active Manager on the affected server and requests that all databases be immediately dismounted.
5. For each affected database copy:
6. The PAM requests the database copy status from all servers in the DAG.
7. The PAM receives a response from all reachable and active DAG members.
8. The PAM tries to determine the best log source among all responding servers by querying the most recent log generation number from each of the responders.
9. Each of the servers responds with the log generation number.
10. The PAM retrieves the current search index catalog status from the cluster database.
11. Based on the log generation number and catalog health of each database copy, the PAM selects the best copies to activate.
12. The PAM updates the mounted location of the database in the cluster database.

13. The PAM initiates database failover by communicating with the Active Manager on one or more other servers.

14. The Active Manager on the selected servers requests that the Microsoft Exchange Replication service attempt to copy the last logs from the previous server and set the mountable flag.

15. When the database is mountable, the Active Manager on the servers mounts the databases.

For more information about Active Manager's best copy selection process, see Active Manager.

Datacenter Failovers

Significant changes have been made since Exchange 2010 regarding site resilience configuration. With namespace simplification, consolidation of server roles, separation of Client Access services and DAG recovery (in Exchange Server, the namespace does not need to move with the DAG), and changes around load balancing, Exchange Server provides site resilience options like the ability to use a single global namespace. If you have more than two locations in which to deploy messaging service components, Exchange Server also provides the ability to configure the messaging service for automatic failover in response to failures that required manual intervention in previous versions.

Exchange leverages fault tolerance built into the namespace through multiple IP addresses, load balancing, and, if necessary, the ability to take servers in and out of service. Exchange Server makes it possible to leverage the clients' ability to cache multiple IP addresses returned from a DNS server in response to a name resolution request. Clients with the ability to cache multiple IP addresses (which includes almost all HTTP-based clients in Exchange Server, such as Outlook, Outlook Anywhere, EAS, EWS, Outlook on the web, EAC, RPS, etc.), all have the ability to use those multiple IP addresses, and this provides failover on the client side. You can configure DNS to hand multiple IP addresses to a client during name resolution. The client asks for mail.contoso.com and gets back two IP addresses, or four IP addresses, for example. However many IP addresses the client gets back will be used reliably by the client. This makes the client a lot better off because if one of the IP addresses fails, the client has one or more others to try to connect to. If a client tries one and it fails, it waits around 20 seconds and then tries the next one in the list. Thus, if you lose connectivity to your primary Client Access services (CAS) array, and you have a second published IP address for a second CAS array, recovery for the clients happens automatically (and in about 21 seconds).

Modern HTTP clients (operating systems and Web browsers that are ten years old or less) simply work with this redundancy automatically. The HTTP stack can accept multiple IP addresses for an FQDN, and if the first IP it tries fails hard (e.g., cannot connect), it will try the next IP in the list. In a soft failure (connect lost after session established, perhaps due to an intermittent failure in the service where, for example, a device is dropping packets and needs to be taken out of service), the user might need to refresh their browser.

With the proper configuration, failover can happen at the client level and clients will be automatically redirected to a second datacenter where Client Access services is running, and the servers that are running Client Access services will proxy the communication back to the user's Mailbox server, which remains unaffected by the outage (because you don't do a switchover). Instead of working to recover service, the service recovers itself and you can focus on fixing the core issue (e.g., replacing a failed load balancer).

Since you can failover the namespace between datacenters, all that is needed to achieve a datacenter failover is a mechanism for failover of the Mailbox role across datacenters. To get automatic failover for the DAG, you simply architect a solution where the DAG is evenly split between two datacenters, and then place the witness server in a third location so that it can be arbitrated by DAG members in either datacenter, regardless of the state of the network between the datacenters that contain the DAG members. The key is that third location is isolated from network failures that affect the locations containing the DAG members.

If you only have two datacenters and would like to be able to configure automatic failover, you can utilize Microsoft Azure as your third location. You will need to create an Azure virtual network and connect it to your two datacenters using a multi-point VPN. You will then be able to place your witness server on a Microsoft Azure virtual machine. For more information, see Using a Microsoft Azure VM as a DAG witness server.

Datacenter switchovers

8/3/2020 • 21 minutes to read • Edit Online

In a site resilient configuration, automatic recovery in response to a site-level failure can occur within a DAG, allowing the messaging system to remain in a functional state. This configuration requires at least three locations, as it requires deploying DAG members in two locations and the DAG's witness server in a third location.

If you don't have three locations, or even if you do have three locations but you want to control datacenter-level recovery actions, you can configure a DAG for manual recovery in the event of a site-level failure. In that event, you would perform a process called a *datacenter switchover*. As with many disaster recovery scenarios, prior planning and preparation for a datacenter switchover can simplify your recovery process and reduce the duration of your outage.

There are four basic steps that you complete to perform a datacenter switchover, after making the initial decision to activate the second datacenter:

1. **Terminate a partially running datacenter**: This step involves terminating Exchange services in the primary datacenter, if any services are still running. This is particularly important for the Mailbox server role because it uses an active/passive high availability model. If services in a partially failed datacenter aren't stopped, it's possible for problems from the partially failed datacenter to negatively affect the services during a switchover back to the primary datacenter.

   > **IMPORTANT**
   >
   > If network or Active Directory infrastructure reliability has been compromised as a result of the primary datacenter failure, we recommend that all messaging services be off until these dependencies are restored to healthy service.

2. **Validate and confirm the prerequisites for the second datacenter**: This step can be performed in parallel with step 1 because validation of the health of the infrastructure in the second datacenter is largely independent of the first datacenter. Each organization typically requires its own method for performing this step. For example, you may decide to complete this step by reviewing health information collected and filtered by an infrastructure monitoring application, or by using a tool that's unique to your organization's infrastructure. This is a critical step, because you don't want to activate the second datacenter when its infrastructure is unhealthy and unstable.

3. **Activate the Mailbox servers**: This step begins the process of activating the second datacenter. This step can be performed in parallel with step 4 because the Microsoft Exchange services can handle database outages and recover. Activating the Mailbox servers involves a process of marking the failed servers from the primary datacenter as unavailable followed by activation of the servers in the second datacenter. The activation process for Mailbox servers depends on whether the DAG is in database activation coordination (DAC) mode. See Datacenter Activation Coordination mode for more information.

   If the DAG is in DAC mode, you can use the Exchange site resilience cmdlets to terminate a partially failed datacenter (if necessary) and activate the Mailbox servers. For example, in DAC mode, this step is performed by using the Stop-DatabaseAvailabilityGroup cmdlet. In some cases, the servers must be marked as unavailable twice (once in each datacenter). Next, the Restore-DatabaseAvailabilityGroup cmdlet is run to restore the remaining members of the database availability group (DAG) in the second datacenter by reducing the DAG members to those that are still operational, thereby reestablishing quorum. If the DAG isn't in DAC mode, you must use the Windows Failover Cluster tools to activate the Mailbox servers. After either process is complete, the database copies that were previously passive in the second datacenter can become active and be mounted. At this point, Mailbox server recovery is complete.

4. **Activate Client Access services**: This involves using the URL mapping information and the Domain Name System (DNS) change methodology to perform all required DNS updates. The mapping information describes what DNS changes to perform. The amount of time required to complete the update depends on the methodology used and the Time to Live (TTL) settings on the DNS record (and whether the deployment's infrastructure honors the TTL).

Users should start to have access to messaging services sometime after steps 3 and 4 are completed. Steps 3 and 4 are described in greater detail later in this topic.

Terminating a Partially Failed Datacenter

If any DAG members in the failed datacenter are still running, they should be terminated.

When the Exchange DAG is in DAC mode, you can disable the servers in a failed datacenter with a single command. This will allow you to mount the databases in another datacenter even if the DAG doesn't have quorum (more than half the members of the DAG available).

When the DAG is in DAC mode, the specific actions to terminate any surviving DAG members in the primary datacenter are as follows:

1. The DAG members in the primary datacenter must be marked as stopped in the primary datacenter. *Stopped* is a state of Active Manager that prevents databases from mounting, and Active Manager on each server in the failed datacenter is put into this state by using the Stop-DatabaseAvailabilityGroup cmdlet. The *ActiveDirectorySite* parameter of this cmdlet can be used to mark all of the servers in the primary datacenter as stopped with a single command. This step may not be possible depending on the failure. This step should be taken if the state of the datacenter permits it. The **Stop-DatabaseAvailabilityGroup** cmdlet should be run against all servers in the primary datacenter. If the Mailbox server is unavailable but Active Directory is operating in the primary datacenter, the **Stop-DatabaseAvailabilityGroup** command with the *ConfigurationOnly* parameter must be run against all servers in this state in the primary datacenter, or the Mailbox server must be turned off. Failure to either turn off the Mailbox servers in the failed datacenter or to successfully perform the **Stop-DatabaseAvailabilityGroup** command against the servers will create the potential for split-brain syndrome to occur across the two datacenters. You may need to individually turn off computers through power management devices to satisfy this requirement.

2. The second datacenter must now be updated to represent which primary datacenter servers are stopped. This is done by running the same **Stop-DatabaseAvailabilityGroup** command with the *ConfigurationOnly* parameter using the same *ActiveDirectorySite* parameter and specifying the name of the Active Directory site in the failed primary datacenter. The purpose of this step is to inform the servers in the second datacenter about which mailbox servers are available to use when restoring service.

When the DAG isn't in DAC mode, the specific actions to terminate any surviving DAG members in the primary datacenter are as follows:

1. The DAG members in the primary datacenter must be forcibly evicted from the DAG's underlying cluster by running the following commands on each member:

   ```
   net stop clussvc
   ```

   ```
   cluster <DAGName> node <DAGMemberName> /forcecleanup
   ```

2. The DAG members in the second datacenter must now be restarted and then used to complete the eviction process from the second datacenter. Stop the Cluster service on each DAG member in the second datacenter by running the following command on each member:

```
net stop clussvc
```

3. On a DAG member in the second datacenter, force a quorum start of the Cluster service by running the following command:

   ```
   net start clussvc /forcequorum
   ```

4. Open the Failover Cluster Management tool and connect to the DAG's underlying cluster. Expand the cluster, and then expand **Nodes**. Right-click each node in the primary datacenter, select **More Actions**, and then select **Evict**. When you're done evicting the DAG members in the primary datacenter, close the Failover Cluster Management tool.

If any Unified Messaging services are in use in the failed datacenter, they must be disabled to prevent call routing to the failed datacenter. You can disable Unified Messaging services by using the Disable-UMService cmdlet (for example, `Disable-UMService EX1`). Alternatively, if you're using a Voice over IP (VoIP) gateway, you can also remove the server entries from the VoIP gateway, or change the DNS records for the failed servers to point to the IP address of the servers in the second datacenter if your VoIP gateway is configured to route calls using DNS.

> **NOTE**
>
> Unified Messaging is not available in Exchange 2019

Activating Mailbox Servers

The steps needed to activate Mailbox servers during a datacenter switchover also depend on whether the DAG is in DAC mode. Before activating the DAG members in the second datacenter, we recommend that you validate that the infrastructure services in the second datacenter are ready for messaging service activation.

When the DAG is in DAC mode, the steps to complete activation of the mailbox servers in the second datacenter are as follows:

1. The Cluster service must be stopped on each DAG member in the second datacenter. You can use the **Stop-Service** cmdlet to stop the service (for example, `Stop-Service ClusSvc`), or use `net stop clussvc` from an elevated command prompt.

2. The Mailbox servers in the standby datacenter are then activated by using the Restore-DatabaseAvailabilityGroup cmdlet. The Active Directory site of the standby datacenter is passed to the **Restore-DatabaseAvailabilityGroup** cmdlet to identify which servers to use to restore service and to configure the DAG to use an alternate witness server. If the alternate witness server wasn't previously configured, you can configure it by using the *AlternateWitnessServer* and *AlternateWitnessDirectory* parameters of the Restore-DatabaseAvailabilityGroup cmdlet. If this command succeeds, the quorum criteria are shrunk to the servers in the standby datacenter. If the number of servers in that datacenter is an even number, the DAG will switch to using the alternate witness server as identified by the setting on the DAG object.

3. The databases can now be activated. Depending on the specific configuration used by the organization, this may not be automatic. If the servers in the standby datacenter have an activation blocked setting, the system won't do an automatic failover from the primary datacenter to the standby datacenter of any database. If no failover restrictions are present for any of the database copies in the standby datacenter, the system will activate copies in the second datacenter assuming they are healthy. If databases are configured with an activation blocked setting that requires explicit manual action, there are two choices for action:

4. Clear the setting that blocks activation. This will make the system return to its default behavior, which is to

activate any available copy.

5. Leave the setting unchanged and use the Move-ActiveMailboxDatabase cmdlet to complete the database activation in the second datacenter. To complete this step using the **Move-ActiveMailboxDatabase** cmdlet when activation blocked is set, you must explicitly identify the target of the move.

6. The last step is to review all error and warning messages from the tasks. Any indicated warnings should be followed up and corrected. The task design model for these commands is to only fail if they can't achieve the fundamental goal of their design. For example, the **Restore-DatabaseAvailabilityGroup** cmdlet will fail if it can't shrink the quorum of the DAG to allow a server in the second datacenter to be restarted for servicing without causing a quorum outage. However, each task's output is also used to identify the issues that require administrator follow-up. You're strongly encouraged to save all task output and review it for follow-up actions.

When the DAG isn't in DAC mode, the steps to complete activation of the mailbox servers in the second datacenter are as follows:

1. The quorum must be modified based on the number of DAG members in the second datacenter.

2. If there's an odd number of DAG members, change the DAG quorum model from a Node a File Share Majority to a Node Majority quorum by running the following command:

   ```
   cluster <DAGName> /quorum /nodemajority
   ```

3. If there's an even number of DAG members, reconfigure the witness server and directory by running the following command in the Exchange Management Shell:

   ```
   Set-DatabaseAvailabilityGroup <DAGName> -WitnessServer <ServerName>
   ```

4. Start the Cluster service on any remaining DAG members in the second datacenter by running the following command:

   ```
   net start clussvc
   ```

5. Perform server switchovers to activate the mailbox databases in the DAG by running the following command for each DAG member:

   ```
   Move-ActiveMailboxDatabase -Server <DAGMemberinPrimarySite> -ActivateOnServer <DAGMemberinSecondSite>
   ```

6. Mount the mailbox databases on each DAG member in the second site by running the following command:

   ```
   Get-MailboxDatabase <DAGMemberinSecondSite> | Mount-Database
   ```

Activating Client Access services

Clients connect to service endpoints (for example Outlook on the web, Autodiscover, Exchange ActiveSync, Outlook Anywhere, POP3, IMAP4, and the RPC Client Access services array) to access Exchange services and data. Therefore, activating Client Access services involves changing the mapping of the DNS records for these service endpoints from IP addresses in the primary datacenter to the IP addresses in the second datacenter that are configured as the new service endpoints. Depending on your DNS configuration, the DNS records that need to be modified may or may not be in the same DNS zone.

Activating Client Access services

Clients will then automatically connect to the new service endpoints in one of two ways:

- Clients will continue to try to connect, and should automatically connect after the TTL has expired for the original DNS entry, and after the entry is expired from the client's DNS cache. Users can also run the `ipconfig /flushdns` command from a command prompt to manually clear their DNS cache.
- Clients starting or restarting will perform a DNS lookup on startup and will get the new IP address for the service endpoint, which will be an Exchange server running Client Access services, or a Client Access services array, in the second datacenter.

Assuming that all appropriate configuration changes have been completed to define and configure the services in the second datacenter to function as they were in the primary datacenter, and assuming that the established DNS configuration is correct, no further changes should be needed to activate Client Access services.

Activating Transport services

Clients and other servers that submit messages typically identify those servers using DNS. Activating transport services in the second datacenter involves changing DNS records to point to the IP addresses of the Mailbox servers in the second datacenter. Clients and sending servers will then automatically connect to the servers in the second datacenter in one of two ways:

- Clients will continue to try to connect, and should automatically connect after the TTL has expired for the original DNS entry, and after the entry is expired from the client's DNS cache. Users can also run the `ipconfig /flushdns` command from a command prompt to manually clear their DNS cache.
- Clients starting or restarting will perform a DNS lookup on startup and will get the new IP address for the SMTP endpoint, which will be a Mailbox server in the second datacenter.

Assuming that all appropriate configuration changes have been completed to define and configure the services in the second datacenter to function as they were in the primary datacenter, and assuming that the established DNS configuration is correct, no further changes should be needed to activate transport services.

Activating Unified Messaging services in Exchange 2016

> **NOTE**
>
> Unified Messaging is not available in Exchange 2019.

Unified Messaging (UM) services in Exchange 2016 connect to an organization's PBX system and phone lines. The logical connection between the PBX system and the Unified Messaging service is provided by an IP gateway. IP gateways include high availability functionality and are able to switch between multiple Unified Messaging services when a failure is detected.

If there are Unified Messaging services in the second datacenter that were in a disabled state because they are dedicated to the site resilience solution, they can be enabled by using the Enable-UMService cmdlet (for example, `Enable-UMService EX4`).

Assuming the IP gateways are associated with Unified Messaging services by using DNS servers, activating Unified Messaging services therefore involves changing DNS records to point to the new IP addresses that will be configured for the Unified Messaging service in the second datacenter. Assuming that all appropriate configuration changes have been completed to define and configure the services in the second datacenter to function as they were in the primary datacenter, and assuming that the established DNS configuration is correct, no further changes should be needed to activate Unified Messaging services.

If the IP gateway in use doesn't support the use of DNS names to resolve the Unified Messaging services, additional configuration steps will be necessary to manually point the IP gateway to the IP addresses of the Unified

Messaging services in the second datacenter.

Activating Edge Transport Servers

The steps to activate the Edge Transport server role will vary, depending on the specific configuration. Edge Transport servers in two datacenters can be configured in either an active/passive or an active/active configuration. In an active/passive configuration, the Edge Transport server in the second datacenter is idle until the second datacenter is activated. In an active/active configuration, Edge Transport servers in both datacenters are delivering mail at all times.

In an active/active configuration, no steps are necessary to activate the second datacenter's Edge Transport servers because they are already running. In an active/passive configuration, the DNS MX resource record for each SMTP domain needs to be updated as part of the switchover from the primary datacenter to the standby datacenter. Although the active/active configuration provides a simple datacenter switchover solution, it has the drawback of requiring careful load monitoring to make sure that after the datacenter switchover, the Edge Transport servers in the second datacenter can provide sufficient capacity to support the increased load now flowing through it, as a result of the Edge Transport servers in the primary datacenter being unavailable.

Even with an active/active configuration, it may be appropriate to update the MX resource records for your Edge Transport servers during a datacenter switchover. Allowing the MX resource record for the failed datacenter to continue to point at the failed datacenter means that when the datacenter starts recovering, it could start experiencing connection attempts to its Edge Transport servers. This could happen while the Edge Transport services are in an unstable state (for example, because dependent services in the datacenter are being restored).

Assuming the DNS records are under the control of the organization, activating Edge Transport servers involves updating the MX resource record for each SMTP domain hosted by the server.

> **NOTE**
>
> If the MX resource record used by your organization isn't hosted by a DNS server under your organization's control, you might consider referencing a CNAME record in the MX resource record and using a CNAME record under the organization's control that can then be updated.

DNS updates enable incoming traffic, and outgoing traffic is handled by the activation of the mailbox databases in a site that has functioning Edge Transport servers:

- When incoming SMTP connections are initiated using the updated name resolution information, SMTP clients will connect to the Edge Transport servers in the second datacenter. Traffic will be appropriately routed by the Edge Transport server, and no further changes are required.
- When outgoing SMTP connections are initiated, they will try the locally available Edge Transport server, and those messages will be queued or immediately sent based on the status of the receiving server.

Restoring Service to the Primary Datacenter

Generally, datacenter failures are either temporary or permanent. With a permanent failure, such as an event that has caused the permanent destruction of a primary datacenter, there's no expectation that the primary datacenter will be activated. However, with a temporary failure (for example, an extended power loss or extensive but repairable damage), there's an expectation that the primary datacenter will eventually be restored to full service.

The process of restoring service to a previously failed datacenter is referred to as a *switchback*. The steps used to perform a datacenter switchback are similar to the steps used to perform a datacenter switchover. A significant distinction is that datacenter switchbacks are scheduled, and the duration of the outage is often much shorter.

It's important that switchback not be performed until the infrastructure dependencies for Exchange have been reactivated, are functioning and stable, and have been validated. If these dependencies aren't available or healthy, it's likely that the switchback process will cause a longer than necessary outage, and the process could fail altogether.

Mailbox Server Role Switchback

The Mailbox server role should be the first role that's switched back to the primary datacenter. The following steps detail the Mailbox server role switchback process:

1. As part of the datacenter switchover process, the Mailbox servers in the primary datacenter were put into a stopped state. When the environment (such as primary datacenter, Exchange dependencies, and wide area network (WAN) connectivity) is ready, the first step is to put the Mailbox servers in the restored primary datacenter into a started state and incorporate them into the DAG. The way in which this is done depends on whether the DAG is in DAC mode.

    a. If the DAG is in DAC mode, you can reincorporate the DAG members in the primary site by using the Start-DatabaseAvailabilityGroup cmdlet. Then, to make sure that the proper quorum model is being used by the DAG, run the Set-DatabaseAvailabilityGroup cmdlet against the DAG without specifying any parameters.

    b. If the DAG isn't in DAC mode, you can reincorporate the DAG members by using the Add-DatabaseAvailabilityGroupServer cmdlet.

2. After the Mailbox servers in the primary datacenter have been incorporated into the DAG, they will need some time to synchronize their database copies. Depending on the nature of the failure, the length of the outage, and actions taken by an administrator during the outage, this may require reseeding the database copies. For example, if during the outage, you remove the database copies from the failed primary datacenter to allow log file truncation to occur for the surviving active copies in the second datacenter, reseeding will be required. Each database can individually proceed from this point forward. After a replicated database copy in the primary datacenter is healthy, it can proceed to the next step.

    > **NOTE**
    >
    > This process doesn't require that all databases be moved at the same time. You are encouraged to move the majority of your organization's databases at one time, but some databases many linger in the second datacenter if there are issues associated with the database copies in the primary datacenter.

3. After a majority of the databases are in a healthy state in the primary datacenter, the switchback outage can be scheduled. When the scheduled time arrives, the following steps must be taken:

    a. During the datacenter switchover process, the DAG was configured to use an alternate witness server. The DAG must be reconfigured to use a witness server in the primary datacenter. If you're using the same witness server and witness directory that was used prior to the primary datacenter outage, you can run the `Set-DatabaseAvailabilityGroup -Identity DAGName` command. If you plan on using a witness server or witness directory that is different from the original witness server and directory, use the Set-DatabaseAvailabilityGroup command to configure the witness server and witness directory parameters with the appropriate values.

    b. The databases being reactivated in the primary datacenter should be dismounted in the second datacenter. You can use the Dismount-Database cmdlet to dismount the databases.

    c. After the databases have been dismounted, the URLs of the servers running Client Access services should be moved from the second datacenter to the primary datacenter. This is accomplished by changing the DNS record for the URLs to point to the Client Access services server or array in the primary datacenter. This will result in the system acting as though a database failover has occurred for each database being moved.

IMPORTANT

Don't proceed to the next step until the URLs for the servers running Client Access services have been moved and the DNS TTL and cache entries have expired. Activating the databases in the primary datacenter prior to moving the URLs to the primary datacenter will result in an invalid configuration (for example, a mounted database that has no Client Access services running in its Active Directory site).

4. Because each database in the primary datacenter is in a healthy state, it can be activated in the primary datacenter by performing database switchovers. This is accomplished by using the Move-ActiveMailboxDatabase cmdlet for each database that will be activated.

5. After each database is moved to the primary datacenter, it can be mounted by using the Mount-Database cmdlet.

After one or more databases are active and mounted in the primary datacenter, switchback procedures for the other server roles can be performed.

Client Access services switchback

As part of the switchover process, the internal and external DNS records used by clients, other servers, and IP gateways to resolve the service endpoints for Client Access services, Transport and Unified Messaging services, and Edge Transport servers were modified to point to the corresponding endpoints in the second datacenter. The switchback process for the other server roles involves modifying those records to point to the restored service endpoints in the primary datacenter.

As with the DNS changes that were made during the switchover to the second datacenter, clients, servers, and IP gateways will continue to try to connect, and should automatically connect after the TTL has expired for the original DNS entry, and after the entry is expired from their DNS cache.

Reestablishing Site Resilience

After switchback to the primary datacenter is completed successfully, you can reestablish site resilience for the primary datacenter by verifying the health and status of each mailbox database copy in the second datacenter. In addition, if any database copies in the second datacenter were originally blocked for activation, you can reconfigure those settings at this time.

Perform a server switchover

8/3/2020 • 2 minutes to read • Edit Online

A server switchover is part of preparing for a scheduled outage for the current Mailbox server.

What do you need to know before you begin?

- Estimated time to complete: 30 seconds per database
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Database availability groups" entry in the High availability and site resilience permissions topic.

> **TIP**
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the EAC to perform a server switchover

You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the"Mailbox database copies" entry in the High availability and site resilience permissions topic.

1. In the EAC, go to **Servers** > **servers**.
2. Select the Mailbox server you want to switchover.
3. In the details pane, select **Server Switchover**.
4. On the **Server Switchover** page, do one of the following:
5. Accept the default setting of **Automatically choose a target server** (in which case, the system automatically selects the best Mailbox server for each database being switched over), and then click **save**.
6. Click **Use the specified server as the target for switchover**, click **Browse** to select a Mailbox server, and then click **save**.
7. When the switchover has completed, click **close** to exit the **Server Switchover** page.

Use the Exchange Management Shell to perform a server switchover

This example performs a server switchover for the server MBX1. The system automatically selects the best Mailbox server for each active database on MBX1.

```
Move-ActiveMailboxDatabase -Server MBX1
```

This example performs a server switchover of the Mailbox server MBX4. When the command completes, MBX5 hosts the active copy of the databases that were previously active on MBX4.

```
Move-ActiveMailboxDatabase -Server MBX4 -ActivateOnServer MBX5
```

For detailed syntax and parameter information, see Move-ActiveMailboxDatabase.

Backup, restore, and disaster recovery

8/3/2020 • 12 minutes to read • Edit Online

Data protection planning is a complex process that relies on many decisions that you make during the planning phase of your deployment. As part of your planning, it's important that you understand the ways in which data can be protected, and to determine which method best suits your organization's needs.

Traditionally, backups have been used for the following scenarios:

- **Disaster recovery**: In the event of a hardware or software failure, multiple database copies in a DAG enable high availability with fast failover and little or no data loss. This eliminates downtime and the resulting lost productivity that's a significant cost of recovering from a past point-in-time backup to disk or tape. DAGs can be extended to multiple sites and can provide resilience against disk, server, network, and datacenter failures.

- **Recovery of accidentally deleted items**: Historically, in a situation where a user deleted items that later needed to be recovered, it involved finding the backup media on which the data that needed to be recovered was stored, and then somehow obtaining the desired items and providing them to the user. With the Recoverable Items folder in Exchange 2016 and Exchange 2019, and the Hold Policy that can be applied to it, it's possible to retain all deleted and modified data for a specified period of time, so recovery of these items is easier and faster. This reduces the burden on Exchange administrators and the IT help desk by enabling end users to recover accidentally deleted items themselves, thereby reducing the complexity and administrative costs associated with single item recovery. For more information, see Messaging policy and compliance in Exchange Server and Data loss prevention in Exchange Server.

- **Long-term data storage**: Backups have also been used as an archive, and typically tape is used to preserve point-in-time snapshots of data for extended periods of time as governed by compliance requirements. The new archiving, multiple-mailbox search, and message retention features in Exchange Server provide a mechanism to efficiently preserve data in an end-user accessible manner for extended periods of time. This eliminates expensive restores from tape, and increases productivity. For more information, see In-Place Archiving in Exchange Server, In-Place eDiscovery in Exchange Server, and In-Place Hold and Litigation Hold in Exchange Server.

- **Point-in-time database snapshot**: If a past point-in-time copy of mailbox data is a requirement for your organization, Exchange provides the ability to create a lagged database copy in a DAG environment. This can be useful in the rare event that store logical corruption replicates to multiple database copies in the DAG, resulting in a need to return to a previous point in time. It may also be useful if an administrator accidentally deletes mailboxes or user data. Recovery from a lagged copy can be faster than restoring from a backup because lagged copies don't require a time-consuming copy process from the backup server to the Exchange server. This can significantly lower total cost of ownership by reducing downtime.

Because there are native Exchange Server features that meet each of these scenarios in an efficient and cost effective manner, you may be able to reduce or eliminate the use of traditional backups in your environment.

Exchange Native Data Protection

Microsoft's preferred architecture for Exchange Server 2016 and Exchange Server 2019 leverages a concept known as Exchange Native Data Protection. Exchange Native Data Protection relies on built-in Exchange features to protect your mailbox data, without the use of backups (although you can still use those features and make backups). Exchange 2016 and Exchange 2019 include several features that, when deployed and configured correctly, can provide native data protection that eliminates the need to make traditional backups of your data. Using the high availability features built into Exchange Server to minimize downtime and data loss in the event of a disaster can also reduce the total cost of ownership of the messaging system. By combining these features with other built-in features, such as Legal Hold, you can reduce or eliminate your use of traditional point-in-time backups and reduce the associated costs.

In addition to determining whether Exchange Server enables you to move away from traditional point-in-time backups, we recommend that you evaluate the cost of your current backup infrastructure. Consider the cost of end-user downtime and data loss when attempting to recover from a disaster using your existing backup infrastructure. Also, include hardware, installation, and license costs, as well as the management cost associated with recovering data and maintaining the backups. Depending on the requirements of your organization, it's quite likely that a pure Exchange 2016 or Exchange 2019 environment with at least three mailbox database copies will provide lower total cost of ownership than one with backups.

There are several issues that you should consider before using the features built into Exchange Server as a replacement for traditional backups. There may also be considerations unique to your organization. Consider the following issues, and note that this isn't an exhaustive list:

- You should determine how many copies of the database need to be deployed. We strongly recommend deploying a minimum of three (non-lagged) copies of a mailbox database before eliminating traditional forms of protection for the database, such as Redundant Array of Independent Disks (RAID) or traditional VSS-based backups.

- You should clearly define the recovery time objective and recovery point objective goals, and you should establish that using a combined set of built-in features in lieu of traditional backups to enable you to meet these goals.

- You should determine how many copies of each database are needed to cover the various failure scenarios against which your system is designed to protect.

- You should determine whether eliminating the use of a DAG or some of its members captures sufficient costs to support a traditional backup solution. If so, you should determine whether that solution improves your recovery time objective or recovery point objective service level agreements (SLAs).

- You should determine whether you can afford to lose a point-in-time copy if the DAG member hosting the copy experiences a failure that affects the copy or the integrity of the copy.

- Exchange Server allows you to deploy much larger mailboxes, with a recommended maximum mailbox database size of 2 terabytes (when two or more highly available mailbox database copies are being used). Based on the larger mailboxes that most organizations are likely to deploy, you should determine your recovery point objective if you have to replay a large number of log files when activating a database copy or a lagged database copy.

- You should determine how you'll detect and prevent logical corruption in an active database copy from replicating to the passive copies of the database. This includes determining the recovery plan for this situation and how frequently this scenario has occurred in the past. If logical corruption occurs frequently in your organization, we recommend that you factor that scenario into your design by using one or more lagged copies, with a sufficient replay lag window to allow you to detect and act on logical corruption when it occurs, but before that corruption is replicated to other database copies.

One of the functions performed at the end of a successful full or incremental backup is the truncation of transaction log files that are no longer needed for database recovery. If backups aren't being taken, log truncation won't occur. To prevent a buildup of log files, you enable circular logging for your replicated databases. When you combine circular logging with continuous replication, you have a new type of circular logging called continuous replication circular logging (CRCL), which is different from Extensible Storage Engine (ESE) circular logging. Whereas ESE circular logging is performed and managed by the Microsoft Exchange Information Store service, CRCL is performed and managed by the Microsoft Exchange Replication service. When enabled, ESE circular logging doesn't generate additional log files and instead overwrites the current log file when needed. However, in a continuous replication environment, log files are needed for log shipping and replay. As a result, when you enable CRCL, the current log file isn't overwritten and closed log files are generated for the log shipping and replay process.

Specifically, the Microsoft Exchange Replication service manages CRCL so that log continuity is maintained and logs aren't deleted if they're still needed for replication. The Microsoft Exchange Replication service and the Microsoft Exchange Information Store service communicate by using remote procedure calls (RPCs) regarding which log files can be deleted.

For truncation to occur on highly available (non-lagged) mailbox database copies, the following must be true:

- The log file has been backed up, or CRCL is enabled.
- The log file is below the checkpoint.
- The other non-lagged copies of the database agree with deletion.
- The log file has been inspected by all lagged copies of the database.

For truncation to occur on lagged database copies, the following must be true:

- The log file is below the checkpoint.
- The log file is older than ReplayLagTime + TruncationLagTime.
- The log file is deleted on the active copy of the database.

Supported Backup Technologies

Exchange Server supports only Exchange-aware, VSS-based backups. Exchange Server includes a plug-in for Windows Server Backup that enables you to make and restore VSS-based backups of Exchange data. To back up and restore Exchange Server, you must use an Exchange-aware application that supports the VSS writer for Exchange Server, such as Windows Server Backup (with the VSS plug-in), Microsoft System Center 2012 - Data Protection Manager, or a third-party Exchange-aware VSS-based application.

For detailed steps about how to back up and restore Exchange data using Windows Server Backup, see Using Windows Server Backup to back up and restore Exchange data.

Exchange Server VSS Writer

Earlier versions of Exchange included two VSS writers: one inside the Microsoft Exchange Information Store service (store.exe) and one inside the Microsoft Exchange Replication service (msexchangerepl.exe). Back in Exchange 2013, the VSS writer functionality previously found in the Microsoft Exchange Information Store service was moved to the Microsoft Exchange Replication service. This architecture remains the same in Exchange 2016 and Exchange 2019. This writer, named Microsoft Exchange Writer, is used by Exchange-aware VSS-based applications to back up active and passive database copies, and to restore backed up database copies. Although the writer runs in the Microsoft Exchange Replication service, it requires the Microsoft Exchange Information Store service to be running for the writer to be advertised. As a result, both services are required to back up or restore Exchange databases.

Exchange Server Recovery

Almost all of the configuration settings for Mailbox servers and Client Access services are stored in Active Directory. As with previous versions of Exchange, Exchange 2016 and Exchange 2019 include a Setup parameter for recovering lost servers. This parameter, */m:RecoverServer*, is used to rebuild and re-create a lost server by using the settings and configuration information stored in Active Directory. However, be aware that there are several settings which are not restored, such as changes to local web.config and other configuration files. In addition, custom registry entries are not restored. We recommend that you use a reliable change management process to track and recreate these changes.

For detailed steps about how to perform a server recovery of a lost Exchange server, see Recover an Exchange Server. For detailed steps about how to recover a lost server that's a member of a database availability group (DAG), see Recover a database availability group member server.

Unified Contact Store Recovery

When Microsoft Lync Server 2013 or Skype for Business Server 2015 is used in an Exchange 2016 or Exchange 2019 environment, the user's Lync/Skype for Business contact information is stored in a special contact folder in the user's mailbox. This is referred to as the unified contact store (UCS). If you restore a UCS-migrated mailbox, the instant messaging contact list for the target user may be affected. If the user was migrated after the last backup, restoring the mailbox will result in a complete loss of the user's contact list. In less severe cases, modifications to the contact list made by the user since the last backup will be lost. To mitigate this potential data loss, ensure the user is migrated back to the instant messaging server prior to restoring the mailbox.

Recovery Database

A recovery database is a special kind of mailbox database that allows you to mount a restored mailbox database and extract data from the restored database as part of a recovery operation. You can use the New-MailboxRestoreRequest cmdlet to extract data from a recovery database. After extraction, the data can be exported to a folder or merged into an existing mailbox. Recovery databases enable you to recover data from a backup or copy of a database without disturbing user access to current data.

Using a recovery database for a Mailbox database from any previous version of Exchange isn't supported. In addition, the target mailbox used for data merges and extraction must be in the same Active Directory forest as the database mounted in the recovery database.

For more information, see Recovery databases. For detailed steps about how to create a recovery database, see Create a recovery database. For detailed steps about how to use a recovery database, see Restore data using a recovery database.

Database Portability

Database portability is a feature that enables an Exchange mailbox database to be moved to and mounted on any other Exchange Mailbox server in the same organization. By using database portability, reliability is improved by removing several error-prone, manual steps from the recovery processes. In addition, database portability reduces the overall recovery times for various failure scenarios.

For detailed steps to use database portability, see Move a mailbox database using database portability.

Dial Tone Portability

Dial tone portability is a feature that provides a limited business continuity solution for failures that affect a mailbox database, a server, or an entire site. Dial tone portability enables a user to have a temporary mailbox for sending and receiving e-mail while the original mailbox is being restored or repaired. The temporary mailbox can be on the same Exchange Mailbox server or on any other Exchange Mailbox server in your organization. This allows an alternative server to host the mailboxes of users who were previously on a server that's no longer available. Clients that support Autodiscover, such as Microsoft Outlook, are automatically redirected to the new server without having to manually update the user's desktop profile. After the user's original mailbox data has been restored, an administrator can merge a user's recovered mailbox and the user's dial tone mailbox into a single, up-to-date mailbox.

The process for using dial tone portability is called a *dial tone recovery*. A dial tone recovery involves creating an empty database on a Mailbox server to replace a failed database. This empty database, referred to as a *dial tone database*, allows users to send and receive e-mail while the failed database is recovered. After the failed database is recovered, the dial done database and the recovered database are swapped, and then the data from the dial tone database is merged into the recovered database.

For more information, see Dial tone portability. For detailed steps to perform a dial tone recovery, see Perform a dial tone recovery.

Using Windows Server Backup to back up and restore Exchange data

8/3/2020 • 3 minutes to read • Edit Online

Microsoft's preferred architecture for Exchange Server leverages a concept known as Exchange Native Data Protection. Exchange Native Data Protection relies on native Exchange features to protect your mailbox data, without the use of traditional backups. But if you want to create backups, Exchange includes a plug-in for Windows Server Backup (WSB) that enables you to create Exchange-aware Volume Shadow Copy Service (VSS)-based backups of Exchange data. To take Exchange-aware backups, you must have the WSB feature installed.

The plug-in, WSBExchange.exe, runs as a service named Microsoft Exchange Server Extension for Windows Server Backup (the short name for this service is WSBExchange). This service is automatically installed and configured for manual startup on all Mailbox servers. The plug-in enables WSB to create Exchange-aware VSS backups.

Before using WSB to back up Exchange data, we recommend that you familiarize yourself with the following features and options for the plug-in:

- Backups taken with WSB occur at the volume level, and the only way to perform an application-level backup or restore is to select an entire volume. To back up a database and its log stream, you must back up the entire volume containing the database and logs, not just the individual folders. You can't back up any data without backing up the entire volume containing the data.

- The backup must be run locally on the server being backed up, and you can't use the plug-in to take remote VSS backups. There is no remote administration of WSB or the plug-in. You can, however, use Remote Desktop Services or Terminal Services to remotely manage backups.

- The backup can be created on a local drive or on a remote network share.

- Only full backups should be taken. Log truncation will occur only after a successful completion of a VSS full backup of a volume or folders containing an Exchange database.

- When restoring data, it's possible to restore only Exchange data. This data can be restored to its original location or to an alternate location. If you restore the data to its original location, WSB and the plug-in automatically handle the recovery process, including dismounting any existing database and replaying logs into the restored database.

- The restore process doesn't support the Exchange recovery database (RDB). If you want to use an RDB, you must restore the data to an alternate location and then manually copy or move the restored data from that location into the RDB folder structure.

- When restoring Exchange data, all backed up databases must be restored together. You can't restore a single database.

- Bare metal restores are supported when using WSB; however, the recommended recovery approach for Exchange servers is to recover the Exchange server and then restore the data. If you're using a third-party backup application (e.g., non-Microsoft), then support for bare metal restores of Exchange may be available from your backup application vendor.

The following table describes the supportability of the backup and recovery options available for Exchange Server with WSB.

IF YOU...	THEN...
Back up the full server...	A VSS copy backup will be performed, and the transaction logs for the databases on the server will not be truncated.
Perform a custom backup and select one or more volumes to back up...	A VSS full backup can be selected, allowing the transaction logs for the databases on the selected volumes to be truncated at the completion of a successful backup.
Perform a custom backup and select one or more folders to back up...	A VSS full backup can be selected and the log files will be truncated; however, restoration of the backup will be limited to file restore, as an Application level restore will not be available as an option.

For detailed steps to back up Exchange using WSB, see Use Windows Server Backup to back up Exchange.

For detailed steps to restore data from a backup taken with WSB, see Use Windows Server Backup to restore a backup of Exchange.

Use Windows Server Backup to back up Exchange

8/3/2020 • 3 minutes to read • Edit Online

You can use Windows Server Backup to back up and restore Exchange databases. Exchange includes a plug-in for Windows Server Backup that allows you to make Volume Shadow Copy Service (VSS)-based backups of Exchange data.

What do you need to know before you begin?

- Estimated time to complete: 1 minute, plus the time it takes to back up the data
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Mailbox recovery" entry in the Recipients Permissions topic.
- The Windows Server Backup feature must be installed on the local computer.
- During the backup operation, a consistency check of the Exchange data files is run to make sure that the files are in a good state and can be used for recovery. If the consistency check succeeds, Exchange data is available for recovery from that backup. If the consistency check fails, the Exchange data isn't available for recovery. Windows Server Backup runs the consistency check on the snapshot taken for the backup. As a result, before copying files from the snapshot to backup media, the consistency of the backup is known, and the user is notified of the consistency check results.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use Windows Server Backup to back up Exchange

1. Start Windows Server Backup.
2. Select **Local Backup**.
3. In the Actions pane, click **Backup Once...** to start the Backup Once Wizard.
4. On the **Backup Options** page, select **Different options**, and then click **Next**.
5. On the **Select Backup Configuration** page, select **Custom**, and then click **Next**.
6. On the **Select Items for Backup** page, click **Add Items** to select the volume(s) to be backed up, and then click **OK**.

   > **NOTE**
   >
   > Choose volumes and not individual folders. The only way to perform an application-level backup or restore is to select an entire volume.

7. Click **Advanced Settings**. On the **Exclusions** tab, click **Add Exclusion** to add any files or file types you want to exclude from the backup.

NOTE

By default, volumes that contain operating system components or applications are included in the backup and can't be excluded.

8. On the **VSS Settings tab**, select **VSS full Backup**, and then click **OK**, and then click **Next**.

9. On the **Specify Destination Type** page, select the location where you want to store the backup, and then click **Next**.

   - If you choose **Local drives**, the **Select Backup Destination** page appears. Select an option from the **Backup destination** dropdown, and then click **Next**.

   - If you choose **Remote shared folder**, the **Specify remote folder** page appears. Specify a UNC path for the backup files, configure access control settings. Choose **Do not inherit** if you want the backup to be accessible only through a specific account. Provide a username and password for an account that has write permissions on the computer hosting the remote folder, and then click **OK**. Alternatively, choose **Inherit** if you want the backup to be accessible by everyone who has access to the remote folder. Click **Next**.

10. On the **Confirmation** page, review the backup settings, and then click **Backup**.

11. On the **Backup Progress** page, you can view the status and progress of the backup operation.

12. Click **Close** to exit the **Backup Progress** page at any time. Any backup in progress will continue to run in the background.

How do you know this worked?

To verify that you've successfully backed up the data, do any of the following:

- On the server on which Windows Server Backup was run, the last backup status will be displayed, which should say Successful. You can also verify that the backup completed successfully by viewing the Windows Server Backup logs.

- Open Event Viewer and verify that a backup completion event was logged in the Application event log.

- Run the following command in the Exchange Management Shell to verify that each database on the selected volume(s) was backed up successfully:

  ```
  Get-MailboxDatabase -Server <ServerName> -Status | Format-List Name,*FullBackup
  ```

  The *SnapshotLastFullBackup* and *LastFullBackup* properties of the database indicate when the last successful backup was taken, and if it was a VSS full backup.

Use Windows Server Backup to restore a backup of Exchange

8/3/2020 • 3 minutes to read • Edit Online

You can use Windows Server Backup to back up and restore Exchange databases. Exchange includes a plug-in for Windows Server Backup that allows you to make and restore Volume Shadow Copy Service (VSS)-based backups of Exchange data. For additional information, see Using Windows Server Backup to back up and restore Exchange data.

What do you need to know before you begin?

- Estimated time to complete: 1 minute, plus the time it takes to restore the data
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Mailbox recovery" entry in the Recipients Permissions topic.
- The Windows Server Backup feature must be installed on the local computer.
- When restoring a database to its original location, the database can remain in a dirty shutdown state and be mountable by the system. When restoring to an alternate location (for example, in preparation to use a recovery database), the database must be manually brought into a clean shutdown state by using Exchange Server Database Utilities (Eseutil.exe).

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use Windows Server Backup to restore a backup of Exchange

1. Start Windows Server Backup.
2. Select **Local Backup**.
3. In the Actions pane, click **Recover...** to start the Recovery Wizard.
4. On the **Getting Started** page, do either of the following:
   - If the data being recovered was backed up on the local server, select **This server (ServerName)**, and then click **Next**.
   - If the data being recovered is from another server, or if the backup being recovered is located on another computer, select **Another server**, and then click **Next**. On the **Specify location type** page, select **Local drives** or **Remote shared folder**, and then click **Next**. If you select **Local drives**, select the drive containing the backup on the **Select backup location** page, and then click **Next**. If you select **Remote shared folder**, enter the UNC path for the backup data on the **Specify remote folder** page, and then click **Next**.
5. On the **Select Backup Date** page, select the date and time of the backup that you want to recover, and then click **Next**.
6. On the **Select Recovery Type** page, select **Applications**, and then click **Next**.

> **NOTE**
>
> If **Applications** is not available as a selection, it indicates that the backup selected for restore was a folder-level backup, and not a volume level backup. You must perform backups at the volume level when backing up Exchange data with Windows Server Backup.

7. On the **Select Application** page, verify that Exchange is selected in the **Applications** field. Click **View Details** to view the application components of the backups. If the backup that you're recovering is the most recent, the **Do not perform a roll-forward recovery of the application database** check box is displayed. Select this check box if you want to prevent Windows Server Backup from rolling forward the database being recovered by committing all uncommitted transaction logs. Click **Next**.

8. On the **Specify Recovery Options** page, specify where you want to recover the data, and then click **Next**:

   - Choose **Recover to original location** if you want to restore the Exchange data directly to its original location. If you use this option, you can't choose which databases are restored; all backed up databases on the volume will be restored to their original locations.

   - Choose **Recover to another location** if you want to restore individual databases and their files to a specified location. Click **Browse** to specify the alternate location. If you use this option, you can choose which databases are restored. After being restored, the data files can then be moved into a recovery database, manually moved back to their original location, or mounted somewhere else in the Exchange organization using database portability. When you restore a database to an alternate location, the restored database will be in a dirty shutdown state. After the restore process has completed, you will need to manually put the database into a clean shutdown state using Eseutil.exe.

9. On the **Confirmation** page, review the recovery settings, and then click **Recover**.

10. On the **Recovery Progress** page, you can view the status and progress of the recovery operation.

11. Click **Close** when the recovery operation has completed.

How do you know this worked?

The **Recovery Progress** page will indicate whether or not the recovery process completed successfully. To further verify that you've successfully restored the data, do any of the following:

- Examine the target directory of the backup and verify that the restored data exists.

- On the server on which Windows Server Backup was run, verify that the job completed successfully by viewing the backup logs.

- Open Event Viewer and verify that a restore completion event was logged in the Application event log.

Recover Exchange servers

8/3/2020 • 5 minutes to read • Edit Online

You can recover a lost Exchange server by using the */Mode:RecoverServer* switch in unattended mode (from the command line) of Exchange Setup. Since most Exchange server settings are stored in Active Directory, the `Setup.exe /Mode:RecoverServer` command uses that information during the installation of Exchange on a new server with the same name.

Recovering a lost Exchange server is often accomplished by using new hardware. However, you can also use an existing server that doesn't already have Exchange installed on it.

This topic shows you how to recover a lost Exchange server that isn't a member of a database availability group (DAG). For detailed steps about how to recover a server that was a member of a DAG, see Recover a database availability group member server.

Looking for other management tasks related to backing up and restoring data? Check out Backup, restore, and disaster recovery.

What do you need to know before you begin?

- Estimated time to complete: 20 minutes
- The account that you'll use to do the server recovery requires the following permissions:
  - Domain Admins security group membership.
  - Exchange Organization Management role group membership.
- If Exchange is installed in a location other than the default location of %ProgramFiles%\Microsoft\Exchange Server\V15, you must include the */TargetDir:<Path>* switch in the `Setup.exe /Mode:RecoverServer` command to specify the location of the Exchange program (binary) files. If you don't use the */TargetDir* switch, the Exchange files will be installed in the default location when you recover the Exchange server.

  To find the install location of Exchange on the lost Exchange server, do the following steps:

  1. Open ADSIEDIT.MSC or LDP.EXE.
  2. Go 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**
  3. Right-click the Exchange server object, and then click **Properties**.
  4. Find the **msExchInstallPath** attribute. This attribute stores the current installation path.
- If you do not have the installation media for the Cumulative Update (CU) version that was installed on the server to be recovered, you can recover a server using the latest available Cumulative Update. Only the last two CUs are available for download. For more information, see Updates for Exchange Server.
- The target server must use the same version of Windows Server as the lost server. For example, you can't recover a lost Exchange 2016 server that was running Windows 2012 R2 on a new server that's running Windows 2016, or vice-versa.
- The same disk drive letters that were used for mounted databases on the lost server must also exist on the target server.

- The target server should have the same general performance characteristics and hardware configuration as the lost server.

- The */Mode:RecoverServer* switch assigns a self-signed certificate to all Exchange Services that require SSL/TLS. If the server previously used an SSL/TLS certificate that was issued by a different certification authority, you'll need to re-import the certificate and configure the services to use the certificate. Otherwise, users will get a certificate prompt when they try to connect (for example, in Outlook).

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at Exchange Server.

Recover a Lost Exchange Server

1. Reset the computer account for the lost server. For detailed steps, see Reset a Computer Account.

2. Install the proper operating system and name the new server with the same name as the lost server. Recovery won't succeed if the target Windows server doesn't have the same name as the lost Exchange server.

3. Join the server to the same domain as the lost server.

4. Install the necessary prerequisites and operating system components on the target server. For details, see Exchange Server system requirements.

5. On the target server, open File Explorer, right-click on the Exchange ISO image file that you downloaded, and then select **Mount**. Note the virtual DVD drive letter that's assigned.

6. Open a Windows Command Prompt window. For example:

   - Press the Windows key + 'R' to open the **Run** dialog, type cmd.exe, and then press **OK**.

   - Press **Start**. In the **Search** box, type **Command Prompt**, then in the list of results, select **Command Prompt**.

7. In the Command Prompt window, use the following syntax:

   ```
   <Virtual DVD drive letter>:\Setup.exe /IAcceptExchangeServerLicenseTerms /Mode:RecoverServer
   [/TargetDir:<Path>] [/DomainController:<ServerNameOrFQDN>] [/DoNotStartTransport]
   [/EnableErrorReporting]
   ```

   This example uses the Exchange installation files on drive E: to install Exchange in the default location (%ProgramFiles%\Microsoft\Exchange Server\V15) and recover the Exchange server.

   ```
   E:\Setup.exe /IAcceptExchangeServerLicenseTerms /Mode:RecoverServer
   ```

   This is the same example, but a custom location for the Exchange program files is required to match the location on the lost server.

   ```
   E:\Setup.exe /IAcceptExchangeServerLicenseTerms /Mode:RecoverServer /TargetDir:"D:\Program
   Files\Exchange"
   ```

   For more information about the optional switches, see Use unattended mode in Exchange Setup.

8. After Setup has completed, but before you put the recovered server into production, reconfigure any

custom settings that were previously present on the server, and then restart the server.

How do you know this worked?

The successful completion of Setup will be the primary indicator that the recovery was successful. To further verify that you've successfully recovered a lost server, open the Windows Services tool (services.msc) and verify that the Microsoft Exchange services have been installed and are running.

Possible issues with the Scripting Agent

If you previously enabled the Scripting Agent in your Exchange organization, the recovery process might fail. The error will look like this:

```
"Initialization failed: '"Scripting Agent initialization failed: "File is not found: 'C:\Program
Files\Microsoft\Exchange Server\V15\Bin\CmdletExtensionAgents\ScriptingAgentConfig.xml'.""' --->
Microsoft.Exchange.Provisioning.ProvisioningException: "Scripting Agent initialization failed: "File is not
found: 'C:\Program Files\Microsoft\Exchange Server\V15\Bin\CmdletExtensionAgents\ScriptingAgentConfig.xml'.""
---> System.IO.FileNotFoundException: "File is not found: 'C:\Program Files\Microsoft\Exchange
Server\V15\Bin\CmdletExtensionAgents\ScriptingAgentConfig.xml'."
```

If you have other Exchange servers in your organization, you'll need to:

1. Disable the Scripting Agent in the Exchange Management Shell on an existing server:

   ```
   Disable-CmdletExtensionAgent -Identity "Scripting Agent"
   ```

2. Run Exchange Setup in recovery mode as described earlier in this topic.

3. Enable the Scripting Agent in the Exchange Management Shell after the Exchange server recovery is complete:

   ```
   Enable-CmdletExtensionAgent -Identity "Scripting Agent"
   ```

If the recovered Exchange server is the only Exchange server in your organization, you'll need to:

1. Rename the file %ExchangeInstallPath%Bin\CmdletExtensionAgents\ScriptingAgentConfig.**xml.sample** to %ExchangeInstallPath%Bin\CmdletExtensionAgents\ScriptingAgentConfig.**xml**.

   The default value of %ExchangeInstallationPath% is %ProgramFiles%\Microsoft\Exchange Server\V15, but the actual value is wherever you installed Exchange on the server.

2. Re-run Exchange Setup in recovery mode as described earlier in this topic.

Recover a database availability group member server

8/3/2020 • 3 minutes to read • Edit Online

If a Mailbox server that's a member of a database availability group (DAG) is lost or fails, and is unrecoverable and needs replacement, you can perform a server recovery operation.

Microsoft Exchange Server Setup includes the switch */m:RecoverServer* that can be used to perform the server recovery operation. Running Setup with the */m:RecoverServer* switch causes Setup to read the server's configuration information from Active Directory for a server with the same name as the server from which you're running Setup.

After the server's configuration information is gathered from Active Directory, the original Exchange files and services are then installed on the server, and the roles and settings that were stored in Active Directory are then applied to the server.

Looking for other management tasks related to DAGs? Check out Manage database availability groups.

What do you need to know before you begin?

- Estimated time to complete: 30 minutes
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Mailbox database copies" entry in the High availability and site resilience permissions topic.
- If Exchange is installed in a location other than the default location, you must use the */TargetDir* Setup switch to specify the location of the Exchange program files. If you don't use the */TargetDir* switch, the Exchange program files will be installed in the default location (%programfiles%\Microsoft\Exchange Server\V15).

  To determine the install location, follow these steps:

  1. Open ADSIEDIT.MSC or LDP.EXE.
  2. Navigate to the following location: **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**
  3. Right-click the Exchange server object, and then click **Properties**.
  4. Locate the **msExchInstallPath** attribute. This attribute stores the current installation path.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use Setup /m:RecoverServer to recover a server

1. Retrieve any replay lag or truncation lag settings for any mailbox database copies that exist on the server

being recovered by using the Get-MailboxDatabase cmdlet:

```
Get-MailboxDatabase DB1 | Format-List *lag*
```

2. Remove any mailbox database copies that exist on the server being recovered by using the Remove-MailboxDatabaseCopy cmdlet:

```
Remove-MailboxDatabaseCopy DB1\MBX1
```

3. Remove the failed server's configuration from the DAG by using the Remove-DatabaseAvailabilityGroupServer cmdlet:

```
Remove-DatabaseAvailabilityGroupServer -Identity DAG1 -MailboxServer MBX1
```

> **NOTE**
>
> If the DAG member being removed is offline and can't be brought online, you must add the `-ConfigurationOnly` parameter to the preceding command. If you use the `-ConfigurationOnly` switch, you must also manually evict the node from the cluster.

4. Reset the server's computer account in Active Directory. For detailed steps, see Reset a Computer Account.

5. Open a Command Prompt window. Using the original Setup media, run the following command:

```
Setup /m:RecoverServer
```

6. When the Setup recovery process is complete, add the recovered server to the DAG by using the Add-DatabaseAvailabilityGroupServer cmdlet:

```
Add-DatabaseAvailabilityGroupServer -Identity DAG1 -MailboxServer MBX1
```

7. After the server has been added back to the DAG, you can reconfigure mailbox database copies by using the Add-MailboxDatabaseCopy cmdlet. If any of the database copies being added previously had replay lag or truncation lag times greater than 0, you can use the *ReplayLagTime* and *TruncationLagTime* parameters of the Add-MailboxDatabaseCopy cmdlet to reconfigure those settings:

```
Add-MailboxDatabaseCopy -Identity DB1 -MailboxServer MBX1
Add-MailboxDatabaseCopy -Identity DB2 -MailboxServer MBX1 -ReplayLagTime 3.00:00:00
Add-MailboxDatabaseCopy -Identity DB3 -MailboxServer MBX1 -ReplayLagTime 3.00:00:00 -TruncationLagTime
3.00:00:00
```

How do you know this worked?

To verify that you've successfully recovered the DAG member, do the following:

- In the Exchange Management Shell, run the following command to verify the health and status of the recovered DAG member.

```
Test-ReplicationHealth <ServerName>
```

```
Get-MailboxDatabaseCopyStatus -Server <ServerName>
```

All of the replication health tests should pass successfully, and the status of databases and their content indexes should be healthy.

Recovery databases

8/3/2020 • 4 minutes to read • Edit Online

A recovery database (RDB) is a special kind of mailbox database that allows you to mount a restored mailbox database and extract data from the restored database as part of a recovery operation. You can use the New-MailboxRestoreRequest cmdlet to extract data from an RDB. After extraction, the data can be exported to a folder or merged into an existing mailbox. RDBs enable you to recover data from a backup or copy of a database without disturbing user access to current data.

Microsoft Exchange Server supports the ability to restore data directly to a recovery database. Mounting the recovered data as a recovery database allows the administrator to restore individual mailboxes or individual items in a mailbox. Restoring to a recovery database can be accomplished in two ways:

- If a recovery database already exists, the application can dismount the database, restore the data onto the recovery database and log files, and then remount the database.
- The database and log files can be restored to any disk location. Exchange analyzes the restored data and replays the transaction logs to bring the databases up to date, and then a recovery database can be configured to point to already recovered database files.

Difference between a mailbox database and a recovery database

RDBs are different from standard mailbox databases in several respects:

- An RDB is created by using the Exchange Management Shell.
- Mail can't be sent to or from an RDB. All client protocol access to an RDB (including SMTP, POP3, and IMAP4) is blocked. This design prevents using an RDB to insert mail into or remove mail from the messaging system.
- Client MAPI access using Microsoft Outlook or Outlook on the web is blocked. MAPI access is supported for an RDB, but only by recovery tools and applications. Both the mailbox GUID and the database GUID must be specified when using MAPI to log into a mailbox in an RDB.
- Mailboxes in an RDB can't be connected to user accounts. To allow a user to access the data in a mailbox in an RDB, the mailbox must be merged into an existing mailbox, or exported to a folder.
- System and mailbox management policies aren't applied. This design prevents items in an RDB from being deleted by the system during the recovery process.
- Online maintenance isn't performed for RDBs.
- Circular logging can't be enabled for RDBs.
- Only one RDB can be mounted at any time on a Mailbox server. The use of an RDB doesn't count against the database limit per Mailbox server.
- You can't create mailbox database copies of an RDB.
- An RDB can be used as a target for restore operations, but not backup operations.
- A recovered database mounted as an RDB isn't tied to the original mailbox in any way.

Using a recovery database

Before you can use an RDB, there are certain requirements that must be met. An RDB can be used for Exchange 2016 and later mailbox databases only. Mailbox databases from previous versions of Exchange aren't supported. In addition, the target mailbox used for data merges and extraction must be in the same Active Directory forest as the database mounted in the RDB.

An RDB can be used to recover data in several situations, such as:

- **Same server dial tone recovery**: You can perform a recovery from an RDB after the original database has been restored from backup, as part of a dial tone recovery operation.
- **Alternate server dial tone recovery**: You can use an alternate server to host the dial tone database, and then later recover data from an RDB after the original database has been restored from backup.
- **Mailbox recovery**: You can recover an individual mailbox from backup when the deleted mailbox retention period has elapsed. You then extract data from the restored mailbox and copy it to a target folder or merge it with another mailbox.
- **Specific item recovery**: You can restore from backup data that has been deleted or purged from a mailbox.

> **NOTE**
>
> Folder access control lists (ACLs) aren't preserved when recovering content into an active mailbox. Because the recovery process typically involves recovering mailbox data and merging the content back into the original database, there should be no need to recover or copy ACLs.

An RDB is designed for mailbox database recovery under the following conditions and scenarios:

- The logical information about the original database and the mailboxes in that database remains intact and unchanged in Active Directory.
- You need to recover a single mailbox or a single database. Recovery scenarios include:
  - Recovering or repairing a database while a dial tone database is in use, with the goal of merging the two databases.
  - Recovering a database on a server other than the original server for that database. If needed, you can then merge the recovered data back to the original server.
  - Recovering deleted items that users previously deleted from their mailbox, after the deleted item retention period has expired.

RDBs are generally not designed for scenarios in which you have to restore entire servers, when you have to restore multiple databases, or when you're in an emergency situation that requires changing or rebuilding your Active Directory topology.

For detailed steps about how to create an RDB, see Create a recovery database. For detailed steps about how to use an RDB, see Restore data using a recovery database.

Create a recovery database

8/3/2020 • 2 minutes to read • Edit Online

You can use the Exchange Management Shell to create a recovery database, a special kind of mailbox database that's used to mount and extract data from the restored database as part of a recovery operation. After you create a recovery database, you can move a recovered or restored mailbox database into the recovery database, and then use the New-MailboxRestoreRequest cmdlet to extract data from the recovered database. After extraction, the data can then be exported to a folder or merged into an existing mailbox. Using recovery databases, you can recover data from a backup or copy of a database without disrupting user access to current data.

Looking for other management tasks related to recovery databases? Check out Recovery databases.

What do you need to know before you begin?

- Estimated time to complete this task: 1 minute
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Mailbox recovery" entry in the Recipients Permissionstopic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to create a recovery database

This example creates the recovery database RDB1 on the Mailbox server MBX2.

```
New-MailboxDatabase -Recovery -Name RDB1 -Server MBX2
```

This example creates the recovery database RDB2 on the Mailbox server MBX1 using a custom path for the database file and log folder.

```
New-MailboxDatabase -Recovery -Name RDB2 -Server MBX1 -EdbFilePath "C:\Recovery\RDB2\RDB2.EDB" -LogFolderPath
"C:\Recovery\RDB2"
```

For detailed syntax and parameter information, see New-MailboxDatabase.

How do you know this worked?

To verify that you've successfully created a recovery database, do the following:

- In the Exchange Management Shell, run the following command to display configuration information for the recovery database.

  ```
  Get-MailboxDatabase <RecoveryDatabaseName> | Format-List
  ```

Other Tasks

After you create a recovery database, you may also want to restore data using a recovery database. For detailed steps, see Restore data using a recovery database.

Restore data using a recovery database

8/3/2020 • 3 minutes to read • Edit Online

A recovery database (RDB) is a special kind of mailbox database that allows you to mount and extract data from a restored mailbox database as part of a recovery operation. RDBs allow you to recover data from a backup or copy of a database without disrupting user access to current data.

After you create an RDB, you can restore a mailbox database into the RDB by using a backup application or by copying a database and its log files into the RDB folder structure. Then you can use the New-MailboxRestoreRequest cmdlet to extract data from the recovered database. Once extracted, the data can then be exported to a folder or merged into an existing mailbox.

For additional management tasks related to RDBs, see Recovery databases.

What do you need to know before you begin?

- Estimated time to complete this task: 1 minute, plus the time it takes to put the database into a clean shutdown state and to extract the data.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Mailbox recovery" entry in the Recipients Permissions topic.
- Some backup applications have the ability to restore Exchange data directly to a recovery database. Windows Server Backup can restore only file-level backups to a recovery database. It cannot be used to restore application-level backups to a recovery database.
- The database and log files containing the recovered data must be restored or copied into the RDB folder structure.
- The database must be in a clean shutdown state. Because an RDB is an alternate restore location for all databases, all restored databases will be in a dirty shutdown state. You must use **Eseutil /R** to put restored databases into a clean shutdown state.

Use the Exchange Management Shell to recover data using a recovery database

1. Copy a recovered database and its log files, or restore a database and it log files, to the location you will use for your recovery database.

2. Use Eseutil to bring that database into a clean shutdown state. In the following example, EXX is the log generation prefix for the database (for example, E00, E01, E02, and so on).

   ```
   Eseutil /R EXX /l <RDBLogFilePath> /d <RDBEdbFolder>
   ```

   The following example illustrates a log generation prefix of E01 and a recovery database and log file path of E:\Databases\RDB1:

   ```
   Eseutil /R E01 /l E:\Databases\RDB1 /d E:\Databases\RDB1
   ```

3. Create a recovery database. Give the recovery database a unique name, but use the name and path of the database file for the EdbFilePath parameter, and the location of the recovered log files for the LogFolderPath

parameter.

```
New-MailboxDatabase -Recovery -Name <RDBName> -Server <ServerName> -EdbFilePath <RDBPathandFileName> -
LogFolderPath <LogFilePath>
```

The following example illustrates creating a recovery database that will be used to recover DB1.edb and its log files, which are located at E:\Databases\RDB1.

```
New-MailboxDatabase -Recovery -Name <RDBName> -Server <ServerName> -EdbFilePath
"E:\Databases\RDB1\DB1.EDB" -LogFolderPath "E:\Databases\RDB1"
```

4. Restart the Microsoft Exchange Information Store service:

```
Restart-Service MSExchangeIS
```

5. Mount the recovery database:

```
Mount-database <RDBName>
```

6. Verify that the mounted database contains the mailbox(es) you want to restore:

```
Get-MailboxStatistics -Database <RDBName> | Format-Table DisplayName,MailboxGUID -AutoSize
```

7. Use the New-MailboxRestoreRequest cmdlet to restore a mailbox or items from the recovery database to a production mailbox.

   The following example restores the source mailbox that has the MailboxGUID 1d20855f-fd54-4681-98e6-e249f7326ddd on mailbox database DB1 to the target mailbox with the alias Morris.

```
New-MailboxRestoreRequest -SourceDatabase DB1 -SourceStoreMailbox 1d20855f-fd54-4681-98e6-e249f7326ddd
-TargetMailbox Morris
```

   The following example restores the content of the source mailbox that has the display name Morris Cornejo on mailbox database DB1 to the archive mailbox for Morris@contoso.com.

```
New-MaiboxRestoreRequest -SourceDatabase DB1 -SourceStoreMailbox "Morris Cornejo" -TargetMailbox
Morris@contoso.com -TargetIsArchive
```

8. Periodically check the status of the Mailbox restore request using Get-MailboxRestoreRequest.

   Once the restore has a status of Completed, remove the restore request using Remove-MailboxRestoreRequest. For example:

```
Get-MailboxRestoreRequest -Status Completed | Remove-MailboxRestoreRequest
```

How do you know this worked?

To verify that you have successfully recovered the mailbox data, open the target mailbox using Outlook or Outlook Web App and verify that the recovered data is present.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Move a mailbox database using database portability

8/3/2020 • 2 minutes to read • Edit Online

Database portability can help reduce overall recovery times for some failure scenarios. By using database portability, reliability is improved by removing several error-prone, manual steps from the recovery processes. Note that Mailbox databases from previous versions of Exchange can't be moved to a Mailbox server running Exchange 2016 or Exchange 2019.

> **NOTE**
>
> When using database portability to recover a mailbox database, the operating system version and the Exchange Server version on the source and target Exchange servers must be the same. For example, if an Exchange 2016 mailbox database was previously mounted on a server running Windows Server 2016, database portability will only work when migrating the database to a server also running Windows Server 2016 and Exchange 2016.

What do you need to know before you begin?

- Estimated time to complete: 5 minutes, plus the time it takes to restore the data, move the database files, and wait for Active Directory replication to complete.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Mailbox recovery" entry in the Recipients Permissions topic.
- You can't use the EAC to move user mailboxes to a recovered or dial tone database using database portability.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to move user mailboxes to a recovered or dial tone database using database portability

1. Verify that the database to be moved is in a clean shutdown state. If the database isn't in a clean shutdown state, perform a soft recovery.

   > **NOTE**
   >
   > When you perform a soft recovery, any uncommitted log files are committed to the database. If you don't have all of the required log files, you can't complete the soft recovery process. Proceed to step 2.

   To commit all uncommitted log files to the database, from a command prompt, run the following command.

   ```
   ESEUTIL /R <Enn>
   ```

> **NOTE**
>
> <E *nn*> specifies the log file prefix for the database into which you intend to replay the log files. The log file prefix specified by <E *nn*> is a required parameter for Eseutil /r.

2. Create a database on a server using the following syntax:

   ```
   New-MailboxDatabase -Name <DatabaseName> -Server <ServerName> -EdbFilePath <DatabaseFileNameandPath> -
   LogFolderPath <LogFilesPath>
   ```

3. Set the *This database can be over written by restore* attribute using the following syntax:

   ```
   Set-MailboxDatabase <DatabaseName> -AllowFileRestore $true
   ```

4. Move the original database files (.edb file, log files, and Exchange Search catalog) to the database folder you specified when you created the new database above.

5. Mount the database using the following syntax:

   ```
   Mount-Database <DatabaseName>
   ```

6. After the database is mounted, modify the user account settings with the Set-Mailbox cmdlet so that the account points to the mailbox on the new mailbox server. To move all of the users from the old database to the new database, use the following syntax.

   ```
   Get-Mailbox -Database <SourceDatabase> |where {$_.ObjectClass -NotMatch
   '(SystemAttendantMailbox|ExOleDbSystemMailbox)'}| Set-Mailbox -Database <TargetDatabase>
   ```

7. Trigger delivery of any messages remaining in queues using the following syntax.

   ```
   Get-Queue <QueueName> | Retry-Queue -Resubmit $true
   ```

After Active Directory replication is complete, all users can access their mailboxes on the new Exchange server. Most clients are redirected via Autodiscover. Outlook on the web users are also automatically redirected.

How do you know this worked?

To verify that you've successfully moved a mailbox, do the following:

- Open the mailbox using Outlook on the web.

- Open the mailbox using Microsoft Outlook.

Dial tone portability

8/3/2020 • 3 minutes to read • Edit Online

Dial tone portability is a feature of Exchange Server 2016 and Exchange Server 2019 that provides a limited business continuity solution for failures that affect a mailbox database, a server, or an entire site. A temporary mailbox maintains users' ability to send email, and this mailbox can be on the same Exchange Mailbox server or on any other Exchange Mailbox server in your organization, provided they contain databases with the same database schema version. This allows an alternative server to host the mailboxes of users who were previously on a server that is no longer available. Clients that support Autodiscover are automatically redirected to the new server without having to manually update the user's desktop profile. After the user's original mailbox data has been restored, an administrator can merge a user's recovered mailbox and the user's dial tone mailbox into a single, up-to-date mailbox.

The process for using dial tone portability is called a *dial tone recovery*. A dial tone recovery involves creating an empty database on a Mailbox server to replace a failed database. This empty database, referred to as a *dial tone database*, allows users to send and receive email messages while the failed database is recovered.

There are three options for performing a dial tone recovery:

- **Dial tone recovery on the server with the failed database**: If the server hosting the failed database is still functional, we recommend that you perform a dial tone recovery on that server. This means less downtime because you don't need to move database files between servers. In addition, you won't need to reconfigure messaging profiles for clients that don't support Autodiscover.
- **Dial tone recovery using an alternate server for the dial tone database**: If a server fails and needs to be rebuilt, the most efficient way to give users basic mail functionality is to create a dial tone database on another server, and use database portability to move the users' mailbox configuration to that new server. Because this process involves moving the dial tone database back to the original (recovered) server, this option adds more time to the overall recovery process. In addition, this process is more complex than performing a dial tone recovery on the original server. When performing this process, the server hosting the dial tone database must have sufficient resources to support the added load of the additional users. In addition, if the users' client doesn't support Autodiscover, their messaging profile will need to be reconfigured to point to the dial tone server.
- **Dial tone recovery using and staying on an alternate server for the dial tone database**: This is similar to the preceding option, except that you don't revert back to the original server. We recommend this option for situations in which it isn't possible or feasible to recover the failed server. In this scenario, users typically remain on an alternate server after the recovery operation has completed. When performing this process, the server hosting the dial tone database must have sufficient resources to support the added load of the additional users. In addition, if the users' client doesn't support Autodiscover, their messaging profile will need to be reconfigured to point to the dial tone server.

All three options follow the same basic steps:

1. **Create an empty dial tone database to replace the failed database.**

   This new database will allow users who had mailboxes on the failed database to send and receive new messages. Dial tone portability allows you to point a user to a different database without moving the mailbox. If you created the dial tone database on a different server than the server that housed the failed database, you need to move the mailbox configuration to that new server.

2. **Restore the old database.**

Use the backup and recovery software you typically use to restore the failed database. If there is no backup of the failed database, recover the failed database using other means if possible. If you're using the same server for dial tone recovery, you need to restore the database to a recovery database (RDB).

3. **Swap the dial tone database with the restored database.**

   After the failed database is restored, swap it with the dial tone database. This gives the users the ability to send and receive email and access all the data in the restored database. If users were moved to a dial tone database on another server, you need to move the mailbox configuration back to the original server.

4. **Merge the databases.**

   To get the data from the dial tone database into the restored database, you merge the data using the New-MailboxRestoreRequest cmdlet.

For detailed steps about how to perform a dial tone recovery, see Perform a dial tone recovery.

Perform a dial tone recovery

8/3/2020 • 3 minutes to read • Edit Online

The process for using dial tone portability is called a dial tone recovery, which involves creating an empty database on a Mailbox server to replace a failed database. To learn more, see Dial tone portability.

What do you need to know before you begin?

- Estimated time to complete: 5 minutes, plus the time it takes to restore and move the data.
- You must have fewer than the maximum number of databases deployed to create a dial tone database (a maximum of five databases per server for Exchange Standard Edition, a maximum of 100 databases per server for Exchange Enterprise Edition).
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Mailbox recovery" entry in the Recipients Permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to perform a dial tone recovery on a single server

> **NOTE**
>
> You can't use the EAC to perform a dial tone recovery on a single server.

1. Make sure that any existing files for the database being recovered are preserved in case they're needed later for further recovery operations.

2. Use the New-MailboxDatabase cmdlet to create a dial tone database, as shown in this example.

   ```
   New-MailboxDatabase -Name DTDB1 -EdbFilePath D:\DialTone\DTDB1.EDB
   ```

3. Use the Set-Mailbox cmdlet to rehome the user mailboxes hosted on the database being recovered, as shown in this example.

   ```
   Get-Mailbox -Database DB1 | Set-Mailbox -Database DTDB1
   ```

4. Use the Mount-Database cmdlet to mount the database so client computers can access the database and send and receive messages, as shown in this example.

```
    Mount-Database -Identity DTDB1
    ```

5. Create a recovery database (RDB) and restore or copy the database and log files containing the data you want to recover into the RDB. For detailed steps, see Create a recovery database.

6. After the data is copied to the RDB, but before mounting the restored database, copy any log files from the failed database to the recovery database log folder so they can be played against the restored database.

7. Mount the RDB, and then use the Dismount-Database cmdlet to dismount it, as shown in this example.

    ```
    Mount-Database -Identity RDB1
    Dismount-Database -Identity RDB1
    ```

8. After the RDB is dismounted, move the current database and log files within the RDB folder to a safe location. This is done in preparation for swapping the recovered database with the dial tone database.

9. Dismount the dial tone database, as shown in this example. Note that your end users will experience an interruption in service when you dismount this database.

    ```
    Dismount-Database -Identity DTDB1
    ```

10. Move the database and log files from the dial tone database folder into the RDB folder.

11. Move the database and log files from the safe location containing the recovered database into the dial tone database folder, and then mount the database, as shown in this example.

    ```
    Mount-Database -Identity DTDB1
    ```

    This ends the service interruption for your end users. They will be able to access their original production database and send and receive messages.

12. Mount the RDB, as shown in this example.

    ```
    Mount-Database -Identity RDB1
    ```

13. Use the Get-Mailbox and New-MailboxRestoreRequest cmdlets to export the data from the RDB and import it into the recovered database, as shown in this example. This will import all the messages sent and received using the dial tone database into the production database.

    ```
    $mailboxes = Get-Mailbox -Database DTDB1
    ```

    ```
    $mailboxes | %{ New-MailboxRestoreRequest -SourceStoreMailbox $_.ExchangeGuid -SourceDatabase RDB1 -
    TargetMailbox $_ }
    ```

14. After the restore operation is complete, you can dismount and remove the RDB, as shown in this example.

    ```
    Dismount-Database -Identity RDB1
    Remove-MailboxDatabase -Identity RDB1
    ```

For detailed syntax and parameter information, see the following topics:

- New-MailboxDatabase
- Get-Mailbox
- Set-Mailbox
- Mount-Database
- Dismount-Database
- Remove-MailboxDatabase

How do you know this worked?

To verify that you've successfully moved a mailbox, do the following:

- Open the mailbox using Outlook on the web.
- Open the mailbox using Microsoft Outlook.

Managed availability

8/3/2020 • 14 minutes to read • Edit Online

Ensuring that users have a good email experience has always been the primary objective for messaging system administrators. In your Exchange Server organization, all aspects of the system must be actively monitored and any detected issues must be resolved quickly. To achieve this, a feature called *Managed Availability* provides built-in monitoring and recovery actions that preserve the end-user experience.

Managed Availability

Managed availability, also known as *Active Monitoring* or *Local Active Monitoring*, is the integration of built-in monitoring and recovery actions with the Exchange high availability platform. It's designed to detect and recover from problems as soon as they occur and are discovered by the system. Unlike previous external monitoring solutions and techniques for Exchange, managed availability doesn't try to identify or communicate the root cause of an issue. It's instead focused on recovery aspects that address three key areas of the user experience:

- **Availability**: Can users access the service?
- **Latency**: How is the experience for users?
- **Errors**: Are users able to accomplish what they want?

Managed availability provides a native health monitoring and recovery solution. It moves away from monitoring individual separate slices of the system to monitoring the end-to-end user experience, and protecting the end user's experience through recovery-oriented actions.

Managed availability is an internal process that runs on every Exchange server. It polls and analyzes hundreds of health metrics every second. If something is found to be wrong, most of the time it will be fixed automatically. But there will always be issues that managed availability won't be able to fix on its own. In those cases, managed availability will escalate the issue to an administrator by means of event logging.

Managed availability is implemented in the form of two services:

- **Exchange Health Manager Service (MSExchangeHMHost.exe)**: This is a controller process used to manage worker processes. It's used to build, execute, and start and stop the worker process, as needed. It's also used to recover the worker process in case that process fails, to prevent the worker process from being a single point of failure.
- **Exchange Health Manager Worker process (MSExchangeHMWorker.exe)**: This is the worker process responsible for performing run-time tasks within the managed availability framework.

Managed availability uses persistent storage to perform its functions:

- XML files in the \bin\Monitoring\config folder are used to store configuration settings for some of the probe and monitor work items.
- Active Directory is used to store global overrides.
- The Windows registry is used to store run-time data, such as bookmarks, and local (server-specific) overrides.
- The Windows crimson channel event log infrastructure is used to store the work item results.
- Health mailboxes are used for probe activity. Multiple health mailboxes will be created on each mailbox database that exists on the server.

Managed Availability Components

As illustrated in the following drawing, managed availability includes three main asynchronous components that are constantly doing work.

Managed Availability Components

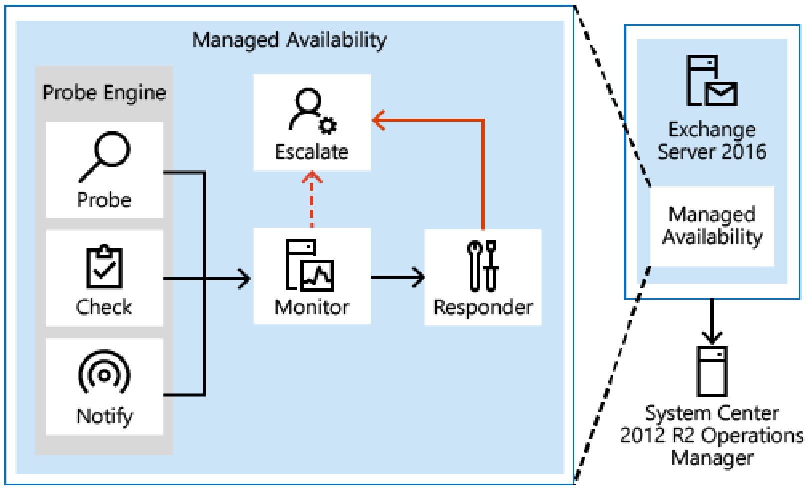

Probes

The first component is called a *Probe*. Probes are responsible for taking measurements on the server and collecting data.

There are three primary categories of probes: recurrent probes, notifications, and checks. Recurrent probes are synthetic transactions performed by the system to test the end-to-end user experience. Checks are the infrastructure that perform the collection of performance data, including user traffic. Checks also measure the collected data against thresholds that are set to determine spikes in user failures, which enable the checks infrastructure to become aware when users are experiencing issues. Finally, the notification logic enables the system to take action immediately, based on a critical event, and without having to wait for the results of the data collected by a probe. These are typically exceptions or conditions that can be detected and recognized without a large sample set.

Recurrent probes run every few minutes and evaluate some aspect of service health. These probes might transmit an email via Exchange ActiveSync to a monitoring mailbox, they might connect to an RPC endpoint, or they might verify Client Access-to-Mailbox connectivity.

All probes are defined on Health Manager service startup in the Microsoft.Exchange.ActiveMonitoring\ProbeDefinition crimson channel. Each probe definitions has many properties, but the most relevant properties are:

- **Name** The name of the probe, which begins with a *SampleMask* of the probe's monitor.
- **TypeName** The code object type of the probe that contains the probe's logic.
- **ServiceName** The name of the health set that contains this probe.
- **TargetResource** The object the probe is validating. This is appended to the name of the probe when it is executed to become a probe result *ResultName*
- **RecurrenceIntervalSeconds** How often the probe executes.
- **TimeoutSeconds** How long the probe will wait before failing.

There are hundreds of recurrent probes. Many of these probes are per-database, so as the number of databases increases, so does the number of probes. Most probes are defined in code and are therefore not directly discoverable.

The basics of a recurrent probe are as follows: start every *RecurrenceIntervalSeconds* and check (or probe) some aspect of health. If the component is healthy, the probe passes and writes an informational event to the Microsoft.Exchange.ActiveMonitoring\ProbeResult channel with a *ResultType* of 3. If the check fails or times out, the probe fails and writes an error event to the same channel. A *ResultType* of 4 means the check failed and a *ResultType* of 1 means that it timed out. Many probes will re-run if they timeout, up to the value of the *MaxRetryAttempts* property.

> **NOTE**
>
> The ProbeResult crimson channel can get very busy with hundreds of probes running every few minutes and logging an event, so there can be a real impact on the performance of your Exchange server if you try expensive queries against the event logs in a production environment.

Notifications are probes that are not run by the health manager framework, but by some other service on the server. These services perform their own monitoring, and then feed their data into the Managed Availability framework by directly writing probe results. You won't see these probes in the ProbeDefinition channel, as this channel only describes probes that will be run by the Managed Availability framework. For example, the ServerOneCopyMonitor Monitor is triggered by probe results written by the MSExchangeDAGMgmt service. This service performs its own monitoring, determines whether there is a problem, and logs a probe result. Most notification probes have the capability to log both a red event that turns the monitor unhealthy and a green event that makes the monitor healthy again.

Checks are probes that only log events when a performance counter passes above or below a defined threshold. They are really a special case of notification probes, as there is a service monitoring the performance counters on the server and logging events to the ProbeResult channel when the configured threshold is met.

To find the counter and threshold that is considered unhealthy, you can look at the monitor for this check. Monitors of the type *Microsoft.Office.Datacenter.ActiveMonitoring.OverallConsecutiveSampleValueAboveThresholdMonitor* or *Microsoft.Office.Datacenter.ActiveMonitoring.OverallConsecutiveSampleValueBelowThresholdMonitor* mean that the probe they watch is a check probe

Monitor

The results of the measurements collected by probes flow into the second component, the *Monitor*. The monitor contains all of the business logic used by the system on the data collected. Similar to a pattern recognition engine, the monitor looks for the various different patterns on all the collected measurements, and then it decides whether something is considered healthy.

Monitors query the data to determine if action needs to be taken based on a predefined rule set. Depending on the rule or the nature of the issue, a monitor can either initiate a responder or escalate the issue to a human via an event log entry. In addition, monitors define how much time after a failure that a responder is executed, as well as the workflow of the recovery action. Monitors have various states. From a system state perspective, monitors have two states:

- **Healthy**: The monitor is operating properly and all collected metrics are within normal operating parameters.

- **Unhealthy**: The monitor isn't healthy and has either initiated recovery through a responder or notified an administrator through escalation.

From an administrative perspective, monitors have additional states that appear in the Exchange Management Shell:

- **Degraded**: When a monitor is in an unhealthy state from 0 through 60 seconds, it's considered Degraded. If a monitor is unhealthy for more than 60 seconds, it is considered Unhealthy.

- **Disabled**: The monitor has been explicitly disabled by an administrator.
- **Unavailable**: The Exchange Health service periodically queries each monitor for its state. If it doesn't get a response to the query, the monitor state becomes Unavailable.
- **Repairing**: An administrator sets the Repairing state to indicate to the system that corrective action is in process by a human, which allows the system and humans to differentiate between other failures that may occur at the same time corrective action is being taken (such as a database copy reseed operation).

Every monitor has a *SampleMask* property in its definition. As the monitor executes, it looks for events in the ProbeResult channel that have a *ResultName* that matches the monitor's *SampleMask*. These events could be from recurrent probes, notifications, or checks. If the monitor's thresholds are achieved, it becomes Unhealthy. From the monitor's perspective, all three probe types are the same as they each log to the ProbeResult channel.

It is worth noting that a single probe failure does not necessarily indicate that something is wrong with the server. It is the design of monitors to correctly identify when there is a real problem that needs fixing. This is why many monitors have thresholds of multiple probe failures before becoming Unhealthy. Even then, many of these problems can be fixed automatically by responders, so the best place to look for problems that require manual intervention is in the Microsoft.Exchange.ManagedAvailability\Monitoring crimson channel. This will include the most recent probe error.

Responders

Finally, there are *Responders*, which are responsible for recovery and escalation actions. As their name implies, responders execute some sort of response to an alert that was generated by a monitor. When something is unhealthy, the first action is to attempt to recover that component. This could include multi-stage recovery actions; for example, the first attempt may be to restart the application pool, the second may be to restart the service, the third attempt may be to restart the server, and the subsequent attempt may be to take the server offline so that it no longer accepts traffic. If the recovery actions are unsuccessful, the system escalates the issue to a human through event log notifications.

Responders take a variety of recovery actions, such as resetting an application worker pool or restarting a server. There are several types of responders:

- **Restart Responder** Terminates and restarts a service.
- **Reset AppPool Responder** Stops and restarts an application pool in Internet Information Services (IIS).
- **Failover Responder** Initiates a database or server failover.
- **Bugcheck Responder** Initiates a bugcheck of the server, thereby causing a server reboot.
- **Offline Responder** Takes a protocol on a server out of service (rejects client requests).
- **Online Responder** Places a protocol on a server back into production (accepts client requests).
- **Escalate Responder** Escalates the issue to an administrator via event logging.

In addition to the above listed responders, some components also have specialized responders that are unique to their component.

All responders include throttling behavior, which provide a built-in sequencing mechanism for controlling responder actions. The throttling behavior is designed to ensure that the system isn't compromised or made worse as a result of responder recovery actions. All responders are throttled in some fashion. When throttling occurs, the responder recovery action may be skipped or delayed, depending on the responder action. For example, when the Bugcheck Responder is throttled, its action is skipped, and not delayed.

Health Sets

From a reporting perspective, managed availability has two views of health, one internal and one external.

The internal view uses *health sets*. Each component in Exchange Server (for example, Outlook on the web, Exchange ActiveSync, the Information Store service, content indexing, transport services, etc.) is monitored by managed availability using probes, monitors, and responders. A group of probes, monitors and responders for a given component is called a *health set*. A health set is a group of probes, monitors, and responders that determine if that component is healthy. The current state of a health set (e.g., whether it is healthy or unhealthy) is determined by using the state of the health set's monitors. If all of a health set's monitors are healthy, then the health set is in a healthy state. If any monitor is not in a healthy state, then the health set state will be determined by its least healthy monitor.

For detailed steps to view server health or health sets state, see Manage health sets and server health.

Health Groups

The external view of managed availability is composed of *health groups*. Health groups are exposed to System Center Operations Manager 2012 R2.

There are four primary health groups:

- **Customer Touch Points** Components that affect real-time user interactions, such as protocols, or the Information Store.
- **Service Components** Components without direct, real-time user interactions, such as the Microsoft Exchange Mailbox Replication service, or the offline address book generation process (OABGen).
- **Server Components** The physical resources of the server, such as disk space, memory and networking.
- **Dependency Availability** The server's ability to access necessary dependencies, such as Active Directory, DNS, etc.

When the Exchange Management Pack is installed, System Center Operations Manager (SCOM) acts as a health portal for viewing information related to the Exchange environment. The SCOM dashboard includes three views of Exchange server health:

- **Active Alerts** Escalation Responders write events to the Windows event log that are consumed by the monitor within SCOM. These appear as alerts in the Active Alerts view.
- **Organization Health** A roll up summary of the overall health of the Exchange organization health is displayed in this view. These rollups include displaying health for individual database availability groups, and health within specific Active Directory sites.
- **Server Health** Related health sets are combined into health groups and summarized in this view.

Overrides

Overrides provide an administrator with the ability to configure some aspects of the managed availability probes, monitors, and responders. Overrides can be used to fine tune some of the thresholds used by managed availability. They can also be used to enable emergency actions for unexpected events that may require configuration settings that are different from the out-of-box defaults.

Overrides can be created and applied to a single server (this is known as a *server override*), or they can be applied to a group of servers (this is known as a *global override*). Server override configuration data is stored in the Windows registry on the server on which the override is applied. Global override configuration data is stored in Active Directory.

Overrides can be configured to last indefinitely, or they can be configured for a specific duration. In addition, global overrides can be configured to apply to all servers, or only servers running a specific version of Exchange.

When you configure an override, it will not take effect immediately. The Microsoft Exchange Health Manager service checks for updated configuration data every 10 minutes. In addition, global overrides will be dependent on Active Directory replication latency.

For detailed steps to view or configure server or global overrides, see Configure managed availability overrides.

Management Tasks and Cmdlets

There are three primary operational tasks that administrators will typically perform with respect to managed availability:

- Extracting or viewing system health
- Viewing health sets, and details about probes, monitors and responders
- Managing overrides

The two primary management tools for managed availability are the Windows Event Log and the Exchange Management Shell. Managed availability logs a large amount of information in the Exchange ActiveMonitoring and ManagedAvailability crimson channel event logs, such as:

- Probe, monitor, and responder definitions, which are logged in the respective *Definition event logs.
- Probe, monitor, and responder results, which are logged in the respective *Results event logs.
- Details about responder recovery actions, including when the recovery action is started, and it is considered complete (whether successful or not), which are logged in the RecoveryActionResults event log.

There are 12 cmdlets used for managed availability, which are described in the following table.

CMDLET	DESCRIPTION
Get-ServerHealth	Used to get raw server health information, such as health sets and their current state (healthy or unhealthy), health set monitors, server components, target resources for probes, and timestamps related to probe or monitor start or stop times, and state transition times.
Get-HealthReport	Used to get a summary health view that includes health sets and their current state.
Get-MonitoringItemIdentity	Used to view the probes, monitors, and responders associated with a specific health set.
Get-MonitoringItemHelp	Used to view descriptions about some of the properties of probes, monitors, and responders.
Add-ServerMonitoringOverride	Used to create a local, server-specific override of a probe, monitor, or responder.
Get-ServerMonitoringOverride	Used to view a list of local overrides on the specified server.
Remove-ServerMonitoringOverride	Used to remove a local override from a specific server.
Add-GlobalMonitoringOverride	Used to create a global override for a group of servers.
Get-GlobalMonitoringOverride	Used to view a list of global overrides configured in the organization.

CMDLET	DESCRIPTION
Remove-GlobalMonitoringOverride	Used to remove a global override.
Set-ServerComponentState	Used to configure the state of one or more server components.
Get-ServerComponentState	Used to view the state of one or more server components.

Manage health sets and server health

8/3/2020 • 2 minutes to read • Edit Online

You can use the built-in health reporting cmdlets to perform a variety of tasks related to managed availability, such as:

- Viewing the health of a server or group of servers
- Viewing a list of health sets
- Viewing a list of probes, monitors, and responders associated with a particular health set
- View a list of monitors and their current health

For more information about health reporting and managed availability, see Managed availability.

What do you need to know before you begin?

- Estimated time to complete each procedure: 2 minutes
- The procedures in this topic require the Exchange Management Shell. For more information, see Open the Exchange Management Shell.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to view server health

You can use the Exchange Management Shell to get a summary of the health of an Exchange server.

Run either of the following commands to view the health sets and health information on an Exchange server:

```
Get-HealthReport -Identity <ServerName>
```

```
Get-ServerHealth -Identity <ServerName> | Format-Table
Server,CurrentHealthSetState,Name,HealthSetName,AlertValue,HealthGroupName -Auto
```

Run any of the following commands to view the health sets on an Exchange server or database availability group:

```
Get-ExchangeServer | Get-HealthReport -RollupGroup
```

```
Get-ExchangeServer | Get-HealthReport -RollupGroup -HealthSetName <HealthSet>
```

```
(Get-DatabaseAvailabilityGroup <DAGName>).Servers | Get-HealthReport -RollupGroup
```

For detailed syntax and parameter information, see Get-HealthReport.

Use the Exchange Management Shell to view a list of health sets

A *health set* is a group of monitors, probes and responders for a component that determine whether the component is healthy or unhealthy.

Run the following command to view the health sets on an Exchange server:

```
Get-HealthReport -Server <ServerName>
```

For detailed syntax and parameter information, see Get-HealthReport.

Use the Exchange Management Shell to view the probes, monitors and responders for a health set

You can use the Exchange Management Shell to view the list of probes, monitors, and responders associated with a health set on an Exchange server.

Run the following command to view the probes, monitors and responders associated with a health set on an Exchange server:

```
Get-MonitoringItemIdentity -Server <ServerName> -Identity <HealthSetName> | Format-Table
Identity,ItemType,Name -Auto
```

For detailed syntax and parameter information, see Get-MonitoringItemIdentity.

Use the Exchange Management Shell to View a List of Monitors and Their Current Health

The health of a monitor is reported by using the "worst of" monitors in the health set. You can view the details of a health set to see which monitors are healthy and which ones are unhealthy.

Run the following command to view a list of the monitors and their current health on an Exchange server:

```
Get-ServerHealth -HealthSet <HealthSetName> -Server <ServerName> | Format-Table Name, AlertValue -Auto
```

For detailed syntax and parameter information, see Get-ServerHealth.

Configure managed availability overrides

8/3/2020 • 4 minutes to read • Edit Online

Managed availability performs continuous probing to detect possible problems with Exchange components or their dependencies, and it performs recovery actions to make sure the end user experience is not impacted due to a problem with any of these components. However, there may be scenarios where the out-of-box settings may not be suitable for your environment. Managed availability probes, monitors, and responders can be customized by creating an override.

There are two types of overrides: local and global. As their names imply, a local override is available only on the server on which it is created, and a global override is used to apply an override to multiple servers. Both types of override can be created for a specific duration or for a specific version of Exchange, but not both at the same time.

> **NOTE**
>
> When you create an override, it doesn't take effect immediately. The Microsoft Exchange Health Management service checks for configuration changes every 10 minutes and loads any detected configuration changes. If you don't want to wait, you can restart the service.

To learn more about managed availability, see Managed availability. For additional management tasks related to managed availability, see Manage health sets and server health.

What do you need to know before you begin?

- Estimated time to complete each procedure: 5 minutes
- The procedures in this topic require the Exchange Management Shell. To open the Exchange Management Shell, see Open the Exchange Management Shell.
- You can only use PowerShell to perform this procedure. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to create local overrides

To create a local override for a specific duration, use the following syntax:

```
Add-ServerMonitoringOverride -Server <ServerName> -Identity <HealthSetName>\<MonitoringItemName>[\
<TargetResource>] -ItemType <Probe | Monitor | Responder | Maintenance> -PropertyName <PropertyName> -
PropertyValue <Value> -Duration <dd.hh:mm:ss>
```

To create a local override for a specific version of Exchange, use the following syntax.

```
Add-ServerMonitoringOverride -Server <ServerName> -Identity <HealthSetName>\<MonitoringItemName>[\
<TargetResource>] -ItemType <Probe | Monitor | Responder | Maintenance> -PropertyName <PropertyName> -
PropertyValue <Value> -Version <15.01.xxxx.xxx>
```

NOTE

When you create the override, the values used in the *Identity* parameter are case-sensitive.

This example adds a local override that disables the responder `ActiveDirectoryConnectivityConfigDCServerReboot` on the server named EXCH03 for 20 days.

```
Add-ServerMonitoringOverride -Server EXCH03 -Identity "AD\ActiveDirectoryConnectivityConfigDCServerReboot" -
ItemType Responder -PropertyName Enabled -PropertyValue 0 -Duration 20.00:00:00
```

How do you know this worked?

To verify that you have successfully created a local override, use the **Get-ServerMonitoringOverride** cmdlet to view the list of local overrides:

```
Get-ServerMonitoringOverride  -Server <ServerIdentity> | Format-List
```

The override should appear in the list.

Use the Exchange Management Shell to remove local overrides

To remove a local override, use the following syntax.

```
Remove-ServerMonitoringOverride -Server <ServerName> -Identity <HealthSetName>\<MonitoringItemName>[\
<TargetResource>] -ItemType <ExistingItemTypeValue> -PropertyName <PropertytoRemove>
```

This example removes the existing local override of the `ActiveDirectoryConnectivityConfigDCServerReboot` responder in the Exchange health set from server EXCH01.

```
Remove-ServerMonitoringOverride -Server EXCH01 -Identity
Exchange\ActiveDirectoryConnectivityConfigDCServerReboot -ItemType Responder -PropertyName Enabled
```

How do you know this worked?

To verify that you have successfully removed a local override, use the **Get-ServerMonitoringOverride** cmdlet to view the list of local overrides:

```
Get-ServerMonitoringOverride  -Server <ServerIdentity> | Format-List
```

The removed override should not appear in the list.

Use the Exchange Management Shell to create global overrides

To create a global override for a specific duration, use the following syntax.

```
Add-GlobalMonitoringOverride -Identity <HealthSetName>\<MonitoringItemName>[\<TargetResource>] -ItemType
<Probe | Monitor | Responder | Maintenance> -PropertyName <PropertytoOverride> -PropertyValue
<NewPropertyValue> -Duration <dd.hh:mm:ss>
```

To create a global override for a specific version of Exchange, use the following syntax.

```
Add-GlobalMonitoringOverride -Identity <HealthSetName>\<MonitoringItemName>[\<TargetResource>] -ItemType
<Probe | Monitor | Responder | Maintenance> -PropertyName <PropertytoOverride> -PropertyValue
<NewPropertyValue> -ApplyVersion <15.01.xxxx.xxx>
```

NOTE

When you create the override, the values used in the *Identity* parameter are case-sensitive.

This example adds a global override that disables the `OnPremisesInboundProxy` probe for 30 days.

```
Add-GlobalMonitoringOverride -Identity "FrontendTransport\OnPremisesInboundProxy" -ItemType Probe -
PropertyName Enabled -PropertyValue 0 -Duration 30.00:00:00
```

This example adds a global override that disables the `StorageLogicalDriveSpaceEscalate` responder for all servers running Exchange version 15.01.0225.042.

```
Add-GlobalMonitoringOverride -Identity "MailboxSpace\StorageLogicalDriveSpaceEscalate" -PropertyName Enabled -
PropertyValue 0 -ItemType Responder -ApplyVersion "15.01.0225.042"
```

How do you know this worked?

To verify that you have successfully created a global override, use the **Get-GlobalMonitoringOverride** cmdlet to view the list of global overrides:

```
Get-GlobalMonitoringOverride
```

The override should appear in the list.

Use the Exchange Management Shell to remove global overrides

To remove a global override, use the following syntax.

```
Remove-GlobalMonitoringOverride -Identity <HealthSetName>\<MonitoringItemName>[\<TargetResource>] -ItemType
<ExistingItemTypeValue> -PropertyName <OverriddenProperty>
```

This example removes the existing global override of the `ExtensionAttributes` property of the `OnPremisesInboundProxy` probe in the `FrontEndTransport` health set.

```
Remove-GlobalMonitoringOverride -Identity FrontEndTransport\OnPremisesInboundProxy -ItemType Probe -
PropertyName ExtensionAttributes
```

How do you know this worked?

To verify that you have successfully removed a global override, use the **Get-GlobalMonitoringOverride** cmdlet to view the list of global overrides:

```
Get-GlobalMonitoringOverride
```

The removed override should not appear in the list.

Server health and performance in Exchange Server

8/3/2020 • 2 minutes to read • Edit Online

Understanding server health and performance is critical to designing and maintaining a high-performance messaging infrastructure. Exchange 2016 and Exchange 2019 continue the features that were introduced in Exchange 2013 to help you manage server health and performance.

Managed availability

Managed availability provides built-in monitoring and recovery actions that preserve the end-user experience. Managed availability is made of two processes: the Exchange Health Manager Service (MSExchangeHMHost.exe) and the Exchange Health Manager Worker process (MSExchangeHMWorker.exe), and the following components:

- **Probe engine**: The *probe engine* takes measurements on the server.
- **Monitoring probe engine**: The *monitoring probe engine* stores the business logic about what constitutes a healthy state. Like a pattern recognition engine, the monitoring probe engine looks for patterns and measurements that differ from a healthy state, and then evaluates whether a component or feature is unhealthy.
- **Responder engine**: When the *responder engine* is alerted about an unhealthy component, it first tries to recover that component. Managed availability enables multi-stage recovery actions. The first attempt may be to restart the application pool, the second attempt may be to restart the corresponding service, and the third attempt may be to restart the server. And, the final attempt may be to put the server offline, so that it no longer accepts traffic. If all of these actions fail, an alert is sent to the help desk.

For more information about managed availability, see Managed availability.

Workload management

Workload management is made of these components:

- *User workload management* is the new name for the user throttling features that were introduced in Exchange 2010. You can customize these setting based on the needs of your environment.
- *System workload management* automatically throttles specific Exchange workloads by monitoring the health of key server resources. These settings should be customized only under the direction of Microsoft Customer Service and Support.

For more information about user workload management, see User workload management in Exchange Server.

User workload management in Exchange Server

8/3/2020 • 3 minutes to read • Edit Online

User workload management allows you to control how Exchange system resources are consumed by users. This feature was available in Exchange 2010 (known as *user throttling*), and was expanded to its current level in Exchange 2013.

A *workload* is a feature, protocol, or service that's been explicitly defined to manage system resources on Exchange servers. Each workload consumes system resources on the Exchange server (for example CPU, memory, network, and disk bandwidth). Examples of workloads include Outlook on the web (formerly known as Outlook Web App), Exchange ActiveSync, mailbox migration, and mailbox assistants.

Control the user consumption of Exchange system resources

By default, the user workload settings allow users to increase their resource consumption for brief periods without experiencing a reduction in bandwidth. Because you can limit user access to resources, there are fewer instances of large resource consumers being locked out. You can further budget user resource consumption by setting a recharge rate for users. The important concepts for user workload management are describe in this list:

- **Burst allowances**: Allows users perform short periods of increased resource consumption without experiencing any throttling.

- **Recharge rate**: Uses a budget system to manage user resource consumption, and specifies the rate at which the user's budget is charged (how much the budget grows by) during the budget time. For example, if the budget time is one hour, a recharge rate value of 600,000 milliseconds indicates that resource budgets for users are recharged at the rate of ten minutes of usage per hour.

- **Traffic shaping (microdelays)**: Works by delaying the user for short periods of time when their resource usage reaches the configured limit over a specific time interval. This delay occurs for very short periods of time (users generally don't notice the delay), and well before the resource consumption causes a significant impact the Exchange server's performance. Traffic shaping preserves the availability of the Exchange server without blocking user productivity, has less user impact than a user lockout, and significantly reduces the chance of a user lockout.

- **Maximum usage**: Temporarily blocks a user who reaches a maximum user resource threshold (the user consumes an unusually high amount of resources over a short time interval). Users who are temporarily blocked from resource usage are unblocked as soon as their resource usage budget allows it (as their budgets are recharged).

You manage user workload settings with these cmdlets in the Exchange Management Shell:

- **View, create, remove, and modify user workload settings**: Get-ThrottlingPolicy, New-ThrottlingPolicy, Remove-ThrottlingPolicy and Set-ThrottlingPolicy.

- **Assign user workload settings to users or computers**: Get-ThrottlingPolicyAssociation and Set-ThrottlingPolicyAssociation

Scopes in user workload settings

By default, there's one throttling policy named `GlobalThrottlingPolicy`. This policy has the scope value Global, which means it applies to all users in the organization. Typically, the settings in the default throttling policy are adequate for users in most Exchange organizations. Instead of customizing the default throttling policy, you can create custom throttling policies that have different settings that the default policy. The scopes that are available in custom throttling policies are:

- **Organization**: The throttling settings apply to *all* users in the organization.
- **Regular**: The throttling settings that apply only to *specific* users in the organization.

The order of precedence for throttling polices are:

1. Throttling policies with the scope value Regular are applied to users before Organization policies and the default throttling policy.
2. Throttling policies with the scope value Organization are applied to users before the default throttling policy.
3. The default throttling policy is applied last, or exclusively to users who don't have Regular or Organization policies assigned to them.

If you create custom throttling policies, the settings should be different than the default throttling policy, and you should plan for the difference in settings from Regular policies to Organization policies to the default policy (for example, least restrictive to most restrictive, or vice-versa).

> **NOTE**
>
> We strongly recommend that you don't modify the default throttling policy, because changes to the default policy could be overwritten by future Exchange updates. Instead, you should create custom throttling policies that contain customized settings.

User throttling in Exchange 2010 coexistence environments

Users with mailboxes on Exchange 2016 servers are throttled using Exchange 2016 throttling features, even if you install Exchange 2016 in an Exchange 2010 organization. This list describes the important considerations for throttling in coexistence environments:

- Exchange 2010 mailboxes remain throttled by Exchange 2010 throttling features when users access their mailboxes through Exchange 2010 Client Access servers.
- When you install Exchange 2016 in an Exchange 2010 organization, Exchange 2016 setup might try to carry some of the Exchange 2010 throttling settings forward. However, the throttling functionality is so different that the effects of any legacy throttling settings will generally not alter how throttling works in Exchange 2016.

Antispam and antimalware protection in Exchange Server

8/3/2020 • 2 minutes to read • Edit Online

Antispam and antimalware protection are included in Exchange Server 2016 and Exchange Server 2019.

- Antispam protection is provided by the same built-in transport agents that were introduced in Exchange Server 2010. These agents are enabled by default on Edge Transport servers, and you can enable many of them on Exchange Mailbox servers.
- Antimalware protection is provided by the Malware agent that was introduced in Exchange Server 2013. The Malware agent is available and enabled by default on Exchange Mailbox servers.

The following table contains links to topics that provide overview information and configuration steps for customizing the built-in spam and malware filtering settings for your organization.

TOPIC	DESCRIPTION
Antispam protection in Exchange Server	Describes the built-in antispam protection features in Exchange Server, and how to configure the antispam protection options.
Running Windows antivirus software on Exchange servers	Describes considerations for running Windows antivirus programs on Exchange servers.

If you're looking for information about antispam features in Microsoft 365, see Anti-spam protection in EOP.

Antispam protection in Exchange Server

8/3/2020 • 5 minutes to read • Edit Online

Spammers, or malicious senders, use a variety of techniques to send unwanted email into your organization. No single tool or process can eliminate all spam. However, Microsoft Exchange provides a layered, multifaceted approach to reducing these unwanted messages. Exchange uses transport agents to provide antispam protection, and the built-in agents that are available in Exchange Server 2016 and Exchange Server 2019 are relatively unchanged from Exchange Server 2010. In Exchange 2016 and Exchange 2019, configuration and management of these agents is available only in the Exchange Management Shell.

For more antispam features and easier management, you can purchase Exchange Online Protection (EOP), which is part of Microsoft 365 and Office 365. To learn more about Microsoft 365 or Office 365 antispam protection, see Anti-spam protection in EOP.

Antispam agents on Mailbox servers

Typically, you enable the antispam agents on a Mailbox server if your organization doesn't have an Edge Transport server, or if it doesn't do other antispam filtering on incoming messages. For more information, see Enable antispam functionality on Mailbox servers.

Like all transport agents, each antispam agent is assigned a priority value. A lower value indicates a higher priority, so typically, an antispam agent with priority 1 acts on a message before an antispam agent with priority 9. However, the SMTP event in the transport pipeline where the antispam agent is registered is also very important in determining the order that antispam agent acts on messages. A low priority antispam agent that's registered early in the transport pipeline acts on messages before a high priority antispam agent that's registered later in the transport pipeline.

Based on the default priority value of the agent and the SMTP event where the agent is registered, this is the order that the antispam agents are applied to messages on Mailbox servers:

1. **Sender Filter agent**: Sender filtering compares the sending server to a list of senders or sender domains that are prohibited from sending messages to your organization. For more information, see Sender filtering.

2. **Sender ID agent**: Sender ID relies on the IP address of the sending server and the Purported Responsible Address (PRA) of the sender to determine whether the sending email address is spoofed. For more information, see Sender ID.

3. **Content Filter agent**: Content filtering agent assigns a spam confidence level (SCL) to each message based on data from legitimate and spam messages. For more information, see Content filtering.

   Spam quarantine is a component of the Content Filter agent that reduces the risk of losing legitimate messages that are incorrectly classified as spam. Spam quarantine provides a temporary storage location for suspicious messages so an administrator can review the messages. For more information, see Spam quarantine in Exchange Server.

   Content filtering also uses the safelist aggregation feature. Safelist aggregation collects safe list data that users configure in Microsoft, Outlook, and Outlook on the web and makes this information available to the Content Filter agent. For more information, see Safelist aggregation.

4. **Protocol Analysis agent (sender reputation)**: The Protocol Analysis agent is the agent that provides sender reputation. Sender reputation uses several tests to calculate a sender reputation level (SRL) on incoming messages that determines the action to take on those messages. For more information, see

Sender reputation and the Protocol Analysis agent.

Antispam agents on Edge Transport servers

If your organization has an Edge Transport server installed in the perimeter network, all of the antispam agents that are available on a Mailbox server are installed and enabled by default on the Edge Transport server. However, the following antispam agents are available only on Edge Transport servers:

- **Connection Filtering agent**: Connection filtering uses an IP block list, IP allow list, IP block list providers, and IP allow list providers to determine whether a connection should be blocked or allowed. For more information, see Connection filtering on Edge Transport servers.

- **Recipient Filter agent**: Recipient filtering uses a recipient block list to identify messages that aren't allowed to enter the organization. The recipient filter also uses the local recipient directory to reject messages sent to invalid recipients. For more information, see Recipient filtering on Edge Transport servers.

  > **NOTE**
  >
  > Although the Recipient Filter agent is available on Mailbox servers, you shouldn't configure it. When recipient filtering on a Mailbox server detects one invalid or blocked recipient in a message that contains other valid recipients, the message is rejected. The Recipient Filter agent is enabled when you install the antispam agents on a Mailbox server, but it isn't configured to block any recipients.

- **Attachment Filtering agent**: Attachment filtering blocks messages or attachments based on the attachment file name, extension, or MIME content type. For more information, see Attachment filtering on Edge Transport servers.

Based on the default priority value of the antispam agent, and the SMTP event in the transport pipeline where the agent is registered, this is the order that the antispam agents are applied to messages on Edge Transport servers:

1. Connection Filtering agent
2. Sender Filter agent
3. Recipient Filter agent
4. Sender ID agent
5. Content Filter agent
6. Protocol Analysis agent (sender reputation)
7. Attachment Filtering agent

Antispam stamps

Antispam stamps are applied to messages and are used by the antispam agents. You can view the antispam stamps to help you diagnose spam-related problems. For more information, see Antispam stamps.

Strategy for antispam approach

Antispam is a balancing act between blocking unwanted messages and allowing legitimate messages. If you configure the antispam features too aggressively, you'll likely block too many legitimate messages (false positives). If you configure the antispam features too loosely, you likely allow too much spam into your organization.

These are some best practices to consider when configuring the built-in antispam features in Exchange:

- Reject messages that are identified by the Connection Filtering agent, Recipient Filter agent, and Sender

Filter agent rather than quarantining the messages or applying antispam stamps. This approach is recommended for these reasons:

  - Messages that are identified by the default settings of the connection filtering, recipient filtering, or sender filtering typically don't require further tests to determine if they're unwanted. For example, if you configured sender filtering to block specific senders, there's no reason to continue to process messages from those senders. (If you didn't want the messages rejected, you wouldn't have put them on the blocked senders list).

  - Configuring a more aggressive level for the antispam agents that encounter messages early in the transport pipeline saves processing, bandwidth, and disk resources. The farther in transport pipeline a message travels, the greater number of variables that the remaining antispam features need to evaluate to successfully identify the message as spam. Reject obvious messages early so you can process ambiguous messages later.

- You need to monitor the effectiveness of the antispam features at their current configuration levels. Monitoring allows you to react to trends and increase or decrease the aggressiveness of the settings. You should start with the default settings to minimize the number of false positives. As you monitor the amount of spam and false positives, you can increase the aggressiveness of the settings based on the type of spam and spam attacks that your organization experiences.

See also

Anti-spam protection in EOP

Enable antispam functionality on Mailbox servers

8/3/2020 • 5 minutes to read • Edit Online

The following antispam agents are available in the Transport service on Exchange 2016 and Exchange 2019 Mailbox servers, but they aren't installed by default:

- Content Filter agent
- Sender Filter agent
- Sender ID agent
- Protocol Analysis agent for sender reputation

You can install these antispam agents on a Mailbox server by using an Exchange Management Shell script, which is important if these agents are your only defense to help prevent spam. Typically, you don't need to install the antispam agents on a Mailbox server when your organization uses other types of antispam filtering on incoming mail.

> **NOTE**
>
> Although the Recipient Filter agent is available on Mailbox servers, you shouldn't configure it. When recipient filtering on a Mailbox server detects one invalid or blocked recipient in a message that contains other valid recipients, the message is rejected. The Recipient Filter agent is enabled when you install the antispam agents on a Mailbox server, but it isn't configured to block any recipients. For more information, see Recipient filtering procedures on Edge Transport servers.

What do you need to know before you begin?

- Estimated time to complete this task: 15 minutes
- You can only use PowerShell to perform this procedure. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- The Connection Filtering agent and the Attachment Filtering agent aren't available on Mailbox servers. They're only available on Edge Transport servers, and they're installed and enabled there by default. However, the Malware agent is installed and enabled by default on Mailbox servers. For more information, see Antimalware protection in Exchange Server.
- If you have other Exchange antispam agents operating on the messages before they reach the Mailbox server (for example, an Edge Transport server in the perimeter network), the antispam agents on the Mailbox server recognize the antispam X-header values that already exist in messages, and those messages pass through without being scanned again.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Transport configuration" entry in the Mail flow permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Step 1: Run the Install-AntispamAgents.ps1 PowerShell script

Run the following command in the Exchange Management Shell on the Mailbox server:

```
& $env:ExchangeInstallPath\Scripts\Install-AntiSpamAgents.ps1
```

How do you know this step worked?

You know this step worked if the script runs without errors and asks you to restart the Microsoft Exchange Transport service. The output looks like this:

```
WARNING: Please exit Windows PowerShell to complete the installation.
WARNING: The following service restart is required for the change(s) to take effect : MSExchangeTransport
WARNING: The following service restart is required for the change(s) to take effect : MSExchangeTransport
Identity                                            Enabled         Priority
--------                                            -------         --------
Content Filter Agent                                True            8
WARNING: Please exit Windows PowerShell to complete the installation.
WARNING: The following service restart is required for the change(s) to take effect : MSExchangeTransport
WARNING: The following service restart is required for the change(s) to take effect : MSExchangeTransport
Sender Id Agent                                     True            9
WARNING: Please exit Windows PowerShell to complete the installation.
WARNING: The following service restart is required for the change(s) to take effect : MSExchangeTransport
WARNING: The following service restart is required for the change(s) to take effect : MSExchangeTransport
Sender Filter Agent                                 True            10
WARNING: Please exit Windows PowerShell to complete the installation.
WARNING: The following service restart is required for the change(s) to take effect : MSExchangeTransport
WARNING: The following service restart is required for the change(s) to take effect : MSExchangeTransport
Recipient Filter Agent                              True            11
WARNING: Please exit Windows PowerShell to complete the installation.
WARNING: The following service restart is required for the change(s) to take effect : MSExchangeTransport
WARNING: The following service restart is required for the change(s) to take effect : MSExchangeTransport
Protocol Analysis Agent                             True            12
WARNING: The agents listed above have been installed. Please restart the Microsoft Exchange Transport
service for
changes to take effect.
```

Step 2: Restart the Microsoft Exchange Transport service

Run the following command in the Exchange Management Shell on the Mailbox server:

```
Restart-Service MSExchangeTransport
```

How do you know this step worked?

You know this step worked if the Microsoft Exchange Transport service restarts without errors. The output looks like this:

```
WARNING: Waiting for service 'Microsoft Exchange Transport (MSExchangeTransport)' to start...
WARNING: Waiting for service 'Microsoft Exchange Transport (MSExchangeTransport)' to start...
WARNING: Waiting for service 'Microsoft Exchange Transport (MSExchangeTransport)' to start...
WARNING: Waiting for service 'Microsoft Exchange Transport (MSExchangeTransport)' to start...
WARNING: Waiting for service 'Microsoft Exchange Transport (MSExchangeTransport)' to start...
WARNING: Waiting for service 'Microsoft Exchange Transport (MSExchangeTransport)' to start...
```

Step 3: Specify the internal SMTP servers in your organization

You need to specify the IP addresses of every internal SMTP server that should be ignored by the Sender ID agent. In fact, you need to specify the IP address of at least one internal SMTP server. If the Mailbox server where you're running the antispam agents is the only SMTP server in your organization, specify the IP address of that computer.

To add the IP addresses of internal SMTP servers without affecting any existing values, run the following command in the Exchange Management Shell on the Mailbox server:

```
Set-TransportConfig -InternalSMTPServers @{Add="<ip address1>","<ip address2>"...}
```

This example adds the internal SMTP server addresses 10.0.1.10 and 10.0.1.11 to the transport configuration of your organization.

```
Set-TransportConfig -InternalSMTPServers @{Add="10.0.1.10","10.0.1.11"}
```

How do you know this step worked?

To verify that you have successfully specified the IP address of at least one internal SMTP server, run the following command in the Exchange Management Shell on the Mailbox server, and verify that the IP address of at least one valid internal SMTP server is displayed.

```
Get-TransportConfig | Format-List InternalSMTPServers
```

Step 4: Next steps

- The Content Filter agent, Sender ID agent, Sender Filter agent, and Protocol Analysis (sender reputation) agent should now be installed and running on the Mailbox server. To verify this, run the following commands in the Exchange Management Shell on the Mailbox server:

  ```
  Get-TransportAgent
  ```

  ```
  Get-ContentFilterConfig | Format-Table Name,Enabled; Get-SenderFilterConfig | Format-Table
  Name,Enabled; Get-SenderIDConfig | Format-Table Name,Enabled; Get-SenderReputationConfig | Format-
  Table Name,Enabled
  ```

- To see detailed information about the configuration of each agent, run the following commands:

  ```
  Get-ContentFilterConfig | Format-List *Enabled,RejectionResponse,*Postmark*,Bypassed*,Quarantine*;
  ```

  ```
  Get-SenderFilterConfig | Format-List *Enabled,*Block*
  ```

```
Get-SenderIDConfig | Format-List *Enabled*,*Action,Bypassed*
```

```
Get-SenderReputationConfig | Format-List *Enabled*,*Proxy*,*Block*,*Ports*
```

- To configure each agent, see the following topics:
  - Content filtering procedures
  - Safelist aggregation procedures
  - Configure Content Filtering to Use Safe Domain Data
  - Exchange spam confidence level (SCL) thresholds
  - Sender filtering procedures
  - Sender ID procedures
  - Sender reputation procedures
- By default, the Content Filter agent, the Sender Filter agent, and the Sender ID agent record their activities in the antispam agent log on the Mailbox server. You can verify that these antispam agents are working when information is written to the log. To see the location and configuration of the log, run the following command in the Exchange Management Shell on the Mailbox server:

  ```
  Get-TransportService | Format-List AgentLog*
  ```

For instructions on how to configure the log, see Configure antispam Agent Logging.

Antispam stamps

8/3/2020 • 6 minutes to read • Edit Online

Antispam stamps in Exchange Server apply diagnostic metadata, or stamps, such as sender-specific information, puzzle validation results, and content filtering results, to messages as they pass through the antispam features that filter inbound messages from the Internet. You can use antispam stamps to see the results of antispam filtering on a message, and to diagnose spam-related problems. The antispam features and stamps are basically unchanged from Exchange Server 2010. There are four major Exchange antispam stamps:

- The phishing confidence level (PCL) stamp
- The Sender ID stamp
- The spam confidence level (SCL) stamp
- The antispam report stamp

The antispam stamps are added to messages as X-header fields in the message header. You can view antispam stamps on a message by using Outlook. For more information, see View antispam stamps in Outlook.

The phishing confidence level stamp

The PCL stamp indicates the likelihood that a message is a phishing message based on its content. The PCL stamp is applied when the message is processed by the Content Filter agent. For more information about content filtering, see Content filtering.

The PCL values are described in the following table.

PCL VALUE	VERDICT	DESCRIPTION
1 through 3	`Neutral`	The message content isn't likely to be phishing.
4 through 8	`Suspicious`	The message content is likely to be phishing.

The PCL value appears in the **X-MS-Exchange-Organization-PCL:** X-header, and the PCL verdict appears in the antispam report stamp as `PCL:PhishingLevel <Verdict>`. Outlook uses the PCL stamp to block the content of suspicious messages.

The Sender ID stamp

The Sender ID stamp is based on the sender policy framework (SPF) that authorizes the use of domains in email. The Sender ID agent determines the Sender ID status for the message. These status values are described in the following table.

STATUS	DESCRIPTION
`Pass`	Both the IP address and Purported Responsible Address (PRA) passed the Sender ID verification check.

STATUS	DESCRIPTION
`Neutral`	Published Sender ID data is explicitly inconclusive.
`SoftFail`	The IP address for the PRA may be in the not permitted set.
`Fail`	The IP Address is not permitted. No PRA is found in the incoming mail or the sending domain does not exist.
`None`	No published SPF data exists in the sender's DNS.
`TempError`	A temporary DNS failure occurred, such as an unavailable DNS server.
`PermError`	The DNS record is invalid, such as an error in the record format.

The Sender ID stamp is displayed in the **X-MS-Exchange-Organization-SenderIdResult:** X-header, and also in the antispam report stamp as `SenderIDStatus <Status>`. The SPF result is displayed in the **Received-SPF** header.

For more information, see the following topics:

- Sender ID
- Sender Policy Framework: SPF Record Syntax

The spam confidence level stamp

> **NOTE**
>
> In November, 2016, Microsoft stopped producing spam definition updates for the SmartScreen filters in Exchange and Outlook. The existing SmartScreen spam definitions were left in place, but their effectiveness will likely degrade over time. For more information, see Deprecating support for SmartScreen in Outlook and Exchange.

The SCL stamp displays the rating of the message based on its content. The Content Filter agent uses Microsoft SmartScreen technology to assess the contents of a message, and to assign an SCL rating to each message. The SCL values are described in the following table.

SCL VALUE	DESCRIPTION
0 through 9	0 indicates an extremely low probability that the message is spam. 9 indicates an extremely high probability that the message is spam.
-1	The message bypassed antispam scanning (for example, the message was from an internal sender).

The SCL value appears in the **X-MS-Exchange-Organization-SCL:** X-header.

The actions that Exchange and Outlook take based on the SCL value depend on your SCL threshold settings. For more information, see Exchange spam confidence level (SCL) thresholds.

The antispam report stamp

The antispam report stamp is a summary of the antispam filter results that have been applied to the message. The Content Filter agent applies this stamp to the message in the **X-MS-Exchange-Organization-Antispam-Report**: X-header. The anti spam report uses the following syntax:

```
X-MS-Exchange-Organization-Antispam-Report: DV:<DATVersion>;CW:CustomList;PCL:PhishingVerdict
<verdict>;P100:PhishingBlock;PP:Presolve;SID:SenderIDStatus <status>;TIME:
<SendReceiveDelta>;MIME:MimeCompliance;OrigIP:<SourceIPAddress>
```

The antispam filter information that can appear in the antispam report stamp is described in the following table. Note that the antispam report stamp only contains results and conclusions from antispam filters that were applied to the message. so the antispam report stamp usually doesn't contain all of the possible stamps and values.

STAMP	DESCRIPTION
DV	The DAT version (DV) stamp indicates the version of the spam definition file that was used when scanning the message.
SA	The signature action (SA) stamp indicates that the message was either recovered or deleted because of a signature that was found in the message.
SV	The signature DAT version (SV) stamp indicates the version of the signature file that was used when scanning the message.
CW	The custom weight (CW) stamp indicates that the message contains an unapproved word or phrase and that the SCL value, or weight, of that unapproved word or phrase was applied to the final SCL score: • Unapproved phrases, or Block phrases, have maximum weight and change the SCL score to 9. • Approved words or phrases, or Allow phrases, have minimum weight and change the SCL score to 0. For more information about how to add approved and unapproved words or phrases to the Content Filtering agent, see Content filtering procedures.
PP	The presolved puzzle (PP) stamp indicates that if a sender's message contains a valid, solved computational postmark (based on Outlook E-mail Postmark validation functionality), it's unlikely that the sender is a malicious sender. In this case, the Content Filter agent would reduce the SCL rating. The Content Filter agent doesn't change the SCL rating if the E-mail Postmark validation feature is enabled and either of the following conditions is true: • An inbound message doesn't contain a computational postmark header. • The computational postmark header isn't valid. For more information about the postmark validation feature, see Content filtering.
TIME:TimeBasedFeatures	Indicates that there was a significant time delay between the time that the message was sent and the time that the message was received. The TIME stamp is used to determine the final SCL rating for the message.
OrigIP	Indicates the IP address of the source messaging server.
MIME:MIMECompliance	Indicates that the email message isn't MIME compliant.

STAMP	DESCRIPTION
P100:PhishingBlock	Indicates that the message contains a URL that's present in a phishing definition file.
IPOnAllowList	Indicates that the sender's IP address is on the IP Allow list. For more information about the IP Allow list, see IP Allow list.
MessageSecurityAntispamBypass	Indicates that the message wasn't filtered for content and that the sender has been granted permission to bypass the antispam filters.
SenderBypassed	Indicates that the Content Filter agent doesn't process any content filtering for messages that are received from this sender. For more information, see Content filtering procedures.
AllRecipientsBypassed	Indicates that one of the following conditions was met for all recipients listed in the message: • The *AntispamBypassedEnabled* parameter on the recipient's mailbox is set to `$true`. For more information, see Use the Exchange Management Shell to configure a mailbox to bypass Exchange antispam filtering. • The message sender is in the recipient's Safe Senders List. For more information about the Safe Senders List, see Use the Exchange Management Shell to configure the safelist collection on a mailbox. • The Content Filter agent doesn't process any content filtering for messages that are sent to this recipient. For more information about recipient exceptions, see Use the Exchange Management Shell to configure recipient and sender exceptions for content filtering.

View antispam stamps in Outlook

8/3/2020 • 2 minutes to read • Edit Online

The built-in antispam agents in Exchange Server apply diagnostic metadata, or stamps, as X-headers to messages as they enter your organization. For more information about these stamps, see Antispam stamps. You can use Microsoft Outlook to view the antispam X-header fields in messages to help you diagnose spam-related problems.

What do you need to know before you begin?

- Estimated time to complete this procedure: less than 5 minutes
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Mailbox access" entry in the Mail flow permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use Outlook 2010 or later to view antispam stamps

1. Open Outlook on a client computer, and in the **Mail** view, double-click a message to open it.
2. In the **Tags** section of the Ribbon, click the **Message Options** icon to display the message **Properties** dialog box.
3. In the **Properties** dialog box, in the **Internet headers** section, use the scroll bar to view the antispam X-headers. The header fields to look for are:
   - **X-MS-Exchange-Organization-SenderIdResult:**
   - **X-MS-Exchange-Organization-SCL:**
   - **X-MS-Exchange-Organization-PCL:**
   - **X-MS-Exchange-Organization-Antispam-Report:**

It can be easier for you to find these values by selecting all of the text in the **Internet headers** field (CTRL key + A), copying the text (CTRL key + C, or right-click and choose **Copy**), and pasting the text into Notepad.

Here's an example of the values that you might find in a suspicious messages:

```
X-MS-Exchange-Organization-SenderIdResult:Fail
X-MS-Exchange-Organization-SCL:6
X-MS-Exchange-Organization-PCL:7X-MS-Exchange-Organization-Antispam-Report:
DV:3.3.15608.880;SID:SenderIDStatus Fail;PCL:PhishingLevel
SUSPICIOUS;CW:CustomList;PP:Presolved;TIME:TimeBasedFeatures;OrigIP:10.1.1.1
```

Configure Exchange antispam settings on mailboxes

8/3/2020 • 18 minutes to read • Edit Online

In Exchange Server, you can configure specific antispam settings on individual mailboxes that are different than the antispam settings that are applied to the rest of the mailboxes in your organization. The antispam settings that are available on mailboxes are basically unchanged from Exchange 2010.

What do you need to know before you begin?

- Estimated time to complete each procedure: 5 minutes
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Antispam features" entry in the Antispam and antimalware permissions topic, and the "Antispam" entry in the Recipients Permissions topic.
- By default, antispam features aren't enabled in the Transport service on a Mailbox server. Typically, you only enable the antispam features on a Mailbox server if your Exchange organization doesn't do any prior antispam filtering before accepting incoming messages. For more information, see Enable antispam functionality on Mailbox servers.
- You can only use PowerShell to perform this procedure. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to enable or disable the junk email rule in a mailbox

By default, the junk email rule (a hidden Inbox rule named Junk E-mail Rule) is enabled in every mailbox, and controls the following Exchange antispam features:

- **Message delivery to the Junk Email folder based on the SCL Junk Email folder threshold**: When a message is assigned a spam confidence level (SCL) value by Exchange, and the SCL value is greater than the SCL Junk Email folder threshold value that's configured for the Exchange organization (the default value is 4) or directly on the mailbox (the default value is not configured), the junk email filter rule moves the message to the Junk Email folder.
- **Message delivery to the Junk Email folder based on the safelist collection on the mailbox**: The entries in the Safe Senders list, Safe Recipients list, and Block Senders list that are configured on the mailbox determine whether the junk email rule delivers the message to the Inbox or the Junk Email folder. Users can configure the safelist collection for their own mailbox in Microsoft Outlook or Outlook on the web. Administrators can configure the safelist collection for a mailbox by using the **Set-MailboxJunkEmailConfiguration** cmdlet.

When the junk email rule is enabled in the mailbox, Exchange is able to deliver messages to the Junk Email folder (based on the Blocked Senders list or SCL Junk Email folder threshold), and prevent messages from being delivered to the Junk Email folder (based on the Safe Senders list). This value corresponds to the Outlook on the web setting: **Automatically filter junk email**.

When the junk email rule is disabled on the mailbox, Exchange can't deliver messages to the Junk Email folder based on the SCL Junk Email folder threshold or the safelist collection on the mailbox. This value corresponds to the Outlook on the web setting: **Don't move email to my Junk Email folder**.

To enable or disable the junk email rule on a mailbox, use the following syntax:

```
Set-MailboxJunkEmailConfiguration <MailboxIdentity> -Enabled <$true | $false>
```

This example disables the junk email rule on Ori Epstein's mailbox.

```
Set-MailboxJunkEmailConfiguration "Ori Epstein" -Enabled $false
```

This example disables the junk email rule on all user mailboxes in the Organizational Unit named North America in the consoto.com domain.

```
Get-Mailbox -RecipientTypeDetails UserMailbox -OrganizationalUnit "contoso.com/North America" | Set-
MailboxJunkEmailConfiguration -Enabled $false
```

This example disables the junk email rule on all user mailboxes in the mailbox database named MDB 01.

```
Get-Mailbox -RecipientTypeDetails UserMailbox -Database "MDB 01" | Set-MailboxJunkEmailConfiguration -Enabled
$false
```

This example disables the junk email rule on all user mailboxes in the organization.

```
$All = Get-Mailbox -RecipientTypeDetails UserMailbox -ResultSize Unlimited; $All | foreach {Set-
MailboxJunkEmailConfiguration $_.Name -Enabled $false}
```

For more information, see Set-MailboxJunkEmailConfiguration.

Notes:

- You can only use the **Set-MailboxJunkEmailConfiguration** cmdlet to disable the junk email rule on a mailbox that's been opened in Outlook (in Cached Exchange mode) or Outlook on the web. If the mailbox hasn't been opened, you'll receive the error:

  ```
  The Junk Email configuration couldn't be set. The user needs to sign in to Outlook Web App before they
  can modify their Safe Senders and Recipients or Blocked Senders lists.
  ```

  If you want to suppress this error for bulk operations, you can add `-ErrorAction SlientlyContinue` to the **Set-MailboxJunkEmailConfiguration** command.

- Disabling the junk email rule on the mailbox prevents the rule from moving messages to the Junk Email folder. However, the Outlook Junk Email Filter can also determine whether a message is spam, and is able to use the safelist collection to move messages to the Inbox or the Junk Email folder. For more information, see the About junk email settings in Outlook section in this topic.

How do you know this worked?

To verify that you have successfully enabled or disabled the junk email rule on a mailbox, use any of the following procedures:

- Replace *<MailboxIdentity>* with the identity of the mailbox, and run the following command to verify the

Enabled property value:

```
Get-MailboxJunkEmailConfiguration <MailboxIdentity> | Format-List Enabled
```

- For bulk operations, use the same filter that identified the mailboxes, and replace the **Set-MailboxJunkEmailConfiguration** command with `Get-MailboxJunkEmailConfiguration | Format-Table -Auto Identity,Enabled`. For example:

```
Get-Mailbox -RecipientTypeDetails UserMailbox -OrganizationalUnit "contoso.com/North America" | Get-
MailboxJunkEmailConfiguration | Format-Table -Auto Identity,Enabled
```

- Replace *<MailboxIdentity>* with the identity of the mailbox, and run the following command to verify the **Enabled** property value of the junk email rule.

```
Get-InboxRule "Junk E-mail Rule" -Mailbox <MailboxIdentity> -IncludeHidden
```

Use the Exchange Management Shell to configure the safelist collection on a mailbox

The safelist collection on a mailbox includes the Safe Senders list, the Safe Recipients list, and the Blocked Senders list. By default, users can configure the safelist collection on their own mailbox in Outlook or Outlook on the web. Administrators can use the corresponding parameters on the **Set-MailboxJunkEmailConfiguration** cmdlet to configure the safelist collection on a user's mailbox. These parameters are described in the following table.

PARAMETER ON SET-MAILBOXJUNKEMAILCONFIGURATION	OUTLOOK WEB APP SETTING
BlockedSendersAndDomains	**Move email from these senders or domains to my Junk Email folder**
ContactsTrusted	**Trust email from my contacts**
TrustedListsOnly	**Don't trust email unless it comes from someone in my Safe Senders and Recipients list**
TrustedSendersAndDomains *TrustedRecipientsAndDomains*	**Don't move email from these senders or domains to my Junk Email folder**

To configure the safelist collection on a mailbox, use the following syntax:

```
Set-MailboxJunkEmailConfiguration <MailboxIdentity> -BlockedSendersAndDomains <EmailAddressesOrDomains |
$null> -ContactsTrusted <$true | $false> -TrustedListsOnly <$true | $false> -TrustedSendersAndDomains
<EmailAddressesOrDomains | $null>
```

To enter multiple values and overwrite any existing entries for the *BlockedSendersAndDomains* and *TrustedSendersAndDomains* parameters, use the following syntax: `"<EmailAddressOrDomain1>","<EmailAddressOrDomain2>"...`. To add or remove one or more values without affecting other existing entries, use the following syntax:

```
@{Add="<EmailAddressOrDomain1>","<EmailAddressOrDomain2>"... ; Remove="<EmailAddressOrDomain3>","
<EmailAddressOrDomain4>...}
```

This example configures the following settings for the safelist collection on Ori Epstein's mailbox:

- Adds the value shopping@fabrikam.com to the Blocked Senders list.
- Removes the value chris@fourthcoffee.com from the Safe Senders list and the Safe Recipients list.
- Configures contacts in the Contacts folder to be treated as trusted senders.

```
Set-MailboxJunkEmailConfiguration "Ori Epstein" -BlockedSendersAndDomains @{Add="shopping@fabrikam.com"} -
TrustedSendersAndDomains @{Remove="chris@fourthcoffee.com"} -ContactsTrusted $true
```

This example empties the Blocked Senders list for all user mailboxes in the Organizational Unit named North America in the contoso.com domain.

```
Get-Mailbox -RecipientTypeDetails UserMailbox -OrganizationalUnit "contoso.com/North America" | Set-
MailboxJunkEmailConfiguration -BlockedSendersAndDomains $null
```

This example adds michelle@tailspintoys.com to the Safe Senders list and Safe Recipients list on all user mailboxes in the mailbox database named MDB 01.

```
Get-Mailbox -RecipientTypeDetails UserMailbox -Database "MDB 01" | Set-MailboxJunkEmailConfiguration -
TrustedSendersAndDomains @{Add="michelle@tailspintoys.com"}
```

This example removes the domain contoso.com from the Blocked Senders list in all user mailboxes in the organization.

```
$All = Get-Mailbox -RecipientTypeDetails UserMailbox -ResultSize Unlimited; $All | foreach {Set-
MailboxJunkEmailConfiguration $_.Name -BlockedSendersAndDomains @{Remove="contoso.com"}}
```

For more information, see Set-MailboxJunkEmailConfiguration.

Notes:

- You can only use the **Set-MailboxJunkEmailConfiguration** cmdlet to configure the safelist collection on a mailbox that's been opened in Outlook (in Cached Exchange mode) or Outlook on the web. If the mailbox hasn't been opened, you'll receive the error:

  ```
  The Junk Email configuration couldn't be set. The user needs to sign in to Outlook Web App before they
  can modify their Safe Senders and Recipients or Blocked Senders lists.
  ```

  If you want to suppress this error for bulk operations, you can add `-ErrorAction SlientlyContinue` to the **Set-MailboxJunkEmailConfiguration** command.

- Disabling the junk email rule in the mailbox prevents the rule from moving messages to the Junk Email folder or keeping messages out of the Junk Email folder based on the safelist collection. However, even with the junk email rule disabled, you can still configure the safelist collection, and the Outlook Junk Email Filter is able to use the safelist collection to move messages to the Inbox or the Junk Email folder. For more information, see the About junk email settings in Outlook section in this topic.

- The safelist aggregation feature of the Content Filter agent is able to share the safelist collection of mailboxes with the built-in Exchange antispam agents. For more information, see Safelist aggregation.

- You can't directly modify the Safe Recipients list by using the **Set-MailboxJunkEmailConfiguration** cmdlet. You modify the Safe Senders list, and those changes are synchronized to the Safe Recipients list.

- The Outlook Junk Email Filter has additional safelist collection settings (for example, **Automatically add people I email to the Safe Senders list**, and separate configuration of the Safe Senders list and Safe Recipients list). For more information, see Use Junk Email Filters to control which messages you see.

How do you know this worked?

To verify that you have successfully configured the safelist collection on a mailbox, use any of following procedures:

- Replace *<MailboxIdentity>* with the identity of the mailbox, and run the following command to verify the property values:

  ```
  Get-MailboxJunkEmailConfiguration <MailboxIdentity> | Format-List trusted*,contacts*,blocked*
  ```

  If the list of email addresses is too long, use this syntax:

  ```
  (Get-MailboxJunkEmailConfiguration <MailboxIdentity>).BlockedSendersAndDomains
  ```

- For bulk operations, specify the filter that you used to configure the safelist collection, and replace the **Set-MailboxJunkEmailConfiguration** command with `Get-MailboxJunkEmailConfiguration | Format-List Identity,trusted*,contacts*,blocked*`. For example:

  ```
  Get-Mailbox -RecipientTypeDetails UserMailbox -OrganizationalUnit "contoso.com/North America" | Get-MailboxJunkEmailConfiguration | Format-List Identity,trusted*,contacts*,blocked*
  ```

Use the Exchange Management Shell to control the availability of junk email settings in Outlook on the web

Administrators can control whether users are allowed to enable or disable the junk email rule, or configure the safelist collection on their own mailboxes in Outlook on the web. This setting doesn't enable or disable the junk email rule in the mailbox; it controls the availability of the junk email settings in Outlook on the web for the mailbox.

To use Outlook on the web mailbox policies to allow or prevent users from configuring the junk email settings on their own mailbox, use the following syntax:

```
Set-OwaMailboxPolicy <OWAMailboxPolicyIdentity> -JunkEmailEnabled <$true | $false>
```

This example prevents all mailboxes that are assigned the Outlook on the web mailbox policy named Default from configuring their junk email settings in Outlook on the web.

```
Set-OwaMailboxPolicy Default -JunkEmailEnabled $false
```

For more information, see Set-OwaMailboxPolicy.

To use Outlook on the web virtual directories to allow or prevent users from configuring the junk email settings on their own mailbox in Outlook on the web, use the following syntax:

```
Set-OwaVirtualDirectory <OWAVirtualDirectoryIdentity> -JunkEmailEnabled <$true | $false>
```

This example prevents all users that connect to the Outlook on the web virtual directory named owa (Default Web Site) on the server named Mailbox01 from configuring their junk email settings.

```
Set-OwaVirtualDirectory "Mailbox01\owa (Default Web Site)" -JunkEmailEnabled $false
```

Note: To apply changes to the Outlook on the web virtual directories, you need to restart Internet Information

Services (IIS) by running the commands `Stop-Service WAS -Force` and `Start-Service W3SVC`.

For more information, see Set-OwaVirtualDirectory.

How do you know this worked?

To verify that you have successfully configured the availability of junk email settings in Outlook on the web, use either of the following procedures:

- For Outlook on the web mailbox policies, run the following command to verify the **JunkEmailEnabled** property value:

  ```
  Get-OwaMailboxPolicy | Format-Table -Auto Name,JunkEmailEnabled
  ```

- For Outlook on the web virtual directories, run the following command to verify the **JunkEmailEnabled** property value:

  ```
  Get-OwaVirtualDirectory | Format-Table -Auto Name,JunkEmailEnabled
  ```

Use the Exchange Management Shell to configure the SCL thresholds on a mailbox

The SCL thresholds are a feature of the Content Filter agent that allows you to escalate the actions that are taken on messages based on their SCL value. For more information, see Exchange spam confidence level (SCL) thresholds.

When you configure an SCL threshold on a mailbox (the value is not blank), the setting on the mailbox overrides the corresponding SCL threshold setting on the Content Filter agent or on the Exchange organization. The SCL thresholds that are available on the mailbox are described in the following table:

SCL THRESHOLD	SCL VALUE COMPARISON OPERATOR	ACTION	AVAILABLE ON THE CONTENT FILTER AGENT?	COMMENTS
Delete	Greater than or equal to	Silently deletes the message (no NDR).	Yes	If this threshold is enabled, the SCL value should be greater than all others.
Reject	Greater than or equal to	Rejects the message with an NDR.	Yes	The SCL value should be less than the delete value, but greater than the quarantine or Junk Email folder values. By default, this threshold is enabled on the Content Filter agent, and has the default value 7.

SCL THRESHOLD	SCL VALUE COMPARISON OPERATOR	ACTION	AVAILABLE ON THE CONTENT FILTER AGENT?	COMMENTS
Quarantine	Greater than or equal to	Redirects the message to the spam quarantine mailbox. For more information about the configuring the spam quarantine mailbox, see Configure a spam quarantine mailbox.	Yes	If this threshold is enabled, the SCL value should be less than the delete or reject values, but greater than the Junk Email folder value.
Junk Email folder	Greater than	Delivers the message to the Junk Email folder in the mailbox. This action is controlled by the junk email rule that's enabled by default in every mailbox. For more information, see the Use the Exchange Management Shell to enable or disable the junk email rule in a mailbox section in this topic.	No You enable or disable the SCL threshold on the mailbox. You configure the SCL threshold value on the Exchange organization, or on the mailbox.	The SCL value should be less than all others. By default, this threshold is enabled on the Exchange organization, and has the default value 4. Because the junk email rule is enabled by default in all mailboxes, messages that arrive in the mailbox with an SCL value of 5 or higher are moved to the Junk Email folder.

To configure the SCL threshold settings on a mailbox, use the following syntax.

```
Set-Mailbox <MailboxIdentity> -SCLDeleteEnabled <$true | $false | $null> -SCLDeleteThreshold <0-9 | $null> -
SCLRejectEnabled <$true | $false | $null> -SCLRejectThreshold <0-9 | $null> -SCLQuarantineEnabled <$true |
$false | $null> -SCLQuarantineThreshold <0-9 | $null> -SCLJunkEnabled <$true | $false | $null> -
SCLJunkThreshold <0-9 | $null>
```

This example disables the SCL Junk email threshold on mailbox of the user named Jeff Phillips.

```
Set-Mailbox "Jeff Phillips" -SCLJunkEnabled $false
```

This example disables the SCL Junk email threshold on all user mailboxes in the Organizational Unit named North America in the consoto.com domain.

```
Get-Mailbox -RecipientTypeDetails UserMailbox -OrganizationalUnit "contoso.com/North America" | Set-Mailbox -
SCLJunkEnabled $false
```

This example disables the SCL Junk email threshold on all user mailboxes in the mailbox database named MDB 01.

```
Get-Mailbox -RecipientTypeDetails UserMailbox -Database "MDB 01" | Set-Mailbox -SCLJunkEnabled $false
```

This example disables the SCL Junk email threshold on all user mailboxes in the organization.

```
$All = Get-Mailbox -RecipientTypeDetails UserMailbox -ResultSize Unlimited; $All | foreach {Set-Mailbox
$_.Name -SCLJunkEnabled $false}
```

Notes:

- To remove the specific SCL thresholds on the mailbox so the SCL threshold is controlled by the Content Filter agent (delete, reject, or quarantine) or the Exchange organization (Junk Email folder), use the value `$null`.
- If you disable the SCL Junk Email folder threshold on the mailbox (*SCLJunkEnabled* is `$false`), but the junk email rule is still enabled in the mailbox, Exchange can still deliver messages to the Junk Email folder based on the Blocked Senders list of the mailbox. Furthermore, even if you disable the Junk E-mail Rule on the mailbox, Outlook (in Cached Exchange mode) can still move messages to the Junk Email folder based on its own determination of whether the message is spam or the Blocked Senders list.

How do you know this worked?

To verify that you have successfully configured the SCL thresholds on a mailbox, use any of the following procedures:

- For a single mailbox, replace *<MailboxIdentity>* with the identity of the mailbox, and run the following command to verify the property values:

  ```
  Get-Mailbox <MailboxIdentity> | Format-List SCL*
  ```

- For bulk operations, specify the filter that you used to configure the SCL thresholds, and replace the **Set-Mailbox** command with `Format-List Name,SCL*`. For example:

  ```
  Get-Mailbox -RecipientTypeDetails UserMailbox -OrganizationalUnit "contoso.com/North America" | Format-
  List Name.SCL*
  ```

Use the Exchange Management Shell to configure the SCL Junk Email folder threshold value for all mailboxes in your organization

The SCL Junk Email folder threshold that's configured on the Exchange organization causes the junk email rule to deliver messages to the Junk Email folder of a mailbox when all of the following conditions are true:

- The message is assigned an SCL value by Exchange (typically, by the Content Filter agent).
- The junk email rule in the mailbox is enabled. It's enabled by default, but it isn't fully functional until the mailbox has been opened in Outlook (in Cached Exchange mode) or Outlook on the web.
- An SCL Junk Email folder threshold isn't configured on the mailbox (by default, it's not configured).
- The SCL value of the message is greater than the SCL Junk Email folder threshold that's configured for the Exchange organization. The default value is 4, which means that messages with an SCL value of 5 or higher are moved to the Junk Email folder by the junk email rule.

To configure the SCL Junk Email folder threshold for all mailboxes in your organization, use the following syntax:

```
Set-OrganizationConfig -SCLJunkThreshold <0-9>
```

This example sets the organization's SCL Junk Email folder threshold value to 5, which means messages with an SCL value of 6 or higher are moved to the Junk Email folder by the junk email rule..

```
Set-OrganizationConfig -SCLJunkThreshold 5
```

Notes:

- You can override the SCL Junk Email folder threshold value on a mailbox by configuring an SCL Junk Email folder threshold on the mailbox. For more information, see the Use the Exchange Management Shell to configure the SCL thresholds on a mailbox section in this topic.

- If you disable the junk email rule in the mailbox, the value of the SCL Junk Email folder threshold for the Exchange organization (or on the mailbox) is meaningless, because the junk email rule is required for Exchange to deliver messages to the Junk Email folder. For more information, see the Use the Exchange Management Shell to enable or disable the junk email rule in a mailbox section in this topic.

How do you know this worked?

To verify that you have successfully configured the SCL Junk Email folder threshold value for all mailboxes in your organization, run the following command to verify the **SCLJunkThreshold** property value:

```
Get-OrganizationConfig | Format-List SCLJunkThreshold
```

Use the Exchange Management Shell to configure a mailbox to bypass Exchange antispam filtering

You can configure messages that are sent to specific mailboxes to bypass all Exchange antispam filters. You can use this setting when Exchange antispam filters are enabled in your organization, but you want to exempt messages that are sent to specific mailboxes from antispam filtering. You can configure this setting for mailboxes with a very low tolerance for false positives (for example, sales or support mailboxes where you can't risk blocking any legitimate messages).

To configure a mailbox to bypass antispam filtering, use the following syntax:

```
Set-Mailbox <MailboxIdentity> -AntispamBypassEnabled <$true | $false>
```

This example exempts messages that are sent to the mailbox named Customer Support from Exchange antispam filtering.

```
Set-Mailbox "Customer Support" -AntispamBypassEnabled $true
```

How do you know this worked?

To verify that you have successfully configured a mailbox to bypass antispam filtering, use the following procedures:

- Replace <MailboxIdentity> with the identity of the mailbox, and run the following command to verify the **AntispamBypassEnabled** property value:

  ```
  Get-Mailbox <MailboxIdentity> | Format-List AntispamBypassEnabled
  ```

- To find all mailboxes in your organization that are configured to bypass antispam filtering, run the following command:

  ```
  Get-Mailbox -ResultSize Unlimited | where {$_.AntispamBypassEnabled -eq $true} | Format-Table
  Name,AntispamBypassEnabled
  ```

About junk email settings in Outlook

To enable, disable, and configure the client-side Junk Email Filter settings that are available in Outlook, use Group Policy. For more information, see Administrative Template files (ADMX/ADML) and Office Customization Tool for Microsoft 365 Apps for enterprise, Office 2019, and Office 2016 and How to deploy junk email settings, such as the Safe Senders list, by using Group Policy.

When the Outlook Junk Email Filter is set to **No automatic filtering** in **Junk** > **Junk E-Mail Options** > **Options**, Outlook doesn't attempt to classify messages as spam, but still uses the safelist collection (the Safe Senders list, Safe Recipients list, and Blocked Senders list) to move messages to the Junk Email folder.

When the Outlook Junk Email Filter is set to **Low** or **High**, the Outlook Junk Email Filter uses its own SmartScreen filter technology to identify and move spam to the Junk Email folder. This spam classification is separate from the SCL Junk Email threshold that's configured on the Exchange organization or on the mailbox. In fact, Outlook ignores the SCL value that's set on a message by Exchange (for all SCL values other than -1), and uses its own criteria to determine whether the message is spam (although the spam verdict from Exchange and Outlook might be the same).

NOTE

In November, 2016, Microsoft stopped producing spam definition updates for the SmartScreen filters in Exchange and Outlook. The existing SmartScreen spam definitions were left in place, but their effectiveness will likely degrade over time. For more information, see Deprecating support for SmartScreen in Outlook and Exchange.

So, the Outlook Junk Email Filter is able to use the mailbox's safelist collection and its own spam classification to move messages to the Junk Email folder, even if the junk email rule and/or the SCL Junk Email threshold are disabled in the mailbox. The difference is whether the junk email rule on the server or the Junk Email Filter in the Outlook client moves the message to the Junk Email folder.

Outlook and Outlook on the web both support the safelist collection. The safelist collection is saved in the Exchange mailbox, so changes to the safelist collection in Outlook appear in Outlook on the web, and vice-versa. The safelist aggregation feature of the Content Filter agent shares these lists with the built-in Exchange antispam agents. For more information, see Safelist aggregation.

Content filtering

8/3/2020 • 6 minutes to read • Edit Online

NOTE

In November, 2016, Microsoft stopped producing spam definition updates for the SmartScreen filters in Exchange and Outlook. The existing SmartScreen spam definitions were left in place, but their effectiveness will likely degrade over time. For more information, see Deprecating support for SmartScreen in Outlook and Exchange.

Content filtering evaluates inbound email messages by assessing the probability that the messages are legitimate or spam. Unlike other filtering technologies, the content filtering uses characteristics from a statistically significant sample of legitimate messages and spam to make its determination. Content filtering in Exchange Server is provided by the Content Filter agent, and is basically unchanged from Exchange Server 2010. Updates to the Content Filter agent are available periodically through Microsoft Update.

By default, the Content Filter agent is enabled on Edge Transport servers, but you can enable it on Mailbox servers. For more information, see Enable antispam functionality on Mailbox servers.

For more information about how to configure the Content Filter agent, see Content filtering procedures.

Using the Content Filter agent

The Content Filter agent assigns a spam confidence level (SCL) to each message by giving it a rating between 0 and 9. A higher number indicates that a message is more likely to be spam. Based on this rating, you can configure the agent to take the following actions:

- **Delete**: The message is silently dropped without a non-delivery report (also known as an NDR, delivery status notification, DSN, or *bounce message*).
- **Reject**: The message is rejected with an NDR.
- **Quarantine**: The message is sent to the spam quarantine mailbox. For more information about the spam quarantine mailbox, see Spam quarantine in Exchange Server.

For example, you may decide that messages with an SCL rating of 7 or higher should be deleted, messages with an SCL rating of 6 should be rejected, and that messages with a SCL rating of 5 should be quarantined.

You can adjust the SCL threshold behavior by assigning different SCL ratings to each of these actions. For more information about how to adjust the SCL threshold to suit your organization's requirements, see Exchange spam confidence level (SCL) thresholds.

NOTE

Messages that are over 11 MB aren't scanned by the Intelligent Message Filter. Instead, they pass through the Content Filter agent without being scanned.

Allow phrases and Block phrases

You can customize how the Content Filter agent assigns SCL values by configuring custom words or phrases the agent will use to apply filter processing. Approved words or phrases are configured with Allow phrases, and unapproved words or phrases with Block phrases. When the Content Filter agent detects an Allow phrase in an inbound message, the agent automatically assigns an SCL value of 0 to the message. Alternatively, when the

Content Filter agent detects a Block phrase in an inbound message, the agent assigns an SCL rating of 9. You can create up to 800 custom words or phrases in any combination of uppercase and lowercase letters. However, the case is ignored by the Content Filter agent.

Outlook Email Postmark validation

The Content Filter agent also includes Outlook Email Postmark validation. This validation is applied to outbound messages to help messaging systems distinguish legitimate email from spam, and to help reduce false positives. In spam filtering, a *false positive* occurs when a spam filter incorrectly identifies a legitimate message as spam. When Outlook Email Postmark validation is enabled, the Content Filter agent parses the inbound message for a computational postmark header. The presence of a valid, solved computational postmark header in the message indicates the client computer that generated the message solved the computational postmark, so the Content Filter agent is likely to lower the message's SCL rating.

Although computers don't require significant processing time to solve individual computational postmarks, processing postmarks for millions of spam messages will be prohibitive to a malicious sender. If a sender's message contains a valid, solved computational postmark, it's unlikely that the sender is malicious, so the Content Filter agent would lower the SCL rating. If the postmark validation feature is enabled and the computational postmark header in an inbound message is invalid or missing, the Content Filter agent won't change the SCL rating.

Bypassing the recipient, sender, and sender domain

In some organizations, all email messages to certain aliases must be accepted, which can cause problems if your organization manages a significant volume of spam. You can configure exceptions to content filtering for specific recipients, senders, and sender domains.

For example, a company named Woodgrove Bank has an alias named customerloans@woodgrovebank.com that provides email support to external loan customers, so the Exchange administrators configure Block phrases to filter messages that are typically used in spam sent by unscrupulous loan agencies. To prevent potentially legitimate messages from being rejected, the administrators set exceptions to content filtering by entering a list of recipient email addresses in the Content Filter agent configuration.

Safelist aggregation

Safelist aggregation is a set of antispam functionality that's shared across Outlook and Exchange. As its name suggests, it collects data from the antispam safe lists that Outlook users configure, and makes this data available to the antispam agents on the Exchange server. The Content Filter agent uses the Outlook Safe Senders Lists, Safe Recipients Lists, and trusted contacts to optimize spam filtering. Email messages from these contacts are identified as safe by the Content Filter agent. Sender filtering and the Sender Filter agent uses the Outlook Blocked Senders list to perform per-recipient sender filtering. For more information, see Safelist aggregation.

Configuring the Content Filter agent

You configure the Content Filter agent by using the Exchange Management Shell. For more information, see Content filtering procedures.

The Content Filter agent depends on updates to determine whether a message is spam. These updates contain data about phishing web sites, Microsoft SmartScreen spam heuristics, and other Intelligent Message Filter updates. These updates generally contain about 6 MB of data that's useful for longer periods of time than other antispam update data.

Content filter updates are available from Microsoft Update. The content filter update data is updated and available for download every two weeks.

Using the SCL value in mail flow rules on Edge Transport servers

On Edge Transport servers, the Edge Rule agent acts on messages before the SCL value is added by the Content

Filter agent. If you want to use the *SCLOver* mail flow rule (also known as a transport rule) condition, you need to configure the Content Filter agent to run before the Edge Rule agent by changing the transport agent priorities. For more information, see Make message SCL values available to mail flow rules on Edge Transport servers.

Notes:

- Although the Content Filter agent runs on other SMTP events, the SCL value is stamped on the message by the instance of the Content Filter agent that's registered on the `OnEndOfData` SMTP event.
- If you configure the Content Filter agent to act on messages before the Edge Rule agent on an Edge Transport server, the server might incur additional processing costs, because messages that would normally be rejected by other mail flow rules are received and evaluated by the Content Filter agent before they are rejected by the Edge Rule agent, Also, you won't be able to configure a mail flow rule to stamp a message that has an SCL value of `-1`, which tells the Content Filter agent to ignore the message.

For more information about transport agents and transport agent priority, see Understanding Transport Agents.

Content filtering procedures

8/3/2020 • 7 minutes to read • Edit Online

Content filtering evaluates incoming messages to determine if a message is legitimate or spam. For more information about content filtering and the Content Filter agent, see Content filtering.

You can configure many aspects of content filtering. For example:

- Enable or disable content filtering on messages from internal (authenticated) and external (unauthenticated) sources (it's enabled by default for incoming messages from external sources).
- Configure exceptions to content filtering for specific senders, recipients, and source domains.
- Configure allowed phrases and blocked phrases to look for in messages.
- Configure the spam confidence level (SCL) thresholds that tell what content filtering should do to messages (delete, reject, or quarantine)

What do you need to know before you begin?

- Estimated time to complete each procedure: less than 5 minutes
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Antispam feature" entry in the Antispam and antimalware permissions topic.
- You can only use PowerShell to perform this procedure. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- By default, antispam features aren't enabled in the Transport service on a Mailbox server. Typically, you only enable the antispam features on a Mailbox server if your Exchange organization doesn't do any prior antispam filtering before accepting incoming messages. For more information, see Enable antispam functionality on Mailbox servers.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to enable or disable content filtering

To disable content filtering, run the following command:

```
Set-ContentFilterConfig -Enabled $false
```

To enable content filtering, run the following command:

```
Set-ContentFilterConfig -Enabled $true
```

> **NOTE**
>
> When you disable content filtering, the underlying Content Filter agent is still enabled. To disable the Content Filter agent, run the command: `Disable-TransportAgent "Content Filter Agent"`.

How do you know this worked?

To verify that you have successfully enabled or disabled content filtering, run the following command to verify the **Enabled** property value:

```
Get-ContentFilterConfig | Format-List Enabled
```

Use the Exchange Management Shell to enable or disable content filtering for external messages

By default, content filtering functionality is enabled for external messages.

To disable content filtering for external messages, run the following command:

```
Set-ContentFilterConfig -ExternalMailEnabled $false
```

To enable content filtering for external messages, run the following command:

```
Set-ContentFilterConfig -ExternalMailEnabled $true
```

How do you know this worked?

To verify that you have successfully enabled or disabled content filtering for external messages, run the following command to verify the **ExternalMailEnabled** property value:

```
Get-ContentFilterConfig | Format-List ExternalMailEnabled
```

Use the Exchange Management Shell to enable or disable content filtering for internal messages

As a best practice, you don't need to apply antispam filters to messages from trusted partners or from inside your organization. There's always a chance that the filters will detect false positives. To reduce the chance that filters will mishandle legitimate email messages, you should typically configure antispam agents to only run on messages from untrusted and unknown sources.

To enable content filtering for internal messages, run the following command:

```
Set-ContentFilterConfig -InternalMailEnabled $true
```

To disable content filtering for internal messages, run the following command:

```
Set-ContentFilterConfig -InternalMailEnabled $false
```

How do you know this worked?

To verify that you have successfully enabled or disabled content filtering for internal messages, run the following command to verify the **InternalMailEnabled** property value:

```
Get-ContentFilterConfig | Format-List InternalMailEnabled
```

Use the Exchange Management Shell to configure recipient and sender exceptions for content filtering

You can specify recipient and sender exceptions that replace the existing values, or you can add or remove specific sender and recipient exceptions without affecting the other existing values.

To replace the existing values, use the following syntax:

```
Set-ContentFilterConfig -BypassedRecipients <recipient1,recipient2...> -BypassedSenders <sender1,sender2...>
-BypassedSenderDomains <domain1,domain2...>
```

This example configures the following exceptions in content filtering:

- The recipients laura@contoso.com and julia@contoso.com aren't checked by content filtering.
- The senders steve@fabrikam.com and cindy@fabrikam.com aren't checked by content filtering.
- All senders in the domain nwtraders.com and all subdomains aren't checked by content filtering.

```
Set-ContentFilterConfig -BypassedRecipients laura@contoso.com,julia@contoso.com -BypassedSenders
steve@fabrikam.com,cindy@fabrikam.com -BypassedSenderDomains *.nwtraders.com
```

To add or remove entries without modifying other existing values, use the following syntax:

```
Set-ContentFilterConfig -BypassedRecipients @{Add="<recipient1>","<recipient2>"...; Remove="<recipient1>","
<recipient2>"...} -BypassedSenders @{Add="<sender1>","<sender2>"...; Remove="<sender1>","<sender2>"...} -
BypassedSenderDomains @{Add="<domain1>","<domain2>"...; Remove="<domain1>","<domain2>"...}
```

This example configures the following exceptions in content filtering:

- Add tiffany@contoso.com and chris@contoso.com to the list of existing recipients who aren't checked by content filtering.
- Add joe@fabrikam.com and michelle@fabrikam.com to the list of existing senders who aren't checked by content filtering.
- Add blueyonderairlines.com to the list of existing domains whose senders aren't checked by content filtering.
- Remove the domain woodgrovebank.com and all subdomains from the list of existing domains whose senders aren't checked by content filtering.

```
Set-ContentFilterConfig -BypassedRecipients @{Add="tiffany@contoso.com","chris@contoso.com"} -BypassedSenders
@{Add="joe@fabrikam.com","michelle@fabrikam.com"} -BypassedSenderDomains @{Add="blueyonderairlines.com";
Remove="*.woodgrovebank.com"}
```

How do you know this worked?

To verify that you have successfully configured the recipient and sender exceptions, run the following command to verify the property values:

```
Get-ContentFilterConfig | Format-List Bypassed*
```

Use the Exchange Management Shell to configure allowed and blocked phrases for content filtering

To add allowed and blocked words and phrases, use the following syntax:

```
Add-ContentFilterPhrase -Influence GoodWord -Phrase <Phrase> -Influence BadWord -Phrase <Phrase>
```

This example allows all messages that contain the phrase "customer feedback".

```
Add-ContentFilterPhrase -Influence GoodWord -Phrase "customer feedback"
```

This example blocks all messages that contain the phrase "stock tip".

```
Add-ContentFilterPhrase -Influence BadWord -Phrase "stock tip"
```

To remove allowed or blocked phrases, use the following syntax:

```
Remove-ContentFilterPhrase -Phrase <Phrase>
```

This example removes the phrase "stock tip":

```
Remove-ContentFilterPhrase -Phrase "stock tip"
```

How do you know this worked?

To verify that you have successfully configured the allowed and block phrases, run the following command to verify the property values:

```
Get-ContentFilterPhrase | Format-Table -Auto Influence,Phrase
```

Use the Exchange Management Shell to configure SCL thresholds for content filtering

To configure the spam confidence level (SCL) thresholds and actions, use the following syntax:

```
Set-ContentFilterConfig -SCLDeleteEnabled <$true | $false> -SCLDeleteThreshold <Value> -SCLRejectEnabled
<$true | $false> -SCLRejectThreshold <Value> -SCLQuarantineEnabled <$true | $false> -SCLQuarantineThreshold
<Value>
```

Notes:

- The Delete action takes precedence over the Reject action, and the Reject action takes precedence over the Quarantine action. Therefore, the SCL threshold for the Delete action should be greater than the SCL threshold for the Reject action, which in turn should be greater than the SCL threshold for the Quarantine action. Only the Reject action is enabled by default, and it has the SCL threshold value 7.

- The Quarantine action requires a spam quarantine mailbox. For more information, see Configure a spam quarantine mailbox.

This example configures the following values for the SCL thresholds:

- The **Delete** action is enabled and the corresponding SCL threshold is set to 9.
- The **Reject** action is enabled and the corresponding SCL threshold is set to 8.
- The **Quarantine** action is enabled and the corresponding SCL threshold is set to 7.

```
Set-ContentFilterConfig -SCLDeleteEnabled $true -SCLDeleteThreshold 9 -SCLRejectEnabled $true -
SCLRejectThreshold 8 -SCLQuarantineEnabled $true -SCLQuarantineThreshold 7
```

How do you know this worked?

To verify that you have successfully configured the SCL thresholds, run the following command to verify the property values:

```
Get-ContentFilterConfig | Format-List SCL*
```

Use the Exchange Management Shell to configure the rejection response for content filtering

When the Reject action is enabled, you can customize the rejection response that's sent to the message sender. The rejection response can't exceed 240 characters.

To configure a custom rejection response, use the following syntax:

```
Set-ContentFilterConfig -RejectionResponse "<Custom Text>"
```

This example configures the Content Filter agent to send a customized rejection response.

```
Set-ContentFilterConfig -RejectionResponse "Your message was rejected because it appears to be SPAM."
```

How do you know this worked?

To verify that you have successfully configured the rejection response, run the following command to verify the property values:

```
Get-ContentFilterConfig | Format-List *Reject*
```

Use the Exchange Management Shell to enable or disable Outlook Email Postmarking

Outlook Email Postmarking validation is a computational proof that Microsoft Outlook applies to outgoing messages to help messaging systems distinguish legitimate email from junk email (reduce false positives). Postmarking was first introduced in Outlook 2007, and is enabled in Outlook by default.

To disable Outlook Email Postmarking, run the following command:

```
Set-ContentFilterConfig -OutlookEmailPostmarkValidationEnabled $false
```

To enable Outlook Email Postmarking, run the following command:

```
Set-ContentFilterConfig -OutlookEmailPostmarkValidationEnabled $true
```

How do you know this worked?

To verify that you have successfully configured Outlook Email Postmarking, run the following command to verify the **OutlookEmailPostmarkValidationEnabled** property value:

```
Get-ContentFilterConfig | Format-List OutlookEmailPostmarkValidationEnabled
```

Safelist aggregation

8/3/2020 • 5 minutes to read • Edit Online

In Exchange Server, *safelist aggregation* refers to sender and recipient email addresses that are collected from all users' Junk Email options in Microsoft Outlook, Outlook on the web, or the **Set-MailboxJunkEmailConfiguration** cmdlet, and shared with the built-in Exchange antispam agents. Safelist aggregation is basically unchanged from Exchange Server 2010.

When you enable and configure safelist aggregation, Exchange can take the following actions based on the safelist aggregation data:

- Deliver incoming messages from senders that have been identified as safe without additional antispam processing (which could potentially identify the messages as spam).
- Block incoming messages from senders that have been identified as malicious.

To configure safelist aggregation, see Safelist aggregation procedures.

In the context of spam filtering, a *false-positive* is a legitimate message that's identified as spam. For organizations that filter hundreds of thousands of messages from the Internet every day, even a small percentage of false-positives means that users might not receive many legitimate messages. Safelist aggregation is likely the most effective way to reduce false-positives.

Information stored in the user's safelist collection

A *safelist collection* is the combined data from the user's Safe Senders list, Safe Recipients list, Blocked Senders list, and (optionally) external contacts. This data is stored in Outlook and in the Exchange mailbox. For more information about adding and removing entries from a user's safelist collection, see Use the Exchange Management Shell to configure the safelist collection on a mailbox.

The following information is stored in a user's safelist collection:

- **Safe senders**: The SMTP email address in the **From:** field.
- **Safe recipients**: The SMTP email address in the **To:** field.
- **Blocked senders**: Just like safe senders, users can block unwanted senders by adding them to their Blocked Senders list.
- **Safe domain**: This is part of the Safe Senders list, but instead of an SMTP email address (masato@contoso.com), the domain of the sender is specified (lcontoso.com).

  Note: By default, Exchange doesn't include safe domains during safelist aggregation. However, you can configure safelist aggregation to include the safe domain data. For more information, see Configure Content Filtering to Use Safe Domain Data.

- **External contacts**: Two types of external contact information can be included in the safelist collection:
  - **Recipients that the user has sent mail to**: These email address are added to the Safe Senders list if the user selects **Automatically add people I e-mail to the Safe Senders list** in the Junk Email options in Outlook.
  - **Contacts in the user's Contacts folder**: These email address are added to the Safe Senders list if the user selects **Also trust e-mail from my Contacts** in the Junk Email options in Outlook, Outlook on the web, or the **Set-MailboxJunkEmailConfiguration** cmdlet.

How Exchange uses the safelist collection

The safelist collection is stored on the user's Mailbox server. A user can have up to 1,024 unique entries in a safelist collection. Exchange has a mailbox assistant, called the Junk Email Options mailbox assistant, that monitors changes to the safelist collection for mailboxes on the server. It then replicates these changes to Active Directory, where the safelist collection is stored on each user object. The safelist collection is optimized for minimized storage and replication. If you have a subscribed Edge Transport server in your perimeter network, the Microsoft Exchange EdgeSync service replicates the safelist collection to the Active Directory Lightweight Directory Services (AD LDS) instance on the Edge Transport server.

The following Exchange antispam agents use the safelist collection:

- The Content Filter agent uses the Safe Senders list data to deliver messages from those senders without additional (unnecessary) processing.
- The Sender Filter agent uses the Blocked Senders list data to reject or delete messages from those senders. For more information, see Sender filtering procedures.

Note:Although the Safe Recipients list can be included in safelist aggregation, the Content Filter agent doesn't act on safe recipient data.

Hashing of safelist collection entries

Safelist collection entries are hashed (SHA-256) one way before they are stored as array sets across three user object attributes, **msExchSafeSenderHash**, **msExchSafeRecipientHash**, and **msExchBlockedSendersHash**, as a binary large object. When data is hashed, an output of fixed length is produced, and the output is likely to be unique. For hashing of safelist collection entries, a 4-byte hash is produced. When a message is received from the Internet, Exchange hashes the sender's email address and compares it to the hashes that are stored on behalf of the destination mailbox. If the sender matches the safe senders hash, the message bypasses content filtering. If the sender matches the blocked senders hash, the message is blocked.

One-way hashing of safelist collection entries performs the following important functions:

- **Minimizes storage and replication space**: Most of the time, hashing reduces the size of the data. Therefore, saving and transmitting a hashed version of a safelist collection entry conserves storage space and replication time. For example, a user who has 200 entries in his or her safelist collection would create about 800 bytes of hashed data stored and replicated in Active Directory.
- **Renders user safelist collections unusable by malicious users**: Because one-way hash values are impossible to reverse-engineer into the original SMTP address or domain, the safelist collections don't yield usable email addresses for malicious users who might compromise an Exchange server.

Enabling safelist aggregation

Safelist aggregation is enabled by default. The safelist collection data is written to Active Directory by the Junk Email Options mailbox assistant. Unlike previous versions of Exchange, you don't need to manually run the **Update-SafeList** cmdlet to hash and write the safelist collection data to Active Directory.

You can still manually run safelist aggregation by using the **Update-Safelist** cmdlet. However, you need to be aware of the replication traffic that might be generated when you run this command. Running **Update-Safelist** on multiple mailboxes where safelists are heavily used might generate a significant amount of network traffic. We recommend that if you run the command on multiple mailboxes, you should run the command during off-peak, non-business hours.

The **Update-SafeList** cmdlet reads the safelist collection from the user's mailbox, hashes each entry, sorts the entries for easy search, and then converts the hash to a binary attribute. Finally, the **Update-SafeList** cmdlet compares the binary attribute that was created to any value stored on the attribute. If the two values are identical, the **Update-SafeList** cmdlet doesn't update the user attribute value with the safelist aggregation data. If the two attribute values are different, the **Update-SafeList** cmdlet updates the safelist aggregation value.

For more information about using **Update-SafeList**, see Safelist aggregation procedures.

Safelist aggregation procedures

8/3/2020 • 5 minutes to read • Edit Online

In Exchange Server, *safelist aggregation* refers to sender and recipient data that's collected from all users' Junk Email options in Microsoft Outlook, Outlook on the web, or the **Set-MailboxJunkEmailConfiguration** cmdlet and shared with the built-in Exchange antispam agents. For more information, see Safelist aggregation.

You can use the procedures in this topic to:

- Configure limits on the number of safe senders and blocked senders that are stored for specific mailboxes.

- Manually run safelist aggregation

- Verify that safelist aggregation is working correctly.

For more information about adding and removing entries from a user's safelist collection, see Use the Exchange Management Shell to configure the safelist collection on a mailbox.

What do you need to know before you begin?

- Estimated time to complete each procedure: 5 minutes

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Recipient Provisioning Permissions" section in the Recipients Permissions topic, and the "Antispam features" section in the Antispam and antimalware permissions topic.

- You can only use PowerShell to perform this procedure. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.

- By default, antispam features aren't enabled in the Transport service on a Mailbox server. Typically, you only enable the antispam features on a Mailbox server if your Exchange organization doesn't do any prior antispam filtering before accepting incoming messages. For more information, see Enable antispam functionality on Mailbox servers.

- Be aware of the replication traffic that might be generated when you run the **Update-SafeList** cmdlet. Running the command on multiple mailboxes where safelists are heavily used might generate a significant amount of network traffic. We recommend that if you run the command on multiple mailboxes, you should run the command during off-peak, non-business hours.

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to configure the mailbox safelist collection limits

You can configure the maximum number of safe senders and blocked senders a user can configure. By default, users can configure up to 5,000 safe senders and 500 blocked senders.

To configure the maximum number of safe senders and blocked senders, use the following syntax:

```
Set-Mailbox <MailboxIdentity> -MaxSafeSenders <Integer> -MaxBlockedSenders <Integer>
```

This example configures the mailbox john@contoso.com to have a maximum of 2,000 safe senders and 200 blocked senders.

```
Set-Mailbox john@contoso.com -MaxSafeSenders 2000 -MaxBlockedSenders 200
```

How do you know this worked?

To verify that you have successfully configured the mailbox safelist collection limits, replace *<MailboxIdentity>* with the identity of the mailbox, and run the following command to verify the mailbox property values.

```
Get-Mailbox <MailboxIdentity> | Format-List Name,Max*Senders
```

Use the Exchange Management Shell to manually run safelist aggregation

Safelist aggregation is done automatically, so you don't need to schedule or manually run the **Update-Safelist** cmdlet. However, you may want to occasionally run this cmdlet to test safelist aggregation.

To manually run safelist aggregation, use the following syntax:

```
Update-Safelist <MailboxIdentity> [-Type <SafeSenders | SafeRecipients | Both>] [-IncludeDomains]
```

This example writes the Safe Senders List for the mailbox john@contoso.com to Active Directory.

```
Update-Safelist john@contoso.com
```

For detailed syntax and parameter information, see Update-SafeList.

Notes:

- You don't need to use the *Type* parameter because:
  - The default value is `SafeSenders`.
  - The Content Filter agent doesn't use Safe Recipients list data, so the `SafeRecipients` or `Both` values are unnecessary.
- By default, safelist aggregation doesn't include domain entries from the Safe Senders list (just email addresses), but you can configure it to include domain entries from the safelist collection. For more information, see Configure Content Filtering to Use Safe Domain Data.

How do you know this worked?

To verify that you have successfully configured safelist aggregation, perform the following steps:

Step 1: Use the Exchange Management Shell to verify the Content Filter agent is enabled on the Exchange server

Run the following command:

```
Get-ContentFilterConfig | Format-List Enabled
```

If the output shows the **Enabled** property to be `True`, content filtering is enabled. If it isn't, run the following command to enable content filtering and the Content Filter agent on the Exchange server:

```
Set-ContentFilterConfig -Enabled $true
```

Step 2: (Optional) Use ADSI Edit to verify replication of the safelist aggregation data to Edge Transport servers

This step is only required if you run the Content Filter agent on a subscribed Edge Transport server in your perimeter network.

You can view the user objects in the Active Directory Lightweight Directory Services (AD LDS) instance on the Edge Transport server to:

- Verify that the safelist collection data is updated for the user objects.
- Verify that the Microsoft Exchange EdgeSync service has replicated the data to the AD LDS instance.

There are three safelist collection attributes for each user object:

- **msExchSafeRecipientsHash**: Stores the hash of the Safe Recipients List collection for the user.
- **msExchSafeSendersHash**: Stores the hash of the Safe Senders List collection for the user.
- **msExchBlockedSendersHash**: Stores the hash of the Blocked Senders List collection for the user.

If a hexadecimal string, such as `0xac 0xbd 0x03 0xca`, is present on the attribute, the user object was updated. If the attribute has a value of `<Not Set>`, the attribute wasn't updated.

You can search for and view the attributes by using ADSI Edit on the Edge Transport server (run ADSIEdit.msc).

Step 3: Send a test message to verify safelist aggregation is working

To test whether safelist aggregation is functioning, you need to send yourself a message from a safe sender that would otherwise be blocked by content filtering (for example, the message contains a blocked phrase). If safelist aggregation is functioning, the message should arrive in your Inbox.

1. Open your Exchange mailbox in Outlook, and add an external email address (associated with an account that you can access) to your Safe Senders List. For more information, see Add names to the Junk Email Filter lists.

2. Use the **Update-SafeList** cmdlet to manually replicate the safelist collection from your mailbox to Active Directory:

   ```
   Update-Safelist <YourMailboxIdentity>
   ```

3. Optional: if you're running the Content Filter agent on a subscribed Edge Transport server in the perimeter network, run the **Start-EdgeSynchronization** cmdlet to force EdgeSync replication.

4. Add a specific word as a blocked phrase to your content filtering configuration. For example:

   ```
   Add-ContentFilterPhrase -Influence BadWord -Phrase "SafeList aggregation test"
   ```

   For details, see Use the Exchange Management Shell to configure allowed and blocked phrases for content filtering.

5. From the external email account in step 1, send a message to your Exchange mailbox that includes the blocked phrase that you configured in step 4.

   If the message is successfully delivered to your Inbox, safelist aggregation is working correctly.

Spam quarantine in Exchange Server

8/3/2020 • 2 minutes to read • Edit Online

Many organizations are bound by legal or regulatory requirements to preserve or deliver all legitimate email messages. In Exchange Server, spam quarantine is a feature of the Content Filter agent that reduces the risk of losing legitimate incoming email messages by providing a temporary storage location for messages that are identified as spam. Spam quarantine is basically unchanged from Exchange Server 2010.

Messages that are identified as spam by the Content Filter agent are wrapped in a non-delivery report (also known as an NDR, delivery status notification, DSN, or bounce message) and delivered to the designated spam quarantine mailbox inside the organization. Administrators can use Microsoft Outlook to review the messages in the spam quarantine mailbox and take appropriate action. For example, you can delete messages, or release legitimate messages to their intended recipients. In addition, you can configure the spam quarantine mailbox to automatically delete messages after a designated time period.

To use the spam quarantine, follow these steps:

1. Verify content filtering is enabled.
2. Create a dedicated mailbox for spam quarantine.
3. Specify the spam quarantine mailbox.
4. Configure the SCL quarantine threshold.
5. Manage the spam quarantine mailbox.
6. Adjust the SCL quarantine threshold as needed.

For detailed instructions, see Configure a spam quarantine mailbox.

More information

The Content Filter agent evaluates incoming messages and applies a spam confidence level (SCL) to each message. The SCL is a numeric value from 0 through 9, where 0 is considered very unlikely to be spam, and 9 is considered very likely to be spam. You can configure the Content Filter agent to take progressively more serious action based on a higher SCL value. For example:

- **SCL is 8 or higher**: Silently delete the message.
- **SCL is 7**: Reject the message with an NDR.
- **SCL is 6**: Quarantine the message.
- **SCL is 5**: Deliver the message to the user's Junk Email folder.
- **SCL is 4 or lower**: Deliver the message to the user's Inbox.

For more information, see Exchange spam confidence level (SCL) thresholds.

As you monitor the spam quarantine mailbox, you can view the results of antispam filtering by inspecting the antispam stamps (X-header fields) that were applied to the message. For more information, see View antispam stamps in Outlook. You can then adjust the SCL thresholds to more accurately filter the spam that's coming into your organization. For example:

- Too many legitimate messages are sent to the spam quarantine mailbox (too many false positives).

- Too many obvious spam messages are sent to the quarantine mailbox (not enough spam is rejected or deleted).

To release a false positive from the spam quarantine to the intended recipient, see the following topics:

- Configure Outlook to show the original sender in the spam quarantine mailbox
- Release quarantined messages from the spam quarantine mailbox

Configure a spam quarantine mailbox

8/3/2020 • 4 minutes to read • Edit Online

Messages determined to be spam by the Content Filter agent can be directed to a spam quarantine mailbox. If the spam confidence level (SCL) quarantine threshold is enabled, all messages that are quarantined are wrapped as non-delivery reports (also known as NDRs, delivery status notifications, DSN, or bounce messages) and are delivered to the spam quarantine mailbox that you specify. Administrators can review quarantined messages and release them to their intended recipients by using Microsoft Outlook.

What do you need to know before you begin?

- Estimated time to complete this task: 30 minutes.
- By default, antispam features aren't enabled in the Transport service on a Mailbox server. Typically, you only enable the antispam features on a Mailbox server if your Exchange organization doesn't do any prior antispam filtering before accepting incoming messages. For more information, see Enable antispam functionality on Mailbox servers.
- The person that's responsible for the spam quarantine mailbox can view potentially private and sensitive messages, and then send mail on behalf of anybody in the Exchange organization.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Step 1: Verify content filtering is enabled

You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Antispam features" entry in the Antispam and antimalware permissions topic.

1. Run the following command to verify that the Content Filter agent is installed and enabled on the Exchange server:

```
Get-TransportAgent "Content Filter Agent"
```

2. Run the following command to verify content filtering is enabled:

```
Get-ContentFilterConfig | Format-List Enabled
```

For more information, see Content filtering procedures.

Step 2: Create a dedicated mailbox for spam quarantine

To create a spam quarantine mailbox, follow these steps:

- **Create a dedicated Exchange database**: We recommend that you create a dedicated database for the spam quarantine mailbox. The spam quarantine mailbox should have a large database, because if the storage quota limit is reached, messages will be lost. For more information, see Manage mailbox databases in Exchange Server.

- **Create a dedicated mailbox and user account**: We recommend that you create a dedicated mailbox and user account for the spam quarantine mailbox. For more information, see Create user mailboxes in Exchange Server.

  You can apply recipient policies, such as messaging records management, mailbox quotas, and delegation rights, according to your organization's compliance policies and needs. For more information, see Messaging records management in Exchange Server.

  > **NOTE**
  >
  > If a quarantined message is rejected because of a storage quota, the message will be lost. Exchange doesn't generate NDRs for quarantined messages because the quarantined messages are wrapped as NDRs.

- **Configure Outlook**: You need to configure the Outlook delegate access permissions to meet the needs of your organization. In addition, you can configure the Outlook profile to show the original sender, recipient, and SCL value of the message. For more information, see Configure Outlook to show the original sender in the spam quarantine mailbox.

Step 3: Specify the spam quarantine mailbox

You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Antispam features" entry in the Antispam and antimalware permissions topic.

Use the following syntax:

```
Set-ContentFilterConfig -QuarantineMailbox <SmtpAddress>
```

This example sends all messages that exceed the spam quarantine threshold to spamQ@contoso.com.

```
Set-ContentFilterConfig -QuarantineMailbox spamQ@contoso.com
```

How do you know this step worked?

To verify that you have successfully specified the spam quarantine mailbox, run the following command to verify the value of the **QuarantineMailbox** property:

```
Get-ContentFilterConfig | Format-List QuarantineMailbox
```

Step 4: Configure the SCL quarantine threshold

The SCL quarantine threshold is the SCL value that redirects a message to the spam quarantine mailbox. You can set the SCL quarantine threshold to a value from 0 through 9, where 0 is considered less likely to be spam, and 9 is considered most likely to be spam.

For more information about how to adjust SCL thresholds to suit your organization's requirements, and how to configure per-mailbox SCL thresholds, see Use the Exchange Management Shell to configure SCL thresholds for content filtering and Use the Exchange Management Shell to configure the SCL thresholds on a mailbox.

Step 5: Manage the spam quarantine mailbox

When you manage your spam quarantine mailbox, follow these guidelines:

- Use **Resend this message** in Outlook to release quarantined messages to their intended recipients. For more information, see Release quarantined messages from the spam quarantine mailbox.
- Monitor the size of the spam quarantine mailbox. The volume of email messages can change because of a large influx of new employees, the natural trend of larger message sizes, or the threshold value on the SCL quarantine action.
- Monitor the spam quarantine mailbox for false positives. If your spam quarantine mailbox includes many false positives, increase your SCL quarantine threshold. For more information about how to determine why false positives are being delivered to the spam quarantine mailbox, see View antispam stamps in Outlook.
- Use the same Outlook profile to view and release quarantined messages from the spam quarantine mailbox. Applying permissions to a different Outlook profile to release messages isn't supported.

> **IMPORTANT**
>
> NDRs for quarantined messages aren't delivered to the spam quarantine mailbox. NDRs that are identified as spam are deleted, even if their SCL value indicates that they should be quarantined. To track these messages, use the agent log or the message tracking log. For more information, see Antispam Agent Logging.

Step 6: Adjust the SCL quarantine threshold

After you configure the SCL quarantine threshold, periodically monitor the settings and adjust them based on your organization's needs. For example, if too many false positives are delivered to the spam quarantine mailbox, raise the SCL quarantine threshold to a larger value. For more information about how to adjust the SCL quarantine threshold, see Use the Exchange Management Shell to configure SCL thresholds for content filtering.

Configure Outlook to show the original sender in the spam quarantine mailbox

8/3/2020 • 5 minutes to read • Edit Online

Spam quarantine is a feature of the Content Filter agent that reduces the risk of losing legitimate messages. Spam quarantine provides a temporary storage location for messages that are identified as spam and that shouldn't be delivered to a user mailbox inside the organization. For more information, see Spam quarantine in Exchange Server.

When a message meets the spam quarantine threshold, it's wrapped in a non-delivery report (also known as an NDR, delivery status notification, DSN, or bounce message) and delivered to the spam quarantine mailbox. Because the quarantined messages are stored as NDRs, the postmaster address of your organization will be listed as the `From:` address for all messages. However, having the original sender address, the original recipient address, and the original spam confidence level (SCL) in the field list would make it easier to locate the message you want to recover.

By default, you can't add these fields in the message view in Microsoft Outlook. You need to create an Outlook form that adds the original sender, original recipient, and original SCL as optional fields that you can select. After you create this custom form, you can configure Outlook to display these fields in the message view.

What do you need to know before you begin?

- Estimated time to complete this procedure: 15 minutes.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Mailbox access" entry in the Mail flow permissions topic.
- This procedure requires that you've configured the quarantine mailbox. For more information, see Configure a spam quarantine mailbox.
- You need to configure an Outlook profile that you use to access the spam quarantine mailbox. For more information about configuring and using multiple Outlook profiles, see Overview of Outlook e-mail profiles.

  You can use MAPI utilities (for example, OutlookSpy or MFCMAPI) to find the MAPI properties in a quarantined message that contain the original sender, recipient, and SCL values. If you find that other MAPI properties give better results than the ones identified in this topic, you can use them in the custom Outlook form.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Step 1: Use Notepad to create a custom Outlook form

1. Open Notepad, and copy the following code into the document.

```
[Description]
MessageClass=IPM.Note
CLSID={00020D31-0000-0000-C000-000000000046}
DisplayName=Quarantine Extension Form
Category=Standard
Subcategory=Form
Comment=This form allows the Original Sender (ReceivedRepresentingEmailAddress), Original Recipient
(To), and Original SCL (OriginalScl) values to be viewed as columns.
LargeIcon=IPML.ico
SmallIcon=IPMS.ico
Version=3.0
Locale=enu
Hidden=1
Owner=Microsoft Corporation
Contact=Your Name
[Platforms]
Platform1=Win16
Platform2=NTx86
Platform9=Win95
[Platform.Win16]
CPU=ix86
OSVersion=Win3.1
[Platform.NTx86]
CPU=ix86
OSVersion=WinNT3.5
[Platform.Win95]
CPU=ix86
OSVersion=Win95
[Properties]
Property01=ReceivedRepresentingEmailAddress
Property02=DisplayTo
Property03=OriginalScl
[Property.ReceivedRepresentingEmailAddress]
Type=31
NmidInteger=0x0078
DisplayName=ReceivedRepresentingEmailAddress
[Property.DisplayTo]
Type=31
NmidInteger=0x0E04
DisplayName=DisplayTo
[Property.OriginalScl]
Type=3
NmidPropset={41F28F13-83F4-4114-A584-EEDB5A6B0BFF}
NmidString=OriginalScl
DisplayName=OriginalScl
[Verbs]
Verb1=1
[Verb.1]
DisplayName=&Open
Code=0
Flags=0
Attribs=2
[Extensions]
Extensions1=1
[Extension.1]
Type=31
NmidPropset={00020D0C-0000-0000-C000-000000000046}
NmidInteger=1
Value=1000000000000000
```

2. Save the file in your Office Forms folder using the following values:

   - Path: `<OfficeInstallPath>\<OfficeVersion>\Forms\<LCID>`
   - <OfficeInstallPath>:

  - For 32-bit versions of Office on 32-bit versions of Microsoft Windows, or 64-bit versions of Office on 64-bit versions of Windows, the default path is `C:\Program Files\Microsoft Office\root`.

  - For 32-bit versions of Office on 64-bit versions of Windows, the default path is `C:\Program Files (x86)\Microsoft Office\root`.

- **<OfficeVersion>**

  - **Outlook 2010**: `Office14`

  - **Outlook 2013**: `Office15`

  - **Outlook 2016**: `Office16`

- **<LCID>**: This is your locale ID (LCID) value. For example, the LCID for US English is 1033. For more information, see Language identifiers and OptionState Id values in Office.

- **Name**: For the rest of this procedure, assume the file is named `QTNE.cfg`. The name of the file isn't important, but be sure to save the file as QTNE.cfg and not QTNE.cfg.txt.

For example, for a 32-bit US English version of Outlook 2016 installed on a 64-bit version of Windows, save the file as:

```
"C:\Program Files (x86)\Microsoft Office\root\Office16\Forms\1033\QTNE.cfg"
```

> **NOTE**
>
> If Windows User Access Control (UAC) prevents you from saving the file in the correct location, save it first to a temporary location, and then copy it.

Step 2: Configure Outlook 2010 or later to use the custom Outlook form

1. Open the spam quarantine mailbox in Outlook on a client computer, and click **File** > **Options** > **Advanced**.

2. In the **Developers** section, click **Custom Forms**.

3. In the **Options** dialog box that opens, click **Manage Forms**.

4. In the **Forms Manager** dialog box that opens, click **Install**. Browse to the location of the `QTNE.cfg` file, select it, and click **Open**. In the **Form Properties** dialog box, review the information, and then click **OK** to install the Quarantine Extension Form in your Personal Forms library.

5. Back in the **Forms Manager** dialog box, click **Close**. Click **OK** twice to close the remaining dialog boxes and return to the main Outlook interface.

6. In the **Mail** view of the Inbox, click **View** > **Add columns**.

7. In the **Show Columns** dialog box that opens, in the **Select available columns from** drop-down list, scroll to the end of the list and select **Forms**.

8. In the **Select Enterprise forms for this folder** dialog box that opens, in the **Selected Forms** field, select **Message** and click **Remove**. In the **Personal Forms** field, select **Quarantine Extension Form**, and then click **Add**. When you're finished, click **Close**.

9. Back in the **Show Columns** dialog box, in the **Available Columns** section, select one or more of the following fields and click **Add** after each field you select.

- **ReceivedRepresentingEmailAddress**: Original sender
- **DisplayTo**: Original recipient (note that this appears as **To** after you add it)
- **OriginalScl**: Original SCL

  Use the **Move Up** or **Move Down** buttons to position the columns in the view. For best results, position the new fields after the **Attachment** field, and before the **From** field. When you're finished, click **OK** twice to return to the main Outlook interface.

How do you know this worked?

You know this procedure worked if you can see the original sender, original recipient, or original SCL values for quarantined messages in the spam quarantine mailbox using Outlook.

Release quarantined messages from the spam quarantine mailbox

8/3/2020 • 2 minutes to read • Edit Online

After you configure the spam quarantine mailbox in Exchange Server, you can use **Resend this message** in Microsoft Outlook to release quarantined messages to their intended recipients. To configure the spam quarantine mailbox, see Configure a spam quarantine mailbox.

What do you need to know before you begin?

- Estimated time to complete this procedure: less than 5 minutes
- **Resend this message** isn't available in Outlook on the web. You need to configure an Outlook profile that you use to access the spam quarantine mailbox. For more information about configuring and using multiple Outlook profiles, see Overview of Outlook email profiles.
- To make it easier to locate the message you that want to recover, you can create a custom Outlook form to show the original sender and recipients in the message view. For detailed steps, see Configure Outlook to show the original sender in the spam quarantine mailbox.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Mailbox access" entry in the Mail flow permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use Outlook 2010 or later to release a message from the spam quarantine mailbox

1. Open the spam quarantine mailbox in Outlook on a client computer.
2. In the **Mail** view, find the message you want to recover in the **Inbox**, and then double-click the message to open it.
3. In the **Move** section of the Ribbon, click **Actions** > **Resend this Message**.
4. When the message opens, click **Send** to resend the message to the intended recipient.

Note: **Resend this message** doesn't work on multiple messages. You need to resend them one at a time.

How do you know this worked?

To verify that you have successfully released the message from the spam quarantine mailbox, contact the recipient and verify that they received the message.

Exchange spam confidence level (SCL) thresholds

8/3/2020 • 11 minutes to read • Edit Online

NOTE

In November, 2016, Microsoft stopped producing spam definition updates for the SmartScreen filters in Exchange and Outlook. The existing SmartScreen spam definitions were left in place, but their effectiveness will likely degrade over time. For more information, see Deprecating support for SmartScreen in Outlook and Exchange.

In Exchange Server, you can define specific actions for messages according to spam confidence level (SCL) thresholds. For example, you can define different thresholds for rejecting, deleting, or quarantining messages on an Exchange server that's running the Content Filter agent. These SCL thresholds and actions are basically unchanged from Exchange Server 2010

The Content Filter agent assigns an SCL rating to messages late in the antispam cycle, after the other antispam agents have processed inbound messages. Many of the other antispam agents that process inbound messages before the Content Filter agent are absolute in how they act on a message. For example, the Connection Filtering agent on an Edge Transport server rejects messages from IP addresses based on a real-time block list. Similarly, the Sender Filter agent blocks messages based on a list of blocked senders, and the Recipient Filter agent blocks messages based on a list of blocked recipients. By processing messages first, the other antispam agents greatly reduce the number of messages (in most cases, blatantly unwanted messages) that need to be processed by the Content Filter agent. For more information about the order that antispam agents process messages in, see Antispam protection in Exchange Server.

Because content filtering isn't an exact science, it's important to be able to adjust the actions of the Content Filter agent based on different SCL values. By carefully monitoring and adjusting the SCL thresholds, you can minimize the following conditions:

- The number of messages in and the size of the spam quarantine mailbox.
- The number of legitimate email messages that are mistakenly quarantined or placed in the user's Junk Email folder (false positives).
- The number of offensive spam email messages that reach the user's mailbox (messages shouldn't even reach the Junk Email folder).
- The number of spam email messages that reach the user's Inbox.

The combination of this SCL thresholds in the Content Filter agent and the SCL Junk Email folder threshold on the user's mailbox helps you implement a more comprehensive and precise antispam strategy, which can help you reduce the overall cost of deploying and maintaining an antispam solution across your Exchange organization.

SCL threshold actions

By adjusting SCL threshold actions, you can escalate the content filtering action taken on messages that have a greater probability of being spam. To understand this functionality, it's helpful to understand the different SCL threshold actions and how they're implemented:

- **SCL delete threshold**: When the message's SCL value is greater than or equal to the SCL delete threshold, the Content Filter agent silently deletes the message. There's no protocol-level communication that tells the source messaging server or sender that the message was deleted. If the message's SCL value is lower than the SCL delete threshold, the Content Filter agent compares the SCL value to the SCL reject threshold.

- **SCL reject threshold**: When the message's SCL value is greater than or equal to the SCL reject threshold, but less than the SCL delete threshold, the Content Filter agent rejects the message and sends a rejection response to the sending system. You can customize the rejection response. In some cases, a non-delivery report (also known as an NDR, delivery status notification, DSN, or bounce message) is sent to the original sender of the message. If the message's SCL value is lower than the SCL reject threshold, the Content Filter agent compares the SCL value to the SCL quarantine threshold.

- **SCL quarantine threshold**: When the message's SCL value is greater than or equal to the SCL quarantine threshold, but less than the SCL reject threshold, the Content Filter agent sends the message to the spam quarantine mailbox. For more information about the configuring the spam quarantine mailbox, see Configure a spam quarantine mailbox.

  Administrators need periodically review the spam quarantine mailbox to verify that too much obvious spam isn't unnecessarily quarantined (the SCL quarantine threshold is too high), and that too much legitimate email isn't quarantined (the SCL quarantine threshold is too low). To view the results of antispam tests on quarantined messages, see View antispam stamps in Outlook.

  If the message's SCL value is lower than the SCL quarantine threshold, the message is delivered to the appropriate Mailbox server, where the organization's or mailbox's SCL Junk Email folder threshold is evaluated.

- **SCL Junk Email folder threshold**: If the message's SCL value is greater than the SCL Junk Email folder threshold that's configured for the organization or on the mailbox, the message is delivered to the Junk Email folder. If the message's SCL value is equal to or lower than the Junk Email folder threshold, the message is delivered to the Inbox.

  Unlike the other SCL thresholds that are controlled by the Content Filter agent, the SCL Junk Email folder threshold is controlled by the junk email rule (a hidden Inbox rule named Junk E-mail Rule) that's enabled by default in every mailbox. The Content Filter agent assigns the SCL value to a message, but the Junk E-mail Rule is responsible for delivering the message to the Junk Email folder. For more information, see Use the Exchange Management Shell to enable or disable the junk email rule in a mailbox.

The Content Filter agent and the Junk Email folder process the SCL threshold value differently. The Content Filter agent uses greater than or equal to for the SCL threshold value, but the Junk Email folder uses greater than. For example, if you configure the Content Filter agent with an SCL delete threshold of 8, all messages with an SCL of 8 or higher are silently deleted. However, if you configure the Junk Email folder with an SCL threshold of 4, all messages with an SCL of 5 or higher are moved to the Junk Email folder, while messages with an SCL of 4 or lower are delivered to the Inbox.

Scope of SCL thresholds

You can configure the SCL thresholds in the following locations:

- **Server configuration**: The SCL delete, reject, and quarantine thresholds on the Content Filter agent.

- **Organization configuration**: The SCL Junk Email folder threshold value on the Exchange organization.

- **Mailbox configuration**: The SCL thresholds on specific mailboxes.

SCL thresholds on the Content Filter agent

You use the **Set-ContentFilterConfig** cmdlet to configure the SCL delete, reject, and quarantine thresholds on an Edge Transport server or Mailbox server where you're running the Content Filter agent. Over time, as you analyze the spam functionality and metrics provided by the antispam logging and reporting features, you can make additional adjustments to these SCL thresholds as needed.

The SCL threshold parameters that are available on the **Set-ContentFilterConfig** cmdlet are described in the following table.

PARAMETER	DESCRIPTION
SCLDeleteEnabled	Enables and disables the SCL delete threshold. Valid values are `$true` or `$false`. The default value is `$false`, which means the SCL delete threshold isn't enabled by default. You set the SCL delete threshold value with the *SCLDeleteThreshold* parameter.
SCLDeleteThreshold	The SCL value that's used when the SCL delete threshold is enabled. A message with an SCL value that's greater than or equal to this value is silently deleted. The maximum value is 9, which is also the default value. When you enable the SCL delete threshold, this value should be greater than all other SCL thresholds.
SCLRejectEnabled	Enables and disables the SCL reject threshold. Valid values are `$true` or `$false`. The default value is `$true`, which means the SCL reject threshold is enabled by default. You set the SCL reject threshold value with the *SCLRejectThreshold* parameter.
SCLRejectThreshold	The SCL value that's used when the SCL reject threshold is enabled. A message with an SCL value that's greater than or equal to this value is rejected, and an NDR is sent to the sender. The maximum value is 9, and the default value is 7. When you enable the SCL reject threshold, this value should be less than the SCL delete threshold, but greater than the SCL quarantine and Junk Email folder thresholds.
SCLQuarantineEnabled	Enables and disables the SCL quarantine threshold. Valid values are `$true` or `$false`. The default value is `$false`, which means the SCL quarantine threshold isn't enabled by default. You set the SCL quarantine threshold value with the *SCLQuarantineThreshold* parameter. For more information about the configuring the spam quarantine mailbox that's required to quarantine messages, see Configure a spam quarantine mailbox.
SCLQuarantineThreshold	The SCL value that's used when the SCL quarantine threshold is enabled. A message with an SCL value that's greater than or equal to this value is redirected to the spam quarantine mailbox. The maximum value is 9, which is also the default value. When you enable the SCL quarantine threshold, this value should be less than the SCL reject threshold, but greater than the SCL Junk Email folder threshold (the *SCLJunkThreshold* parameter on the **Set-OrganizationConfig** or **Set-Mailbox** cmdlets).

For examples of configuring the SCL thresholds on the Content Filter agent, see Use the Exchange Management Shell to configure SCL thresholds for content filtering.

SCL thresholds on the organization

You use the *SCLJunkThreshold* parameter **Set-OrganizationConfig** cmdlet to set the SCL Junk Email folder threshold value for all mailboxes in the organization. This is the only SCL threshold that you can configure at the organization level. The SCL value is typically assigned to messages by the Content Filter agent.

The *SCLJunkThreshold* parameter on the **Set-OrganizationConfig** cmdlet is described in the following table.

PARAMETER	DESCRIPTION
SCLJunkThreshold	The SCL value that's used when the junk email rule is enabled in the mailbox, and an SCL Junk Email folder threshold isn't configured on the mailbox. A message with an SCL value that's greater than this value is moved to the Junk Email folder by the junk email rule. The maximum value is 9, and the default value is 4, which means that messages with an SCL value of 5 or higher are moved to the Junk Email folder, and messages with an SCL value of 4 or lower are delivered to the Inbox.

Notes:

- The SCL Junk Email folder threshold is enabled by default, because the junk email rule is enabled by default in all mailboxes. If the junk email rule is disabled in the mailbox, the SCL Junk Email folder threshold (for the organization or the mailbox) is disabled for the mailbox.
- You can disable the junk email rule in a mailbox by using the *Enabled* parameter on the **Set-MailboxJunkEmailConfiguration** cmdlet, but only after the mailbox has been opened in Outlook (in Cached Exchange mode) or Outlook on the web.
- You can control the availability of the junk email settings in Outlook on the web, which prevents users from enabling or disabling the junk email rule in their own mailbox.

For more information, see the Configure Exchange antispam settings on mailboxes topic.

SCL thresholds on a mailbox

You can use the **Set-Mailbox** cmdlet to configure all SCL thresholds on a mailbox. The same SCL parameters that are available on the **Set-ContentFilterConfig** cmdlet are also available on the **Set-Mailbox** cmdlet:

- *SCLDeleteEnabled*
- *SCLDeleteThreshold*
- *SCLRejectEnabled*
- *SCLRejectThreshold*
- *SCLQuarantineEnabled*
- *SCLQuarantineThreshold*

Unlike the SCL threshold parameters on **Set-ContentFilterConfig**, the parameters on the **Set-Mailbox** cmdlet also accept the value `$null` (the value is blank), which is the default value for all SCL thresholds on the mailbox. This blank default value indicates that no SCL thresholds are configured on the mailbox, so the Content Filter agent uses its SCL threshold settings for messages that are sent to the mailbox.

If you configure an SCL threshold on a mailbox (the value isn't blank), the setting override the corresponding SCL threshold on the Content Filter agent for messages that are sent to the mailbox. The SCL thresholds that you configure on the mailbox are stored in Active Directory, and are replicated to subscribed Edge Transport servers by the Microsoft Exchange EdgeSync service.

The results are similar for the *SCLJunkThreshold* parameter that's available on **Set-OrganizationConfig** and **Set-Mailbox**: the SCL Junk Email folder threshold value that you configure on the mailbox (the value isn't blank) overrides the SCL value on the organization for messages that are sent to the mailbox.

NOTE

SCL thresholds on a mailbox are not enforced for messages that are received from distribution groups.

The SCL threshold setting that's unique to a mailbox is the ability to enable or disable the SCL Junk Email folder threshold. The *SCLJunkEnabled* parameter is only available on the **Set-Mailbox** cmdlet, and is described in the following table.

PARAMETER	DESCRIPTION
SCLJunkEnabled	Enables and disables the SCL Junk Email folder threshold on the mailbox. Valid values are `$true`, `$false`, or `$null` (blank). The default value is blank (`$null`), which means the SCL Junk Email folder threshold isn't configured on the mailbox, and is controlled by whether the junk email rule is enabled or disabled in the mailbox. The default SCL Junk Email folder threshold value is set by the *SCLJunkThreshold* parameter on the **Set-OrganizationConfig** cmdlet. You can override this value for the mailbox by using the *SCLJunkThreshold* parameter on the **Set-Mailbox** cmdlet.

Notes:

- You can disable the SCL Junk Email folder threshold on a mailbox by disabling the junk email rule in the mailbox. However, disabling the rule also prevents the rule from using the mailbox's safelist collection (Safe Senders list, Safe Recipients list, Blocked Senders list) to move messages to the Junk Email folder, or keep messages out of the Junk Email folder.

- Even if the junk email rule is disabled in the mailbox, and the SCL Junk Email folder threshold is disabled on the mailbox, the client-side Outlook Junk Email Filter can still move messages to the Junk Email folder.

For more information, see the Configure Exchange antispam settings on mailboxes topic.

Monitoring the SCL thresholds

You can use several built-in scripts that are located in the `%ExchangeInstallPath%Scripts` folder, such as **Get-AntispamSCLHistogram.ps1**, for gathering filtering result data. If the data indicates that you need to make immediate adjustments, reconfigure the SCL thresholds. Otherwise, collect data and analyze the spam reporting to determine whether adjustments are required.

Sender filtering

8/3/2020 • 2 minutes to read • Edit Online

Sender filtering compares a list of blocked senders that's maintained by the Exchange administrator to the value of the **MAIL FROM** command in SMTP connections to determine what to do with inbound email messages from those blocked senders. Sender filtering in Exchange Server is provided by the Sender Filter agent, and is basically unchanged from Exchange Server 2010.

You can configure the Sender Filter agent block single senders (for example, kim@contoso.com), whole domains (contoso.com), or domains and all subdomains (*.contoso.com). You can control whether the agent inspects messages from internal sources, external sources, or both. You can also configure the action to take on messages from blocked senders:

- **Reject**: The Sender Filter agent rejects the SMTP request with a `554 5.1.0 Sender Denied` SMTP session error and closes the connection.
- **Stamp status**: The Sender Filter agent accepts the message and updates the message to indicate that it came from a blocked sender. The Content Filter agent uses this information when it calculates the spam confidence level (SCL) of the message. For more information about content filtering and the Content Filter agent, see Content filtering.

By default, the Sender Filter agent is enabled on Edge Transport servers, but you can enable it on Mailbox servers. For more information, see Enable antispam functionality on Mailbox servers.

For more information about how to configure the Sender Filter agent, see Sender filtering procedures.

> **IMPORTANT**
>
> The **MAIL FROM:** SMTP headers can be spoofed, so you shouldn't rely exclusively on the Sender Filter agent. Instead, you should use both the Sender Filter agent and the Sender ID agent. The Sender ID agent uses the originating IP address of the sending server to verify that the domain in the **MAIL FROM:** SMTP header matches the domain that's registered. For more information about the Sender ID agent, see Sender ID.

Using the Sender Filter agent to block messages

By default, the Sender Filter agent is configured to only inspect messages from external sources. *External sources* are defined as unauthenticated sources. You can configure the Sender Filter agent to inspect messages from internal (authenticated) sources. However, as best practice, you typically don't need to apply antispam filters to messages from trusted partners or from inside your organization.

You can also configure the Sender Filter agent to block inbound messages that don't specify a sender and domain in the **MAIL FROM** SMTP command. This setting helps to prevent NDR attacks on the Exchange server. Most legitimate SMTP messages come from SMTP servers that provide a sender and domain in the **MAIL FROM** command.

Specify the action for messages from blocked senders

After you've configured the blocked senders and the sources that are monitored by the Sender Filter agent, you need to configure the Sender Filter agent to reject or accept and stamp messages from those senders. We recommend that you reject the messages, because the chance of false positives based on the specific list of blocked senders is much less than other calculated message properties.

There are only two scenarios where a legitimate message might be rejected by the Sender Filter agent:

- You mistype the blocked sender.
- The domain in your Blocked Senders list is later re-registered to a legitimate company.

Sender filtering procedures

8/3/2020 • 6 minutes to read • Edit Online

Sender filtering filters inbound messages by comparing a list of blocked senders to the value of the **MAIL FROM** command in SMTP connections. For more information about sender filtering and the Sender Filter agent, see Sender filtering.

You can configure many aspects of sender filtering. For example:

- Enable or disable sender filtering on inbound messages from internal (authenticated) and external (unauthenticated) sources (it's enabled by default for messages from external sources).
- Configure blocked senders and blocked domains.
- Specify whether to block messages with blank senders.
- Configure the action that sender filtering takes on messages that contain blocked senders or domains.

What do you need to know before you begin?

- Estimated time to complete each procedure: 5 minutes
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Antispam features" entry in the Antispam and antimalware permissions topic.
- You can only use PowerShell to perform this procedure. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- By default, antispam features aren't enabled in the Transport service on a Mailbox server. Typically, you only enable the antispam features on a Mailbox server if your Exchange organization doesn't do any prior antispam filtering before accepting incoming messages. For more information, see Enable antispam functionality on Mailbox servers.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to enable or disable sender filtering

To disable sender filtering, run the following command:

```
Set-SenderFilterConfig -Enabled $false
```

To enable sender filtering, run the following command:

```
Set-SenderFilterConfig -Enabled $true
```

NOTE

When you disable sender filtering, the underlying Sender Filter agent is still enabled. To disable the Sender Filter agent, run the command: `Disable-TransportAgent "Sender Filter Agent"`.

How do you know this worked?

To verify that you have successfully enabled or disabled sender filtering, run the following command to verify the **Enabled** property value:

```
Get-SenderFilterConfig | Format-List Enabled
```

Use the Exchange Management Shell to enable or disable sender filtering for external connections

By default, sender filtering is enabled for external (unauthenticated) SMTP connections.

To disable sender filtering for external connections, run the following command:

```
Set-SenderFilterConfig -ExternalMailEnabled $false
```

To enable sender filtering for external connections, run the following command:

```
Set-SenderFilterConfig -ExternalMailEnabled $true
```

How do you know this worked?

To verify that you have successfully enabled or disabled sender filtering for external SMTP connections, run the following command to verify the **ExternalMailEnabled** property value:

```
Get-SenderFilterConfig | Format-List ExternalMailEnabled
```

Use the Exchange Management Shell to enable or disable sender filtering for internal connections

As a best practice, you don't need to apply antispam filters to messages from trusted partners or from inside your organization. To reduce the chance that filters will mishandle legitimate email messages, you typically configure antispam agents to only run on messages from external sources.

To enable sender filtering for internal (authenticated) SMTP connections, run the following command:

```
Set-SenderFilterConfig -InternalMailEnabled $true
```

To disable sender filtering for internal connections, run the following command:

```
Set-SenderFilterConfig -InternalMailEnabled $false
```

How do you know this worked?

To verify that you have successfully enabled or disabled sender filtering for internal SMTP connections, run the following command to verify the **InternalMailEnabled** property value:

```
Get-SenderFilterConfig | Format-List InternalMailEnabled
```

Use the Exchange Management Shell to configure blocked senders and domains for sender filtering

You can specify blocked senders and domains that replace the existing values, or you can add or remove specific blocked senders and domains without affecting the other existing values.

To replace the existing values, use the following syntax:

```
Set-SenderFilterConfig -BlockedSenders <sender1,sender2...> -BlockedDomains <domain1,domain2...> -
BlockedDomainsAndSubdomains <domain1,domain2...>
```

This example configures the Sender Filter agent to block messages from kim@contoso.com and john@contoso.com, messages from the fabrikam.com domain, and messages from northwindtraders.com and all its subdomains.

```
Set-SenderFilterConfig -BlockedSenders kim@contoso.com,john@contoso.com -BlockedDomains fabrikam.com -
BlockedDomainsAndSubdomains northwindtraders.com
```

To add or remove entries without modifying other existing values, use the following syntax:

```
Set-SenderFilterConfig -BlockedSenders @{Add="<sender1>","<sender2>"...; Remove="<sender1>","<sender2>"...} -
BlockedDomains @{Add="<domain1>","<domain2>"...; Remove="<domain1>","<domain2>"...} -
BlockedDomainsAndSubdomains @{Add="<domain1>","<domain2>"...; Remove="<domain1>","<domain2>"...}
```

This example configures the Sender Filter agent with the following information:

- Add chris@contoso.com and michelle@contoso.com to the list of existing senders who are blocked.
- Remove tailspintoys.com from the list of existing sender domains that are blocked.
- Add blueyonderairlines.com to the list of existing sender domains and subdomains that are blocked.

```
Set-SenderFilterConfig -BlockedSenders @{Add="chris@contoso.com","michelle@contoso.com"} -BlockedDomains
@{Remove="tailspintoys.com"} -BlockedDomainsAndSubdomains @{Add="blueyonderairlines.com"}
```

How do you know this worked?

To verify that you have successfully configured blocked senders, run the following command to verify the property values:

```
Get-SenderFilterConfig | Format-List Blocked*
```

Use the Exchange Management Shell to configure sender filtering to block messages with blank senders

To enable or disable blocking messages that have blank senders, use the following syntax:

```
Set-SenderFilterConfig -BlankSenderBlockingenabled <$true | $false>
```

This example configures the Sender Filter agent to block messages that don't specify a sender in the **MAIL FROM:** SMTP command:

```
Set-SenderFilterConfig -BlankSenderBlockingEnabled $true
```

How do you know this worked?

To verify that you have successfully enabled or disabled blocking messages with blank senders, run the following command to verify the property value:

```
Get-SenderFilterConfig | Format-List BlankSenderBlockingEnabled
```

Use the Exchange Management Shell to configure the action for sender filtering

Typically, you want to reject messages from blocked senders or domains, and this is the default action. However, you can configure sender filtering to allow these message into your organization for further analysis by other antispam agents.

To configure the action that sender filtering takes on messages from blocked senders or domains, use the following syntax:

```
Set-SenderFilterConfig -Action <Reject | StampStatus>
```

This example configures the Sender Filter agent to allow messages from blocked senders or domains. The Sender Filter agent updates the message to indicate that it came from a blocked sender. This information is used in the calculation of the message's spam confidence level (SCL).

```
Set-SenderFilterConfig -Action StampStatus
```

This example configures the Sender Filter agent to reject messages from blocked senders or domains. The Sender Filter agent rejects the SMTP request with a `554 5.1.0 Sender Denied` SMTP session error and closes the connection.

```
Set-SenderFilterConfig -Action Reject
```

How do you know this worked?

To verify that you have successfully configured the action for sender filtering, run the following command to verify the **Action** property value:

```
Get-SenderFilterConfig | Format-List Action
```

Use the Exchange Management Shell to configure the action for sender filtering for blocked senders from SafeList aggregation

SafeList aggregation adds blocked senders that are defined by your users in Microsoft Outlook or Outlook on the web to the Blocked Senders list that's used by the Sender Filter agent. For more information, see Safelist aggregation.

To configure the action that sender filtering takes on messages that contain blocked senders that are defined by SafeList aggregation, use the following syntax:

```
Set-SenderFilterConfig -RecipientBlockedSenderAction <Delete | Reject>
```

This example configures the Sender Filter agent to silently drop messages that contain blocked senders that are defined by SafeList aggregation.

```
Set-SenderFilterConfig -RecipientBlockedSenderAction Delete
```

This example configures the Sender Filter agent to reject messages that contain blocked senders that are defined by SafeList aggregation with a non-delivery report (also known as an NDR, delivery status notification, DSN or bounce message).

```
Set-SenderFilterConfig -RecipientBlockedSenderAction Reject
```

How do you know this worked?

To verify that you have successfully configured the action for sender filtering for blocked senders from SafeList aggregation, run the following command to verify the **RecipientBlockedSenderAction** property value:

```
Get-SenderFilterConfig | Format-List RecipientBlockedSenderAction
```

Sender ID

8/3/2020 • 4 minutes to read • Edit Online

Sender ID is used to detect *spoofing*. A spoofed email message is modified to appear as if it originates from a sender other than the actual sender of the message. In the past, it was relatively easy to send spoofed email messages, because the sender's email address in the message header wasn't validated. Sender ID uses the **RECEIVED** SMTP header and a query to the DNS records for the sender's domain to determine if the sender's email address is spoofed. Sender ID in Exchange Server is provided by the Sender ID agent, and is basically unchanged from Exchange Server 2010.

By default, the Sender ID agent is enabled on Edge Transport servers, but you can enable it on Mailbox servers. For more information, see Enable antispam functionality on Mailbox servers.

For more information about how to configure the Sender ID agent, see Sender ID procedures.

Using Sender ID to combat spoofing

When the Exchange server receives an inbound message, the Sender ID agent verifies the sender's IP address by querying the DNS records for the sender's domain. This check confirms that the message was received from an authorized IP address for the sender's domain. The IP address of the authorized sending server is referred to as the *purported responsible address* (PRA).

Administrators publish sender policy framework (SPF) records in DNS that identify the authorized outbound messaging servers for the domain. If an SPF record is available in DNS for the sender's domain, the Sender ID agent parses the SPF record to determine if the source IP address is authorized to send email for the domain that's specified in the sender's email address. For more information about what an SPF record contains and how to create an SPF record, see Sender Policy Framework: SPF Record Syntax.

Sender ID status values

The Sender ID agent generates a Sender ID status for the message. The Sender ID status can be set to one of the following values:

- **Pass**: Both the IP address and the PRA passed the Sender ID verification check.
- **Neutral**: The published Sender ID data is explicitly inconclusive.
- **Soft fail**: The IP address for the PRA might be in the not permitted set.
- **Fail**: The IP Address is not permitted. No PRA is found in the incoming mail, or the sender's domain doesn't exist.
- **None**: No published SPF data exists in DNS for the sender's domain.
- **TempError**: A temporary DNS failure occurred, such as an unavailable DNS server.
- **PermError**: The DNS record is invalid, such as an error in the record format.

Note:: If the source IP address is missing, the Sender ID status can't be set. Exchange continues to process the message without including a Sender ID status, and the message isn't returned or rejected. In this scenario, the Sender ID status isn't set, and an application event is logged.

The Sender ID status is added to the message metadata, and is later converted to a MAPI property. The junk email filter in Outlook uses this MAPI property during the calculation of the spam confidence level (SCL).

Outlook neither displays the Sender ID status, nor flags a message as junk based solely on the Sender ID value. Instead, Outlook uses the Sender ID status value only during the calculation of the SCL for the message.

For more information about how the Sender ID status is displayed in messages, see Antispam stamps.

Sender ID options for handling spoofed mail and unreachable DNS servers

You can configure the actions to take when the Sender ID agent identifies messages that contain spoofed senders (the Sender ID status is `Fail`), and when a DNS server can't be reached (the Sender ID status is `TempError`):

- **Stamp status**: The Sender ID agent stamps the Sender ID status in the metadata of the message, and allows the delivery of the message to continue. This is the default option.
- **Reject**: The Sender ID agent rejects the message with a 5 *xx* level SMTP error response, which includes text that corresponds to the Sender ID status.
- **Delete**: The Sender ID agent silently deletes the message without an SMTP error response. The Exchange server sends a fake **OK** SMTP command to the source server, and then deletes the message. Because the source server assumes the message was sent, it doesn't try to resend the message in the same session.

For more information about how to configure the action to take for spoofed mail and unreachable DNS servers, see Sender ID procedures.

Updating your organization's Internet facing DNS to support Sender ID

The effectiveness of Sender ID depends on specific DNS data. The more organizations that configure SPF records for their domains, the more effectively Sender ID is able to identify spoofed messages.

To support the Sender ID infrastructure, you need to create SPF records for the domains that your organization sends messages from. For more information about how to create and deploy SPF records, see Sender Policy Framework: SPF Record Syntax.

Specifying recipients and sender domains to exclude from Sender ID filtering

You can exclude specific recipients and sender domains from Sender ID filtering by using the **Set-SenderIdConfig** cmdlet in the Exchange Management Shell. For more information, see Sender ID procedures.

Sender ID procedures

8/3/2020 • 4 minutes to read • Edit Online

Sender ID detects spoofed email messages by using the Sender Policy Framework (SPF) record in DNS to compare the source IP address with the domain in the sender's email address. For more information about Sender ID and the Sender ID agent, see Sender filtering

What do you need to know before you begin?

- Estimated time to complete each procedure: 5 minutes
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Antispam features" entry in the Antispam and antimalware permissions topic.
- You can only use PowerShell to perform this procedure. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- By default, antispam features aren't enabled in the Transport service on a Mailbox server. Typically, you only enable the antispam features on a Mailbox server if your Exchange organization doesn't do any prior antispam filtering before accepting incoming messages. For more information, see Enable antispam functionality on Mailbox servers.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to enable or disable Sender ID

To disable Sender ID, run the following command:

```
Set-SenderIDConfig -Enabled $false
```

To enable Sender ID, run the following command:

```
Set-SenderIDConfig -Enabled $true
```

NOTE

When you disable Sender ID, the underlying Sender ID agent is still enabled. To disable the Sender ID agent, run the command: `Disable-TransportAgent "Sender ID Agent"`.

How do you know this worked?

To verify that you have successfully enabled or disabled Sender ID, run the following command to verify the `Enabled` property value:

```
Get-SenderIDConfig | Format-List Enabled
```

Use the Exchange Management Shell to enable or disable Sender ID for external connections

By default, Sender ID is enabled for external (unauthenticated) SMTP connections.

To disable sender filtering for external connections, run the following command:

```
Set-SenderIDConfig -ExternalMailEnabled $false
```

To enable Sender ID for external connections, run the following command:

```
Set-SenderIDConfig -ExternalMailEnabled $true
```

How do you know this worked?

To verify that you have successfully enabled or disabled Sender ID for external SMTP connections, run the following command to verify the **ExternalMailEnabled** property value:

```
Get-SenderFilterConfig | Format-List ExternalMailEnabled
```

Use the Exchange Management Shell to enable or disable Sender ID for internal connections

As a best practice, you don't need to apply antispam filters to messages from trusted partners or from inside your organization. To reduce the chance that filters will mishandle legitimate email messages, you typically configure antispam agents to only run on messages from external sources.

To enable Sender ID for internal (authenticated) SMTP connections, run the following command:

```
Set-SenderIDConfig -InternalMailEnabled $true
```

To disable Sender ID for internal connections, run the following command:

```
Set-SenderIDConfig -InternalMailEnabled $false
```

How do you know this worked?

To verify that you have successfully enabled or disabled Sender ID for internal SMTP connections, run the following command to verify the **InternalMailEnabled** property value:

```
Get-SenderIDConfig | Format-List InternalMailEnabled
```

Use the Exchange Management Shell to configure the Sender ID action for spoofed messages

To configure the Sender ID action for spoofed messages, use the following syntax:

```
Set-SenderIDConfig -SpoofedDomainAction <StampStatus | Reject | Delete>
```

This example configures the Sender ID agent to reject any messages with a 5 *xx* SMTP error response when sender's domain has an SPF record, and the IP address of the source server isn't listed as an authoritative server for the domain (the Sender ID status is `Fail`).

```
Set-SenderIDConfig -SpoofedDomainAction Reject
```

How do you know this worked?

To verify that you have successfully configured the Sender ID action for spoofed messages, run the following command to verify the **SpoofedDomainAction** property value:

```
Get-SenderIDConfig | Format-List SpoofedDomainAction
```

Use the Exchange Management Shell to configure the Sender ID action for transient errors

To configure the Sender ID action for transient errors, use the following syntax:

```
Set-SenderIDConfig -TempErrorAction <StampStatus | Reject | Delete>
```

This example configures the Sender ID agent to stamp the messages when the Sender ID status can't be determined due to a temporary DNS server error (the Sender ID status is `TempError`). The message will be processed by other antispam agents and the Content Filter agent will use the mark when determining the SCL value for the message.

```
Set-SenderIDConfig -TempErrorAction StampStatus
```

Note that `StampStatus` is the default value for the *TempErrorAction* parameter.

How do you know this worked?

To verify that you have successfully configured the Sender ID action for transient errors, run the following command to verify the **TempErrorAction** property value:

```
Get-SenderIDConfig | Format-List TempErrorAction
```

Use the Exchange Management Shell to configure recipient and sender domain exceptions

To replace the existing values, run the following command:

```
Set-SenderIDConfig -BypassedRecipients <recipient1,recipient2...> -BypassedSenderDomains <domain1,domain2...>
```

This example configures the Sender ID agent to bypass the Sender ID check for messages sent to kim@contoso.com and john@contoso.com, and to bypass the Sender ID check for messages sent from the fabrikam.com domain.

```
Set-SenderIDConfig -BypassedRecipients kim@contoso.com,john@contoso.com -BypassedSenderDomains fabrikam.com
```

To add or remove entries without modifying other existing values, use the following syntax:

```
Set-SenderIDConfig -BypassedRecipients @{Add="<recipient1>","<recipient2>"...; Remove="<recipient1>","
<recipient2>"...} -BypassedSenderDomains @{Add="<domain1>","<domain2>"...; Remove="<domain1>","<domain2>"...}
```

This example configures the Sender ID agent with the following settings:

- Add chris@contoso.com and michelle@contoso.com to the list of existing recipients who bypass the Sender ID check.
- Remove tailspintoys.com from the list of existing domains that bypass the Sender ID check.

```
Set-SenderIDConfig -BypassedRecipients @{Add="chris@contoso.com","michelle@contoso.com"} -
BypassedSenderDomains @{Remove="tailspintoys.com"}
```

How do you know this worked?

To verify that you have successfully configured recipient and sender domain exceptions, run the following command to verify the property values:

```
Get-SenderIDConfig | Format-List BypassedRecipients,BypassedSenderDomains
```

Sender reputation and the Protocol Analysis agent

8/3/2020 • 8 minutes to read • Edit Online

Sender reputation is part of the Exchange antispam functionality that blocks messages according to many characteristics of the sender. Sender reputation relies on persisted data about the sender to determine the action to take on inbound messages. The Protocol Analysis agent is the underlying agent for sender reputation functionality.

For more information about how to configure sender reputation and the Protocol Analysis agent, see Sender reputation procedures.

By default, the Protocol Analysis agent is enabled on Edge Transport servers, but you can enable it on Mailbox servers. For more information, see Enable antispam functionality on Mailbox servers.

Calculating the sender reputation level (SRL)

A sender reputation level (SRL) is calculated from the following statistics:

- **HELO/EHLO analysis**: The HELO and EHLO SMTP commands are intended to provide the domain name, such as Contoso.com, or IP address of the sending SMTP server to the receiving SMTP server. Malicious users, or *spammers*, frequently forge the HELO/EHLO statement in various ways. For example, they type an IP address that doesn't match the IP address from which the connection originated. Spammers also put domains that are known to be locally supported at the receiving server in the HELO statement in an attempt to appear as if the domains are in the organization. In other cases, spammers change the domain that's passed in the HELO statement. The typical behavior of a legitimate user may be to use a different, but relatively constant, set of domains in their HELO statements.

  Therefore, analysis of the HELO/EHLO statement on a persender basis may indicate that the sender is likely to be a spammer. For example, a sender that provides many different unique HELO/EHLO statements in a specific time period is more likely to be a spammer. Senders who consistently provide an IP address in the HELO statement that doesn't match the originating IP address as determined by the Connection Filtering agent are also more likely to be spammers. Remote senders who consistently provide a local domain name in the HELO statement that's in the same organization as the Exchange server are also more likely to be spammers.

- **Reverse DNS lookup**: Sender reputation also verifies that the originating IP address from which the sender transmitted the message matches the registered domain name that the sender submits in the HELO or EHLO SMTP command.

  Sender reputation performs a reverse DNS query by submitting the originating IP address to DNS. The result that's returned by DNS is the domain name that's registered by using the domain naming authority for that IP address. Sender reputation compares the domain name that's returned by DNS to the domain name that the sender submitted in the HELO/EHLO SMTP command. If the domain names don't match, the sender is likely to be a spammer, and the overall SRL rating for the sender is increased.

  The Sender ID agent performs a similar task, but the success of the Sender ID agent relies on legitimate senders to update their DNS infrastructure to identify all the email-sending SMTP servers in their organization. By performing a reverse DNS lookup, you can help identify potential spammers.

- **Analysis of SCL ratings on messages from a particular sender**: When the Content Filter agent processes a message, it assigns a spam confidence level (SCL) rating to the message. The SCL rating is a number from 0 through 9. A higher SCL rating indicates that a message is more likely to be spam. Data about each sender and the SCL ratings that their messages yield is persisted for analysis by sender reputation. Sender reputation calculates statistics about a sender according to the ratio between all messages from that sender that had a low SCL rating in the past and all messages from that sender that had a high SCL rating in the past. Additionally, the number of messages that have a high SCL rating that the sender has sent in the last day is applied to the overall SRL.

- **Sender open proxy test**: An *open proxy* is a proxy server that accepts connection requests from anyone anywhere and forwards the traffic as if it originated from the local hosts. Proxy servers relay TCP traffic through firewall hosts to provide user applications transparent access across the firewall. Because proxy protocols are lightweight and independent of user application protocols, proxies can be used by many different services. Proxies can also be used to share a single Internet connection by multiple hosts. Proxies are usually set up so that only trusted hosts inside the firewall can cross through the proxies. A legitimate sender may be an open proxy because of an unintentional misconfiguration or malware.

  Open proxies provide an ideal way for malicious users to hide their true identities and launch denial of service attacks (DoS) or send spam. As more proxy servers are configured to be open by default, open proxies have become more common. Additionally, malicious users can use multiple open proxies together to hide the sender's originating IP address.

  When sender reputation performs an open proxy test, it does so by formatting an SMTP request in an attempt to connect back to the Exchange server from the open proxy. If an SMTP request is received from the proxy, sender reputation verifies that the proxy is an open proxy and updates the open proxy test statistic for that sender.

Sender reputation weighs each of these statistics and calculates an SRL for each sender. The SRL is a number from 0 through 9 that predicts the probability that a specific sender is a spammer or otherwise malicious user. A value of 0 indicates that the sender isn't likely to be a spammer; a value of 9 indicates that the sender is likely to be a spammer.

You can configure a block threshold from 0 through 9 at which sender reputation issues a request to the Sender Filter agent, and, therefore, blocks the sender from sending a message into the organization. When a sender is blocked, the sender is added to the Blocked Senders list for a configurable time period. How blocked messages are handled depends on the configuration of the Sender Filter agent. The following actions are the options for handling blocked messages:

- **Reject**: Messages are returned in a non-delivery report (also known as an NDR, delivery status notification, DSN, or bounce message)

- **Delete**: Messages are silently deleted without an NDR.

- **Accept**: Messages are accepted and marked as coming from a blocked sender

For more information about the Sender Filter agent, see Sender filtering.

If a sender is included in the IP Block list or Microsoft IP Reputation Service, sender reputation issues an immediate request to the Sender Filter agent to block the sender. To take advantage of this functionality, you need to enable and configure the Microsoft Exchange Antispam Update Service.

By default, sender reputation sets a rating of 0 for senders that haven't been analyzed. After a sender has sent 20 or more messages, sender reputation calculates an SRL that's based on the statistics described earlier in this topic.

When to use the SRL

Sender reputation acts on messages during two phases of the SMTP session:

- **At the MAIL FROM: SMTP command**: Sender reputation acts on a message only if the message was blocked or otherwise acted on by the Connection Filtering agent, Sender Filter agent, Recipient Filter agent, or Sender ID agent. In this case, sender reputation retrieves the sender's current SRL rating from the sender profile that's persisted about that sender on the Exchange server. After this rating is retrieved and evaluated, the Exchange server configuration dictates the behavior that occurs at a particular connection according to the block threshold.

- **After the "end of data" SMTP command**: The end of data transfer (**EOD**) SMTP command is given when all the actual message data is sent. At this point in the SMTP session, many of the antispam agents have processed the message. As a by-product of antispam processing, the statistics that sender reputation relies on are updated. Therefore, sender reputation has the data to calculate or recalculate an SRL rating for the sender.

Configuring the detection of open proxy servers

When sender reputation calculates an SRL, sender reputation tries to connect to the sender's originating IP address by using a variety of common proxy protocols, such as SOCKS4, SOCKS5, HTTP, Telnet, Cisco, and Wingate. Sender reputation formats a protocol-specific request in an attempt to connect back to the Exchange server from the open proxy server by using an SMTP request. If an SMTP request is received from the proxy server, sender reputation verifies that the proxy server is an open proxy server and adjusts the SRL rating according to this result. By default, the detection of open proxy servers is enabled in sender reputation.

For more information about how to configure the detection of open proxy servers, see Sender reputation procedures.

Setting the SRL block threshold

The SRL is a number from 0 through 9 that predicts the probability that a specific sender is a spammer or otherwise malicious user. You need to set an SRL threshold for sender blocking to specify the SRL value that causes sender reputation to block a sender. By default, the SRL block threshold is 7, which means senders that have an SRL of 7, 8 or 9 are blocked.. You should monitor the effectiveness of sender reputation and the Protocol Analysis agent at the default level.

On an Edge Transport server, if the SRL block threshold is met or exceeded by a particular sender, sender reputation adds the sender to the IP Block list on the Connection Filtering agent. Sometimes, spammers send batches of spam from a single sender. In this scenario, if sender reputation calculates an SRL that exceeds the SRL block threshold, the sender is added to the Sender Block List for a configurable duration of time. The default duration is 24 hours. After 24 hours, the sender is removed from the Sender Block List and can send messages again.

When a sender is added to the IP Block list, sender reputation deletes the profile for the sender. Sender reputation deletes the profile because the blocked sender's existing profile indicates that the sender's SRL exceeds the SRL block threshold. This would cause the blocked sender to be added to the IP Block list again as soon as the duration for sender blocking ends.

For more information about how to configure sender blocking, see Sender reputation procedures.

Sender reputation procedures

8/3/2020 • 5 minutes to read • Edit Online

Sender reputation and the Protocol Anaysis agent block unwanted messages according to various characteristics of the sender. Sender reputation relies on persisted data about the sender to determine what action, if any, to take on an inbound message. For more information, see Sender reputation and the Protocol Analysis agent.

What do you need to know before you begin?

- Estimated time to complete each procedure: 5 minutes
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Antispam features" entry in the Antispam and antimalware permissions topic.
- You can only use PowerShell to perform this procedure. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- By default, antispam features aren't enabled in the Transport service on a Mailbox server. Typically, you only enable the antispam features on a Mailbox server if your Exchange organization doesn't do any prior antispam filtering before accepting incoming messages. For more information, see Enable antispam functionality on Mailbox servers.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to enable or disable sender reputation

To disable sender reputation, run the following command:

```
Set-SenderReputationConfig -Enabled $false
```

To enable sender reputation, run the following command:

```
Set-SenderReputationConfig -Enabled $true
```

NOTE

The Protocol Analysis agent is the underlying agent for sender reputation functionality. When you disable sender reputation, the Protocol Analysis agent is still enabled. To disable the Protocol Analysis agent, run the command: `Disable-TransportAgent "Protocol Analysis Agent"`.

How do you know this worked?

To verify that you have successfully enabled or disabled sender reputation, run the following command to verify the **Enabled** property value:

```
Get-SenderReputationConfig | Format-List Enabled
```

Use the Exchange Management Shell to enable or disable sender reputation for external messages

By default, sender reputation is enabled for external messages (messages from external sources).

To disable sender reputation for external messages, run the following command:

```
Set-SenderReputationConfig -ExternalMailEnabled $false
```

To enable sender reputation for external messages, run the following command:

```
Set-SenderReputationConfig -ExternalMailEnabled $true
```

How do you know this worked?

To verify that you have successfully enabled or disabled sender reputation for external messages, run the following command to verify the **ExternalMailEnabled** property value:

```
Get-SenderReputationConfig | Format-List ExternalMailEnabled
```

Use the Exchange Management Shell to enable or disable sender reputation for internal messages

As a best practice, you don't need to apply antispam filters to messages from trusted partners or from inside your organization. There's always a chance that the filters will detect false positives. To reduce the chance that filters will mishandle legitimate email messages, you should typically configure antispam agents to only run on messages from untrusted and unknown sources.

To enable sender reputation for internal messages, run the following command:

```
Set-SenderReputationConfig -InternalMailEnabled $true
```

To disable sender reputation for internal messages, run the following command:

```
Set-SenderReputationConfig -InternalMailEnabled $false
```

How do you know this worked?

To verify that you have successfully enabled or disabled sender reputation for internal messages, run the following command to verify the **InternalMailEnabled** property value:

```
Get-SenderReputationConfig | Format-List InternalMailEnabled
```

Use the Exchange Management Shell to configure sender blocking in

sender reputation

Sender blocking uses the calculated sender reputation level (SRL) of the sender and a specified SRL threshold to temporarily block the sender. To configure the sender blocking in sender reputation, use the following syntax:

```
Set-SenderReputationConfig -SenderBlockingEnabled <$true | $false> -SrlBlockThreshold <0 - 9> [-
SenderBlockingPeriod <0 - 48>]
```

This example lowers the sender reputation level (SRL) block threshold to 6 (which means senders with an SRL of 6, 7, 8, or 9 are blocked), and blocks the offending senders for 36 hours:

```
Set-SenderReputationConfig -SrlBlockThreshold 6 -SenderBlockingPeriod 36
```

This example disables sender blocking.

```
Set-SenderReputationConfig -SenderBlockingEnabled $false
```

Notes:

- The default value of the *SenderBlockingEnabled* parameter is `$true`.
- The default value of the *SenderBlockingPeriod* parameter is 24.
- The default value of the *SrlBlockThreshold* parameter is 7.
- You can't disable sender blocking and open proxy server detection at the same time. One must be enabled when the other is disabled, or they both can be enabled.

How do you know this worked?

To verify that you have successfully configured sender blocking in sender reputation, run the following command to verify the property values:

```
Get-SenderReputationConfig | Format-List *block*
```

Use the Exchange Management Shell to configure open proxy server detection in sender reputation

By default, sender reputation uses open proxy server detection as one of the criteria to calculate the SRL of the source server. In open proxy server detection, the Exchange server tries to send a test message from the source messaging server. If the test message is successfully delivered back to the Exchange server, it indicates the source server is configured as an open proxy server (intentionally or unintentionally).

Open proxy server detection uses the protocols and TCP ports that are described in the following table, so these outbound ports need to be open in your firewall:

PROTOCOLS	PORTS
SOCKS4, SOCKS5	1081, 1080
Wingate, Telnet, Cisco	23
HTTP CONNECT, HTTP POST	6588, 3128, 80

Also, if your organization uses a proxy server to control outbound Internet traffic, you need to configure sender reputation to use your proxy server to access the Internet. Specifically, you need to define the proxy server name, type, and TCP port that sender reputation requires to access the Internet.

To configure open proxy server detection in sender reputation, use the following syntax:

```
Set-SenderReputationConfig -OpenProxyDetectionEnabled <$true | $false> [-ProxyServerName <String> -
ProxyServerPort <Port> -ProxyServerType <None | Socks4 | Socks5 | HttpConnect | HttpPost | Telnet | Cisco |
Wingate>]
```

This example configures sender reputation to connect to the Internet through the proxy server named SERVER01 that uses the HTTP CONNECT protocol on port 80.

```
Set-SenderReputationConfig -ProxyServerName SERVER01 -ProxyServerPort 80 -ProxyServerType HttpConnect
```

This example disables open proxy server detection in sender reputation.

```
Set-SenderReputationConfig -OpenProxyDetectionEnabled $false
```

Notes:

- The default value of the *OpenProxyDetectionEnabled* parameter is `$true`.
- The default value of the *ProxyServerName* parameter is blank (`$null`).
- The default value of the *ProxyServerPort* parameter is 0.
- The default value of the *ProxyServerType* parameter is `None`.
- You can't disable open proxy server detection and sender blocking at the same time. One must be enabled when the other is disabled, or they both can be enabled.

How do you know this worked?

To verify that you have successfully configured open proxy server detection in sender reputation, run the following command to verify the property values:

```
Get-SenderReputationConfig | Format-List *proxy*
```

See also

Get-SenderReputationConfig

Set-SenderReputationConfig

Attachment filtering on Edge Transport servers

8/3/2020 • 3 minutes to read • Edit Online

In Exchange Server, you can use attachment filtering on Edge Transport servers to control the attachments that users receive in email messages. Attachment filtering is performed by the Attachment Filtering agent, which is available only on Edge Transport servers, and is basically unchanged from Exchange Server 2010.

To configure the attachment filtering options, see Attachment filtering procedures on Edge Transport servers.

Types of attachment filtering

You can use the following types of attachment filtering to control attachments that enter or leave your organization through an Edge Transport server:

- **Filtering based on file name or file name extension**: You specify the exact file name or file name extension that you want to filter. For example, `BadFileName.exe` or `*.exe`.
- **Filtering based on file MIME content type**: You specify the MIME content type value that you want to filter. The MIME content type value indicates what the attachment is: for example, a JPEG image, an executable file, or a Microsoft Excel file. Content types are expressed as *<type>/ <subtype>*. For example, a JPEG image file is expressed as `image/jpeg`.

  To view a complete list of file name extensions and content types that attachment filtering can detect, run the following command in the Exchange Management Shell on the Edge Transport server:

  ```
  Get-AttachmentFilterEntry | Format-Table -Auto Type,Name
  ```

After you define the files to look for, you can configure the action to take on messages that contain these attachments. You can't specify different actions for different types of attachments. You configure one of the following actions for all the messages that match any of the attachment filters:

- **Reject (block) the message**: he message is blocked. The sender receives a non-delivery report (also known as an NDR, delivery status notification, DSN, or bounce message) that explains that the message wasn't delivered because it contained an unacceptable attachment. You can customize the text in the NDR. The default text is: `Message rejected due to unacceptable attachments`.
- **Strip the attachment but allow the message through**: The attachment is removed from the message. However, the message itself and any other attachments that don't match the filter are allowed through. If an attachment is stripped, it's replaced with a text file that explains why the attachment was removed. This is the default action.
- **Silently delete the message**: The message is deleted. Neither the sender nor the recipient receives notification.

Notes:

- You can't retrieve messages that have been blocked or attachments that have been stripped. When you configure attachment filters, carefully examine all possible file name matches and verify that legitimate attachments won't be affected by the filter.
- If you remove attachments from digitally signed, encrypted, or rights-protected messages, you invalidate the digital signature, which makes encrypted and rights-protected messages unreadable. A way to avoid this

problem for outbound messages is to sign or encrypt the messages after they've been processed by the Attachment Filtering agent.

For more information, see Attachment filtering procedures on Edge Transport servers.

Default attachments in attachment filtering

The default attachments that are defined in attachment filtering are described in the following table.

TYPE	NAME
`ContentType`	`application/hta`
`ContentType`	`application/javascript`
`ContentType`	`application/msaccess`
`ContentType`	`application/prg`
`ContentType`	`application/x-javascript`
`ContentType`	`application/x-msdownload`
`ContentType`	`message/partial`
`ContentType`	`text/javascript`
`ContentType`	`text/scriptlet`
`ContentType`	`x-internet-signup`
`FileName`	`*.ade`
`FileName`	`*.adp`
`FileName`	`*.app`
`FileName`	`*.asx`
`FileName`	`*.bas`
`FileName`	`*.bat`
`FileName`	`*.chm`
`FileName`	`*.cmd`
`FileName`	`*.com`
`FileName`	`*.cpl`

TYPE	NAME
`FileName`	`*.crt`
`FileName`	`*.csh`
`FileName`	`*.exe`
`FileName`	`*.fxp`
`FileName`	`*.hlp`
`FileName`	`*.hta`
`FileName`	`*.inf`
`FileName`	`*.ins`
`FileName`	`*.isp`
`FileName`	`*.js`
`FileName`	`*.jse`
`FileName`	`*.ksh`
`FileName`	`*.lnk`
`FileName`	`*.mda`
`FileName`	`*.mdb`
`FileName`	`*.mde`
`FileName`	`*.mdt`
`FileName`	`*.mdw`
`FileName`	`*.mdz`
`FileName`	`*.msc`
`FileName`	`*.msi`
`FileName`	`*.msp`
`FileName`	`*.mst`
`FileName`	`*.ops`

TYPE	NAME
FileName	*.pcd
FileName	*.pif
FileName	*.prf
FileName	*.prg
FileName	*.ps1
FileName	*.ps1xml
FileName	*.ps11
FileName	*.ps11xml
FileName	*.ps2
FileName	*.ps2xml
FileName	*.psc1
FileName	*.psc2
FileName	*.reg
FileName	*.scf
FileName	*.scr
FileName	*.sct
FileName	*.shb
FileName	*.shs
FileName	*.url
FileName	*.vb
FileName	*.vbe
FileName	*.vbs
FileName	*.wsc
FileName	*.wsf

TYPE	NAME
`FileName`	`*.wsh`
`FileName`	`*.xnk`

Attachment filtering procedures on Edge Transport servers

8/3/2020 • 5 minutes to read • Edit Online

Attachment filtering in Exchange Server is provided by the Attachment Filter agent that's available only on Edge Transport servers. Attachment filtering can help prevent files in email messages from entering your organization. You can configure one or more attachment filter entries to filter attachments either by content type or by file name.

What do you need to know before you begin?

- Estimated time to complete each procedure: 5 minutes.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Antispam features" entry in the Antispam and antimalware permissions and the "Transport agents" entry in the Mail flow permissions topic.
- Configuration changes that you make to attachment filtering on an Edge Transport server are made only to the local computer. If you have multiple Edge Transport servers in your perimeter network, you need to configure attachment filtering on each Edge Transport server separately.
- You can only use PowerShell to perform this procedure. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to enable or disable attachment filtering

When you enable or disable the Attachment Filtering agent, the change takes effect after you restart the Microsoft Exchange Transport service. When you restart the Microsoft Exchange Transport service on an Edge Transport server, mail flow on the server is temporarily interrupted.

To disable attachment filtering, run the following command:

```
Disable-TransportAgent "Attachment Filtering Agent"
```

To enable attachment filtering, run the following command:

```
Enable-TransportAgent "Attachment Filtering Agent"
```

After you enable or disable attachment filtering, restart the Microsoft Exchange Transport service by running the following command:

```
Restart-Service MSExchangeTransport
```

How do you know this worked?

To verify that you successfully enabled or disabled attachment filtering, run the following command to verify the **Enabled** property value:

```
Get-TransportAgent "Attachment Filtering Agent"
```

Use the Exchange Management Shell to view and find attachment filtering entries

Attachment filtering entries define the message attachments that you want to keep out of your organization. To view the attachment filtering entries that are used by the Attachment Filtering agent, run the following command:

```
Get-AttachmentFilterEntry | Format-Table -Auto Type,Name
```

To find a specific MIME content type entry, use the following syntax:

```
Get-AttachmentFilterEntry ContentType:<MIMEContentType>
```

For example, to see if there's a MIME content type entry for JPEG images, run the following command:

```
Get-AttachmentFilterEntry ContentType:image/jpeg
```

If you receive the error, `Couldn't find the specified identity.`, then the MIME content type isn't defined in the attachment filtering entries.

To view a specific file name or file name extension entry, use the following syntax:

```
Get-AttachmentFilterEntry FileName:<FileName or FileNameExtension>
```

For example, to see if there's a file name extension entry for JPEG attachments, run the following command:

```
Get-AttachmentFilterEntry FileName:*.jpg
```

If you receive the error, `Couldn't find the specified identity.`, then the file name or file name extension isn't defined in the attachment filtering entries.

For more information, see Get-AttachmentFilterEntry.

Use the Exchange Management Shell to add attachment filtering entries

To add an attachment filtering entry that filters attachments by MIME content type, use the following syntax:

```
Add-AttachmentFilterEntry -Name <MIMEContentType> -Type ContentType
```

The following example adds a MIME content type entry that filters JPEG images.

```
Add-AttachmentFilterEntry -Name image/jpeg -Type ContentType
```

To add an attachment filtering entry that filters attachments by file name or file name extension, use the following syntax:

```
Add-AttachmentFilterEntry -Name <FileName or FileNameExtension> -Type FileName
```

The following example filters attachments that have the .jpg file name extension.

```
Add-AttachmentFilterEntry -Name *.jpg -Type FileName
```

For more information, see Add-AttachmentFilterEntry.

How do you know this worked?

To verify that you successfully added an attachment filtering entry, send a test message that contains the prohibited attachment from an external mailbox to an internal recipient and verify that the message and the attachment are processed as you expect.

Use the Exchange Management Shell to remove attachment filtering entries

To remove an attachment filtering entry that filters attachments by MIME content type, use the following syntax:

```
Remove-AttachmentFilterEntry ContentType:<ContentType>
```

The following example removes the MIME content type entry for JPEG images.

```
Remove-AttachmentFilterEntry ContentType:image/jpeg
```

To remove an attachment filtering entry that filters attachments by file name or file name extension, use the following syntax:

```
Remove-AttachmentFilterEntry FileName:<FileName or FileNameExtension>
```

The following example removes the file name entry for the .jpg file name extension.

```
Remove-AttachmentFilterEntry FileName:*.jpg
```

For more information, see Remove-AttachmentFilterEntry.

How do you know this worked?

To verify that you successfully removed an attachment filtering entry, send a test message that contains the allowed attachment from an external mailbox to an internal recipient, and verify that the message was successfully delivered with the attachment.

Use the Exchange Management Shell to view the attachment filtering action

To view the attachment filtering action that's used when a prohibited attachment is detected in a message, run the following command:

```
Get-AttachmentFilterListConfig | Format-List Action,AdminMessage,RejectResponse,ExceptionConnectors
```

Use the Exchange Management Shell to configure the attachment filtering action

To configure the attachment filtering action that's used when a prohibited attachment is detected in a message, use the following syntax:

```
Set-AttachmentFilterListConfig [-Action <Reject | Strip | SilentDelete>] [-RejectResponse "<Message text>"] [-
AdminMessage "<Replacement file text>"] [-ExceptionConnectors <ConnectorGUID>]
```

This example makes the following changes to the attachment filtering configuration:

- Reject (block) messages that have prohibited attachments. Note that you can't specify different actions for different types of attachments.
- Use a custom response for rejected messages.

```
Set-AttachmentFilterListConfig -Action Reject -RejectResponse "This message contains a prohibited attachment.
Your message can't be delivered. Please resend the message without the attachment."
```

For more information, see Set-AttachmentFilterListConfig.

How do you know this worked?

To verify that you successfully configured the attachment filtering action, send a test message that contains a prohibited attachment from an external mailbox to an internal recipient and verify that the message and the attachment are processed as you expect.

Connection filtering on Edge Transport servers

8/3/2020 • 8 minutes to read • Edit Online

Connection filtering is an antispam feature in Exchange Server that allows or blocks email based on the message source. Connection filtering is performed by the Connection Filtering agent that's available only on Edge Transport servers, and is basically unchanged from Exchange Server 2010. The Connection Filtering agent relies on the IP address of the connecting mail server to determine what action, if any, to take on an inbound message.

By default, the Connection Filtering agent is the first antispam agent to evaluate an inbound message on an Edge Transport server. The source IP address of the SMTP connection is checked against the allowed and blocked IP addresses. If the source IP address is specifically allowed, the message is sent to the recipients in your organization without additional processing by other antispam agents. If the source IP address is specifically blocked, the SMTP connection is dropped. If the source IP address isn't specifically allowed or blocked, the message flows through the other antispam agents on the Edge Transport server.

Connection filtering compares the IP address of the source mail server to the values in the IP Allow list, the IP Block list, IP Allow list providers, and IP Block list providers. You need to configure at least one of these four IP address data stores for connection filtering to function. If you don't specify any IP address data, you should disable the Connection Filtering agent. For more information, see Connection filtering procedures on Edge Transport servers.

IP Block list

The IP Block list contains the IP addresses of email servers that you want to block. You manually maintain the IP addresses in the IP Block list. You can add individual IP addresses or IP address ranges. You can specify an expiration time that specifies how long the IP address entry will be blocked. When the expiration time is reached, the IP address entry in the IP Block list is disabled.

If the Connection Filtering agent finds the source IP address on the IP Block list, the SMTP connection will be dropped after all the **RCPT TO** headers (envelope recipients) in the message are processed.

IP addresses can also be automatically added to the IP Block list by the Sender Reputation feature of the Protocol Analysis agent. For more information, see Sender reputation and the Protocol Analysis agent.

IP Allow list

The IP Allow list contains the IP addresses of email servers that you want to designate as trustworthy sources of email. Email from mail servers that you specify in the IP Allow list is exempt from processing by other Exchange antispam agents.

You manually maintain the IP addresses in the IP Allow list. You can add individual IP addresses or IP address ranges. You can specify an expiration time that specifies how long the IP address entry will be allowed. When the expiration time is reached, the entry in the IP Allow list is disabled.

IP Block List providers

IP Block List providers are frequently referred to as *real-time block lists*, or RBLs. IP Block List providers compile lists of mail server IP addresses that send spam. Many IP Block List providers also compile lists of mail server IP addresses that could be used for spam. Examples include mail servers that are configured for open relay, Internet service providers (ISPs) that assign dynamic IP addresses, and ISPs that allow SMTP mail server traffic from dial-up accounts.

When you configure connection filtering to use an IP Block List provider, the Connection Filtering agent compares the IP address of the connecting mail server to the list of IP addresses at the IP Block List provider. If there's a match, the message isn't allowed in your organization. You can configure connection filtering to use multiple IP Block List providers, and you assign different priority values to each provider.

The Connection Filtering agent checks the source IP address at the IP Allow list and the IP Block list. If the IP address doesn't exist on either list, the Connection Filtering agent queries the IP Block List provider according to the priority value that you assigned. If the IP address is defined at an IP Block List provider, the Edge Transport server waits for and processes the **RCPT TO** header, responds to the sending mail server with an `SMTP 550` error, and closes the connection. The connection isn't immediately dropped so that the connection attempt can be logged, and because you can specify recipients that are exempt from having messages blocked by any IP Block list providers.

If the IP address isn't defined at any of the IP Block List providers, the Content Filtering agent hands the message off to the next transport agent on the Edge Transport server.

For each IP Block List provider, you can customize the `SMTP 550` error that's returned to the sender when a message is blocked. You should identify the IP Block List provider that identified the message source as spam. If a legitimate source mail server is erroneously identified as a spam source, the administrator can then contact the IP Block List provider and take the steps necessary to remove the mail server from the IP Block List provider.

IP Block List providers can return different codes to identify why an IP address is defined in their lists. Most IP Block List providers return bitmask or absolute value data types. Within these data types, the IP Block List provider can use multiple values to classify the IP address by threat type.

There are issues to consider when using IP Block list providers:

- Outages or delays at the IP Block list provider service can cause delays in the processing of messages on the Edge Transport server. You should always select reliable IP Block list providers.
- Source servers that you know to be legitimate can be erroneously identified as spam sources. For example, the mail server can be unintentionally configured to act as an open relay. You should always select IP Block list providers that provide clear procedures for evaluation and removal from their services.

Bitmask and absolute value examples

This section shows an example of the status codes returned by most Block List providers. For details about the status codes that the provider returns, see the documentation from the specific provider.

For bitmask data types, the IP Block List provider service returns a status code of 127.0.0. *x*, where the integer *x* is any one of the values listed in the following table.

Values and status codes for bitmask data types

VALUE	STATUS CODE
1	The IP address is on an IP Block list.
2	The SMTP server is configured to act as an open relay.
4	The IP address supports a dial-up IP address.

For absolute value types, the IP Block List provider returns explicit responses that define why the IP address is defined in their block lists. The following table shows examples of absolute values and the explicit responses.

Values and status codes for absolute value data types

VALUE	EXPLICIT RESPONSE
127.0.0.2	The IP address is a direct spam source.
127.0.0.4	The IP address is a bulk mailer.
127.0.0.5	The remote server sending the message is known to support multistage open relays.

IP Allow List providers

IP Allow List providers are also known as *safe lists*. IP Allow List providers are configured just like IP Block List providers, but the results are the opposite: they define mail server IP addresses that are definitely not associated with spam activity. If the IP address of the connecting mail server is defined at an IP Allow List provider, the message is exempt from processing by other Exchange antispam agents. For this reason, IP Block List providers are used much more frequently than IP Allow List providers. Choose your IP Allow List providers carefully.

Test IP Block List providers and IP Allow List providers

After you configure connection filtering to use an IP Block List provider or an IP Allow List provider, you can run tests to verify that the providers are working correctly. Most providers provide test IP addresses that you can use to test their services. When you test a provider, the Connection Filtering agent issues a DNS query that should result in a specific response from the provider. For more information about how to test IP addresses against an IP Block List provider service or an IP Allow List provider service, see Connection filtering procedures on Edge Transport servers.

Configure connection filtering on Edge Transport servers that aren't directly connected to the Internet

You can use connection filtering on Edge Transport servers that don't directly receive email from the Internet. In this scenario, the Edge Transport server is behind another mail server that receives and processes messages directly from the Internet. For example, your organization might send email traffic through an antispam server, service, or appliance before the messages reach the Edge Transport server. In this scenario, the Connection Filtering agent needs to extract the correct source IP address from the message. To do this, the Connection Filtering agent needs to parse the **Received** header field values in the message header and compare those values to the known IP addresses of the mail server that sits between the Edge Transport server and the Internet.

Every mail server that accepts and relays an SMTP message along the delivery path adds its own **Received** header field in the message header. The **Received** header typically contains the domain name and IP address of the mail server that processed the message.

If the Edge Transport server doesn't accept messages directly from the Internet, you need to use the *InternalSMTPServers* parameter on the **Set-TransportConfig** cmdlet on an Exchange Mailbox server to identify the IP address of the mail server that sit between the Edge Transport server and the Internet. The IP address data is replicated to Edge Transport servers by EdgeSync. When messages are received by the Edge Transport server, the Connection Filtering agent assumes an IP address in a **Received** header field that doesn't match a value specified by the *InternalSMTPServers* parameter is the source IP address that needs to be checked. Therefore, you need specify all internal SMTP servers in order for connection filtering to function correctly.

Connection filtering procedures on Edge Transport servers

8/3/2020 • 18 minutes to read • Edit Online

Connection filtering is an antispam feature that's provided by the Connection Filtering agent, which is available only on Edge Transport servers in Exchange Server. Connection filtering enables the following features:

- IP Block list
- IP Block List providers
- IP Allow list
- IP Allow List providers

Each of these features can be enabled or disabled separately.

For more information about connection filtering, see Connection filtering on Edge Transport servers.

What do you need to know before you begin?

- Estimated time to complete each procedure: 5 minutes.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Antispam features" entry in the Antispam and antimalware permissions topic.
- You can only use PowerShell to perform this procedure. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to enable or disable connection filtering

To completely enable or disable connection filtering, you enable or disable the Connection Filtering agent. The change takes effect after you restart the Microsoft Exchange Transport service. When you restart the Microsoft Exchange service on an Edge Transport server, mail flow on the server is temporarily interrupted.

To disable connection filtering, run the following command:

```
Disable-TransportAgent "Connection Filtering Agent"
```

To enable connection filtering, run the following command:

```
Enable-TransportAgent "Connection Filtering Agent"
```

To make the change take effect, restart the Microsoft Exchange Transport service by running the following command:

```
Restart-Service MSExchangeTransport
```

How do you know this worked?

To verify that you successfully enabled or disabled connection filtering, run the following command to verify the **Enabled** property value.

```
Get-TransportAgent "Connection Filtering Agent" | Format-List Enabled
```

IP Block list procedures

These procedures apply to the IP Block list that you manually configure. They don't apply to IP Block List providers.

Use the **IPBlockListConfig** cmdlets to view and configure how connection filtering uses the IP Block list. Use the **IPBlockListEntry** cmdlets to view and configure the IP addresses in the IP Block list.

Use the Exchange Management Shell to view the configuration of the IP Block list

To view the configuration of the IP Block list, run the following command:

```
Get-IPBlockListConfig | Format-List *Enabled,*Response
```

Use the Exchange Management Shell to enable or disable the IP Block list

To disable the IP Block list, run the following command:

```
Set-IPBlockListConfig -Enabled $false
```

To enable the IP Block list, run the following command:

```
Set-IPBlockListConfig -Enabled $true
```

For more information, see Set-IPBlockListConfig.

How do you know this worked?

To verify that you successfully enabled or disabled the IP Block list, run the following command to verify the **Enabled** property value.

```
Get-IPBlockListConfig | Format-List Enabled
```

Use the Exchange Management Shell to configure the IP Block list

To configure the IP Block list, use the following syntax:

```
Set-IPBlockListConfig [-ExternalMailEnabled <$true | $false>] [-InternalMailEnabled <$true | $false> -
MachineEntryRejectionResponse "<Custom response text>"] [-StaticEntryRejectionResponse "<Custom response
text>"]
```

This example configures the IP Block list with the following settings:

- The IP Block list filters incoming connections from internal and external mail servers. By default, connections are filtered from external mail servers only (*ExternalMailEnabled* is set to `$true`, and *InternalMailEnabled* is set to `$false`). Non-authenticated connections and authenticated connections from external partners are considered external.

- The custom response text for connections that were filtered by IP addresses that were automatically added to the IP Block list by the sender reputation feature of the Protocol Analysis agent is set to the value "Connection from IP address {0} was rejected by sender reputation."

- The custom response text for connections that were filtered by IP addresses that were manually added to the IP Block list is set to the value "Connection from IP address {0} was rejected by connection filtering."

```
Set-IPBlockListConfig -InternalMailEnabled $true -MachineEntryRejectionResponse "Connection from IP address {0} was rejected by sender reputation." -StaticEntryRejectionResponse "Connection from IP address {0} was rejected by connection filtering."
```

For more information, see Set-IPBlockListConfig.

How do you know this worked?

To verify that you successfully configured the IP Block list, run the following command to verify the property values.

```
Get-IPBlockListConfig | Format-List *MailEnabled,*Response
```

Use the Exchange Management Shell to view IP Block list entries

To view all IP Block list entries, run the following command:

```
Get-IPBlockListEntry
```

Note that each IP Block list entry is identified by an integer value. The identity integer is assigned in ascending order when you add entries to the IP Block list and the IP Allow list.

To view a specific IP Block list entry, use the following syntax:

```
Get-IPBlockListEntry <-Identity IdentityInteger | -IPAddress IPAddress>
```

For example, to view the IP Block list entry that contains the IP address 192.168.1.13, run the following command:

```
Get-IPBlockListEntry -IPAddress 192.168.1.13
```

For more information, see Get-IPBlockListEntry.

> **NOTE**
>
> When you use the *IPAddress* parameter, the resulting IP Block list entry can be an individual IP address, an IP address range, or a Classless InterDomain Routing (CIDR) IP. To use the *Identity* parameter, you specify the integer value that's assigned to the IP Block list entry.

Use the Exchange Management Shell to add IP Block list entries

To add IP Block list entries, use the following syntax:

```
Add-IPBlockListEntry <-IPAddress IPAddress | -IPRange IP range or CIDR IP> [-ExpirationTime <DateTime>] [-
Comment "<Descriptive Comment>"]
```

This example adds the IP Block list entry for the IP address range 192.168.1.10 through 192.168.1.15 and configures the IP Block list entry to expire on July 4, 2018 at 15:00.

```
Add-IPBlockListEntry -IPRange 192.168.1.10-192.168.1.15 -ExpirationTime "7/4/2018 15:00"
```

For more information, see Add-IPBlockListEntry.

How do you know this worked?

To verify that you successfully added an IP Block list entry, run the following command and verify that the new IP Block list entry is displayed.

```
Get-IPBlockListEntry
```

Use the Exchange Management Shell to remove IP Block list entries

To remove IP Block list entries, use the following syntax:

```
Remove-IPBlockListEntry <IdentityInteger>
```

This example removes the IP Block list entry that has the *Identity* value 3.

```
Remove-IPBlockListEntry 3
```

This example removes the IP Block list entry that contains the IP address 192.168.1.12 without using the *Identity* integer value. Note that the IP Block list entry can be an individual IP address or an IP address range.

```
Get-IPBlockListEntry -IPAddress 192.168.1.12 | Remove-IPBlockListEntry
```

For more information, see Remove-IPBlockListEntry.

How do you know this worked?

To verify that you successfully removed an IP Block list entry, run the following command and verify that the IP Block list entry you removed is gone.

```
Get-IPBlockListEntry
```

IP Block List provider procedures

These procedures apply to IP Block List providers. They don't apply to the IP Block list.

Use the **IPBlockListProvidersConfig** cmdlets to view and configure how connection filtering uses all IP Block List providers. Use the **IPBlockListProvider** cmdlets to view, configure, and test IP Block List providers.

Use the Exchange Management Shell to view the configuration of all IP Block List providers

To view how connection filtering uses all IP Block List providers, run the following command:

```
Get-IPBlockListProvidersConfig | Format-List *Enabled,Bypassed*
```

For more information, see Get-IPBlockListProvidersConfig.

Use the Exchange Management Shell to enable or disable all IP Block List providers

To disable all IP Block List providers, run the following command:

```
Set-IPBlockListProvidersConfig -Enabled $false
```

To enable all IP Block List providers, run the following command:

```
Set-IPBlockListProvidersConfig -Enabled $true
```

For more information, see Set-IPBlockListProvidersConfig.

How do you know this worked?

To verify that you enabled or disabled all IP Block List providers, run the following command to verify the value of the **Enabled** property:

```
Get-IPBlockListProvidersConfig | Format-List Enabled
```

Use the Exchange Management Shell to configure all IP Block List providers

To configure how connection filtering uses all IP Block List providers, use the following syntax:

```
Set-IPBlockListProvidersConfig [-BypassedRecipients <recipient1,recipient2...>] [-ExternalMailEnabled <$true |
$false>] [-InternalMailEnabled <$true | $false>]
```

This example configures all IP Block List providers with the following settings:

- IP Block List providers filter incoming connections from internal and external mail servers. By default, connections are filtered from external mail servers only (*ExternalMailEnabled* is set to `$true`, and *InternalMailEnabled* is set to `$false`). Non-authenticated connections and authenticated connections from external partners are considered external.
- Messages sent to the internal recipients chris@fabrikam.com and michelle@fabrikam.com are excluded from filtering by IP Block List providers. Note that if you want to add recipients to the list without affecting existing recipients, use the syntax, `@{Add="<recipient1>","<recipient2>"...}`.

```
Set-IPBlockListProvidersConfig -BypassedRecipients chris@fabrikam.com,michelle@fabrikam.com -
InternalMailEnabled $true
```

For more information, see Set-IPBlockListProvidersConfig.

How do you know this worked?

To verify that you successfully configured all IP Block List providers, run the following command to verify the property values:

```
Get-IPBlockListProvidersConfig | Format-List *MailEnabled,Bypassed*
```

Use the Exchange Management Shell to view IP Block List providers

To view the summary list of all the IP Block List providers, run the following command:

```
Get-IPBlockListProvider
```

To view the details of a specific provider, use the following syntax:

```
Get-IPBlockListProvider <IPBlockListProviderIdentity>
```

This example show the details of the provider named Contoso IP Block List Provider.

```
Get-IPBlockListProvider "Contoso IP Block List Provider" | Format-List
Name,Enabled,Priority,LookupDomain,*Match,*Response
```

For more information, see Get-IPBlockListProvider.

Use the Exchange Management Shell to add an IP Block List provider

To add an IP Block List provider, use the following syntax:

```
Add-IPBlockListProvider -Name "<Descriptive Name>" -LookupDomain <FQDN> [-Priority <Integer>] [-Enabled <$true
| $false>] [-AnyMatch <$true | $false>] [-BitmaskMatch <IPAddress>] [-IPAddressesMatch
<IPAddressStatusCode1,IPAddressStatusCode2...>] [-RejectionResponse "<Custom Text>"]
```

This example creates an IP Block List provider named "Contoso IP Block List Provider" with the following options:

- **FQDN to use the provider**: rbl.contoso.com
- **Bitmask code to use from the provider**: 127.0.0.1

```
Add-IPBlockListProvider -Name "Contoso IP Block List Provider" -LookupDomain rbl.contoso.com -BitmaskMatch
127.0.0.1
```

> **NOTE**
>
> When you add a new IP Block List provider, it's enabled by default (the value of *Enabled* is `$true`), and the priority value is incremented (the first entry has the *Priority* value 1).

For more information, see Add-IPBlockListProvider.

How do you know this worked?

To verify that you successfully added an IP Block List provider, run the following command and verify that the new IP Block List provider is displayed.

```
Get-IPBlockListProvider
```

Use the Exchange Management Shell to enable or disable an IP Block List provider

To enable or disable a specific IP Block List provider, use the following syntax:

```
Set-IPBlockListProvider <IPBlockListProviderIdentity> -Enabled <$true | $false>
```

This example disables the provider named Contoso IP Block List Provider.

```
Set-IPBlockListProvider "Contoso IP Block List Provider" -Enabled $false
```

This example enables the provider named Contoso IP Block List Provider.

```
Set-IPBlockListProvider "Contoso IP Block List Provider" -Enabled $true
```

For more information, see Set-IPBlockListProvider.

How do you know this worked?

To verify that you successfully enabled or disabled an IP Block List provider, run the following command to verify the value of the **Enabled** property:

```
Get-IPBlockListProvider | Format-Table -Auto Name,LookupDomain,Priority,Enabled
```

Use the Exchange Management Shell to configure an IP Block List provider

The configuration options that are available on the **Set-IPBlockListProvider** cmdlet are identical to those on the **Add-IPBlockListProvider** cmdlet.

To configure an existing IP Block List provider, use the following syntax:

```
Set-IPBlockListProvider <IPBlockListProviderIdentity> -Name "<Descriptive Name>" -LookupDomain <FQDN> [-
Priority <Integer>] [-AnyMatch <$true | $false>] [-BitmaskMatch <IPAddress>] [-IPAddressesMatch
<IPAddressStatusCode1,IPAddressStatusCode2...>] [-RejectionResponse "<Custom Text>"]
```

For example, to add the IP address status code 127.0.0.1 to the list of existing status codes for the provider named Contoso IP Block List Provider, run the following command:

```
Set-IPBlockListProvider "Contoso IP Block List Provider" -IPAddressesMatch @{Add="127.0.0.1"}
```

For more information, see Set-IPBlockListProvider.

How do you know this worked?

To verify that you successfully configured an IP Block List provider, run the following command to verify the property values. Be sure to replace *<IPBlockListProviderIdentity>* with the name of the IP Block List provider.

```
Get-IPBlockListProvider <IPBlockListProviderIdentity> | Format-List
```

Use the Exchange Management Shell to test an IP Block List provider

To test an IP Block List provider, use the following syntax:

```
Test-IPBlockListProvider <IPBlockListProviderIdentity> -IPAddress <IPAddressToTest>
```

This example tests the provider named Contoso IP Block List Provider by looking up the IP address 192.168.1.1.

```
Test-IPBlockListProvider "Contoso IP Block List Provider" -IPAddress 192.168.1.1
```

For more information, see Test-IPBlockListProvider.

Use the Exchange Management Shell to remove an IP Block List provider

To remove an IP Block List provider, use the following syntax:

```
Remove-IPBlockListProvider <IPBlockListProviderIdentity>
```

This example removes the IP Block List provider named Contoso IP Block List Provider.

```
Remove-IPBlockListProvider "Contoso IP Block list Provider"
```

For more information, see Remove-IPBlockListProvider.

How do you know this worked?

To verify that you successfully removed an IP Block List provider, run the following command and verify that the IP Block List provider you removed is gone.

```
Get-IPBlockListProvider
```

IP Allow list procedures

These procedures apply to the IP Allow list that you manually configure. They don't apply to IP Allow List providers.

Use the **IPAllowListConfig** cmdlets to view and configure how connection filtering uses the IP Allow list. Use the **IPAllowListEntry** cmdlets to view and configure the IP addresses in the IP Allow list.

Use the Exchange Management Shell to view the configuration of the IP Allow list

To view the configuration of the IP Allow list, run the following command.

```
Get-IPAllowListConfig | Format-List *Enabled
```

For more information, see Get-IPAllowListConfig.

Use the Exchange Management Shell to enable or disable the IP Allow list

To disable the IP Allow list, run the following command:

```
Set-IPAllowListConfig -Enabled $false
```

To enable the IP Allow list, run the following command:

```
Set-IPAllowListConfig -Enabled $true
```

How do you know this worked?

To verify that you successfully enabled or disabled the IP Allow list, run the following command to verify the value of the **Enabled** property:

```
Get-IPAllowListConfig | Format-List Enabled
```

Use the Exchange Management Shell to configure the IP Allow list

To configure the IP Allow list, use the following syntax:

```
Set-IPAllowListConfig [-ExternalMailEnabled <$true | $false>] [-InternalMailEnabled <$true | $false>
```

This example configures the IP Allow list to filter incoming connections from internal and external mail servers. By default, connections are filtered from external mail servers only (*ExternalMailEnabled* is set to `$true`, and *InternalMailEnabled* is set to `$false`). Non-authenticated connections and authenticated connections from external partners are considered external.

```
Set-IPAllowListConfig -InternalMailEnabled $true
```

For more information, see Set-IPAllowListConfig.

How do you know this worked?

To verify that you successfully configured the IP Allow list, run the following command to verify the property values:

```
Get-IPAllowListConfig | Format-List *MailEnabled
```

Use the Exchange Management Shell to view IP Allow list entries

To view all IP Allow list entries, run the following command:

```
Get-IPAllowListEntry
```

Note that each IP Allow list entry is identified by an integer value. The identity integer is assigned in ascending order when you add entries to the IP Block list and the IP Allow list.

To view a specific IP Allow list entry, use the following syntax:

```
Get-IPAllowListEntry <-Identity IdentityInteger | -IPAddress IPAddress>
```

For example, to view the IP Allow list entry that contains the IP address 192.168.1.13, run the following command:

```
Get-IPAllowListEntry -IPAddress 192.168.1.13
```

For more information, see Get-IPAllowListEntry.

> **NOTE**
>
> When you use the *IPAddress* parameter, the resulting IP Allow list entry can be an individual IP address, an IP address range, or a Classless InterDomain Routing (CIDR) IP. To use the *Identity* parameter, you specify the integer value that's assigned to the IP Allow list entry.

Use the Exchange Management Shell to add IP Allow list entries

To add IP Allow list entries, use the following syntax:

```
Add-IPAllowListEntry <-IPAddress IPAddress | -IPRange IP range or CIDR IP> [-ExpirationTime <DateTime>] [-
Comment "<Descriptive Comment>"]
```

This example adds the IP Allow list entry for the IP address range 192.168.1.10 through 192.168.1.15 and configures the IP Allow list entry to expire on July 4, 2018 at 15:00.

```
Add-IPAllowListEntry -IPRange 192.168.1.10-192.168.1.15 -ExpirationTime "7/4/2018 15:00"
```

For more information, see Add-IPAllowListEntry.

How do you know this worked?

To verify that you successfully added an IP Allow list entry, run the following command and verify that the new IP Allow list entry is displayed.

```
Get-IPAllowListEntry
```

Use the Exchange Management Shell to remove IP Allow list entries

To remove IP Allow list entries, use the following syntax:

```
Remove-IPAllowListEntry <IdentityInteger>
```

This example removes the IP Allow list entry that has the *Identity* value 3.

```
Remove-IPAllowListEntry 3
```

This example removes the IP Allow list entry that contains the IP address 192.168.1.12 without using the *Identity* integer value. Note that the IP Allow list entry can be an individual IP address or an IP address range.

```
Get-IPAllowListEntry -IPAddress 192.168.1.12 | Remove-IPAllowListEntry
```

For more information, see Remove-IPAllowListEntry.

How do you know this worked?

To verify that you successfully removed an IP Allow list entry, run the following command and verify that the IP Allow list entry you removed is gone.

```
Get-IPAllowListEntry
```

IP Allow List provider procedures

These procedures apply to IP Allow List providers. They don't apply to the IP Allow list.

Use the **IPAllowListProvidersConfig** cmdlets to view and configure how connection filtering uses all IP Allow List providers. Use the **IPAllowListProvider** cmdlets to view, configure, and test IP Allow List providers.

Use the Exchange Management Shell to view the configuration of all IP Allow List providers

To view how connection filtering uses all IP Allow List providers, run the following command:

```
Get-IPAllowListProvidersConfig | Format-List *Enabled
```

For more information, see Get-IPAllowListProvidersConfig.

Use the Exchange Management Shell to enable or disable all IP Allow List providers

To disable all IP Allow List providers, run the following command:

```
Set-IPAllowListProvidersConfig -Enabled $false
```

To enable all IP Allow List providers, run the following command:

```
Set-IPAllowListProvidersConfig -Enabled $true
```

For more information, see Set-IPAllowListProvidersConfig.

How do you know this worked?

To verify that you enabled or disabled all IP Allow List providers, run the following command to verify the **Enabled** property value:

```
Get-IPAllowListProvidersConfig | Format-List Enabled
```

Use the Exchange Management Shell to configure all IP Allow List providers

To configure how connection filtering uses all IP Allow List providers, use the following syntax:

```
Set-IPAllowListProvidersConfig [-ExternalMailEnabled <$true | $false>] [-InternalMailEnabled <$true | $false>]
```

This example configures all IP Allow List providers to filter incoming connections from internal and external mail servers. By default, connections are filtered from external mail servers only (*ExternalMailEnabled* is set to `$true`, and *InternalMailEnabled* is set to `$false`). Non-authenticated connections and authenticated connections from external partners are considered external.

```
Set-IPAllowListProvidersConfig -InternalMailEnabled $true
```

For more information, see Set-IPAllowListProvidersConfig.

How do you know this worked?

To verify that you successfully configured all IP Allow List providers, run the following command to verify the property values:

```
Get-IPAllowListProvidersConfig | Format-List *MailEnabled
```

Use the Exchange Management Shell to view IP Allow List providers

To view the summary list of all the IP Allow List providers, run the following command.

```
Get-IPAllowListProvider
```

To view the details of a specific provider, use the following syntax:

```
Get-IPAllowListProvider <IPAllowListProviderIdentity>
```

This example show the details of the provider named Contoso IP Allow List Provider.

```
Get-IPAllowListProvider "Contoso IP Allow List Provider" | Format-List
Name,Enabled,Priority,LookupDomain,*Match
```

For more information, see Get-IPAllowListProvider.

Use the Exchange Management Shell to add an IP Allow List provider

To add an IP Allow List provider, use the following syntax:

```
Add-IPAllowListProvider -Name "<Descriptive Name>" -LookupDomain <FQDN> [-Priority <Integer>] [-Enabled <$true
| $false>] [-AnyMatch <$true | $false>] [-BitmaskMatch <IPAddress>] [-IPAddressesMatch
<IPAddressStatusCode1,IPAddressStatusCode2...>]
```

This example creates an IP Allow List provider named "Contoso IP Allow List Provider" with the following options:

- **FQDN to use the provider**: allow.contoso.com
- **Bitmask code to use from the provider**: 127.0.0.1

```
Add-IPAllowListProvider -Name "Contoso IP Allow List Provider" -LookupDomain allow.contoso.com -BitmaskMatch
127.0.0.1
```

> **NOTE**
>
> When you add a new IP Allow List provider, it's enabled by default (the value of *Enabled* is `$true`), and the priority value is incremented (the first entry has the *Priority* value 1).

For more information, see Add-IPAllowListProvider.

How do you know this worked?

To verify that you successfully added an IP Allow List provider, run the following command and verify that the new IP Allow List provider is displayed.

```
Get-IPAllowListProvider
```

Use the Exchange Management Shell to enable or disable an IP Allow List provider

To enable or disable a specific IP Allow List provider, use the following syntax:

```
Set-IPAllowListProvider <IPAllowListProviderIdentity> -Enabled <$true | $false>
```

This example disables the provider named Contoso IP Allow List Provider.

```
Set-IPAllowListProvider "Contoso IP Allow List Provider" -Enabled $false
```

This example enables the provider named Contoso IP Allow List Provider.

```
Set-IPAllowListProvider "Contoso IP Allow List Provider" -Enabled $true
```

For more information, see Set-IPAllowListProvider.

How do you know this worked?

To verify that you successfully enabled or disabled an IP Allow List provider, run the following command to verify the **Enabled** property value:

```
Get-IPAllowListProvider | Format-Table -Auto Name,LookupDomain,Priority,Enabled
```

Use the Exchange Management Shell to configure an IP Allow List provider

The configuration options that are available on the **Set-IPAllowListProvider** cmdlet are identical to those on the **Add-IPAllowListProvider** cmdlet.

To configure an existing IP Allow List provider, use the following syntax:

```
Set-IPAllowListProvider <IPAllowListProviderIdentity> -Name "<Descriptive Name>" -LookupDomain <FQDN> [-
Priority <Integer>] [-AnyMatch <$true | $false>] [-BitmaskMatch <IPAddress>] [-IPAddressesMatch
<IPAddressStatusCode1,IPAddressStatusCode2...>]
```

For example, to add the IP address status code 127.0.0.1 to the list of existing status codes for the provider named Contoso IP Allow List Provider, run the following command:

```
Set-IPAllowListProvider "Contoso IP Allow List Provider" -IPAddressesMatch @{Add="127.0.0.1"}
```

For more information, see Set-IPAllowListProvider.

How do you know this worked?

To verify that you successfully configured an IP Allow List provider, run the following command. Be sure to replace *<IPAllowListProviderIdentity>* with the name of the IP Allow List provider.

```
Get-IPAllowListProvider <IPAllowListProviderIdentity> | Format-List
```

Use the Exchange Management Shell to test an IP Allow List provider

To test an IP Allow List provider, use the following syntax:

```
Test-IPAllowListProvider <IPAllowListProviderIdentity> -IPAddress <IPAddressToTest>
```

This example tests the provider named Contoso IP Allow List Provider by looking up the IP address 192.168.1.1.

```
Test-IPAllowListProvider "Contoso IP Allow List Provider" -IPAddress 192.168.1.1
```

For more information, see Test-IPAllowListProvider.

Use the Exchange Management Shell to remove an IP Allow List provider

To remove an IP Allow List provider, use the following syntax:

```
Remove-IPAllowListProvider <IPAllowListProviderIdentity>
```

This example removes the IP Allow List provider named Contoso IP Allow List Provider.

```
Remove-IPAllowListProvider "Contoso IP Allow List Provider"
```

For more information, see Remove-IPAllowListProvider.

How do you know this worked?

To verify that you successfully removed an IP Allow List provider, run the following command and verify that the IP Allow List provider you removed is gone.

```
Get-IPAllowListProvider
```

Recipient filtering on Edge Transport servers

8/3/2020 • 4 minutes to read • Edit Online

Recipient filtering is an antispam feature in Exchange Server that relies on the **RCPT TO** SMTP header to determine what action, if any, to take on an inbound message. Recipient filtering is performed by the Recipient Filter agent, and is basically unchanged from Exchange Server 2010.

For more information about how to configure the Recipient Filter agent, see Recipient filtering procedures on Edge Transport servers.

The Recipient Filter agent blocks messages according to the characteristics of the intended recipient in the organization. The Recipient Filter agent can help you prevent the acceptance of messages in the following scenarios:

- **Nonexistent recipients**: You can prevent delivery to recipients that aren't in the organization's address book. For example, you may want to stop delivery to frequently misused account names, such as administrator@contoso.com or support@contoso.com.
- **Restricted distribution groups**: You can prevent delivery of Internet mail to distribution groups that should be used only by internal users.
- **Mailboxes that should never receive messages from the Internet**: You can prevent delivery of Internet mail to a specific mailbox or alias that's typically used inside the organization, such as Helpdesk.

The Recipient Filter agent acts on recipients from one or both of the following data sources:

- **Recipient Block list**: An administrator-defined list of recipients who should never receive messages from the Internet.
- **Recipient Lookup**: Queries Active Directory to verify that the recipient exists in the organization. On an Edge Transport server, Recipient Lookup requires access to Active Directory information that's provided by EdgeSync to the local instance of Active Directory Lightweight Directory Services (AD LDS). For more information, see Edge Subscriptions.

When you enable the Recipient Filter agent, one of the following actions is taken on inbound messages according to the characteristics of the recipients. These recipients are indicated by the **RCPT TO** header.

- If the inbound message contains a recipient that is on the Recipient Block list, the Exchange server sends a `550 5.1.1 User unknown` SMTP session error to the sending server.
- If the inbound message contains a recipient that doesn't match any recipients in Recipient Lookup, the Exchange server sends a `550 5.1.1 User unknown` SMTP session error to the sending server.
- If the recipient isn't on the Recipient Block list and the recipient is found in Recipient Lookup, the Exchange server sends a `250 2.1.5 Recipient OK` SMTP response to the sending server, and the next antispam agent in the chain processes the message.

> **NOTE**
>
> Although the Recipient Filter agent is available on Mailbox servers, you shouldn't configure it. When recipient filtering on a Mailbox server detects one invalid or blocked recipient in a message that contains other valid recipients, the message is rejected. The Recipient Filter agent is enabled when you install the antispam agents on a Mailbox server, but it isn't configured to block any recipients. For more information, see Enable antispam functionality on Mailbox servers.

Configuring recipient lookup

One of the most effective ways to reduce spam is to validate recipients before accepting inbound messages from the Internet. You enable the blocking of messages sent to recipients who don't exist in the Exchange organization, and the blocking of specific recipients using the **Set-RecipientFilterConfig** cmdlet in the Exchange Management Shell. For more information, see Recipient filtering procedures on Edge Transport servers.

Tarpitting functionality

Recipient Lookup functionality enables the sending server to determine whether an email address is valid or invalid. As mentioned earlier, when the recipient of an inbound message is a known recipient, the Exchange server sends back a `250 2.1.5 Recipient OK` SMTP response to the sending server. This functionality provides an ideal environment for a *directory harvest attack*, where a spammer uses an automated program to collect email addresses that return a `250 2.1.5 Recipient OK` SMTP response.

To combat directory harvest attacks, Exchange includes tarpitting functionality. *Tarpitting* is the practice of artificially delaying server responses for specific SMTP communication patterns that indicate high volumes of mail, so that the cost of sending spam increases for the spammer.

If tarpitting isn't configured, the Exchange server immediately returns a `550 5.1.1 User unknown` SMTP session error to the sender when a recipient isn't located in Recipient Lookup. Alternatively, if tarpitting is configured, the Exchange server waits a specified number of seconds before it returns the `550 5.1.1 User unknown` error. This pause in the SMTP session makes automating a directory harvest attack more difficult and less cost-effective for the spammer. By default, tarpitting is configured for 5 seconds on Receive connectors.

To configure the delay before SMTP returns the `550 5.1.1 User unknown` error, you set the tarpitting interval using the *TarpitInterval* parameter on the **Set-ReceiveConnector** cmdlet. For more information, see Message throttling on Receive connectors.

Multiple namespaces

The Recipient Filter agent performs recipient lookups only for authoritative domains. If your organization accepts and forwards messages on behalf of another domain that's configured as an internal relay or external relay domain, the Recipient Filter agent doesn't perform a recipient lookup on recipients in those domains. However, if the recipient is specified in the Recipient Block list, the recipient will still be blocked by the Recipient Filter agent.

Note that you can also configure accepted domains locally on an Edge Transport server. If the domain is configured as internal relay or external relay domain, the Recipient Filter agent on the Edge Transport server also doesn't perform a recipient lookup on recipients in those domains.

Recipient filtering procedures on Edge Transport servers

8/3/2020 • 4 minutes to read • Edit Online

Recipient filtering is provided by the Recipient Filter agent. When recipient filtering is enabled on an Exchange server, it filters inbound messages that come from the Internet but aren't authenticated. These messages are handled as external messages. For more information about recipient filtering and the Recipient Filter agent, see Recipient filtering on Edge Transport servers.

Recipient filtering on Edge Transport servers

NOTE

Although the Recipient Filter agent is available on Mailbox servers, you shouldn't configure it. When recipient filtering on a Mailbox server detects one invalid or blocked recipient in a message that contains other valid recipients, the message is rejected. If you install the antispam agents on a Mailbox server, the Recipient Filter agent is enabled by default. However, it isn't configured to block any recipients. For more information, see Enable antispam functionality on Mailbox servers.

What do you need to know before you begin?

- Estimated time to complete each procedure: less than 5 minutes

- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Antispam features" entry in the Antispam and antimalware permissions topic.

- You can only use PowerShell to perform this procedure. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.

- Although the Recipient Filter agent is available on Mailbox servers, you shouldn't configure it. When recipient filtering on a Mailbox server detects one invalid or blocked recipient in a message that contains other valid recipients, the message is rejected. The Recipient Filter agent is enabled when you install the antispam agents on a Mailbox server, but it isn't configured to block any recipients. For more information, see Enable antispam functionality on Mailbox servers.

- The *AddressBookEnabled* parameter on the **Set-AcceptedDomain** cmdlet enables or disables recipient filtering for recipients in an accepted domain. By default, recipient filtering is enabled for authoritative domains, and disabled for internal relay domains and external relay domains. To view the status of the *AddressBookEnabled* parameter for the accepted domains in your organization, run the command: `Get-AcceptedDomain | Format-List Name,AddressBookEnabled`.

- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

TIP

Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to enable or disable recipient filtering

To disable recipient filtering, run the following command:

```
Set-RecipientFilterConfig -Enabled $false
```

To enable recipient filtering, run the following command:

```
Set-RecipientFilterConfig -Enabled $true
```

> **NOTE**
>
> When you disable recipient filtering, the underlying Recipient Filter agent is still enabled. To disable the Recipient Filter agent, run the command: `Disable-TransportAgent "Recipient Filter Agent"`.

How do you know this worked?

To verify that you've successfully enabled or disabled recipient filtering, run the following command to verify the **Enabled** property value:

```
Get-RecipientFilterConfig | Format-List Enabled
```

Use the Exchange Management Shell to enable or disable recipient filtering for external connections

By default, recipient filtering is enabled for external (unauthenticated) SMTP connections.

To disable recipient filtering for external connections, run the following command:

```
Set-RecipientFilterConfig -ExternalMailEnabled $false
```

To enable recipient filtering for external connections, run the following command:

```
Set-RecipientFilterConfig -ExternalMailEnabled $true
```

How do you know this worked?

To verify that you've successfully enabled or disabled recipient filtering for external SMTP connections, run the following command to verify the **ExternalMailEnabled** property value:

```
Get-RecipientFilterConfig | Format-List ExternalMailEnabled
```

Use the Exchange Management Shell to enable or disable recipient filtering for internal connections

As a best practice, you don't need to apply antispam filters to messages from trusted partners or from inside your organization. To reduce the chance that filters will mishandle legitimate email messages, you typically configure antispam agents to only run on messages from external sources.

To enable recipient filtering for internal (authenticated) SMTP connections, run the following command:

```
Set-RecipientFilterConfig -InternalMailEnabled $true
```

To disable recipient filtering for internal connections, run the following command:

```
Set-RecipientFilterConfig -InternalMailEnabled $false
```

How do you know this worked?

To verify that you've successfully enabled or disabled recipient filtering for internal SMTP connections, run the following command to verify the **InternalMailEnabled** property value:

```
Get-RecipientFilterConfig | Format-List InternalMailEnabled
```

Use the Exchange Management Shell to enable or disable the Recipient Block list

To enable the Recipient Block list, run the following command:

```
Set-RecipientFilterConfig -BlockListEnabled $true
```

To disable the Recipient Block list, run the following command:

```
Set-RecipientFilterConfig -BlockListEnabled $false
```

How do you know this worked?

To verify that you've successfully enabled or disabled the Recipient Block list, run the following command to verify the **BlockListEnabled** property value:

```
Get-RecipientFilterConfig | Format-List BlockListEnabled
```

Use the Exchange Management Shell to configure the Recipient Block list

To replace the existing values, use the following syntax:

```
Set-RecipientFilterConfig -BlockedRecipients <recipient1,recipient2...>
```

This example configures the Recipient Block list with the valuesmark@contoso.com and kim@contoso.com:

```
Set-RecipientFilterConfig -BlockedRecipients mark@contoso.com,kim@contoso.com
```

To add or remove entries without modifying other existing values, use the following syntax:

```
Set-RecipientFilterConfig -BlockedRecipients @{Add="<recipient1>","<recipient2>"...; Remove="<recipient1>","<recipient2>"...}
```

This example adds chris@contoso.com to the list of recipients, and removes michelle@contoso.com from the list of recipients in the Recipient Block list:

```
Set-RecipientFilterConfig -BlockedRecipients @{Add="chris@contoso.com"; Remove="michelle@contoso.com"}
```

How do you know this worked?

To verify that you've successfully configured the Recipient Block list, run the following command to verify the **BlockedRecipients** property value:

```
Get-RecipientFilterConfig | Format-List BlockedRecipients
```

Use the Exchange Management Shell to enable or disable Recipient Lookup

To enable Recipient Lookup to block messages to recipients that don't exist in your organization, run the following command:

```
Set-RecipientFilterConfig -RecipientValidationEnabled $true
```

To disable Recipient Lookup, run the following command:

```
Set-RecipientFilterConfig -RecipientValidationEnabled $false
```

Note: Recipient Lookup on an Edge Transport server requires an Edge subscription. For more information, see Edge Subscriptions.

How do you know this worked?

To verify that you've successfully enabled or disabled Recipient Lookup, run the following command to verify the **RecipientValidationEnabled** property value:

```
Get-RecipientFilterConfig | Format-List RecipientValidationEnabled
```

Antimalware protection in Exchange Server

8/3/2020 • 7 minutes to read • Edit Online

Antimalware protection in Exchange Server 2016 helps combat viruses and spyware in your email messaging environment. *Viruses* infect other programs and data, and they spread throughout your computer looking for programs to infect. *Spyware* gathers personal information (for example, sign-in information and personal data) and sends it back to its author.

The antimalware protection in Exchange Server was introduced in Exchange 2013, and is provided by the Transport agent named Malware Agent. The agent scans messages as they travel through the Transport service on a Mailbox server. You configure malware filtering by using:

- **Antimalware policies**: Specify inbound and outbound scanning and notification options for malware filtering. There's a default policy that applies to all recipients in the Exchange organization, and you can create addtional policies that are applied in a specific order.
- **Antimalware server settings**: Specify the error and retry actions, and the engine and definition update settings for malware filtering. The Malware agent uses Internet access on TCP port 80 (HTTP) to check for engine and definition updates every hour.
- **Antimalware scripts**: Enable or disable malware filtering on the server, and manually download engine and definition updates.

For procedures related to malware filtering, see Procedures for antimalware protection in Exchange Server. For more information about the antispam features in Exchange Server, see Antispam protection in Exchange Server.

Antimalware policies

Antimalware policies control the actions and notification options for malware detections. The important settings in antimalware policies are:

- **Action**: Specifies what to do when a message is found to contain malware. The options are:
  - Delete the message (this is the default value).
  - Replace all attachments with a text file that contains this default text:

    > Malware was detected in one or more attachments included with this email. All attachments have been deleted.

  - Replace all attachments with a text file that contains the custom text you specify.
- **Notifications**: When an antimalware policy is configured to delete messages, you can choose whether to send a notification message to the sender. You can send notification messages based on whether the sender is internal or external. The default notification message has these properties:
  - **From**: Postmaster postmaster@ *<defaultdomain>*.com
  - **Subject**: Undeliverable message
  - **Message text**: This message was created automatically by mail delivery software. Your email message was not delivered to the intended recipients because malware was detected.

  You can customize the message properties for internal and external notifications. You can also specify additional recipients (administrators) to receive notifications for undeliverable messages from internal or external senders.

- **Recipient filters**: For custom antimalware policies, you can specify recipient conditions and exceptions that determine who the policy applies to. You can use these properties for conditions and exceptions:
  - By recipient
  - By accepted domain
  - By group membership

  You can only use a condition or exception once, but the condition or exception can contain multiple values. Multiple values of the same condition or exception use OR logic (for example, *<recipient1>* or *<recipient2>*). Different conditions or exceptions use AND logic (for example, *<recipient1>* and *<member of group 1>*).

- **Priority**: If you create multiple custom antimalware policies, you can specify the order that they're applied.

Antimalware policies in the Exchange admin center vs the Exchange Management Shell

The basic elements of an antimalware policy are:

- **The malware filter policy**: Specifies the action and notification options for malware filtering.
- **The malware filter rule**: Specifies the priority and recipient filters (who the policy applies to) for a malware filter policy.

The difference between these two elements isn't obvious when you manage antimalware polices in the Exchange admin center (EAC):

- When you create an antimalware policy in the EAC, you're actually creating a malware filter rule and the associated malware filter policy at the same time using the same name for both.
- When you modify an antimalware policy in the EAC, settings related to the name, priority, enabled or disabled, and recipient filters modify the malware filter rule. Other settings (actions and notification options) modify the associated malware filter policy.
- When you remove an antimalware policy from the EAC, the malware filter rule and the associated malware filter policy are removed.

In the Exchange Management Shell, the difference between malware filter policies and malware filter rules is apparent. You manage malware filter policies by using the ***-MalwareFilterPolicy** cmdlets, and you manage malware filter rules by using the ***-MalwareFilterRule** cmdlets.

- In the Exchange Management Shell, you create the malware filter policy first, then you create the malware filter rule that identifies the policy that the rule applies to.
- In the Exchange Management Shell, you modify the settings in the malware filter policy and the malware filter rule separately.
- When you remove a malware filter policy from the Exchange Management Shell, the corresponding malware filter rule isn't automatically removed, and vice versa.

Default antimalware policy

Every Mailbox server has a built-in antimalware policy named Default that has these properties:

- The malware filter policy named Default is applied to all recipients in the Exchange organization, even though there's no malware filter rule (recipient filters) associated with the policy.
- The policy named Default has the custom priority value Lowest that you can't modify (the policy is always

applied last). Any custom antimalware policies that you create always have a higher priority than the policy named Default.

- The policy named Default is the default policy (the **IsDefault** property has the value `True`), and you can't delete the default policy.

Antimalware server settings

You can use the **Get-MalwareFilteringServer** and **Set-MalwareFilteringServer** cmdlets in the Exchange Management Shell to view and configure the update, timeout, and download settings for the Malware agent on the Mailbox server. For procedures that use these cmdlets, see Use the Exchange Management Shell to bypass malware filtering on Mailbox servers and Use the Exchange Management Shell to configure malware filtering to rescan messages that were already scanned by EOP.

Antimalware scripts

Exchange includes two Exchange Management Shell scripts that you can use to manage malware filtering:

- `Disable-Antimalwarescanning.ps1` disables the Malware agent, and malware engine and definition updates on the Mailbox server.

- `Enable-Antimalwarescanning.ps1` enables the Malware agent, enables malware engine and definition updates, and runs engine and definition updates on the Mailbox server.

- `Update-MalwareFilteringServer.ps1` manually runs malware engine and definition updates on the Mailbox server.

For more information about using these scripts, see Use the Exchange Management Shell to enable or disable malware filtering on Mailbox servers and Download antimalware engine and definition updates.

Antimalware protection options in Exchange Server

This list describes the antimalware options for Exchange:

- **Built-in antimalware protection**: You can use the built-in antimalware protection in Exchange to help you combat malware. You can use it by itself, or you can pair it with other antimalware solutions to provide a layered defense against malware.

- **Exchange Online Protection (EOP)**: You can pay for a subscription to EOP, which is the antimalware solution that's used in Microsoft 365 and Office 365. EOP leverages partnerships with several antimalware engines to provide efficient, cost effective, and multi-layered antimalware protection. The advantages of paring the built-in antimalware protection with EOP are:

  - EOP uses multiple antimalware engines, while the built-in antimalware protection uses a single engine.

  - EOP has reporting capabilities, including malware statistics.

  - EOP provides the message trace feature for self-troubleshooting mail flow problems including malware detections.

    For more information about EOP, see Anti-malware protection in EOP.

- **Third-party antimalware protection**: You can buy a third-party antimalware program.

Antimalware FAQ for Exchange

This section answers the frequently asked questions about built-in malware filtering and scanning in Exchange.

Why did malware that was identified by other antimalware services get past Exchange antimalware filtering?

There are two likely reasons:

- The most likely scenario is the message attachment doesn't actually contain any active malicious code. Some antimalware engines are more aggressive than others, and these engines might stop messages simply because they contain truncated malware payloads that don't actually do anything.
- The malware you received is a new variant, and our antimalware engine hasn't released a pattern file for it (yet).

I received a message with an unfamiliar attachment. Is this malware or can I disregard this attachment?

We strongly advise that you don't open any attachments that you don't recognize. If you would like us to investigate the attachment, submit it to us as described in the next item.

How do I submit known malware, suspicious files, or false positives to Microsoft?

Save a copy of the message and upload the message to the Microsoft Security Intelligence website so we can examine it.

If the sample contains malware, we'll take corrective action to prevent the virus from going undetected. if the sample is clean, we'll take corrective action to prevent the file from being detected as malware.

Where can I get the messages that have been deleted by the malware filter?

You can't. The messages were found to contain active malicious code, so they were deleted.

Can I use mail flow rules to bypass malware filtering?

No, you can't use mail flow rules (also known as transport rules) to bypass the Malware agent. Instead, send the attachment in a password-protected .zip file (password-protected file .zip files are bypassed by malware filtering).

Procedures for antimalware protection in Exchange Server

8/3/2020 • 18 minutes to read • Edit Online

Exchange Server includes the Malware Agent that's installed on Mailbox servers. For more information about malware filtering in Exchange, see Antimalware protection in Exchange Server.

This topic describes the following procedures for managing malware filtering in Exchange:

- Disable or enable malware filtering on a Mailbox server
- Bypass malware filtering on a Mailbox server
- Create antimalware policies
- View antimalware policies
- Modify antimalware policies
- Enable and disable antimalware policies
- Set the priority of antimalware policies
- Remove antimalware policies
- Configure malware filtering to scan messages that were already scanned by Exchange Online Protection (EOP).

What do you need to know before you begin?

- We recommend that you manually download antimalware engine and definition updates on your Exchange server prior to placing it into production. For more information, see Download antimalware engine and definition updates.
- An antimalware policy consists of a malware filter policy and a malware filter rule. Each element controls different settings that don't overlap. The difference between these elements isn't visible in the EAC, but it's obvious in the Exchange Management Shell because you use different cmdlets to manage the settings (***-MalwareFilterPolicy** and ***-MalwareFilterRule**). This topic refers to antimalware policies for procedures in the EAC, and malware filter policies and malware filter rules for procedures in the Exchange Management Shell. For more information, see Antimalware protection in Exchange Server.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Antimalware" entry in the Antispam and antimalware permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to enable or disable malware

filtering on Mailbox servers

Disabling malware filtering on a Mailbox server disables the Malware agent and definition and engine updates.

1. To disable malware filtering on the local Mailbox server, run this command in the Exchange Management Shell:

   ```
   & $env:ExchangeInstallPath\Scripts\Disable-AntimalwareScanning.ps1
   ```

   To enable malware filtering on the local Mailbox server, run this command in the Exchange Management Shell:

   ```
   & $env:ExchangeInstallPath\Scripts\Enable-AntimalwareScanning.ps1
   ```

   If the command was successful, you see this message:

   ```
   Anti-malware scanning is successfully <enabled or disabled>. Please restart MSExchangeTransport for the
   changes to take effect.
   ```

   Note: The enable script also applies malware engine and definition updates as needed.

2. Restart the Exchange Transport service by running this command, which will temporarily interrupt mail flow on the server:

   ```
   Restart-Service MSExchangeTransport
   ```

   The change might take up to 10 minutes to take effect.

How do you know this worked?

To verify that you've successfully enabled or disabled malware filtering on a Mailbox server, run this command in the Exchange Management Shell, and verify the value of the **Enabled** property:

```
Get-TransportAgent "Malware Agent"
```

Use the Exchange Management Shell to bypass malware filtering on Mailbox servers

Bypassing malware filtering allows you to temporarily disable malware filtering on the server without disrupting mail flow (you don't need to restart the Exchange Transport service).

Note: You should *only* bypass malware filtering on a Mailbox server when you're troubleshooting a problem. When you're done, you should turn malware filtering back on.

To bypass or reenable malware filtering on a Mailbox server, use this syntax:

```
Set-MalwareFilteringServer -Identity <ServerIdentity> -BypassFiltering <$true | $false>
```

This example bypasses malware filtering on the server named Mailbox01.

```
Set-MalwareFilteringServer -Identity Mailbox01 -BypassFiltering $true
```

This example reenables malware filtering on the same server.

```
Set-MalwareFilteringServer -Identity Mailbox01 -BypassFiltering $false
```

The change might take up to 10 minutes to take effect.

For detailed syntax and parameter information, see Set-MalwareFilteringServer.

How do you know this worked?

To verify that you've temporarily bypassed or reenabled malware filtering on a Mailbox server, run this command in the Exchange Management Shell, and verify the value of the **BypassFiltering** property:

```
Get-MalwareFilteringServer | Format-List Name,BypassFiltering
```

Create antimalware policies

Use the EAC to create antimalware policies

Creating an antimalware policy in the EAC creates the malware filter rule and the associated malware filter policy at the same time using the same name for both.

1. In the EAC, go to **Protection** > **Malware filter**, and then click **New** +.

2. In the **New anti-malware policy** page that opens, configure these settings:

   - **Name**: Enter a unique, descriptive name for the policy.

   - **Description**: Enter an optional description for the policy.

   - **Malware detection response**: Select one of these options:

     - **Delete the entire message**: Prevents the entire message from being delivered to the intended recipients. This is the default value.

     - **Delete all attachments and use default alert text**: Replaces all message attachments (not just the detected ones) with a text file that contains this default text:

       > Malware was detected in one or more attachments included with this email. All attachments have been deleted.

     - **Delete all attachments and use custom alert text**: Replaces all message attachments (not just the detected ones) with a text file that contains custom text you specify in the **Custom alert text** field.

   > **NOTE**
   >
   > If malware is detected in the **message body** of an inbound or outbound message, the entire message is deleted, regardless of the setting you configure for **Malware detection response**.

   - **Notification**: The settings in this section control notifications when malware filtering deletes the message. The settings don't apply to messages where all attachments are replaced by the default or custom alert text.

     - **Sender Notifications**: Select one or both of these options:

       - **Notify internal senders**: An internal sender is inside the Exchange organization.

       - **Notify external senders**: An external sender is outside the Exchange organization.

     - **Administrator Notifications**: Select one or both of these options:

       - **Notify administrator about undelivered messages from internal senders**: If you select this option, enter a notification email address in the **Administrator email address** field.

       - **Notify administrator about undelivered messages from external senders**: If you select this option, enter a notification email address in the **Administrator email address** field.

     - **Customize Notifications**: These settings replace the default notification text that's used for senders or administrators. For more information about the default values, see Antimalware policies.

       - **Use customized notification text**: If you select this option, you need to use the **From name** and **From address** fields to specify the sender's name and email that's used in the customized notification message.

       - **Messages from internal senders**: If you elected to notify senders or administrators about undeliverable messages from internal senders, you need to use the **Subject** and **Message** fields to specify the subject and message body of the custom notification message.

       - **Messages from external senders**: If you elected to notify senders or administrators about undeliverable messages from external senders, you need to use the **Subject** and **Message** fields to specify the subject and message body of the custom notification message.

   - **Applied to**: The settings in this section identify the internal recipients that the policy applies to.

     - **If**: Click on the **Select one** drop down, and select conditions for the rule:

       - **The recipient is**: Specifies one or more mailboxes, mail users, or mail contacts in the Exchange organization. In the **Select members** dialog box that appears, select one or more recipients from the list, and then click **add ->**. In the **Check names** field, you can use wildcards for multiple email addresses (for example: *@fabrikam.com). When you're finished, click **OK**.

       - **The recipient domain is**: Specifies recipients in one or more of the configured accepted domains in the Exchange organization. In the dialog box that appears, select one or more domains, and then click **add ->**. When you're finished, click **OK**.

       - **The recipient is a member of**: Specifies one or more groups in the Exchange organization. In the **Select members** dialog box that appears, select one or more groups from the list, and then click **add ->**. When you're finished, click **OK**.

     You can only use one a condition once, but you can specify multiple values for the condition. To add more conditions, click **Add condition** and select from the remaining options.

     - **Except if**: To add exceptions for the rule, click **Add exception**, click on the **Select one** drop down, and configure an exception for the rule. The settings and behavior is exactly like the conditions.

3. When you're finished, click **Save**.

Use the Exchange Management Shell to create antimalware policies

Creating an antimalware policy in the Exchange Management Shell is a two-step process:

1. Create the malware filter policy.

2. Create the malware filter rule that specifies the malware filter policy that the rule applies to.

Notes:

- You can create a new malware filter rule and assign an existing, unassociated malware filter policy to it. A malware filter rule can't be associated with more than one malware filter policy.
- There are two settings that you can configure on new antimalware policies in the Exchange Management Shell that aren't available in the EAC until after you create the policy:
  - Create the new policy as disabled (*Enabled* `$false` on the **New-MalwareFilterPolicy** cmdlet).
  - Set the priority of the policy during creation (*Priority <Number>*) on the **New-MalwareFilterRule** cmdlet).
- Malware filter policies that you create in the Exchange Management Shell don't appear in the EAC until you assign the malware filter policy to a malware filter rule.
- A setting that's available in the Exchange Management Shell that isn't available in the EAC is the ability to turn malware filtering on or off for inbound messages or outbound messages by using the *BypassInboundMessages* or *BypassOutboundMessages* parameters on the **New-MalwareFilterPolicy** cmdlet.

Step 1: Use the Exchange Management Shell to create a malware filter policy

To create a malware filter policy, use this syntax:

```
New-MalwareFilterPolicy -Name "<PolicyName>" [-Action <DeleteMessage | DeleteAttachmentAndUseDefaultAlert |
DeleteAttachmentAndUseCustomAlert>] [-AdminDisplayName "<OptionalComments>"] [-BypassInboundMessages <$true |
$false>] [-BypassOutboundMessages <$true | $false>] [-CustomNotifications <$true | $false>] [<Inbound
notification options>] [<Outbound notification options>]
```

This example creates a new malware filter policy named Contoso Malware Filter Policy with these settings:

- Block messages that contain malware (we aren't using the *Action* parameter, and the default value is `DeleteMessage`).
- Don't notify the message sender when malware is detected in the message (we aren't using the *EnableExternalSenderNotifications* or *EnableInternalSenderNotifications* parameters, and the default value for both is `$false`).
- Notify the administrator admin@contoso.com when malware is detected in a message from an internal sender.

```
New-MalwareFilterPolicy -Name "Contoso Malware Filter Policy" -EnableInternalSenderAdminNotifications $true -
InternalSenderAdminAddress admin@contoso.com
```

For detailed syntax and parameter information, see New-MalwareFilterPolicy.

Step 2: Use the Exchange Management Shell to create a malware filter rule

To create a malware filter rule, use this syntax:

```
New-MalwareFilterRule -Name "<RuleName>" -MalwareFilterPolicy "<PolicyName>" <Recipient filters> [<Recipient
filter exceptions>] [-Comments "<OptionalComments>"]
```

This example creates a new malware filter rule named Contoso Recipients with these settings:

- The malware filter policy named Contoso Malware Filter Policy is associated with the rule.

- The rule applies to recipients in the contoso.com domain.

```
New-MalwareFilterRule -Name "Contoso Recipients" -MalwareFilterPolicy "Contoso Malware Filter Policy" -
RecipientDomainIs contoso.com
```

For detailed syntax and parameter information, see New-MalwareFilterRule.

How do you know this worked?

To verify that you've successfully created an antimalware policy, do any of these steps:

- In the EAC, go to **Protection** > **Malware filter**. Verify that the rule you created is in the list. Click **Edit** ✏ to verify the settings of the rule.

- In the Exchange Management Shell, replace *<PolicyName>* with the name of the malware filter policy, and run this command to verify the property values:

```
Get-MalwareFilterPolicy -Identity "<PolicyName>" | Format-List
```

- In the Exchange Management Shell, replace *<RuleName>* with the name of the malware filter rule, and run this command to verify the property values:

  ```
  Get-MalwareFilterRule -Identity "<RuleName>" | Format-List
  ```

- Use an European Institute for Computer Antivirus Research (EICAR) test file to verify that the malware filter is working correctly:

1. Open Notepad, and insert this text (and only this text) into an empty file:

   ```
   X5O!P%@AP[4\PZX54(P^)7CC)7}$EICAR-STANDARD-ANTIVIRUS-TEST-FILE!$H+H*
   ```

   Save the file as EICAR.txt in a location that's easy for you to find, and that's excluded from scanning by your computer's antivirus program. The file will be 68 bytes in size.

2. Create an email messages, attach the EICAR.txt file to the message, and send the message to a recipient in your Exchange organization who should be affected by the malware policy.

3. Check the recipient's mailbox to verify that malware filtering acted on the message: the message was deleted, or the message was delivered with the replacement alert text file for the attachment, and the notification messages were delivered to the sender and/or administrators.

4. When you're finished, delete the EICAR.TXT file so other users aren't unnecessarily alarmed.

View antimalware policies

Use the EAC to view antimalware policies

1. In the EAC, go to **Protection** > **Malware filter**.

2. When you select a policy, information about the policy is displayed in the details pane. To see more information about the policy, click **Edit** ✏.

   - The **Enabled** property value, the **Priority** property value, and the settings on the **Applied to** tab are in the malware filter rule.

   - The settings on the **General** and **Settings** tabs are in the malware filter policy.

Use the Exchange Management Shell to view malware filter policies

To return a summary list of all malware filter policies, run this command:

```
Get-MalwareFilterPolicy
```

To return detailed information about a specific malware filter policy, use the this syntax:

```
Get-MalwareFilterPolicy -Identity "<PolicyName>" | Format-List [<Specific properties to view>]
```

This example returns all the property values for the malware filter policy named Executives.

```
Get-MalwareFilterPolicy -Identity "Executives" | Format-List
```

This example returns only the specified properties for the same policy.

```
Get-MalwareFilterPolicy -Identity "Executives" | Format-List
Action,AdminDisplayName,CustomNotifications,Enable*Notifications
```

For detailed syntax and parameter information, see Get-MalwareFilterPolicy.

Use the Exchange Management Shell to view malware filter rules

To return a summary list of all malware filter rules, run this command:

```
Get-MalwareFilterRule
```

To return detailed information about a specific malware filter rule, use this syntax:

```
Get-MalwareFilterRule -Identity "<RuleName>" | Format-List [<Specific properties to view>]
```

This example returns all the property values for the malware filter rule named Executives.

```
Get-MalwareFilterRule -Identity "Executives" | Format-List
```

This example returns only the specified properties for the same rule.

```
Get-MalwareFilterRule -Identity "Executives" | Format-List
Name,Priority,State,MalwareFilterPolicy,*Is,*SentTo,*MemberOf
```

For detailed syntax and parameter information, see Get-MalwareFilterRule.

Modify antimalware policies

No additional settings are available when you modify a malware policy in the EAC or the Exchange Management Shell. They're the same settings that were available when you created the policy.

Use the EAC to modify an antimalware policy

1. In the EAC, go to **Protection** > **Malware filter**.

2. Select the policy, and then click **Edit** ✎. For information about the settings, see the Use the EAC to create antimalware policies section in this topic.

Notes:

- Instead of everything on one page, the settings are divided among the **General**, **Settings**, and **Applied to** tabs. The **Applied to** tab isn't available on the default policy named Default.

- You can't rename the default policy.

Use the Exchange Management Shell to modify a malware filter policy

To modify a malware filter policy, use this syntax:

```
Set-MalwareFilterPolicy -Identity "<PolicyName>" <Settings>
```

For detailed syntax and parameter information, see Set-MalwareFilterPolicy.

Use the Exchange Management Shell to modify a malware filter rule

When you modify a malware filter rule in the Exchange Management Shell, you can't disable or enable the rule (there's no *Enabled* parameter on the **Set-MalwareFilterRule** cmdlet). Instead, you use the **Disable-MalwareFilterRule** and **Enable-MalwareFilterRule** cmdlets as described later in this topic.

To modify a malware filter rule, use this syntax:

```
Set-MalwareFilterRule -Identity "<RuleName>" <Settings>
```

For detailed syntax and parameter information, see Set-MalwareFilterRule.

Enable or disable antimalware policies

By default, antimalware policies are enabled when you create them in the EAC or the Exchange Management Shell, but you can use the Exchange Management Shell to create a disabled malware filter rule (use the **New-MalwareFilterRule** cmdlet and the *Enabled* parameter with the value `$false`).

Use the EAC to enable or disable an antimalware policy

1. In the EAC, go to **Protection** > **Malware filter**.

2. Select the policy from the list, and then configure one of the following settings:

   - **Disable the policy**: Clear the check box in the **Enabled** column.

   - **Enable the policy**: Select the check box in the **Enabled** column.

Use the Exchange Management Shell to enable or disable malware filter rules

To enable or disable a malware filter rule in the Exchange Management Shell, use this syntax:

```
<Enable-MalwareFilterRule | Disable-MalwareFilterRule> -Identity "<RuleName>"
```

This example disables the malware filter rule named Marketing Department.

```
Disable-MalwareFilterRule -Identity "Marketing Department"
```

This example enables same rule.

```
Enable-MalwareFilterRule -Identity "Marketing"
```

For detailed syntax and parameter information, see Enable-MalwareFilterRule and Disable-MalwareFilterRule.

How do you know this worked?

To verify that you've successfully enabled or disabled an antimalware policy, use either of these procedures:

- In the EAC, go to **Protection** > **Malware filter**, and in the list of antimalware policies, verify the status of the check box in the **Enabled** column.
- In the Exchange Management Shell, run this command to see the list of rules and their **State** property values:

  ```
  Get-MalwareFilterRule
  ```

Set the priority of custom antimalware policies

By default, antimalware policies are given a priority that's based on the order they were created in (newer polices are lower priority than older policies). A lower priority number indicates a higher priority for the policy, and policies are processed in priority order (higher priority policies are processed before lower priority policies). No two policies can have the same priority.

Notes:

- In the EAC, you can only change the priority of the antimalware policy after you create it. In the Exchange Management Shell, you can override the default priority when you create the malware filter rule (which can affect the priority of existing rules).
- The default antimalware policy named Default has the priority value Lowest, and you can't change it.

Use the EAC to set the priority of custom antimalware policies

In the EAC, antimalware policies are processed in the order that they're displayed (the first policy has the **Priority** value 0). To change the priority of a policy, move the policy up or down in the list (you can't directly modify the **Priority** number in the EAC).

1. In the EAC, go to **Protection** > **Malware filter**.
2. Select a policy, and then click **Move up** (🡅) or **Move down** (🡇) to move the rule up or down in the list.

Use the Exchange Management Shell to set the priority of custom malware filter rules

The highest priority value you can set on a rule is 0. The lowest value you can set depends on the number of rules. For example, if you have five rules, you can use the priority values 0 through 4. Changing the priority of an existing rule can have a cascading effect on other rules. For example, if you have five rules (priorities 0 through 4), and you change the priority of a rule to 2, the existing rule with priority 2 is changed to priority 3, and the rule with priority 3 is changed to priority 4.

To set the priority of a malware filter rule in the Exchange Management Shell, use the following syntax:

```
Set-MalwareFilterRule -Identity "<RuleName>" -Priority <Number>
```

This example sets the priority of the rule named Marketing Department to 2. All existing rules that have a priority less than or equal to 2 are decreased by 1 (their priority numbers are increased by 1).

```
Set-MalwareFilterRule -Identity "Marketing Department" -Priority 2
```

Note: To set the priority of a new rule when you create it, use the *Priority* parameter on the **New-MalwareFilterRule** cmdlet.

How do you know this worked?

To verify that you've successfully modified the priority of an antimalware policy, use either of these procedures:

- In the EAC, go to **Protection** > **Malware filter**, and verify the **Priority** value of the antimalware policies in the list.
- In the Exchange Management Shell, run this command to see the list of rules and their **Priority** property values:

  ```
  Get-MalwareFilterRule
  ```

Remove antimalware policies

Note: You can't remove the default antimalware policy.

Use the EAC to remove antimalware policies

When you use the EAC to remove an antimalware policy, the malware filter rule and the corresponding malware filter policy are both removed.

1. From the EAC, go to **Protection** > **Malware filter**.
2. Select the antimalware policy you want to remove from the list, and then click **Delete** (🗑).

Use the Exchange Management Shell to remove malware filter policies

When you use the Exchange Management Shell to remove a malware filter policy, the corresponding malware filter rule isn't removed.

To remove a malware filter policy in the Exchange Management Shell, use this syntax:

```
Remove-MalwareFilterPolicy -Identity "<PolicyName>"
```

This example removes the malware filter policy named Marketing Department.

```
Remove-MalwareFilterPolicy -Identity "Marketing Department"
```

For detailed syntax and parameter information, see Remove-MalwareFilterPolicy.

Use the Exchange Management Shell to remove malware filter rules

When you use the Exchange Management Shell to remove a malware filter rule, the associated malware filter policy isn't removed.

To remove a malware filter rule in the Exchange Management Shell, use this syntax:

```
Remove-MalwareFilterRule -Identity "<RuleName>"
```

This example removes the malware filter rule named Marketing Department:

```
Remove-MalwareFilterRule -Identity "Marketing Department"
```

For detailed syntax and parameter information, see Remove-MalwareFilterRule.

How do you know this worked?

To verify that you've successfully removed an antimalware policy, use either of these procedures:

- In the EAC, go to **Protection** > **Malware filter**, and verify that the policy you removed is no longer in the list.

- In the Exchange Management Shell, run this command to verify that the malware filter policy you removed is no longer listed:

```
Get-MalwareFilterPolicy
```

- In the Exchange Management Shell, run this command to verify that the malware filter rule you removed is no longer listed:

```
Get-MalwareFilterRule
```

Use the Exchange Management Shell to configure malware filtering to rescan messages that were already scanned by EOP

By default, messages in transit that have been scanned by Exchange Online Protection (EOP) aren't scanned again by the Malware agent in Exchange. But, rescanning these messages can provide another layer of defense against malware.

To enable or disable scanning for malware in messages that have been already been scanned by EOP, use this syntax in the Exchange Management Shell:

```
Set-MalwareFilteringServer -Identity <ServerIdentity> -ForceRescan <$true | $false>
```

This example enables scanning for malware in messages that have already been scanned by EOP on the Mailbox server named Mailbox01.

```
Set-MalwareFilteringServer -Identity Mailbox01 -ForceRescan $true
```

This example disables scanning for malware in messages that have already been scanned by EOP on the same server.

```
Set-MalwareFilteringServer -Identity Mailbox01 -ForceRescan $false
```

How do you know this worked?

To verify that you've configured malware filtering to rescan messages that were already scanned by EOP, run this command in the Exchange Management Shell, and verify the value of the `ForceRescan` property:

```
Get-MalwareFilteringServer | Format-List Name, ForceRescan
```

Download antimalware engine and definition updates

8/3/2020 • 3 minutes to read • Edit Online

Administrators can manually download antimalware engine and definition (signature) updates. We strongly recommend that you download engine and definition updates before you put the Exchange server into production.

What do you need to know before you begin?

- Estimated time to complete: 5 minutes
- You can only use PowerShell to perform this procedure.

  To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.
- To download updates, your computer needs to be able to access the Internet and to establish a connection on TCP port 80 (HTTP). If your organization uses a proxy server for Internet access, see the following section in this topic.
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Antimalware" entry in the Antispam and antimalware permissions topic.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts in the Exchange admin center.

> **TIP**
>
> Having problems? Ask for help in the Exchange forums. Visit the forums at: Exchange Server, Exchange Online, or Exchange Online Protection.

Use the Exchange Management Shell to manually download engine and definition updates

To download engine and definition updates, run the following command:

```
& $env:ExchangeInstallPath\Scripts\Update-MalwareFilteringServer.ps1 -Identity <FQDN of server>
```

This example manually downloads the engine and definition updates on the Exchange server named mailbox01.contoso.com:

```
& $env:ExchangeInstallPath\Scripts\Update-MalwareFilteringServer.ps1 -Identity mailbox01.contoso.com
```

Optionally, you can use the *EngineUpdatePath* parameter to download updates from somewhere other than the default location. You can use this parameter to specify an alternate HTTP address or a UNC path. If you specify a UNC path, the network service must have access to the path.

This example manually downloads engine and definition updates on the Exchange server named mailbox01.contoso.com from the UNC path `\\FileServer01\Data\MalwareUpdates`:

```
& $env:ExchangeInstallPath\Scripts\Update-MalwareFilteringServer.ps1 -Identity mailbox01.contoso.com -
EngineUpdatePath \\FileServer01\Data\MalwareUpdates
```

How do you know this worked?

In order to verify that updates were downloaded successfully, you need to access Event Viewer and view the event log. We recommend that you filter only FIPFS events, as described in the following procedure.

1. From the **Start** menu, click **All Programs** > **Administrative Tools** > **Event Viewer**.
2. In Event Viewer, expand the **Windows Logs** folder, and then click **Application**.
3. In the **Actions** menu, click **Filter Current Log**.
4. In the **Filter Current Log** dialog box, from the **Event sources** drop-down list, select the **FIPFS** check box, and then click **OK**.

If engine updates were downloaded successfully, you will see Event ID 6033, which will appear similar to the following:

```
MS Filtering Engine Update process performed a successful scan engine update.
```

```
Scan Engine: Microsoft
```

```
Update Path: http://forefrontdl.microsoft.com/server/scanengineupdate
```

```
Last Update time: 2012-08-16T13:22:17.000Z
```

```
Engine Version: 1.1.8601.0
```

```
Signature Version: 1.131.2169.0
```

Use the Exchange Management Shell to configure proxy server settings for antimalware updates

If your organization uses a proxy server to control access to the Internet, you need to identify the proxy server so that you can successfully download antimalware engine and definition updates. Proxy server settings that are available using the **Netsh.exe** tool, Internet Explorer connection settings, and the *InternetWebProxy* parameter on the **Set-ExchangeServer** cmdlet don't affect how antimalware updates are downloaded.

To configure the proxy server settings for antimalware updates, perform the following steps.

1. Run the following command:

   ```
   Add-PsSnapin Microsoft.Forefront.Filtering.Management.Powershell
   ```

2. Use the **Get-ProxySettings** and **Set-ProxySettings** cmdlets to view and configure the proxy server settings that are used to download antimalware updates. The **Set-ProxySettings** cmdlet uses the following syntax:

   ```
   Set-ProxySettings -Enabled <$true | $false> -Server <Name or IP address of proxy server> -Port <TCP
   port of proxy server>
   ```

   For example, to configure antimalware updates to use the proxy server at address 172.17.17.10 on TCP port 80, run the following command.

```
Set-ProxySettings -Enabled $true -Server 172.17.17.10 -Port 80
```

To verify the proxy server settings, run the **Get-ProxySettings** cmdlet.

For more information

Procedures for antimalware protection in Exchange Server

Manually update scan engines in Microsoft Exchange Server

Running Windows antivirus software on Exchange servers

8/3/2020 • 8 minutes to read • Edit Online

When you run Windows antivirus programs on Microsoft Exchange servers, you can help enhance the security and health of your Exchange organization. However, if they aren't configured correctly, Windows antivirus programs can cause problems in Exchange Server.

There are two basic components of any Windows antivirus program:

- *Memory-resident scanning* or *real-time protection* monitors all files and processes that are loaded and running in a computer's active memory.
- *File-level scanning* refers to checking files on the hard disk for viruses manually or on a regular schedule. Some antivirus programs start an on-demand scan automatically after the virus signatures are updated to make sure that all files are scanned with the latest signatures.

The biggest potential problem is a Windows antivirus program might lock or quarantine an open log file or database file that Exchange needs to modify. This can cause severe failures in Exchange Server, and it might also generate 1018 event log errors. Therefore, excluding these files from being scanned by the Windows antivirus program is very important.

Another issues to consider is that Windows antivirus programs can't replace email-based antispam and antimalware solutions because Windows antivirus programs that run on Windows servers can't detect viruses, malware, and spam that are distributed only through email.

Recommended exclusions for Windows antivirus programs on Exchange servers

When you deploy a Windows antivirus program on an Exchange server, make sure that the folder exclusions, process exclusions, and file name extension exclusions that are described in these sections are configured for both memory-resident and file-level scanning.

Note: The **%ExchangeInstallPath%** value is typically `C:\Program Files\Microsoft\Exchange Server\V15\` (includes a trailing "\"), the **%SystemRoot%** value is typically `C:\Windows` (doesn't include a trailing "\"), and the **%SystemDrive%** value is typically `C:` (doesn't include a trailing "\").

The locations of many of these Exchange folders are configurable in the Exchange Management Shell. To learn how to open the Exchange Management Shell in your on-premises Exchange organization, see Open the Exchange Management Shell.

Folder exclusions

Exclude the following folders from file-level scanning and memory-resident scanning on Exchange servers.

> **NOTE**
>
> Unified Messaging is not available in Exchange 2019.

FOLDER	CATEGORY	DESCRIPTION	SERVERS	
`%SystemRoot%\Cluster`	DAGs	The cluster quorum database and other files for database availability groups (DAGs).	Mailbox servers	
`%SystemDrive%\DAGFileShareWitnesses\<DAGFQDN>`	DAGs	The witness directory on the witness server that's configured for the DAG. The witness server can be virtually any Microsoft Windows server in the local Active Directory forest that isn't already a member of the DAG. To see the actual location, run the following command: Get-DatabaseAvailabilityGroup <DAGName> | Format-List *Witness*	Any	
`%ExchangeInstallPath%ClientAccess\OAB`	Offline Address Books	Offline Address Book files.	Mailbox servers	
`%ExchangeInstallPath%FIP-FS`	Antimalware and DLP	Content scanning that's used by the Malware agent and data loss prevention (DLP).	Mailbox servers	
`%ExchangeInstallPath%GroupMetrics`	MailTips	Group Metrics files that are used to calculate values for the Large Audience and External Recipients MailTips.	Mailbox servers	

FOLDER	CATEGORY	DESCRIPTION	SERVERS	
`%ExchangeInstallPath%Logging`	Exchange process logs	This folder contains many different types of Exchange logs in subfolders. For example: • Calendar Repair Assistant logs • Managed Folder Assistant logs • IMAP4 protocol logs • POP3 protocol logs To see the actual locations, run the following commands: Get-MailboxServer -Identity <ServerName> | Format-List *LogPath* Get-PopSettings <ServerName> | Format-List LogFileLocation Get-ImapSettings <ServerName> | Format-List LogFileLocation	Mailbox servers	
`%ExchangeInstallPath%Mailbox`	Mailbox databases	Exchange databases, checkpoint files, and log files. By default, these files are located in subfolders based on the name of the database. To see the actual locations, run the following command: Get-MailboxDatabase -Server <ServerName> | Format-List EdbFilePath,LogFolderPath By default, database context index files are located in the same folder as the database files in a subfolder that's named after the GUID of the database.	Mailbox servers	
`%ExchangeInstallPath%TransportRoles\Data\Adam`	EdgeSync	Active Directory Lightweight Directory Services (AD LDS) and log files.	Edge Transport servers	
`%ExchangeInstallPath%TransportRoles\Data\IpFilter`	Connection filtering	IP filter database, checkpoint, and log files.	Edge Transport servers	

FOLDER	CATEGORY	DESCRIPTION	SERVERS	
`%ExchangeInstallPath%TransportRoles\Data\Queue`	Queues	Queue database, checkpoint, and log files.	Mailbox servers Edge Transport servers	
`%ExchangeInstallPath%TransportRoles\Data\SenderReputation`	Sender reputation	Sender Reputation database, checkpoint, and log files.	Edge Transport servers Mailbox servers	
`%ExchangeInstallPath%TransportRoles\Data\Temp`	Content conversion	Content conversion that's done in the transport pipeline.	Mailbox servers Edge Transport servers	
`%ExchangeInstallPath%TransportRoles\Logs`	Transport logs	Mail flow and transport pipeline logs are located in subfolders, for example: • Agent logging • Connectivity logging • Message tracking • Pipeline tracing • Send and Receive connector protocol logging To see the actual locations, run the following commands: Get-TransportService <ServerName> | Format-List *LogPath,*TracingPath Get-FrontEndTransportService <ServerName> | Format-List *LogPath Get-MailboxTransportService <ServerName> | Format-List *LogPath,*TracingPath	Mailbox servers Edge Transport servers (Transport service only)	
`%ExchangeInstallPath%TransportRoles\Pickup`	Pickup directory	The Pickup directory is used by administrators for mail flow testing or by applications that need to create and submit their own message files. To see the actual location, run the following command: Get-TransportService <ServerName> | Format-List PickupDirectoryPath	Mailbox servers Edge Transport servers	

FOLDER	CATEGORY	DESCRIPTION	SERVERS	
`%ExchangeInstallPath%TransportRoles\Replay`	Replay directory	The Replay directory receives messages from foreign gateway servers and can also be used to resubmit messages that administrators export from the queues of Exchange servers. To see the actual location, run the following command: Get-TransportService <ServerName>	Format-List ReplayDirectoryPath	Mailbox servers Edge Transport servers
`%ExchangeInstallPath%UnifiedMessaging\Grammars`	Unified Messaging	Grammar files for different locales, for example en-EN or es-ES.	Exchange 2016 Mailbox servers	
`%ExchangeInstallPath%UnifiedMessaging\Prompts`	Unified Messaging	Voice prompts, greetings, and informational message files.	Exchange 2016 Mailbox servers	
`%ExchangeInstallPath%UnifiedMessaging\Temp`	Unified Messaging	Temporary files generated by Unified Messaging.	Exchange 2016 Mailbox servers	
`%ExchangeInstallPath%UnifiedMessaging\Voicemail`	Unified Messaging	Voice mail files that are temporarily stored.	Exchange 2016 Mailbox servers	
`%ExchangeInstallPath%Working\OleConverter`	Content conversion	Transport Neutral Encoding Format (TNEF), also known as Rich Text Format (RTF), to MIME/HTML conversions.	Mailbox servers Edge Transport servers	
`%SystemDrive%\inetpub\temp\IIS Temporary Compressed Files`	Web components	Internet Information Services (IIS) compression folder that's used with Outlook on the web.	Mailbox servers	

FOLDER	CATEGORY	DESCRIPTION	SERVERS	
`%SystemRoot%\Microsoft.NET\Framework64\v4.0.30319\Temporary ASP.NET Files`	Web components	Temporary files that are used with Exchange services. These files are located in the following subfolders: autodiscover ecp ecp ews mapi mapi_emsmdb microsoft-server-activesync oab owa owa_calendar powershell root rpc	Mailbox servers	
`%SystemRoot%\System32\Inetsrv`	Web components	IIS system files.	Mailbox servers	
`%SystemRoot%\Temp\OICE_<GUID>`	Exchange Search	Temporary files used by the Exchange Search service and Microsoft Filter Pack to perform file conversion in a sandboxed environment.	Mailbox servers	

Process exclusions

Many antivirus programs support the scanning of processes, which can adversely affect Microsoft Exchange if the incorrect processes are scanned. Therefore, you should exclude the following Exchange or related processes from process scanning.

PROCESS	PATH	COMMENTS	SERVERS
ComplianceAuditService.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange Compliance Audit service (MSComplianceAudit)	Mailbox servers
Dsamain.exe	`%SystemRoot%\System32`	Microsoft Exchange ADAM service (ADAM_MSExchange) (Active Directory Lightweight Directory Services (AD LDS) on subscribed Edge Transport servers)	Edge Transport servers
EdgeTransport.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange Transport service worker process	Mailbox servers Edge Transport servers

PROCESS	PATH	COMMENTS	SERVERS
fms.exe	`%ExchangeInstallPath%FIP-FS\Bin`	Content scanning component that's used by the Malware agent and DLP.	Mailbox servers
hostcontrollerservice.exe	`%ExchangeInstallPath%Bin\Search\Ceres\HostController`	Microsoft Exchange Search Host Controller service (HostControllerService)	Mailbox servers
inetinfo.exe	`%SystemRoot%\System32\inetsrv`	Internet Information Services (IIS)	Mailbox servers
Microsoft.Exchange.AntispamUpdateSvc.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange Antispam Update service (MSExchangeAntispamUpdate)	Mailbox servers Edge Transport servers
Microsoft.Exchange.ContentFilter.Wrapper.exe	`%ExchangeInstallPath%TransportRoles\agents\Hygiene`	Content Filter agent	Mailbox servers Edge Transport servers
Microsoft.Exchange.Diagnostics.Service.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange Diagnostics service (MSExchangeDiagnostics)	Mailbox servers Edge Transport servers
Microsoft.Exchange.Directory.TopologyService.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange Active Directory Topology service (MSExchangeADTopology)	Mailbox servers
Microsoft.Exchange.EdgeCredentialSvc.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange Credential service (MSExchangeEdgeCredential)	Edge Transport servers
Microsoft.Exchange.EdgeSyncSvc.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange EdgeSync service (MSExchangeEdgeSync)	Mailbox servers
Microsoft.Exchange.Imap4.exe	`ExchangeInstallPath%FrontEnd\PopImap`	Microsoft Exchange IMAP4 service (MSExchangeImap4)	Mailbox servers
Microsoft.Exchange.Imap4service.exe	`%ExchangeInstallPath%ClientAccess\PopImap`	Microsoft Exchange IMAP4 Backend service (MSExchangeIMAP4BE)	Mailbox servers
Microsoft.Exchange.Notifications.Broker.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange Notifications Broker service (MSExchangeNotificationsBroker)	Mailbox servers
Microsoft.Exchange.Pop3.exe	`%ExchangeInstallPath%FrontEnd\PopImap`	Microsoft Exchange POP3 service (MSExchangePop3)	Mailbox servers
Microsoft.Exchange.Pop3service.exe	`%ExchangeInstallPath%ClientAccess\PopImap`	Microsoft Exchange POP3 Backend service (MSExchangePOP3BE)	Mailbox servers

PROCESS	PATH	COMMENTS	SERVERS
Microsoft.Exchange.ProtectedServiceHost.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange Service Host service (MSExchangeServiceHost)	Mailbox servers Edge Transport servers
Microsoft.Exchange.RPCClientAccess.Service.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange RPC Client Access service (MSExchangeRPC)	Mailbox servers
Microsoft.Exchange.Search.Service.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange Search service (MSExchangeFastSearch)	Mailbox servers
Microsoft.Exchange.Servicehost.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange Service Host service (MSExchangeServiceHost)	Mailbox servers Edge Transport servers
Microsoft.Exchange.Store.Service.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange Information Store service (MSExchangeIS)	Mailbox servers
Microsoft.Exchange.Store.Worker.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange Information Store service worker process	Mailbox servers
Microsoft.Exchange.UM.CallRouter.exe	`%ExchangeInstallPath%FrontEnd\CallRouter`	Microsoft Exchange Unified Messaging Call Router service (MSExchangeUMCR)	Exchange 2016 Mailbox servers
MSExchangeCompliance.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange Compliance Service (MSExchangeCompliance)	Mailbox servers
MSExchangeDagMgmt.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange DAG Management service (MSExchangeDagMgmt)	Mailbox servers
MSExchangeDelivery.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange Mailbox Transport Delivery service (MSExchangeDelivery)	Mailbox servers
MSExchangeFrontendTransport.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange Frontend Transport service (MSExchangeFrontEndTransport)	Mailbox servers
MSExchangeHMHost.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange Health Manager service (MSExchangeHM)	Mailbox servers Edge Transport servers
MSExchangeHMWorker.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange Health Manager service worker process	Mailbox servers Edge Transport servers

PROCESS	PATH	COMMENTS	SERVERS
MSExchangeMailboxAssistants.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange Mailbox Assistants service (MSExchangeMailboxAssistants)	Mailbox servers
MSExchangeMailboxReplication.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange Mailbox Replication service (MSExchangeMailboxReplication)	Mailbox servers
MSExchangeRepl.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange Replication service (MSExchangeRepl)	Mailbox servers
MSExchangeSubmission.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange Mailbox Transport Submission service (MSExchangeSubmission)	Mailbox servers
MSExchangeTransport.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange Transport service (MSExchangeTransport)	Mailbox servers Edge Transport servers
MSExchangeTransportLogSearch.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange Transport Log Search service (MSExchangeTransportLogSearch)	Mailbox servers Edge Transport servers
MSExchangeThrottling.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange Throttling service (MSExchangeThrottling)	Mailbox servers
Noderunner.exe	`%ExchangeInstallPath%Bin\Search\Ceres\Runtime\1.0`	Microsoft Exchange Search service (MSExchangeFastSearch)	Mailbox servers
OleConverter.exe	`%ExchangeInstallPath%Bin`	Converts rich text format (RTF) messages to MIME/HTML for external recipients.	Mailbox servers
ParserServer.exe	`%ExchangeInstallPath%Bin\Search\Ceres\ParserServer`	Microsoft Exchange Search service (MSExchangeFastSearch)	Mailbox servers
Powershell.exe	`C:\Windows\System32\WindowsPowerShell\v1.0`	Exchange Management Shell	Mailbox servers Edge Transport servers
ScanEngineTest.exe	`%ExchangeInstallPath%FIP-FS\Bin`	Content scanning component that's used by the Malware agent and DLP	Mailbox servers
ScanningProcess.exe	`%ExchangeInstallPath%FIP-FS\Bin`	Content scanning component that's used by the Malware agent and DLP	Mailbox servers

PROCESS	PATH	COMMENTS	SERVERS
UmService.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange Unified Messaging service (MSExchangeUM)	Exchange 2016 Mailbox servers
UmWorkerProcess.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange Unified Messaging service worker process	Exchange 2016 Mailbox servers
UpdateService.exe	`%ExchangeInstallPath%FIP-FS\Bin`	Content scanning component that's used by the Malware agent and DLP	Mailbox servers
W3wp.exe	`%SystemRoot%\System32\inetsrv`	Internet Information Services (IIS)	Mailbox servers
wsbexchange.exe	`%ExchangeInstallPath%Bin`	Microsoft Exchange Server Extension for Windows Server Backup (wsbexchange)	Mailbox servers

File name extension exclusions

In addition to excluding specific folders and processes, you should exclude the following Exchange-specific file name extensions in case folder exclusions fail or files are moved from their default locations.

EXTENSIONS	DESCRIPTION	SERVERS
.config	Application-related extensions	Mailbox servers Edge Transport servers
.chk .edb .jfm .jrs .log .que	Database-related extensions	Mailbox servers Edge Transport servers
.dsc .txt	Group Metrics-related extensions	Mailbox servers
.cfg .grxml	Unified Messaging-related extensions	Exchange 2016 Mailbox servers
.lzx	Offline address book-related extensions	Mailbox servers

Unified Messaging in Exchange Server 2016

8/3/2020 • 2 minutes to read • Edit Online

Unified Messaging in Exchange Server 2016 is basically unchanged from Exchange Server 2013. For information about Exchange 2013 Unified Messaging, see Unified Messaging.

About Exchange documentation

8/3/2020 • 2 minutes to read • Edit Online

You're reading a collection of conceptual and procedural topics organized by subject or by technologies used by Microsoft Exchange. You can access each topic directly from the table of contents in the left pane, from a link in another Help topic, from the results of a search, or from your own custom list of favorite topics.

Other information related to Exchange documentation is in Third-party copyright notices.

Where to find Exchange documentation

Exchange documentation is your primary gateway to in-depth technical information about Microsoft Exchange.

The Exchange Team Blog contains technical articles written by the Exchange Team, as well as product announcements and updates. The blog is an excellent way to interact with the Exchange Team. We read and respond to your feedback and comments.

If you're an admin for an Exchange hybrid or Exchange Online deployment, you may also be interested in Manage Microsoft 365 and Office 365.

Additional resources

Looking for more than just documentation? Check out these other Exchange resources:

- Exchange Server Forums: The forum provides a place to discuss Exchange with users and Exchange Team members.
- Exchange and Exchange Online development: You'll find Exchange developer documentation here.
- Support for business: Select **Servers** > **Exchange Server** for support resources for multiple versions of Exchange.
- Accessibility for people with disabilities: This topic provides important information about features, products, and services that help make Microsoft Exchange more accessible for people with disabilities.

Accessibility for people with disabilities

8/3/2020 • 3 minutes to read • Edit Online

Microsoft is committed to making its products and services easier for everyone to use. The following sections provide information about the features, products, and services that make Microsoft Exchange more accessible for people with disabilities:

- Accessibility features of Exchange
- Accessibility features of Exchange Help
- Accessibility products and services from Microsoft

Accessibility features of Exchange

The following features help make Microsoft Exchange more accessible for people with disabilities:

- Keyboard shortcuts in the Exchange admin center
- Keyboard Shortcuts in Outlook on the web

In addition, some accessibility features and utilities of Windows may benefit Exchange users with disabilities. Also, Windows PowerShell size and color changes provide accessibility options when using the Exchange Management Shell. For more information about Windows PowerShell accessibility options, see Accessibility in Windows PowerShell ISE.

Accessibility features of Exchange Help

Every figure in Help for Microsoft Exchange, including screenshots, diagrams, flow charts, and other figures, has associated alternate text. Users who have difficulty viewing figures can pause the cursor on the figure to read the alternate text. The alternate text describes what is illustrated in the figure.

Accessibility products and services from Microsoft

The following sections provide information about the features, products, and services that make Microsoft Windows more accessible for people with disabilities.

> **NOTE**
>
> The information in this section applies only to users who license Microsoft products in the United States. If you obtained this product outside of the United States, visit the Microsoft Accessibility website for a list of telephone numbers and addresses for Microsoft support services. You can contact your subsidiary to find out whether the type of products and services described in this section are available in your area. You can learn more about the accessibility features included in Microsoft products on the Accessibility in Microsoft Products web site.

Accessibility features of Windows

The Windows operating system has many built-in accessibility features that are useful for individuals who have difficulty typing or using a mouse, are blind or have low vision, or who are deaf or hard-of-hearing. The features are installed during Setup. For more information about these features, see Help in Windows and Microsoft Accessibility.

- **Free step-by-step tutorials**: Microsoft offers a series of step-by-step tutorials that provide detailed procedures for adjusting the accessibility options and settings on your computer. This information is presented in a side-by-side format so that you can learn how to use the mouse, the keyboard, or a combination of both.

To find step-by-step tutorials for Microsoft products, see Microsoft Accessibility.

- **Assistive technology products for Windows**: A wide variety of assistive technology products are available to make computers easier to use for people with disabilities. You can search a catalog of assistive technology products that run on Windows at Microsoft Accessibility.

  If you use assistive technology, be sure to contact your assistive technology vendor before you upgrade your software or hardware to check for possible compatibility issues.

Documentation in alternative formats

If you have difficulty reading or handling printed materials, you can obtain the documentation for many Microsoft products in more accessible formats. You can obtain an index of accessible product documentation at Microsoft Accessibility.

In addition, you can obtain additional Microsoft publications from Learning Ally. Learning Ally distributes these documents to registered, eligible members of their distribution service. For information about the availability of Microsoft product documentation and books from Microsoft Press, contact Learning Ally.

> Learning Ally
> 20 Roszel Road
> Princeton, NJ 08540
> Telephone number from within the United States: (800) 221-4792
> Web site: Learning Ally

Customer service for people with hearing impairments

If you're deaf or hard-of-hearing, complete access to Microsoft product and customer services is available through a text telephone (TTY/TDD) service:

- For customer service, contact Microsoft Sales Information Center at (800) 892-5234 between 6:30 A.M. and 5:30 P.M. Pacific Time, Monday through Friday, excluding holidays.

- For technical assistance in the United States, contact Microsoft Product Support Services at (800) 892-5234 between 6:00 A.M. and 6:00 P.M. Pacific Time, Monday through Friday, excluding holidays. In Canada, dial (905) 568-9641 between 8:00 A.M. and 8:00 P.M. Eastern Time, Monday through Friday, excluding holidays.

Microsoft Support Services are subject to the prices, terms, and conditions in place at the time the service is used. For more information, see Microsoft Support.

For more information

For more information about how accessible technology for computers helps to improve the lives of people with disabilities, see Microsoft Accessibility.

Third-party copyright notices

8/3/2020 • 2 minutes to read • Edit Online

Outside In HTML Export © 1991, 2011 Oracle

Platforms Supported - Outside In HTML Export:

Windows (32-bit):

Windows 2000

Windows Server 2003

Windows Vista

Windows Server 2008

Windows XP

Windows 7

Windows Itanium (64 bit):

Windows .NET Server 2003 Enterprise Edition for Itanium

Windows (64 bit):

Windows 2003 x 64 Datacenter

Windows 2003 x 64 Enterprise

Windows 2003 x 64 Standard Windows Server

Windows Server 2008

Windows Server 2008 R2

Windows 7

Keyboard shortcuts in the Exchange admin center

8/3/2020 • 2 minutes to read • Edit Online

Microsoft is committed to making its products and services easier for everyone to use. This topic provides information about the keyboard shortcuts that make Exchange Server and other Microsoft products and services more accessible for people with disabilities.

Keyboard shortcuts in the Exchange admin center in Exchange Server

By using keyboard shortcuts in the Exchange admin center (EAC), you can quickly accomplish the common tasks that are described in the following table. To learn more about the EAC, see Exchange admin center in Exchange Server.

Keyboard shortcuts in the EAC

TO DO THIS	USE THIS KEYBOARD SHORTCUT
Move between areas or between controls in the EAC	Tab Shift-Tab
Move between items in drop-down menus in the EAC	Up arrow key Down arrow key Note that you can't use Tab or SHIFT-Tab to move between items in drop-down menus
Move within lists from one item to another	Up arrow key Down arrow key Page Up Page Down Home End Note that you can also use the Up, Down, Left, and Right arrow keys to: • Move between option buttons. • Move within a group of associated check boxes.
Move within primary property pages from one item to another	Up arrow key Down arrow key Page Up Page Down Home End Tab Shift-Tab You can use Enter or the Spacebar to activate your selection.

TO DO THIS	USE THIS KEYBOARD SHORTCUT
Move within secondary property pages from one item to another	Up arrow key Down arrow key Page Up Page Down Home End Tab Shift-Tab You can use Enter or the Spacebar to activate your selection.

Keyboard shortcuts in other Microsoft products and services

To learn about accessibility features in Microsoft 365 or Office 365, including keyboard shortcuts, visit the Microsoft Accessibility website.

Exchange Server Privacy Statement

8/3/2020 • 11 minutes to read • Edit Online

Microsoft is committed to protecting your privacy, while delivering software that brings you the performance, power, and convenience you desire in your personal computing. This privacy statement applies to Microsoft Exchange Server 2016 and Exchange Server 2019. It focuses on features that communicate with the Internet. It does not apply to any other online or offline Microsoft sites, products, or services.

Exchange 2016 and Exchange 2019 deliver email, calendaring, contact management, and other online collaboration functionalities on your PC's, mobile phones, and web browsers.

This privacy statement addresses the deployment and use of Exchange 2016 and Exchange 2019 in an enterprise network environment. If you use Exchange Server technologies as a service operated by Microsoft or a third party, please refer to the service-specific privacy and security policies provided by Microsoft or the third-party service provider.

IT administrators of Exchange 2016 and Exchange 2019 may choose to enable or disable certain Internet-enabled features in Exchange, or to deploy other privacy impacting technologies, based on legal or compliance considerations or internal policies. You should direct privacy-related requests related to the entity that's providing your access to Exchange 2016 or Exchange 2019. Microsoft is not responsible for the privacy practices of its customers or other third parties.

Collection and Use of Your Information

The information we collect from you will be used by Microsoft and its controlled subsidiaries and affiliates to enable the features you're using and to provide the service(s) or carry out the transaction(s) you have requested or authorized. It may also be used to analyze and improve Microsoft products and services.

We may send standard service communications such as welcome letters, billing reminders, information on technical service issues, and security announcements. Some Microsoft services may send periodic member letters that are considered part of the service. We may occasionally request your feedback, invite you to participate in surveys, or send you promotional mailings to inform you of other products or services available from Microsoft and its affiliates.

In order to offer you a more consistent and personalized experience in your interactions with Microsoft, information collected through one Microsoft service may be combined with information collected through other Microsoft services. We may also supplement the information we collect with information obtained from other companies. For example, we may use services from other companies that enable us to derive a general geographic area based on your IP address in order to customize certain services to your geographic area.

Except as described in this statement, personal information you provide will not be transferred to third parties without your consent.

We occasionally hire other companies to provide limited services on our behalf, such as answering customer questions about products or services or performing statistical analysis of our services. We will provide those companies only the personal information they need to deliver the service, and they are prohibited from using that information for any other purpose.

Microsoft may access or disclose information about you, including the content of your communications, in order to: (a) comply with the law or respond to lawful requests or legal process; (b) protect the rights or property of Microsoft or our customers, including the enforcement of our agreements or policies governing your use of the services; or (c) act on a good faith belief that such access or disclosure is necessary to protect the personal safety of

Microsoft employees, customers, or the public.

Information that is collected by or sent to Exchange 2016 and Exchange 2019 may be stored and processed in the United States or in any other country/region in which Microsoft or its affiliates, subsidiaries, or service providers maintain facilities. Microsoft abides by the safe harbor framework as set forth by the United States Department of Commerce regarding the collection, use, and retention of data from the European Union, the European Economic Area, and Switzerland.

Collection and Use of Information about Your Computer

When you use software with Internet-enabled features, information about your computer ("standard computer information") is sent to the web sites you visit and online services you use. Microsoft uses standard computer information to provide you Internet-enabled services, to help improve our products and services, and for statistical analysis. Standard computer information typically includes information such as your IP address, operating system version, browser version, and regional and language settings. In some cases, standard computer information may also include hardware ID, which indicates the device manufacturer, device name, and version. If a particular feature or service sends information to Microsoft, standard computer information will be sent as well.

The privacy details for other Exchange 2016 or Exchange 2019 features, software, or services listed in this privacy statement describe what additional information is collected and how it is used.

Security of Your Information

Microsoft is committed to helping protect the security of your information. We use a variety of security technologies and procedures to help protect your information from unauthorized access, use, or disclosure. For example, information you provide is stored on computer systems with limited access, which are located in controlled facilities.

Specific Feature: Microsoft Error Reporting

What This Feature Does: Microsoft Error Reporting provides a service that allows you to report problems you may be having to Microsoft and to receive information that may help you avoid or solve such problems.

Information Collected, Processed, or Transmitted: For information about the information collected, processed, or transmitted by Microsoft Error Reporting, see Privacy Statement for the Microsoft Error Reporting Service.

Use of Information: The error reporting data that you submit may be used to solve customer problems and to improve Microsoft software and services.

Choice/Control: You will be offered the opportunity to participate in Microsoft Error Reporting the first time an error is encountered. When you choose to enable it, Microsoft Error Reporting will automatically report problems you encounter to Microsoft. In addition, your IT administrator can choose to enable or disable Microsoft Error Reporting during the Exchange Server setup process for all users.

In rare cases, such as problems that are especially difficult to solve, Microsoft may request additional data, including sections of memory (which may include memory shared by any or all applications running at the time the problem occurred), some registry settings, and one or more files from your computer. Your current documents may also be included. When additional data is requested, you will have an opportunity to view the information contained in the error report before choosing whether or not to send the report to Microsoft.

Important Information: Enterprise customers can use Group Policy to configure how Microsoft Error Reporting works on their computers. Configuration options include the ability to turn off Microsoft Error Reporting. If you're an administrator and want to configure Group Policy for Microsoft Error Reporting, technical details are available at the Group Policy Settings Reference for Windows and Windows Server.

Specific Feature: Online Feedback

What This Feature Does: Online Feedback allows you to provide feedback about products and services directly to Microsoft.

Information Collected, Processed, or Transmitted: If you choose to use Online Feedback, the content of the message and standard computer information will be sent to Microsoft.

Use of Information: The information submitted may be used to improve Microsoft sites, products, or services. The information that we collect from this feature may also be used to request additional information about feedback provided about the product or service.

Choice/Control: Use of Online Feedback is optional.

Specific Feature: Online Help

What This Feature Does: Clicking or otherwise using Help connects to online support materials, providing you with the most up-to-date content available.

Information Collected, Processed, or Transmitted: When you use Help, the request is sent to Microsoft, as well as any rating or feedback provided about the help topics. Any personal information entered into the search or feedback boxes will be sent to Microsoft but will not be used to identify or contact you.

Use of Information: Help uses search information to return the most relevant results, develop new content, and improve the existing content.

Choice/Control: Do not use Help if you do not wish to connect to online support materials.

Specific Feature: Customer Experience Improvement Program

What This Feature Does: If you choose to participate, the Customer Experience Improvement Program (CEIP) collects basic information about your hardware configuration and how you use Microsoft software and services in order to identify trends and usage patterns. CEIP also collects the type and number of errors you encounter, software and hardware performance, and the speed of services. Microsoft does not collect your name, address, or other contact information.

Information Collected, Processed, or Transmitted: CEIP information is automatically sent to Microsoft when the feature is turned on. For more information about the information collected, processed, or transmitted by CEIP, see the Privacy Statement for the Microsoft Customer Experience Improvement Program.

Use of Information: Microsoft uses this information to improve the quality, reliability, and performance of Microsoft software and services.

Choice/Control: CEIP is turned off by default unless your IT administrator has chosen to turn it on for you. You will be prompted to sign up in the Exchange installer. Unless your administrator has restricted your ability to do so, you can change your CEIP settings at any time.

Specific Feature: Bing Maps Extension

What This Feature Does: the Bing Maps extension will appear in Outlook and Outlook on the web (formerly known as Outlook Web App) when Exchange 2016 or Exchange 2019 detects the presence of an address in the body of an email and allow you to query the Bing Maps service for a map of the location.

Information Collected, Processed, or Transmitted: when you click on the Bing Maps extension from the user interface, the information that Exchange determines to be an address will be passed to the Bing Maps service, which will perform a query based on the address and return a map for it. For more information on Bing's privacy practices, see the Bing section in Microsoft Privacy Statement.

Use of Information: The address information is used to display the map for the address.

Choice/Control: This feature can be turned off by the IT administrator, or the end user.

Specific Feature: Offline

What This Feature Does: In Outlook on the web, the Offline features stores contacts, calendar and email information on a user's machine so that it is accessible without a network connection.

Information Collected, Processed, or Transmitted: Information collected and stored includes contacts, calendar and email from the user's mailbox on the Exchange Server. This feature does not transmit information to Microsoft.

Use of Information: Information is stored locally on the user's machine to enable offline access.

Choice/Control: In Outlook on the web, this Offline feature is off by default and the user must enable offline mode. Additionally, the IT administrators can disable the option for this feature so that users are not able to turn it on.

Specific Feature: Sender Photo

What This Feature Does: In Outlook on the web, the Sender Photo feature allows a recipient to see the photo of the sender of an email he or she is viewing, if the sender is from the same organization as the recipient.

Information Collected, Processed, or Transmitted: When Outlook on the web believes a sender's photo may be available, the sender's email address is sent to the Exchange server. Transmission may be sent partially in an unencrypted form. This feature does not transmit information to Microsoft.

Use of Information: The sender's email address will be used to locate the sender's photo from the destination Exchange server.

Choice/Control: IT administrators can turn off this feature.

Specific Feature: Contact Card

What This Feature Does: When browsing emails in Outlook on the web, a user can click on the name of the sender or one of the recipients in a mail to retrieve the contact card for the individual. The contact card displays information about that individual.

Information Collected, Processed, or Transmitted: When a user requests the contact card information for someone by clicking on the person's name from an email, the email address or similar identifier for the person whose information is requested is sent to the Exchange server. Transmission may be sent partially in an unencrypted form. This feature does not transmit information to Microsoft.

Use of Information: The email address of the person whose contact card is requested will be used to locate the person's contact card information from the destination Exchange server.

Choice/Control: Do not click on an individual's name to retrieve the individual's contact card if you do not wish to send the individual's identification to the Exchange server.

Changes to This Privacy Statement

Microsoft is committed to helping protect the security of your information. We use a variety of security technologies and procedures to help protect your information from unauthorized access, use, or disclosure. For example, information you provide is stored on computer systems with limited access, which are located in controlled facilities.

For More Information

Microsoft welcomes your comments regarding this privacy statement and its supplements. If you have questions or believe that we have not adhered to these documents, please contact us using this web form, or by email, at the following address:

Microsoft Privacy

Microsoft Corporation

One Microsoft Way

Redmond, Washington 98052 USA

425-882-8080

To find the Microsoft subsidiary in your country or region, see Microsoft office locations around the world.