



**Ninth Edition (2023)**

# **Office 365 for IT Pros**

*The ultimate guide to planning, deploying, and managing the Microsoft 365 Cloud Productivity Suite*

**Tony Redmond**

with Paul Robichaux, Brian Desmond,  
Gareth Gudger, Christina Wheeler,  
Juan Carlos González, and Ben Lee

Technical Editor: Vasil Michev

# Office 365 for IT Pros (2023 Edition)

## Mastering Microsoft 365 Office Applications

Published by Tony Redmond (<https://office365itpros.com>)

© Copyright 2015-2022 by Tony Redmond.

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

The example companies, organizations, products, domain names, email addresses, logos, people, places, and events depicted herein are fictitious. No association with any real company, organization, people, domain name, email address, logo, person, place, or event is intended or should be inferred. The book expresses the views and opinions of the authors. The information presented in the book is provided without any express, statutory, or implied warranties. The authors cannot be held liable for any damages caused or alleged to be caused either directly or indirectly by this book.

*Although the authors are members of Microsoft's Most Valuable Professional (MVP) program, the content of this book solely represents their views and opinions about Office 365 and any other technologies mentioned in the text and is not endorsed in any way by Microsoft Corporation.*

**Please be respectful of the rights of the authors and do not make and distribute copies of this eBook available to others.**

Ninth (2023) edition. Previous editions appeared under the following titles:

- Office 365 for Exchange Professionals (May 2015 and September 2015).
- Office 365 for IT Pros (3<sup>rd</sup> edition – June 2016).
- Office 365 for IT Pros (4<sup>th</sup> edition – June 2017).
- Office 365 for IT Pros (5<sup>th</sup> edition – July 2018).
- Office 365 for IT Pros (6<sup>th</sup> edition – July 2019).
- Office 365 for IT Pros (7<sup>th</sup> edition – July 2020).
- Office 365 for IT Pros (8<sup>th</sup> edition – July 2021).

The authors issue regular updates for this eBook until they publish a new edition, usually after a year. People who buy the book through [gumroad.com](https://gumroad.com) can use their accounts to download updates as they become available (the View content link in your receipt always accesses the latest available files). Updates are also available for the Kindle edition, but this depends on the willingness of Amazon to make purchasers aware that an update exists. We try to persuade Amazon of this need about once a month. Information about how to access updated files is in our [FAQ](#).

The companion volume for this book has other information that you might find interesting. We include the companion volume with the EPUB/PDF version and is available separately for the Kindle edition.

This is the original version of the 2023 edition published on **1 July 2022** (monthly update #85.5). You can find information about the changes made in each update in [our change log](#) or through our [Facebook page](#).

Tony Redmond took the photo used on the front cover in Dubrovnik, Croatia by in May 2022.

# Table of Contents

<b>Foreword</b> .....	<b>xi</b>
<b>Introduction</b> .....	<b>xiii</b>
What We Cover.....	xiii
PowerShell Examples.....	xiv
The Colored Boxes.....	xiv
Book Updates.....	xv
The Author Team .....	xv
Thanks .....	xvi
Our Sponsor.....	xvi
Comments and feedback.....	xvii
<b>Chapter 1: The Microsoft 365 Ecosystem and Office 365</b> .....	<b>1</b>
Microsoft 365: More than Office 365.....	2
Cloud Infrastructure.....	3
Products and Licenses.....	15
Trial Tenants.....	21
The Commercial Success of the Microsoft Cloud .....	21
Service Level Agreements for Online Services .....	24
Leveraging the Breadth of Office 365.....	27
<b>Chapter 2: Embracing the Microsoft 365 Cloud</b> .....	<b>29</b>
Should You Move to the Cloud?.....	29
What to Do Before You Move .....	30
What To Do While You Move .....	60
What To Do After You Move.....	62
The Cloud Is a Journey.....	63
<b>Chapter 3: Managing Identities</b> .....	<b>64</b>
The Role of Azure Active Directory.....	64
Identity Architectures .....	66
Standalone Identity.....	67
Hybrid Identity Authentication Infrastructure .....	69
Alternate Login ID.....	73
Passwordless Authentication.....	74
Understanding Azure AD Authentication.....	75
Customizing the Tenant Sign-In Page.....	79
Joining Computers to Azure AD .....	81

Guest Access to Azure AD .....	82
Protecting Your Identities.....	88
Controlling Access.....	97
App Registrations and Permissions .....	108
Connecting to LinkedIn .....	110
PowerShell and Azure AD.....	112
<b>Chapter 4: Tenant Management.....</b>	<b>113</b>
Cloud versus On-Premises Management.....	113
Administrative Interfaces .....	115
Managing Licenses, Plans, and Billing .....	123
Managing Integrated Apps.....	124
Managing Feature Releases.....	125
Managing Connectivity .....	127
Protecting Data with Encryption.....	130
Monitoring.....	132
Service Requests.....	140
Customizing the Microsoft 365 Interface.....	145
Reporting .....	147
Backing Up Office 365 .....	152
<b>Chapter 5: Managing User Accounts.....</b>	<b>159</b>
Managing User Accounts in the Microsoft 365 Admin Center.....	159
Enabling Self-Service User Management .....	160
Managing User License Assignments .....	161
Managing User Role Assignments .....	164
Managing Privileged Accounts.....	173
Managing Exchange Online Mailboxes.....	174
Administrator Access to User Mailbox Settings.....	186
Securing the Data of Ex-employees .....	187
<b>Chapter 6: Managing Exchange Online .....</b>	<b>196</b>
Exchange Online.....	196
Native Data Protection .....	198
Managing Mailboxes.....	200
Autodiscover.....	223
Recovering Deleted Mailboxes.....	223
Inactive Mailboxes.....	225
Automatic Mailbox Maintenance .....	231

Archive Mailboxes.....	237
Shared Mailboxes .....	247
Mail Contacts and Mail Users .....	259
Blocking Basic Authentication .....	261
<b>Chapter 7: Mail Flow.....</b>	<b>266</b>
Configuring Mail Flow.....	266
Managing Connectors .....	281
Mail Flow Rules.....	282
Remote Domains .....	290
Device and app mail relay to Exchange Online .....	291
Exchange Online Protection (EOP).....	292
Microsoft Defender for Office 365 (MDO).....	320
Attack Simulation Training.....	328
Investigations .....	333
Explorer .....	334
Monitoring and Troubleshooting Mail Flow.....	336
Transport Limits.....	346
<b>Chapter 8: Managing SharePoint and OneDrive for Business .....</b>	<b>349</b>
SharePoint Online.....	349
OneDrive for Business.....	387
Managing OneDrive for Business.....	396
Migrating to SharePoint Online and OneDrive for Business .....	398
Microsoft Lists.....	400
Delve and Microsoft Search.....	402
How Microsoft Viva and SharePoint Syntex Brings AI to SharePoint .....	408
Other Services Using SharePoint Online and OneDrive for Business .....	412
<b>Chapter 9: Tasks.....</b>	<b>414</b>
Delivering Lightweight Planning.....	414
Planner Basics.....	419
Creating New Plans.....	420
Plan Settings .....	425
Copying a Plan .....	427
Closing or Archiving a Plan.....	428
Planner and Microsoft 365 Compliance.....	429
Deleting Plans and Plan Data.....	429
Creating and Managing Tasks .....	430

Moving Planner Data to a Different Tenant .....	438
Using Planner Offline .....	438
Planner and Guest Users .....	439
Planner, Tasks, and Teams .....	440
Linking Planner and the Microsoft 365 Message Center .....	442
<b>Chapter 10: Managing Video .....</b>	<b>444</b>
The Cloud Video Platform .....	444
Stream User Functionality .....	446
Stream and Teams .....	453
Stream Audit Events .....	458
<b>Chapter 11: Managing Microsoft 365 Groups .....</b>	<b>460</b>
Modern Groups and Distribution Lists .....	460
An Identity and Membership Service .....	460
Group Components .....	464
Implementing Groups .....	468
Azure AD Policy for Groups .....	469
Creating New Microsoft 365 Groups .....	483
Managing Groups with Outlook Clients .....	490
Guest Access to Microsoft 365 Groups .....	492
Controlling Guest Access to Groups .....	502
Removing and Recovering Groups .....	508
Group Expiration Policy .....	512
Dynamic Microsoft 365 Groups .....	516
Groups and Compliance .....	519
Yammer and Groups .....	520
Evaluating Yammer, Groups, and Teams .....	524
Distribution Lists .....	527
Recipient Moderation .....	544
Migration from Email Distribution Lists .....	546
Comparing Groups, Distribution Lists, and Shared Mailboxes .....	550
<b>Chapter 12: Teams Basics .....</b>	<b>551</b>
Workgroup Collaboration .....	551
The Structure of Teams .....	562
Maintaining Team Membership .....	573
Teams Messaging .....	578
Personal (1:1) and Group Chats .....	592

Teams Meetings .....	600
Teams Calling .....	622
Viewing Organizational Information .....	623
Presence and Status.....	625
Files: Linking Teams and SharePoint Online .....	627
Teams for Frontline Workers.....	631
Can Teams Replace Email?.....	632
Teams and Email Interaction .....	636
<b>Chapter 13: Managing Teams.....</b>	<b>640</b>
Keeping Teams in Good Shape .....	640
Creating a Deployment Plan for Teams.....	640
Teams Management.....	641
Creating Teams.....	652
Dynamic Teams.....	660
Hiding Teams from Exchange Online .....	661
Using a Team-Enabled Group as a Distribution List .....	662
Deleting (and Restoring) Channels and Teams.....	663
Channel Moderation.....	664
Managing Settings for a Team .....	665
Guest Access for Teams.....	667
Email Integration for Teams Channels.....	672
Teams and Compliance .....	675
Auditing Teams.....	686
Teams and the Groups Expiration Policy.....	686
Archiving Teams .....	687
Reporting Teams Usage .....	689
Extending Teams.....	691
Teams App Setup Policies .....	695
Teams App Permission Policies.....	697
Teams App Store.....	698
Teams and Bots .....	698
Office Connectors and Teams.....	699
Teams Approvals.....	700
Debugging Teams Clients .....	701
Migration to Teams.....	703
<b>Chapter 14: Managing Teams Calling and Devices.....</b>	<b>704</b>

Teams Calling Fundamentals .....	704
Teams Meeting Enhancements .....	716
Teams Phone .....	720
Teams Devices.....	733
Troubleshooting and Monitoring Calls .....	745
<b>Chapter 15: Managing Clients .....</b>	<b>759</b>
Many Clients, One Service .....	759
Managing the Microsoft 365 Apps Suite.....	762
Using the Microsoft Apps Admin Center .....	771
Managing the Outlook Client Family .....	775
Managing Outlook Web App.....	780
Managing Client Access and Protocols.....	788
Managing OneDrive for Business Clients.....	795
Managing Teams Clients.....	796
Managing Microsoft Authenticator .....	805
Optimizing Microsoft 365 Client Network Access .....	808
<b>Chapter 16: Managing Devices.....</b>	<b>811</b>
Comparing the Three Solutions .....	811
Getting Started with Intune .....	812
Managing Apps .....	815
Managing Devices .....	823
Security by Compliance .....	825
Intune Management .....	827
<b>Chapter 17: Managing Data Governance and Compliance .....</b>	<b>831</b>
Data Governance.....	831
Principles of Data Governance .....	833
Compliance Permissions .....	835
Retention Policies and Publishing Label Policies .....	836
Rules or Principles of Retention .....	838
Retention Policies .....	839
Auto-Label Retention Policies.....	859
Retention and Sensitivity Labels .....	863
Creating New Retention Labels.....	866
Using Retention Labels.....	872
Records Management.....	879
Processing Manual Dispositions .....	882



Event-based Retention .....	886
Removing Retention Labels .....	887
Data Classification Dashboard .....	888
Ingesting Items for Review from External Sources .....	891
Using PowerShell with Retention Labels and Policies .....	892
Moving Data Between Tenants .....	901
Understanding the Exchange Mailbox Lifecycle .....	902
Exchange Mailbox Retention Policies .....	906
<b>Chapter 18: Managing eDiscovery .....</b>	<b>911</b>
Content Searches .....	911
Creating and Running a Content Search .....	914
Auditing of Search Activities .....	931
Data Subject Requests .....	932
Microsoft Purview eDiscovery .....	933
Premium eDiscovery .....	940
Using PowerShell with Content Searches .....	942
Using PowerShell to Manage eDiscovery Cases .....	950
In-place Holds and Litigation Holds .....	953
<b>Chapter 19: Managing Data Loss Prevention .....</b>	<b>955</b>
Leak Prevention with Software .....	955
Microsoft Purview Data Loss Prevention .....	956
Sensitive Information Types .....	957
Microsoft Purview DLP Policies .....	959
Endpoint DLP .....	971
Creating Custom Sensitive Information Types .....	972
Document Fingerprinting .....	976
<b>Chapter 20: Managing Information Protection .....</b>	<b>978</b>
The Need to Protect Data .....	978
Rights Management .....	979
Enabling Rights Management for a Tenant .....	983
Sensitivity Labels .....	985
Protecting SharePoint Online and OneDrive for Business .....	1013
Protecting Email .....	1018
Managing Office 365 Message Encryption .....	1027
Hybrid Protection .....	1036
Protecting Windows Files .....	1037

Managing Sensitivity Labels with PowerShell.....	1039
PowerShell for Rights Management.....	1048
Using Microsoft Defender for Cloud Apps to Protect Office 365 Content.....	1053
Microsoft Information Protection Auditing.....	1054
Cloud Exit for Encrypted Content.....	1055
<b>Chapter 21: Managing Auditing and Reporting .....</b>	<b>1056</b>
Auditing Framework.....	1056
Activity Alerts and Alert Policies.....	1076
Office 365 Cloud App Security.....	1082
Third-Party Auditing Alternatives.....	1087
Exchange Online Administrative Auditing.....	1087
Exchange Online Mailbox Auditing.....	1089
Reporting Workload Activity.....	1094
Communications Compliance.....	1101
Information Barriers.....	1110
<b>Chapter 22: Power Platform .....</b>	<b>1120</b>
In Search of No-Code/Low-Code Automation.....	1120
Power Platform Administration.....	1121
Power Automate.....	1126
Power Apps.....	1135
Power BI.....	1143
Power Pages.....	1146
Power Virtual Agents.....	1147
Teams and the Power Platform.....	1148
<b>Chapter 23: Managing Tenants with PowerShell and the Microsoft Graph .....</b>	<b>1151</b>
The Power of Automation.....	1151
Managing Exchange Online with PowerShell.....	1165
Managing Microsoft 365 Groups with PowerShell.....	1172
Managing Teams with PowerShell.....	1198
Sending Messages with PowerShell.....	1213
Understanding How to Use the Microsoft Graph to Manage Microsoft 365.....	1215
Approaching the Use of Graph API Requests with PowerShell.....	1222
The Microsoft Graph Explorer.....	1230
Using the Microsoft Graph PowerShell SDK.....	1233
Using Azure Automation with Microsoft 365.....	1254
Office Connectors.....	1255

Managing Self-Service Purchases .....	1261
<b>Chapter 24: What don't you know you don't know about Office 365 and Azure AD? .....</b>	<b>1262</b>
Azure Active Directory Recovery .....	1263
Tenant-to-Tenant Migration .....	1268
Practical Auditing in a Hybrid World .....	1275
Group Management .....	1278
About Quest.....	1281
<b>Appendix .....</b>	<b>1282</b>
Annualized Run Rate for the Microsoft Cloud .....	1282
Growth in Office 365 User Numbers .....	1282
Quarterly Performance Against SLA.....	1283

# Foreword



Many people talk about the cloud as though it is made of pixie dust. Toss a few clever phrases around, write a check, and life is as easy as sipping margaritas on the beach. If that sounds like a pile of manure, it's because it is.

The authors of this book write "the Cloud is a Journey." That's why I love this book. Tony Redmond and his fellow authors have been around the block enough times to know manure from masterpieces. They deliver a clear-eyed view of the issues you need to think through moving to the cloud and what to do once you are there. There will be plenty of time for margaritas on the beach, but first, you must use this book to apprentice yourself to masters and learn the art and science of being an Office 365 IT Pro.

Office 365 relieves you from solving hundreds if not thousands of extremely difficult problems. You don't have to plan, size, purchase, and deploy the right server, power, storage, and network systems. You don't have to provision the OS, storage, security, networking, monitoring, management, backup, and other software stacks. You don't have to run the infrastructure 24x7, patch everything on a timely basis, verify that everything continues to work after the patches, and update the systems to take advantage of the latest and greatest set of features. Office 365 does all that for you.

But if Office 365 does all that, do we need IT Pros anymore?

Having spent over four decades in the computer industry, I've lived through, and led through, several major technology transitions. At each of these transitions, some IT Pros fear that their jobs will go away. Jobs change but they rarely go away.

IT Pros used to walk around with an unfolded paper clip in their pocket to toggle dip-switches to configure CD-ROM readers to work on PCs. Then plug-n-play was invented. We got rid of a lot of unfolded paper clips but not many IT Pros. Most learned new skills and prospered.

As soon as one problem goes away, IT Pros move on to work on the next set of issues to help move the business forward. And they became more valuable in the process. There wasn't much business value generated in flipping dip switches.

So it is with Office 365. IT Pros no longer plan, deploy, and operate the infrastructure to run Office applications. That frees them up to focus on things like managing Data Governance and Compliance, eDiscovery, Information Protection, and Data Loss Prevention, all of which are covered in chapters in this book. It should be readily obvious that those topics are far more important and valuable to a business than monitoring disk usage on a server.

Technology transitions like the cloud are not easy. You spend your time becoming really good at something. Perhaps you are the company's go-to person for a certain technology. And then a technology transition comes along, and you must start all over. It's hard, but that is our way. When you choose a career as an IT Pro, you signed up for walking the path of lifelong learning. This book is an excellent companion for that journey.

In this book, you will find many PowerShell examples. PowerShell is a must-have skill for any serious IT Pro. GUIs are pretty and often easy, but they limit your ability to be a highly effective IT Pro. You can't cut and paste mouse clicks. You can't code review mouse clicks. You can't rerun a set of mouse clicks on five thousand user accounts. You can't share your mouse clicks with other people. With PowerShell, you can plan your work, get it reviewed by others, and then perfectly repeat it over and over again. Not with a GUI. With PowerShell, you are part of a community where members help each other. With a GUI, you are on your own.

Lastly, as information expands, expertise narrows. As paradigms shift, expertise expires. Therefore, IT Pros need to be generalists and exhibit a growth mindset. That includes the ability to learn, collaborate, detect, and apply patterns. PowerShell was designed with this in mind. It has a single parser for all cmdlets, so you don't have to relearn how to type a command. It provides a regular verb-noun syntax with the goal of enabling you to think about what you want, type it and get it. You can run interactive commands, simple scripts, or complex scripts. You can run them on Windows, Linux, macOS, and even from the Azure Portal. And it is fun.

Best wishes with your journey to the cloud and Office 365, and congratulations to Tony Redmond, Paul Robichaux, Christina Wheeler, Brian Desmond, Juan Carlos Gonzalez, Ben Lee, and Gareth Gudger for bringing clarity and a wealth of good advice to help you on your way.

Jeffrey Snover

Technical Fellow and Inventor of PowerShell, Microsoft Corporation

# Introduction

Welcome to the Ninth (2023) edition of **Office 365 for IT Pros**, a book focused on the Microsoft 365 cloud productivity suite. The goal of this book is to help tenant administrators, architects, and technologists understand and exploit the capabilities available in the enterprise version of Office 365 and the surrounding Microsoft 365 ecosystem. Office 365 had its 10<sup>th</sup> anniversary on June 28, 2021. We have seen enormous change since its release. It's been a blast tracking, analyzing, and documenting the evolution of Office 365 since we started on the Office 365 for IT Pros path in the summer of 2014.

Some accuse Microsoft of applying branding too liberally with the result that they confuse customers about just what Microsoft 365 means. In the context of this book, Office 365 means the enterprise Office services accessed by tenants via the internet, including Exchange Online, SharePoint Online, OneDrive for Business, Teams, Planner, and Yammer. We also cover the functionality available in the various administrative portals and how to use Power Automate, Power Apps, PowerShell, and the Microsoft Graph to automate administrative processes. Although we cover clients like Outlook, we don't cover applications like Word, Excel, and PowerPoint except in passing.

This is not an official Microsoft publication and Microsoft endorses none of the opinions expressed here in any way. Instead, it's a collection of thoughts, ideas, and perspectives from a team of highly experienced MVPs (members of the Microsoft Most Valuable Professional program).

## What We Cover

This book helps you maximize your use of Office 365. The major topics include:

- Introducing Office 365.
- Cloud Adoption, or what you should do to run an effective cloud onboarding project.
- Managing user identities with Azure Active Directory, including synchronization with on-premises directories and different authentication methods.
- Managing the overall Office 365 framework.
- Understanding the two basic workloads: Exchange Online and SharePoint Online (including OneDrive for Business), including how mail flow works to bring messages into and out of your tenant.
- Working with Microsoft 365 Groups, including how to manage the membership and access service enabled by Groups.
- Understanding the architecture of Teams and how this popular application leverages so many parts of the Microsoft 365 ecosystem.
- Understanding the Teams architecture and how to manage an application built on top of many other Microsoft 365 and Azure components.
- How the new version of Microsoft Stream works.
- Using Microsoft Planner to organize team and group projects effectively.
- Managing Office 365 desktop, mobile, and browser clients, with or without Microsoft Endpoint Manager.
- How data governance and compliance work inside Office 365, including how retention policies and retention labels work across Exchange Online, SharePoint Online, OneDrive for Business, and Teams.
- How Office 365 captures audit records for over 1,500 different events and how to use that audit data to answer questions about how people use (and sometimes abuse) the service.
- Using Power Automate and Power Apps to automate different user and administrative operations.
- Protecting Microsoft 365 content with data loss prevention policies and encryption (sensitivity labels).

- Using PowerShell and the Microsoft Graph APIs to automate common administrative processes within a tenant.

It's a lot of information to cover in a single book. We hope that you'll like what you find here and appreciate the effort we go through to publish the monthly updates to refresh the material and make sure it is as up-to-date as possible.

## Companion Book

Over the years, we have removed content (including entire chapters) that appeared in previous versions and moved it into a companion volume. The intention was to make space for new topics and content that we think is important to tenants today.

We don't often update the material in the companion volume and only plan to fix errors and omissions. We still think that the information in the companion volume is interesting and valuable, but its time in the limelight might have passed or the information is less relevant to current operations and administration.

## Cloud and On-Premises Products

This book is about Office 365. It follows that when we reference "Exchange" or "SharePoint," we mean the cloud version rather than the on-premises code. Where necessary, we are explicit. For instance, we say "Exchange Online" when we mean that a feature belongs to the cloud version. If we say "Exchange" or "SharePoint," the discussion applies to both cloud and on-premises versions. It's easier with applications like Teams and Planner because they don't exist on-premises (and never will).

## PowerShell Examples

Everyone involved in running an Office 365 tenant should have at least a passing acquaintance with PowerShell (and probably the Microsoft Graph APIs). We include many examples of using PowerShell across the book to help readers understand the value of being able to automate operations. Please remember that we write code to illustrate principles, not to create off-the-shelf solutions. Every organization has a certain way of doing things, its coding standards, and code libraries. We hope that our code helps you understand how PowerShell works with Office 365 so that you can take that learning and use it according to the standards laid down for your tenant. Before putting the code into production, you should add better error handling and probably rewrite it to move some code into functions or enable parameter passing, or simply make the code fit better with other scripts used in your tenant. We do not represent any of the code included in the book to be complete answers to a problem. Instead, you should treat our code as a starting point and be prepared to do some additional work to improve, smoothen, and bulletproof the code before using it to process real data.

Readers of previous editions know about our [GitHub repository](#) (see [this article](#) for information). Referencing scripts in GitHub makes the book a little shorter and allows text to flow better on the page. Another advantage is that the code in GitHub is more likely to work because you can download a complete script instead of transcribing code from a book. An increasing number of our scripts use Microsoft Graph API calls to get work done. This reflects the current state of technology, including the fact that some Microsoft 365 workloads don't support a PowerShell module, and we hope that you can use these examples to build solutions to meet administrative or operational needs.

## The Colored Boxes

From time to time we will want to draw your attention to something that we think is important. We use three colors to highlight information. The first type is a note.

**Note:** This is some additional information about a topic that we're discussing. We've included it because we think it adds some value.

We also have some warnings for things that you need to understand.

**Warning:** Warnings or other cautionary notes will appear like this. Try not to ignore these, the lessons were often learned the hard way and we'd hate to see you suffer the same pain.

And there are lots of real-world observations that we think will interest you.

**Microsoft 365 Groups:** Every group is represented as a group object in Azure Active Directory...

## Book Updates

Apart from its multi-platform nature, we also wanted to avoid the static nature that content often takes in traditional technical books. Given that Office 365 changes all the time, it didn't seem to make much sense to say that any text was definitive and that led us to the decision to create a book with content that changes to match what happens inside Office 365.

Due to the fast-changing nature of Office 365, some of the user interface elements illustrated in this book might have changed when you read this book. The same is true for details of how a feature works. Our tenants use the "Targeted Release" option to allow us to see new features before Microsoft releases them to the general Office 365 base. We may cover something here that you might not see yet in your tenant. Coping with a fast-changing (or even ever-changing) environment is just one of the challenges faced by tenant administrators. New features appear, options move around, and options function in a slightly different manner. Things are just very different from the somewhat staid situation that often occurs inside a typical on-premises infrastructure.

Microsoft's documentation has improved dramatically recently but still often fails to keep up to date with the rapid change within Office 365. We do our best to keep an eye on what is happening and what changes, but please forgive us if we overlook some detail that Microsoft recently revealed or updated. Think of this as an opportunity to demonstrate how good a detective you are in seeking the right answer based on the evidence presented in this book, on Microsoft's websites, and in the voluminous amount of text that you will find in blogs scattered around the web. Of course, blog authors are not seers either and their text begins to degrade as soon as it appears, so you must gather evidence and put it in context with what you see in Office 365 at the time when you're trying to solve a problem or get something to work as you believe it should. Welcome to the world of cloud software!

We release updated versions of this book monthly (see [our FAQ](#) for information about downloading updates). Our policy is to fix errors as soon as we find them, or a reader tells us about a problem. We will update (patch) the version that is currently online to give subscribers the opportunity of always being able to fetch the latest content from where they bought the book (Gumroad or Amazon). It's just like fixing bugs in a software program.

## The Author Team

The Office 365 for IT Pros writing team is:

- Tony Redmond.
- Paul Robichaux.
- Ben Lee.



Office 365 for IT Pros

- Brian Desmond.
- Juan Carlos González.
- Christina Wheeler.
- Gareth Gudger.

Our Technical Editor is Vasil Michev. For more information on the team, see [our online bios](#).

## Microsoft MVP Program

Tony, Juan Carlos, Paul, Gareth, Christina, and Vasil are proud members of Microsoft's Most Valuable Professional (MVP) program. See [this page](#) for more information about the MVP Program.



## Thanks

Many people at Microsoft helped us to chase down technical detail or explain how things work inside Office 365. The list of those who have helped is now too long to include here. We extend our thanks to everyone at Microsoft who has helped us chase down a technical issue or answered questions. We owe you a huge debt.

## Our Sponsor



It's hard to find the time to gather information, make sure that it's current, and write it up. An enormous amount of effort has gone into the creation of the original book and the many revisions and rewrites required for this edition. We could not undertake the task without the help and support of our sponsor, Quest Software. We are very grateful for the support extended by Quest and for the support given in the past by our previous sponsors.

Quest helps Office 365 tenants to migrate, manage, and secure Azure Active Directory, Exchange Online, OneDrive for Business, SharePoint Online and Teams. Quest delivers the most comprehensive set of Office 365 and hybrid management solutions, including solutions from [recently acquired Quadrotech](#), [Binary Tree](#) and [Metalogix](#). Quest solutions allow you to approach migrations with confidence, moving workloads to Office 365 with little to no end user disruption. Supported migration scenarios include:

- Active Directory consolidation and restructuring
- Office 365 tenant to tenant (Teams, SharePoint Online, OneDrive for Business, Exchange Online and Azure Active Directory).

## Office 365 for IT Pros

- Google Drive/Box/Dropbox/EFSS to OneDrive for Business, SharePoint Online, or Teams.
- On-premises SharePoint and Exchange to Office 365.

Quest technology enables Office 365 tenants to simplify administration, improve security posture and protect data across Office 365 and Azure Active Directory with solutions for:

- Backup and recovery.
- License management.
- Auditing.
- Group management.
- Reporting, and
- Governance.

For nearly 20 years, customers have used Quest software to migrate, manage and secure Microsoft platforms. Quest continues to set the bar for Microsoft Platform Management, with key products winning industry awards and gaining recognition year after year. Gartner recently named Quest as a Representative Vendor in its 2019 Market Guide for Cloud Office Migration Tools, which provides recommendations for migration scoping, third-party tool evaluation, and how to gain value from such tools post migration. Gartner also cited Quest as the ONLY vendor to deliver all 40 of the 40 essential features and functionalities expected in an Office 365 migration tool.

To learn more about Quest solutions for Office 365, visit <https://www.quest.com/solutions/office-365/>.

## Comments and feedback

Comments about the content as well as pointers to where little errors might have crept into the text are always welcome. Please send your comments, suggestions, and observations to [BookComments@Office365ITPros.com](mailto:BookComments@Office365ITPros.com) or post to our [Facebook page](#).

# Chapter 1: The Microsoft 365 Ecosystem and Office 365

***Tony Redmond***

Microsoft launched Office 365 in June 2011. Much later, it introduced Microsoft 365 by bundling Office 365, Windows 10 Enterprise, and the Enterprise and Mobility Suite. Microsoft then updated the Office 365 plans for small to medium businesses to use Microsoft 365 branding and changed the name of the Office Pro Plus desktop application suite to Microsoft 365 apps for enterprise. Although the marketing team might be happy to have applied Microsoft 365 to as many products and solutions as possible, it created some confusion in the market.

Over a decade later, over 350 million people use Office 365 in 249 countries worldwide. As the cloud Office service, it is the cornerstone of the Microsoft 365 ecosystem. Office 365 is a set of cloud services, referred to as *workloads*. A cloud service is a resource delivered over the Internet, including software resources (Software as a service, or SaaS) or infrastructure (Infrastructure as a Service, or IaaS). Exchange Online and SharePoint Online, the two base Office 365 workloads, are examples of cloud-based services. Some services have roots in on-premises products such as Exchange Server and SharePoint Server. The codebases for these services are very different from their on-premises counterparts, notably because of the engineering effort to run at a massive scale. Other services like Teams and Planner are cloud-only and have no allegiance to any on-premises technology. Indeed, the nature of these services and how they consume other cloud services and microservices (services built for a specific purpose, like Azure Key Vault) means that they will never appear in an on-premises variant.

The basic idea behind cloud services is that customers can transfer the responsibility for running workloads to a cloud provider, who then charges a fixed monthly or annual fee based on some unit of work, such as a user account or mailbox. Customers either start from scratch with a cloud provider or migrate their data across the internet to data centers run by the cloud provider where enough computing, network, and operational capacity exists to handle the work generated by hundreds of thousands of companies. The value proposition is that the massive economy of scale created by cloud providers allows them to deliver the same or better functionality at a lower cost than is possible for on-premises IT to deliver. As more organizations move to the cloud, the economics improve even more (even with Microsoft's first price increase in ten years effective March 1, 2022), which then lets providers enhance and grow their services.

Microsoft 365 is not just one application or workload. Rather, it is a complex ecosystem composed of multiple moving parts connected via components like the Microsoft 365 substrate, common services like Search, Artificial Intelligence, and Machine Learning, and common programming interfaces in the Microsoft Graph APIs. The ecosystem did not develop overnight and is rooted in nearly twenty years of engineering effort. Office 365 is a core part of the Microsoft 365 ecosystem.

It's hard to keep any book about technology completely up to date. We do not attempt to claim that the coverage here is perfect because Microsoft 365 continues to evolve to keep its cloud services "evergreen." A cloud service that stays static is less attractive to its users than one where changes and updates appear all the time. To set some achievable boundaries, this book covers Office 365 in-depth and takes in Microsoft 365 as required.

# Microsoft 365: More than Office 365

Microsoft calls Microsoft 365 “the world’s productivity cloud,” saying that it represents their vision for the future of productivity tools spanning an integrated set of apps and services. The upshot of this marketing activity is that Microsoft liberally applies the Microsoft 365 moniker to a wide range of products offered to consumers and the enterprise. To be clear, this book covers Microsoft 365 as it affects Office 365 for enterprise customers and does not deal with packages offered to consumers or small businesses. Table 1-1 lists the Microsoft 365 packages for the enterprise.

<b>Variant</b>	<b>Target market</b>	<b>Components</b>
<a href="#">Microsoft 365 Enterprise</a>	Companies with more than 300 users.	Windows 11 Enterprise Office 365 (E3 or E5) EMS
<a href="#">Microsoft 365 Business</a>	Small to medium companies (up to 300 users).	Windows 11 Business Microsoft 365 Business Standard EMS
<a href="#">Microsoft 365 Frontline</a>	Customer service and support workers.	Windows 11 Enterprise Office 365 F3 EMS
Microsoft 365 Education	Educational establishments.	Same offerings as for Enterprise, Business, and Frontline
Microsoft 365 Non-Profit	Non-profit organizations.	Same as Microsoft 365 Business
<a href="#">Microsoft 365 Government</a>	U.S. government and state agencies.	Same as Microsoft 365 Enterprise (E3 and E5 bundles)

Table 1-1: Microsoft 365 products

Although the enterprise Microsoft 365 plans include Office 365, customers can still buy Office 365 separately. No difference exists in the functionality available in applications like Exchange Online, SharePoint Online, and Teams in the Microsoft 365 plans over what is available in the Office 365 plans. The differences exist in areas like device and identity management and advanced compliance and data governance features.

At the launch of Microsoft 365, the offering was simply a bundle of several software packages. This isn’t a bad thing, because if Office 365 proves anything, it shows how more functional technology can be when multiple applications are available. Teams, for instance, cannot work if Exchange Online, SharePoint Online, Groups, and OneDrive for Business are unavailable. However, Microsoft has bundled an increasing amount of functionality in the Microsoft 365 plans, especially in the areas of compliance, data governance, and identity management. Two useful tools are available to help understand what’s available, and the costs involved:

- [PDFs with overviews](#) of the Office 365 and Microsoft 365 plans.
- A [Microsoft comparison of the features available to each plan](#), including the costs of subscriptions and add-ons.

Over time, Microsoft introduced components that work across the suite instead of being restricted to individual workloads. Good examples of where this has happened are the management consoles like the Microsoft 365 admin center and the Microsoft Purview Compliance portal. Looking back, the progress to integrate applications enabled Office 365 to progress from being a loose collection of barely cloudified on-premises applications to evolving into an integrated ecosystem. That progression took the best part of six years; it’s happening more rapidly in Microsoft 365.

Evidence exists that an increasing percentage of the user base, particularly in medium to large enterprises (over one thousand seats), see their most cost-efficient licensing arrangement as one built around Microsoft

365. In April 2022, Microsoft said that 45% of Office 365 seats are part of Microsoft 365 plans. Many of these customers buy Microsoft 365 because they want to use Enterprise Mobility and Security and advanced Azure AD features. The influence of Microsoft is seen in the steady growth of Enterprise Mobility and Security seats to 218 million (April 2022), many of whom also use Office 365.

## Cloud Infrastructure

Attention to detail and absolute adherence to well-defined procedures are the hallmarks of how Microsoft runs its cloud infrastructure. Without attention to detail enforced through automated processes and procedures, it would be impossible to manage a service for millions of tenants ranging from small to the very largest companies.

An indication of the scale of Microsoft's cloud businesses is its assertion (June 2019) that it serves [a billion users and twenty million businesses](#). The investment needed to buy land and build data centers, install computing resources, power, and cooling, and automate the management and security of applications and data at the scale of Microsoft 365 is massive. The ability to deliver services to customers at a competitive price is only possible for companies with very deep pockets. Microsoft is one of those companies. Google and Amazon are other examples. Let's look at some aspects of Microsoft's cloud infrastructure.

### Data centers and Regions

Microsoft 365 tenants share a single large logical infrastructure composed of hundreds of thousands of servers spread across multiple Microsoft data centers. Figure 1-1 shows the Office 365 data center deployment in April 2022. By its very nature, this picture is incomplete as it does not show some of the infrastructure that is under development. In addition, it does not convey the deep investment made to create "edge" network termination points set up by Microsoft to bring user traffic quickly into its data centers from all around the world or the internal network that transports tenant data between data centers. In total, the [Microsoft Cloud](#) (Microsoft 365 and Azure) spans well over 200 data centers. Not every Microsoft Cloud data center delivers Office 365 services to customers, but the number of local data centers delivering Office 365 is growing over time.

Microsoft organizes Office 365 data centers into regions. The data center region selected to host the data for new tenants is based on the country (location) selected by the tenant. Since launching Office 365, Microsoft has gradually built out the data center infrastructure. In many cases, new data centers are launched to keep data local ("[in-geo data residency](#)") to accommodate customer choice and satisfy local regulations. Where large regional data centers such as the European Union used to be the focus for Office 365 service delivery, localized service is now available in individual countries like France, Germany, Switzerland, and the U.K. The same is true in Asia-Pacific, where services come from data centers in countries like Japan, South Korea, Singapore, and Australia. Microsoft calls the country-level regions "Go Local," as in "Go Local Japan" or "Go Local Australia." This is an internal Microsoft notation that reflects the purpose of country-level regions.

One big advantage gained by segmenting tenant workload across multiple data center regions is that a fault that affects users in one region seldom spreads to other regions. Known single points of failure that can affect multiple regions do exist, but in general, outages don't affect more than a single workload running in a single region, such as a failure affecting mailboxes running in a single Exchange Online Forest or sites served by a single SharePoint Online farm.

Apart from the ability to serve large customer populations, natural and economic advantages such as ambient temperature (to reduce the need for cooling) or availability of cheap hydroelectric power influence data center placement. Security is of prime concern and Microsoft pays great attention to the physical security of the

buildings (you will not find large signs proclaiming Microsoft's ownership anywhere) as well as cyber-security for the data contained within the buildings.

Because Microsoft's data center infrastructure is constantly growing and expanding, the live location for tenant data also changes. In addition, Microsoft moves workload to rebalance servers across data centers (within the same region) and to make more effective use of available resources. Even though the underlying infrastructure is changing all the time, users can continue to work and access their information from anywhere around the world.



Figure 1-1: Office 365 data center map (source: Microsoft)

## Co-Location of Consumer and Business Services

In addition to its enterprise cloud applications, Microsoft hosts over 400 million Outlook.com users of the consumer email service in the same data centers. Outlook.com mailboxes use Exchange Online with the sole difference between the two services being the feature set exposed to users. Outlook.com mailboxes run on the same server, storage, and network infrastructure as Exchange Online to take advantage of features like Native Data Protection and Exchange Online Protection and the same clients are available for both services. Although the same engineering teams support the two services, the functionality in Outlook.com is much less comprehensive than that available to even the entry-level Exchange Online plan. However, the two services share some features (like the method to connect the Outlook mobile clients) and Microsoft introduces some functionality into one service before they decide to do the same for the other. For example, "Sweep" rules first appeared in Outlook.com and are now available in OWA, while the calendar and calendar sharing features now available in Outlook.com originated in Exchange Online. Taken together, the infrastructure shared by Exchange Online and Outlook.com delivers email service to billions of mailboxes (not all of the mailboxes belong to humans). Microsoft uses the same shared infrastructure approach with OneDrive for Business and the consumer version of OneDrive.

## Number of Servers

As you'd expect, the infrastructure is also massive if looked at in terms of the number of individual servers running different workloads. The last public figure for Exchange Online is 275,000 mailbox servers (September 2020) while that for SharePoint Online is 180,000 servers (May 2019). Add in the servers used for Teams, Planner, Yammer, Azure AD, other Azure services, and supporting services, the number deployed to run Office 365 workloads comfortably surpasses one million physical servers.

Even with so many resources, demand flowing from an event can exert enormous pressure on an online service when user numbers grow dramatically over a short period. During the Covid-19 pandemic, tens of millions of new users were onboarded. In the case of Teams, the number of daily active users increased from 20 million in November 2019 to 75 million in April 2020, creating a surge in user demand that also impacted SharePoint Online, OneDrive for Business, Exchange Online, and other services. To maintain responsiveness for end-users, Microsoft trimmed the functionality of different applications. Background processes didn't run as often and Microsoft adjusted some features to reduce demand, such as Stream dropping the resolution of Teams meeting recordings from 1080p to 720p. Although some interruptions happened, generally Office 365 remained online and handled the swelling load. After Microsoft had the chance to commission new server and storage resources in the data centers, they reversed the feature adjustments.

## Current Data Center Regions

Microsoft publishes information about the [locations of Microsoft 365 data centers](#). The same page includes details of where customer data resides on a country-by-country basis for all the offered services. All data centers deliver core services (Exchange Online, SharePoint Online, and OneDrive for Business). Other services, like Teams and Project Online, might also run in the same data centers. In some cases, other data centers within the region deliver specific services, as in North America where the Planner service runs from data centers in California and Virginia.

The point of delivery for services changes over time and if you are concerned about data sovereignty, you should check with Microsoft to understand exactly where your data are for all applicable applications. Azure AD is another service that can come from another region. For instance, the U.K. region uses Azure AD from the EMEA and U.S. data centers.

## New Data Centers

Once a new data center comes online, a sophisticated migration process moves tenants from other data centers to the new location. The same is true when Microsoft creates a new region. For instance, after the United Kingdom data centers came online, some tenants asked to move their work to those data centers to keep their data remained "in country;" the same happened in France or when Australian and New Zealand tenants moved to the Australian data centers. This work happens behind the scenes (just like regular mailbox moves) so that the eventual switchover is fast and painless. Microsoft has [a documented process](#) to help tenants with specific data residency requirements to request a move for their core data to a new region after it comes online.

The creation of a new data center region normally means that the data center first delivers the base workloads to tenants. It can take some time before the full range of service capabilities is available, including applications (like Teams or Planner) and utilities. Microsoft is currently building out new data center regions in Italy, Poland, and New Zealand.

## Workloads Running Within Data Center Regions

Microsoft distributes work across all data centers within a region to protect data against failure. For instance, the active-active design for Exchange Online Database Availability Groups (DAGs) means that mailbox database copies exist in at least two data centers within a region. In addition, as Microsoft adds data centers to a region, the opportunity exists to spread database copies to those data centers. For instance, new DAGs built for use by Exchange Online in the European Union region might include databases spread across the Amsterdam, Dublin, Helsinki, and Vienna data centers. Spreading data across so many data centers reduces the risk that any individual outage will affect a sizeable number of users. It is a deployment strategy that the

average on-premises administrator could never contemplate because of the investment needed to build out the underlying data centers and network.

Another way to understand what data center region supports data for your tenant is to access the Organization Profile (in the Settings section of the Microsoft 365 admin center). Go to the Data Location section and you will see the region for some, but not all, of the workloads used by the tenant (Figure 1-2). You can confirm these locations against Microsoft's [list of where customer data is stored](#). Microsoft has committed to delivering what it calls the [EU data boundary](#) by the end of 2022, meaning that all services and data used by tenants based in the European Union will remain within the European Union. Customers in Norway and Switzerland will also be able to access the EU data boundary if they wish.

## Data location

As part of our transparency principles, we publish the location where Microsoft stores your customer content. For more information about Microsoft's contractual commitments, see the [Online Services Terms](#).

[Learn more at the Office 365 Trust Center](#)

Service	Data at Rest
 Exchange	European Union
 SharePoint	European Union
 Skype for Business	European Union
 Microsoft Teams	European Union

For applications which you are not subscribed to, please see [Where is my data](#).

Figure 1-2: Data locations for a tenant

Although Azure AD holds most of the information for tenant accounts and configurations in the same data center region as a tenant's data, an exception exists in that Microsoft stores five user-related attributes including the User Principal Name and password hash for tenant accounts in the U.S. This is to make sure that authentication can happen as quickly as possible, no matter where in the world a user is located. For more information on this topic, see Microsoft's [support article on the situation](#) for European customers.

## Sovereign Clouds

A sovereign cloud is a data center region that exists because Microsoft must meet specific requirements imposed by the target customer base or geography. Three sovereign clouds are currently active. The U.S. government cloud (GCC) is a set of special data center regions created to meet the specific security needs of different levels of U.S. government and state authorities; the other is in China. The commercial data center regions, or clouds, are general-purpose in that these regions serve customers from any country and any industry in the coverage area.

## Multi-Geo Microsoft 365

Organizations can choose to distribute Exchange Online, SharePoint Online, OneDrive for Business, Teams, and Microsoft 365 Groups data across different data center regions (multi-geo). This means that the organization has a home region (known as the "central geo") where most of its work runs and one or more satellite geos. Sites including team, communication, and hub sites, created by distributed users are in their



assigned geo-location (otherwise known as the PDL, or “preferred data location”, an attribute of their Azure AD account). Microsoft says that they are exploring how to add multi-geo capabilities to other applications.

The normal use case for multi-geo is an international company that needs to satisfy data sovereignty requirements. For example, a U.S. company might have subsidiaries in France and the UK. Due to the sensitive nature of the work done in EMEA, the company does not want to store the data in the U.S., as would be normal. With multi-geo, they can choose to have user data for the supported workloads stored in local data centers.

Behind the scenes, once a tenant is enabled for multi-geo, Microsoft transfers the data belonging to the selected users to the satellite geos. During the transfer process, users continue to work as normal until the transfer is complete, at which point they switch over. Cross-region synchronization within the tenant’s Azure AD instance ensures that all the users within the tenant see a single worldwide picture and can continue to share work with colleagues without hindrance. The only thing that changes is the location of user data. OneDrive for Business is an exception because the transfer process does not move data belonging to an existing OneDrive account. Instead, administrators must run the *Start-SPOUserAndContentMove* cmdlet to [transfer user data to the new location](#).

Distributing user data across multiple regions causes some technical challenges, one of which is eDiscovery. Organizations need to consider if they wish to run eDiscovery locally or at a global level. See [this page](#) for more information.

Multi-geo capabilities are not a solution to poor network performance. Some people assume that this is the case because “the data is closer to the users.” This is a fallacy because, in most cases, poor network performance is due to issues such as lack of bandwidth or other problems in the link connecting users to Microsoft or poor routing inside the tenant’s internal network. Once inside Microsoft’s data center network, traffic flows from region to region very quickly and users do not see a difference when they move to a local region.

Multi-geo is available for most [general-purpose Microsoft 365 data center regions](#). It is not available for the sovereign clouds. To qualify, organizations must have an enterprise agreement with Microsoft and purchase multi-geo licenses for at least 5% of their accounts (previously, the organization had to license at least 250 accounts). Additional monthly fees of \$2 per user are payable for each user licensed for multi-geo capabilities. The fees cover the connection of the base workloads to different data center regions. Because of the extra cost, the most likely customers for multi-geo are multinational companies with complex data sovereignty needs and relatively large numbers of users. For more information, see the [multi-geo home page](#).

## Relationship with Azure

Microsoft 365 has a strong relationship with Azure, the other major component in Microsoft’s cloud strategy.

- The Office 365 data centers are co-located with Azure data centers.
- Office 365 uses Azure Active Directory as its authentication and identity service.
- Office 365 applications use many Azure services such as Azure Key Vault.

Office 365 applications use Azure Storage heavily and SharePoint Online is its largest single application consumer (using Azure SQL). Teams uses Azure Cosmos DB for its message and media stores. Stream is in the process of moving from Azure storage to SharePoint Online. At this point, the only major Office 365 storage system not running in Azure is Exchange Online mailbox storage, which continues to use physical mailbox servers. According to Microsoft, plans are in place to move Exchange Online to Azure storage.

## Microsoft 365 Substrate

It is an undeniable fact of IT life that it is much easier to process data in a single repository than it is to process data drawn from several. The Microsoft 365 substrate is the single repository for user documents, emails, meetings, tasks, groups, chats, and other data. The Microsoft 365 substrate holds data generated by applications like Exchange Online, SharePoint Online, OneDrive for Business, Yammer, and Teams, either by an application that uses the substrate as its primary repository or as a "digital twin" of data stored elsewhere, which is often a much smaller version holding just the information needed for search and compliance, together with a pointer to the original item.

Exchange Online mailbox databases are the physical implementation of the substrate. Exchange Online mailboxes have stored non-email application data for years. Now, the substrate creates the digital twin items from applications like SharePoint Online, Teams, and Planner and stores the data in hidden mailbox folders. It might seem surprising that Exchange Online databases serve in this role, but the Exchange database engine (ESE) has always been good at managing many different types of data, can scale up efficiently, and runs on low-cost storage. In addition, the Exchange Online Native Data Protection model makes sure that copies of application data exist in four databases, including a lagged database, to ensure robust data availability. ESE also supports "sharding," the ability to connect different chunks of data into a logical whole (for instance, 50 GB mailboxes are combined to form expandable archives), which makes it easy to present data on a per-user or per-tenant basis.

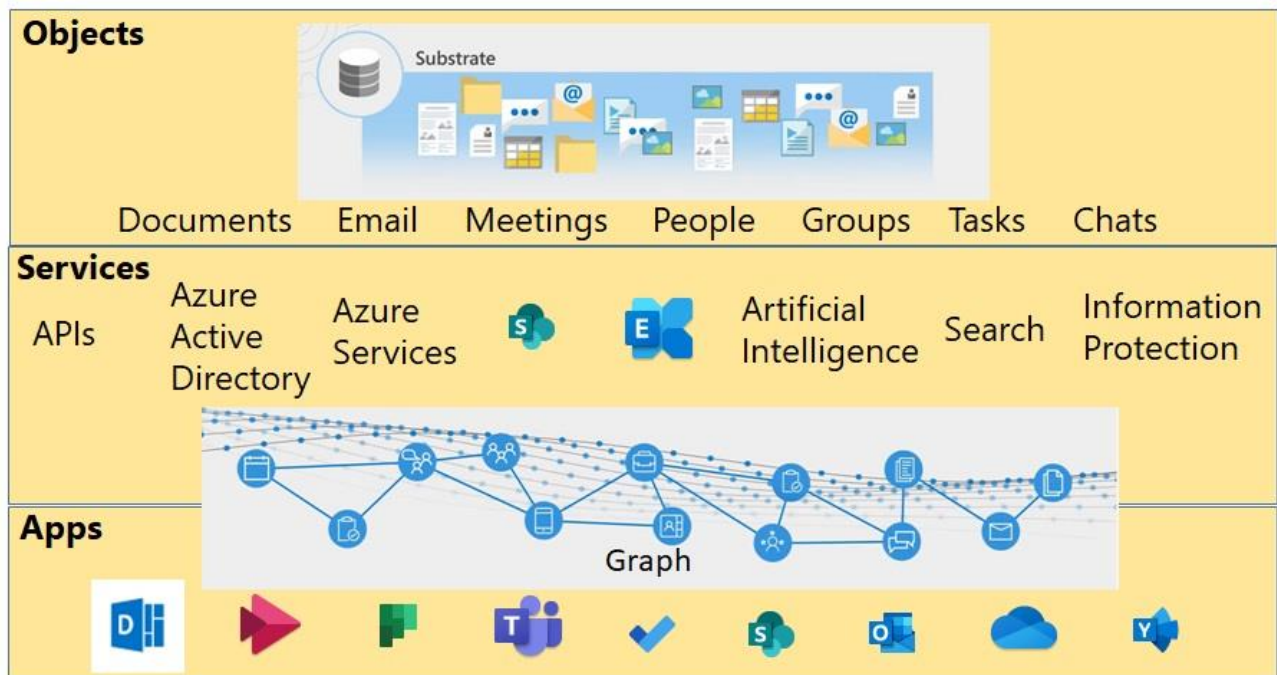


Figure 1-3: The Microsoft 365 substrate

Figure 1-3 shows how the substrate underpins Microsoft 365 with a common data repository for services to process and apps to present to users. SharePoint Online and Exchange Online appear as services because these base workloads deliver services such as email, storage, and document management to other apps in the service.

The existence of the substrate does not mean that applications will migrate to Exchange Online storage. Teams will continue to use Cosmos DB and SharePoint Online will continue to use Azure SQL. What it does mean is that copying data from applications into the substrate gives common services like Microsoft Search, cognitive services, and Artificial Intelligence an integrated platform to work against.

The substrate is the essential foundation for what Microsoft 365 is today. Without the substrate, Microsoft 365 apps couldn't get to data as easily as they can. or integrate with other Microsoft 365 apps or with other services as easily as they now can. It would also be harder for apps to incorporate new technologies like Microsoft Loop components. The substrate also facilitates "graph traversal," the ability to follow information from one graph to another, such as the way that the people card can connect to the LinkedIn graph to retrieve details of a contact's career. And without the substrate, cloud applications would still be like the collection of loosely cloudified on-premises products which formed the original service.

## Tenant Domains

Customers that use Microsoft 365 are tenants of the service. Each tenant occupies a subdomain within the overall Microsoft 365 infrastructure called its service domain. You can think of a tenant as the container for the company within Microsoft 365 and the domain is a sub-domain under the onmicrosoft.com root. For example, the Contoso company might run a tenant domain called contoso.onmicrosoft.com. Each user account has a separate user principal name to sign in (or connect) to applications. For instance, the user principal name *TRedmond@office365itpros.onmicrosoft.com* is an account belonging to the office365itpros tenant.

Tenants often have one or more domains registered in the Internet Domain Name Service (DNS) and want to continue using that domain with Microsoft 365. These domains, sometimes called vanity domains, might be part of the company branding, which is why they must remain in use. To enable this to happen, an organization can register their domains with Microsoft 365 and associate the domains with their tenant. When this happens, user accounts can have user principal names and email addresses belonging to the vanity domains, meaning that we can replace the somewhat clunky service domain addresses like *TRedmond@office365itpros.onmicrosoft.com* with the more elegant and brand-friendly *Tony.Redmond@office365itpros.com*.

Microsoft 365 does not care what name a company gives to its tenant, but you should because you can't change the name of a tenant after its creation. Any change required by corporate restructurings such as a merger, acquisition, or divestiture usually ends up in a tenant-to-tenant migration, an operation that is often expensive and long drawn out. Some flexibility in outward presence is possible by using a new vanity domain for external-facing email addresses and configuring Exchange to handle messages sent to those addresses. However, other applications can surface tenant names in different places (it's now possible to [rename the domain for SharePoint Online](#)), all of which means that some care and attention is necessary to ensure that the tenant name you use is the right one for your company.

## Directories and Identities

Microsoft 365 supports both cloud-pure and hybrid (cloud/on-premises) environments. A key aspect of the support is the ability to reliably authenticate user accounts with Azure AD and on-premises Active Directory. To enable this to happen, a variety of directories and other tools provision, store, maintain, and synchronize identities.

Azure Active Directory is the cornerstone of Microsoft's cloud identity story and acts as the source of authority for cloud user accounts. In hybrid deployments, where the on-premises Active Directory is always the master directory, Azure AD synchronizes with Active Directory to form a seamless view of user accounts and configuration data drawn from both environments. The tool used for this purpose is Azure Active Directory Connect (AAD Connect).

Many of the applications running inside Microsoft 365 need to store information specific to their operation. For instance, Exchange Online uses EXODS, its directory store, to hold information about public folders that

are not mail-enabled. These objects exist only inside Exchange Online and are irrelevant to the other workloads, so there is no reason to store information about them in Azure AD. The same logic applies to workload-specific objects and configuration data used by SharePoint Online (SPODS), Yammer, and Teams. The Identities chapter presents a comprehensive discussion about directory services within Microsoft 365 and the different forms of identities (on-premises, cloud, and hybrid).

## Automation

The Microsoft 365 infrastructure uses a sophisticated workflow engine called "Central Admin" (CA) that is capable of handling more than a hundred million workflow tasks per month. The idea is to automate the common tasks needed to keep services running as much as possible to remove the possibility that human error will compromise systems. The need to avoid human mistakes in the management of cloud systems is seen in the [DNS error that interrupted service in September 2011](#) and the [command typo that knocked out part of the Amazon Web Services](#) infrastructure in March 2017. A smoothly functioning workflow engine also achieves a reliable and robust throughput of actions across the system. Developers create CA tasks as scripted workflows in either C# or PowerShell. CA is responsible for the execution of scheduled tasks to perform actions such as server deployment, database rebalancing within a DAG, and so on. More complex tasks such as the addition of new capacity to the service still require some human intelligence and planning, but the application of a structured model and great attention to detail has enabled Microsoft to reduce the time necessary to complete even very complex tasks down from weeks to days.

Microsoft standardizes server configurations whenever possible. This does not mean that each server uses the same components, as this would be impossible in an industry where components change often. However, it does mean that a server will have the same general characteristics (CPU, disk, memory) and that software is installed in the same way on all servers of a specific type. Low-cost components such as JBOD arrays allow Microsoft to increase the storage available to tenants while still being market competitive. It also means that servers are built from modules to eliminate cabling. Everything is optimized for mass production and servers are integrated into racks at the factory and shipped to the data center ready to be plugged in and brought online.

Exchange is a good example of an application where sustained engineering investment has delivered huge performance improvements and made cloud economics possible. Exchange Online uses JBOD SATA drives to deliver cost-effective storage. Using these disks implies a risk of a higher failure rate than the more expensive "enterprise-class" drives often found in corporate data centers and indeed, across Microsoft's data centers, hard disk failures are the most common event in the tens of thousands of hardware events that are handled monthly. The low cost of storage and Exchange's performance profile makes it possible for Microsoft to offer enterprise users a 100 GB mailbox quota and auto-expanding archives, and to hold a mass of non-user data (such as information about Files usage) in mailboxes. Without low-cost storage, the monthly subscription to a Microsoft 365 enterprise plan would be much dearer.

Software components help to insulate users from the effect of hardware failures. For instance, Exchange's Active Manager will failover a database to a new server quickly if a disk problem is detected. It will also create a new copy of the failed database using the autoreseed feature if replacement disks are available. Across the entire fabric, a CA workflow called "RepairBox" constantly checks for hardware failures and will open support tickets automatically if an issue is detected like a failed disk. A technician can then replace the failed disk (no attempt is made to fix the disk). The same workflow monitors servers for inconsistencies in their state to detect and fix problems with configurations.

Even with such a sophisticated and smooth-running automatic support infrastructure, some problems still occur. For instance, "stragglers" are servers that run out-of-date software versions that might deliver an

inconsistent service to users. The server infrastructure is in a constant state of server refresh to introduce new software builds and new features. As such, with so many servers and so many updates, some server updates do not happen as well as they should, which is the usual reason why a straggler exists.

## Networks

Given that Microsoft 365 is a set of cloud services, it should come as no surprise that the network is a precious resource. Without enough high-quality bandwidth, users will be unable to connect to services, migrations cannot transfer data from on-premises servers, and hybrid connectivity will not work. Microsoft does not control the backbone used by tenants to connect their internal networks to its services as the links making up the backbone are managed by a large set of Internet Service Providers (ISPs) around the world. Although the Internet was originally designed to survive a nuclear holocaust, local failures caused by cable problems, ISP data center issues, and hardware failure can all prevent access to services.

Microsoft cannot control the Internet, but it can control traffic flow within the network that connects its data centers. This is a dedicated and tightly controlled and monitored network. Dark fiber optical connections link the data centers to ensure maximum data flow across the network. Automatic redundancy is deployed so that a temporary outage is contained and automatically addressed. Everything that can be done to ensure that the service is maintained is done, but even so, like all cloud services, the SLA for Office 365 can only be guaranteed at the boundary of the cloud provider's data centers as defined by the edge servers that handle inbound and outbound traffic.

## Monitoring, Telemetry, and Intelligence

The sheer size of Microsoft 365 and the telemetry gathered through user interaction make a colossal set of signals available to Microsoft. For example, in late 2020, Exchange Online processed more than 300 billion messages monthly. Increasingly, Microsoft applies artificial intelligence and machine learning techniques to analyze and make sense of the data. The output is used to guide choices in software engineering investment to create new functionality. For an application like Exchange Online that's been around for a long time and is very mature, the telemetry offers insight into issues that might have lingered without ever being addressed simply because they are not a huge problem at the scale of on-premises servers but are when viewed through the lens of the cloud.

Microsoft uses a "Data Insights Engine" to process the billions of events generated hourly, aggregating and analyzing events to understand how the overall service is running and to detect problems with individual components. The general approach is that if a problem surfaces in many entities or through different signals, then it must be true. By depending on signals from multiple resources you can get close to 100% fidelity when it comes to the automatic detection of problems, or "Red Alerts" as they are known. In addition, by analyzing signals from diverse sources, engineers can focus on where the root cause of the problem is likely to be with a high degree of accuracy and this, in turn, allows the launching of automatic recovery actions with a high degree of confidence that they will fix the problem. Taking a data-driven and analytic approach to the detection and resolution of problems is key to being able to run at scale.

In addition to its signal processing engine, Microsoft uses some much simpler techniques to know when something might be going wrong. For instance, if a spike in page views occurs for the Service Health Dashboard, it might be due to customers checking the dashboard to know whether a problem exists with the service. Such a spike can often be correlated with an output from the signal processing engine but sometimes it leads to a discovery of a problem identified by human beings. Microsoft can record the characteristics of that problem as a recognizable scenario for future automatic identification and resolution. It is also important to say that signals used to identify issues are both active (those generated by specific events) and passive

(those generated by servers and users during normal work). Microsoft uses a process of triangulation to spot abnormalities between the two sets that can point to a developing problem. As in all knowledge, learning how to measure the pulse of cloud operations and figure out what is normal and what's not is an evolving art that requires great dedication and ongoing observation by both automated systems and humans.

## A Constant State of Change

The size and technical complexity of the Microsoft 365 ecosystem create a unique challenge for anyone interacting with the technology. Even more challenging is the way that the ecosystem flexes and changes as different Microsoft engineering groups make software updates available to customers. At Microsoft's FY21 Q3 results briefing in April 2021, [CEO Satya Nadella said](#) that Teams added: "over 300 features over the past year, including more than 100 new capabilities so far in 2021." Adding Exchange Online, SharePoint Online, OneDrive for Business, Yammer, Planner, Azure AD, and solutions like Microsoft Purview Information Protection and Data Lifecycle Management to the mix means that a customer might have to cope with over 500 changes annually.

Inside Microsoft 365, applications introduce new features on an ongoing and constant basis. The changes range from tweaks to a client interface to the introduction of a completely new feature that changes the behavior of an application. This is a major difference between the traditional on-premises model where new software releases often appear on an annual release cycle that customers can then factor into carefully-planned "change windows". Becoming accustomed to the pace of change in the cloud can be quite a challenge for those used to the older way of deploying software updates but it is the method employed by most major cloud services.

The best way for administrators to know what Microsoft is working on is to keep an eye on the online [Microsoft 365 Roadmap](#) (Figure 1-4), which documents planned features coming across Microsoft 365. Although the roadmap sometimes misses an update and you always must keep your eyes peeled on what appears in the service to detect some new functionality, Microsoft refreshes the roadmap regularly and its contents are comprehensive enough to allow tenants to plan for new functionality.

Microsoft organizes the Microsoft 365 roadmap into the following sections:

- **Launched:** Features that Microsoft has deployed to all applicable customers.
- **Rolling Out:** Features that Microsoft is now deploying. As you can imagine with such a large and distributed service, it can take some weeks to deploy new software to every server running in every data center around the world. Microsoft usually posts to [the Microsoft 365 blog](#) or [individual product blogs in the Microsoft Technical Community](#) to inform customers about new features when they begin the roll-out process. The appearance of a post is no guarantee that a new feature will show up in a specific tenant anytime soon as this depends on whether the new feature belongs to the set available in the plans the tenant uses, the length of time that the feature spends in Targeted Release (previously First Release) status, and the time taken for Microsoft to deploy the feature to all applicable tenants after it reached standard release status.
- **In Development:** Features that Microsoft has announced are under development. Microsoft does not commit to delivering any of the features listed here as reasons might emerge to change or cancel a feature before the code reaches tenants.

Each roadmap item has a feature identifier (in the form *ID 61652*). You can match the feature identifier against change notifications announced in the Message Center, part of the Microsoft 365 admin center (see Chapter 4), where a notification includes text such as "This message is associated with Microsoft 365 Roadmap ID 61652". In addition, each item has a date to tell you when Microsoft added it to the roadmap, when the new functionality should be available, and when Microsoft last modified the information for the item. The

integration between Planner and the Message Center is another tool a tenant can use to track change as Microsoft introduces updates into Office 365. See Chapter 9 for more information about the integration with Planner.

It is important to understand that the roadmap gives general guidance as to what might change in the future. The Message Center is a more authoritative view of new developments that apply to your tenant. Roadmap items give a glimpse into the future, but they might not be available for six months or more. Once a development shows up in the Message Center, it's more likely to appear in the following few weeks, meaning that it is time to prepare for change.

The screenshot shows the Microsoft 365 roadmap interface. At the top, there's a navigation bar with 'Microsoft 365' and various links. Below that is a header section with the title 'Microsoft 365 roadmap' and a brief description. A search bar contains the text 'planner'. To the right of the search bar are several filter dropdowns: 'Product', 'Release phase', 'Platform', 'Cloud instance', and 'New or updated'. Below the filters, it says 'Showing 11 updates'. There are three status filters: '3 In development', '0 Rolling out', and '8 Launched'. Below these filters is a list of updates, each with a title and a 'GA' (General Availability) date. The updates listed are:
 

- Microsoft Planner: Rich text and images in Planner task notes (GA: September 2022)
- Microsoft Planner: Recurring tasks (GA: June 2022)
- SharePoint: Planner cards appear within SharePoint team site home page activity feed (GA: May 2022)
- Microsoft Planner: Move tasks to any of your plans (GA: November 2021)
- Microsoft Teams: Tab actions are moving (GA: November 2021)

 At the bottom right of the list, there is a 'Feedback' button.

Figure 1-4: Browsing the Microsoft 365 roadmap

You can apply filters to the roadmap to show updates for specific products. Using filters to navigate the roadmap is useful as the sheer number of documented changes can be overwhelming at first glance. Filters also allow you to zero in on functionality that is most important to your company. This includes application-level features (for example, OneNote or Outlook), service-level features (for example, Exchange Online or SharePoint Online), and even sector capabilities (for example, features due for delivery in GCC).

The roadmap supports a download facility, meaning that you can apply filters to find the set of information you're interested in and then download details of those features to a CSV file, which you can process later with Excel or load into Power BI for further analysis.

**Versions of Software:** Microsoft uses a variety of methods to release software to tenants. If you configure a tenant for *standard release*, it means that you use software that is generally available (GA). This is the most stable version of an application. *Targeted release* means that you can choose for the complete tenant or specific users to see new features sometime before Microsoft makes it generally available. The exact period depends on the app and the complexity of the feature, but it is usually between four and eight weeks. A further delay might then come into play to allow client updates to appear to support new features, especially in desktop clients. All in all, it might take 90 days between Microsoft announcing that a

feature is generally available and everyone in every tenant being able to use the feature. Control over targeted release is through the Organization Profile tab of Org Settings in the Microsoft 365 admin center.

Some product groups run dedicated test programs (*technology adoption programs*, or TAP) to make beta software available for use in customer production environments. TAP software is a targeted release for use by selected customers. Microsoft often talks about *rings* of software releases to describe the progression from initial builds to a targeted release and finally become the standard release. Rings go from the development group (1) to Microsoft (2) to TAP (3) to release (4). Because development never ceases, the only guarantee you have is that the version of software you use today will change over time. The ever-changing nature of the service is why Microsoft calls Office 365 *evergreen software*.

## Command and Control

Microsoft 365 is a massive machine that moves forward at its own pace. Thanks to the roadmap, we know a lot more about what Microsoft is working on for the future than we did in the first few years, but even so, you cannot get away from the fact that customers cede an enormous amount of control to Microsoft when they sign up for cloud services. You accept that Microsoft will deliver a wide range of functionality, but you don't get to vote on what functionality Microsoft will deliver, when users see a new feature, or when changes show up. It just happens. This is a very different experience from the careful control that most IT departments exercise over the computer systems used in-house.

Most of the time this is not a problem and users like to find that new features continually appear. Google proved the attractiveness of an evolving interface to end users when it kept Gmail in what seemed to be a perpetual beta for many years. Even now, new features appear all the time in Gmail and Google Workplace, so Microsoft is simply keeping pace with its competition when it refreshes applications frequently.

However, dealing with a rapid update cadence can be problematic in the following ways:

- **User support:** Originally, two broad categories of users existed – people who access services through desktop clients and seldom make use of browser-based applications such as Yammer and Planner and those who use a browser (the third category is mobile users). In some cases, you don't have a choice because some applications only have a browser interface. Many Exchange users prefer the desktop version of Outlook because it allows them to continue working when offline. To some degree, Outlook insulates users from the ongoing changes that occur within the service. New features only appear after installing a new version of Outlook on a workstation, and even the “click-to-run” version of Outlook is slow to introduce the user interface necessary to support new functionality. On the other hand, those who use OWA to interact with Exchange Online might see new features show up on an ongoing basis. It's worth noting that Microsoft's “One Outlook” initiative intends to bring the different clients closer together, notably by allowing the Outlook desktop clients to share OWA components (this already happens with components like the calendar room finder). Every change in a user interface creates a potential flow of calls to local IT support as users seek information about why the change occurred and what it means to them. This is a very different mode of working from the normal carefully controlled and planned change management practiced by corporate IT departments. On the other hand, the Outlook desktop client is based on what is now an old architecture. More modern desktop clients, like Teams, have auto-update capabilities which means that they pick up new functionality as soon as Microsoft releases updated code.
- **Administration:** Those who manage a tenant work mostly through a set of web portals such as the Microsoft 365 admin center, SharePoint Online admin center, and Teams admin center. Microsoft updates and refreshes these portals over time to accommodate new applications and features, or for their purposes such as rebranding. In addition to the web portals, administrators can use PowerShell



or Graph-based tools to manage applications. Keeping up to date with the release of these features and learning how to manage them with the available tools (also usually dependent on whatever Microsoft provides) can be a challenge for tenant administrators.

- **IT management:** The role of IT management in an on-premises environment is well understood. They accept needs and requirements from the business and work out how best to meet these requests in line with the available budget, the current IT infrastructure, and the knowledge and experience available to the company. The business will continue to generate needs and requirements and Microsoft 365 is now one of the ways that the requests can be satisfied. The big difference is that IT management must accept whatever functionality comes from cloud providers. The task of matching needs and capability is easy if a good match is available within Microsoft 365. Things become a little trickier when functionality changes, as a decision to invest in an on-premises solution or to buy in software from another provider, might be undermined if Microsoft decides to offer equivalent functionality inside Microsoft 365.

The issues listed above are less important to small companies than they are for large companies who tend to have well-entrenched methods to control the introduction of new technology. In the same way, new companies usually embrace new technology faster than older companies do. A further consideration comes from the employee mix, where younger employees more readily embrace rapid change than their older peers do.

Apart from preparing IT Support staff with information about new changes so that they can respond to user queries, it's hard to manage how functionality flexes and changes within the service because of its rapid release cadence. One method that IT departments have used to manage the interaction between Microsoft 365 and end-users is to appoint a change coordinator. This role is to monitor changes, collect information about the changes, and make sure that the information about the changes gets to the right people in a form that makes sense to them and can be used (whenever possible) by the business to gain an advantage. The [Microsoft 365 Roadmap](#) is a prime source of information for this person but they should not limit their scan for information to Microsoft sources. A wealth of tips and advice exists across blogs, social networks, and conferences that need to be mined to form a complete picture of what is going on within Microsoft's cloud ecosystem. And that picture has then to be placed into context with the business goals and needs of the company so that the information is best used.

## Products and Licenses

Every active user in a tenant needs licenses to access functionality. In some cases, that statement is not completely true because situations exist where Microsoft 365 functionality is available to unlicensed users. Microsoft is closing off those holes, so it's best to adopt the attitude that licenses are necessary to use features. A tenant can buy a mixture of products and assign licenses for those products to users based on their individual needs. For example, some users might only need kiosk licenses to gain browser access to applications while others need products that include the Microsoft 365 enterprise apps and sophisticated Teams calling plans. Both sets of licenses exist quite happily within the tenant and administrators can move licenses between users based on the available license pool.

Users can access the applications available through the licenses assigned to their account when they open the [Office portal](#). The portal uses Microsoft Graph queries to display a set of recently used documents tailored for the individual user.

Three terms are important to understand when discussing Microsoft 365 products and licenses:

- **Product names** are how Microsoft refers to licenses. Office 365 E3 is a product, as is Viva Topics or Enterprise Mobility and Security E5.
- Stockkeeping units (**SKUs**) are the internal names used to manage licenses. For instance, the ENTERPRISEPACK SKU refers to Office 365 E3. SKUs also cover add-ons like SharePoint Syntex.
- **Service plans** control access to a licensed feature. You cannot buy a service plan as Microsoft includes them in SKUs. Office 365 E5 includes over 50 service plans. Somewhat confusingly, Microsoft also refers to products like Office 365 E5 which bundles many different service plans into a single offering as “plans.”

Microsoft 365 spans many different products, each of which specifies a range of functionality available to users and administrators. Details of [the current Office 365 enterprise products](#) and [Microsoft 365 enterprise products](#) are available online. Other pages describe offerings targeted at individual professionals and small companies.

## Service Families

Microsoft divides Office 365 into a set of service families, each divided into individual products sold to customers. Table 1-2 lists the service families and plans available in October 2021.

<b>Office 365 Service Family</b>	<b>Available Products</b>
Microsoft 365 Business <i>Available for up to a maximum of 300 users</i>	Microsoft 365 Business Basic Microsoft 365 Business Standard Microsoft 365 Business Premium
Enterprise <i>Available for an unlimited number of users</i>	Office 365 Enterprise F1 and F3 (front line workers) Office 365 Enterprise E1 Office 365 Enterprise E3 Office 365 Enterprise E5
Education <i>Available for an unlimited number of users</i>	Office 365 Education
Nonprofit	Office 365 Nonprofit
Government <i>Available for an unlimited number of users</i>	Office 365 Government E1 Office 365 Government E3 Office 365 Government E5 Office 365 Government F1 + F3
Sovereign Clouds	See <a href="#">this page</a> for information about the plans available in China.

Table 1-2: Service Families and products

Products available in specific markets vary from country to country, as does the pricing. Sometimes differences in tax regimes drive the price charged by Microsoft in a certain country and sometimes the competitive landscape within the country is the primary influence. Customers buy products based on their appropriateness, status, and availability. For instance, you can only buy Education licenses if your organization is a recognized university or another educational establishment. Likewise, nonprofit licenses are only available to organizations that meet Microsoft’s criteria to have nonprofit status. Microsoft makes education and nonprofit licenses available at a large discount compared to enterprise plans.

Government offerings are broadly like the corresponding enterprise plans. However, new applications invariably take longer to show up in the government cloud because of the need to meet specific requirements for cloud services. Teams is a good example. Despite being generally available to enterprise customers since March 2017, the U.S. Government Cloud (GCC) only [qualified Teams for use in July 2018](#). Since then, hundreds of individual Teams features have been through an approval process for deployment in GCC (and the GCC

High and Department of Defense clouds). Some features do not appear in GCC for several months after commercial release. Sometimes it's because of the approval process, and sometimes it's because a feature depends on other software that's not yet approved.

The Microsoft 365 Business products are for small businesses rather than large enterprises. Each includes the basic applications (Exchange Online, SharePoint Online, OneDrive for Business, Teams, Groups, and Planner). The difference between the business and enterprise plans is in the detail and depth of the functionality rather than the individual server applications. For example, if you need access to the widest range of compliance features, you need to buy Microsoft 365 enterprise because those products include that functionality. The assumption here is that individual professionals or small companies probably do not need quite the full range of compliance features.

A challenge that exists when writing about products is that they flex and change over time to reflect competitive pressures, new offerings, and new markets. The best idea for anyone considering moving to the cloud is to do some Internet research to discover what Microsoft offers in your geography and the monthly price per user. You also need to figure out whether the organization qualifies for academic, education, or non-profit licenses. And if you work for a government or state agency, you might find that you qualify to buy government licenses.

## Functionality Available in Office 365 Enterprise

Office 365 enterprise is the major subject of this book. Although E1 is the cheapest Office 365 offering, it still includes the fundamental collaboration capabilities offered by Exchange, SharePoint, and Teams. Office 365 E5 is the most expensive and includes features that every user might not need or want to use. Table 1-3 compares the services in Office 365 E1 plan against the higher-end E3 and E5 products to illustrate the difference in functionality that grows as the price increases. The obvious differences between E1 and E3 are the lack of access to the Microsoft 365 enterprise apps, and the compliance features (mailbox holds, data loss prevention, information protection, encryption, etc.). Office 365 E5 adds more automation and advanced functionality, like Office 365 Cloud App Security and Microsoft Purview eDiscovery Premium. [Microsoft 365 E3 and E5](#) include Office 365 E3 and E5 respectively. The Microsoft 365 E3 and E5 plans also include Windows 11 Enterprise and Enterprise Mobility and Security.

<b>Feature</b>	<b>Enterprise E1 \$8/month</b>	<b>Enterprise E3 \$23/month</b>	<b>Enterprise E5 \$38/month</b>
Office Online (web versions of Word, Excel, etc.)	Yes	Yes	Yes
Office for smartphones and tablets (up to 5 installs per user on PCs/Macs, tablets, and phones)	No	Yes	Yes
Microsoft 365 Apps for enterprise (the click to run Office desktop applications)	No	Yes	Yes
"Basic" Exchange email functionality (mail and calendars)	Yes	Yes	Yes
Advanced Exchange email functionality (100 GB mailbox and 1.5TB archive storage, legal hold, mailbox auditing, Data Loss Prevention)	Yes	Yes	Yes
Teams (conversations and audio/video meetings)	Yes	Yes	Yes
SharePoint Online (team sites)	Yes	Yes	Yes
OneDrive for Business File Storage and Sharing	Yes	Yes	Yes
Yammer Corporate Social Network	Yes	Yes	Yes
Microsoft Stream (Plan 2)	Yes	Yes	Yes
Delve	Yes	Yes	Yes
Basic management (including PowerShell)	Yes	Yes	Yes

Rolling updates	Yes	Yes	Yes
Microsoft Purview Information Protection (rights management)	No	Yes	Yes
Security and Compliance (auditing, search, eDiscovery)	Yes	Yes	Yes
Viva Insights	Yes	Yes	Yes
eDiscovery Premium	No	No	Yes
Advanced data governance (auto-apply, trainable classifiers)	No	No	Yes
Data Loss Prevention (DLP)	No	Yes	Yes (including Teams)
Teams Phone system and audio conferencing	No	No	Yes
Microsoft Defender for Office 365	No	Plan 1	Plan 2
Office 365 Cloud App Security	No	No	Yes
FastTrack onboarding service	Yes	Yes	Yes
Workflow automation with Power Automate (Flow)	Yes	Yes	Yes
Visio web app	Yes	Yes	Yes

Table 1-3: Comparing functionality across Office 365 Enterprise

Cheaper frontline licenses are available to accommodate deskless workers or those who use a shared PC. These plans include email and web-based versions of the Office applications. Mailboxes for F3 users are smaller than other plans (2 GB) and cannot be archive-enabled. Even so, these mailboxes are large enough to meet the needs of many users. Office 365 F1 and F3 include Teams, 2 GB of OneDrive for Business storage, Planner, Yammer, Sway, and Stream. F3 users also get Power Automate, Power Apps, and Exchange Online.

Microsoft 365 is flexible when it comes to altering the number and type of licenses that a tenant owns. Each user has a separate license, so you can mix and match license types within a tenant. The flexibility in licensing means that you can increase or reduce capacity as the number of users grows or declines over time. It also allows you to create a customized mix of licenses in a pool for allocation to users based on their needs. For example, a company of 10,000 users might have some people who only need intermittent access to email and never join online meetings. Kiosk (front-line worker) plans might satisfy these users while the enterprise products serve other users better. Even within the enterprise, a division might exist between users who only need basic functionality and those who need the more extended variety. In some cases, the need for advanced compliance or data governance features drives the choice of Office 365 E5 or Microsoft 365 E5 while factors such as the need to replace an old PBX will convince tenants to buy add-on licenses.

Microsoft publishes [Service Descriptions](#) and [comparisons](#) to guide customers through the functionality available in its current products. As in all negotiations, before you can decide which products are right for your company, you need to understand what functionality you need now, what might be necessary for the future, and the features that you do not need. Once you know the functionality you need, you can discuss licensing requirements and pricing with Microsoft. An independent and more graphical way to view the applications available in the available plans is available in the [Periodic Table of Office 365](#).

Remember that product details change over time, even within specific features. For example, in 2019, Microsoft moved the MyAnalytics (now Viva Insights) and Stream Plan 2 features from Office 365 E5 to E3. Microsoft introduces new features and applications on an ongoing basis. The plan you settled on two years ago might not be the best now. It is sensible to use the annual subscription renewal cycle as a reminder to check the functionality offered in the various products and confirm which is best for your organization.

## Adding Cost

Microsoft is not a charitable organization, and it is in its interest to sell as many licenses for its cloud services to customers as it can. The sources of added revenue for Microsoft are:

- Convincing customers to use higher-priced products. For example, getting organizations to upgrade from Office 365 E3 to Office 365 E5 benefits customers by enabling access to many compliance features. Microsoft gains by charging an extra \$15/month per user.
- Selling plan add-ons. If customers do not want to buy a higher-priced plan, they might be able to buy specific functionality through something like the Viva Topics add-on.
- Expanding to include other cloud services. Tenants have basic access to Azure AD. You might like to buy premium licenses to increase the overall security of the tenant and add functionality through features such as conditional access policies, group access reviews, group expiration policies, and password write-back to an on-premises directory used in hybrid deployments. The same is true of mobile device management, which you can perform at a basic level through the ActiveSync management tools built into Exchange Online but is easier and more functional when you deploy Enterprise Mobility and Security. An alternative to buying separate plans is to upgrade to Microsoft 365 to take advantage of the bundled price for Office 365, Enterprise Mobility and Security, and Windows 11.

Buying options can increase a tenant's monthly bill by a large amount. On the other hand, if the organization needs the functionality and it enables you to decommission older systems (especially on-premises servers), then the cost might be justifiable. Office 365 products include lots of functionality but sometimes you need just a little bit more. For example, assume that you license Office 365 E3 for everyone, but the activity of some accounts might need to be monitored by Office 365 Cloud App Security, which is part of Office 365 E5. You could simply buy some E5 licenses and assign the licenses to those accounts, but it is sometimes possible to buy the specific feature through an add-on. If this is the case, it is usually cheaper to buy exactly what you need rather than to upgrade to the next level. To discover what add-ons are available to you, click **Purchase Services** in the **Billing** section of the Microsoft 365 admin center, which brings you to the [service catalog](#). You can then select whatever add-on you need and decide how many licenses to buy. Microsoft charges for add-ons on a per-user, per month basis. The monthly charge for an add-on varies from country to country.

## Enterprise Mobility & Security and Azure Active Directory Premium

Tenants use Azure AD to store information about user accounts and settings. Although the functionality available in the version of Azure AD is enough to allow users to authenticate and access the apps, some organizations, especially enterprise deployments, pay extra to access the features exposed through Azure AD Premium features. You can buy Azure AD Premium licenses separately or as part of [Microsoft Enterprise Mobility and Security](#) (EMS). From April 2020, the Microsoft 365 Business Premium plan includes Azure AD Premium P1. Table 1-4 lists the functionality available through the EMS plans.

<b>Enterprise Mobility + Security plans</b>	<b>Identity and access management</b>	<b>Managed mobile productivity</b>	<b>Information protection</b>	<b>Identity driven security</b>
<i>EMS Plan E5 (\$16.40/month)</i>	Azure AD Premium P2	Microsoft Endpoint Manager	Azure Information Protection Premium P2	<a href="#">Microsoft Defender for Cloud Apps</a>

<i>EMS Plan E3 (\$10.60/month)</i>	Azure AD Premium P1	<a href="#">Microsoft Endpoint Manager</a>	Azure Information Protection P1	<a href="#">Microsoft Defender for Identity</a>
--	------------------------	--	---------------------------------------	---

Table 1-4: Functionality available in the Enterprise Mobility and Security plans

Table 1-5 lists some of the features supported by Azure AD Premium P1 that are usually of interest to tenants. The [full feature list enabled by Premium licenses is available online](#). Azure AD Premium P1 includes other features, such as [conditional access policies for Exchange Online and SharePoint Online](#), which tenants can deploy to force users to sign in with multi-factor authentication based on a network location. The P2 plan adds identity protection and privileged identity management.

The decision to invest in Enterprise Mobility and Security or Azure AD Premium licenses depends on the value that a tenant can gain from the available functionality. Some tenants will never need any of these features, some will need just one or two features, and some will use all the features. As a decision to use these products is a factor that can influence the overall cost of Microsoft 365 for an organization, it is a factor to include in budget discussions.

<b>Feature</b>	<b>Use</b>
Enhanced multi-factor authentication	Protects other cloud workloads in addition to Office 365.
Password write-back	Enables the write-back of user passwords to an on-premises Active Directory.
Connect health	Delivers information about directory synchronization performed with Azure AD Connect (AAD Connect).
Dynamic groups	Allows groups to have dynamic membership calculated based on queries executed against Azure AD (including dynamic Microsoft 365 Groups).
Group expiration policy	Expires Groups after a predetermined period and allows group owners to renew their groups for a further period.
Advanced reports	Includes reports such as password reset activity and irregular sign-in or anomalous sign-in activity.
Assign licenses via Groups	You can achieve a basic level of automation by using Groups to assign licenses to members of those groups.
Conditional access policies	Control who can access your tenant and the conditions under which they are allowed access.

Table 1-5: Azure AD Premium features of interest to enterprise tenants

## Licensing Requirements for Azure Active Directory

As described above, some Azure AD features need premium licenses. Microsoft says that “a [proper license is required](#) if a user benefits directly or indirectly from any feature covered by that license.” Sometimes Microsoft does not enforce a license requirement to block someone from using a feature and sometimes a partial block is in place. For example, if an administrator account has an Azure AD Premium P1 license, they can use the Azure portal to create dynamic Groups or define a group expiration policy. The Azure portal enforces the license requirement by not revealing the necessary UI unless the user has the correct license. However, no license requirement exists in PowerShell, so you can use it to create dynamic groups or create a policy. It is also true that dynamic groups work even if the accounts that come within the scope of a query do not have a premium license. Microsoft could enforce the requirement at any time in the future, so it is both unwise and contrary to the licensing agreement to take advantage of any gaps you find.

## Buying Through CSPs (Cloud Solution Providers)

Microsoft is not the only company that sells its licenses. You can also buy Microsoft 365 through [Microsoft cloud solution providers](#) to gain the benefit of the insight and knowledge that they have about how to deploy and use Office 365. In many countries, local resellers package Office 365 with other services to suit the local market. You still get the same Office 365 that Microsoft sells along with whatever added benefits the reseller can deliver. The reseller might charge an extra fee on top of the subscriptions levied by Microsoft. The question, therefore, is whether it ever makes sense to buy from a reseller.

The answer depends on exactly what added services you get for the extra fee. For instance, some resellers will take care of the entire migration process for you while others deliver a “white glove” service where they manage any support issues that occur. Dealing with cloud support can be a particularly frustrating area so interacting with a local services company that understands how cloud support works and has direct knowledge of Microsoft 365 to help solve problems without the need for escalation can be very valuable.

## Trial Tenants

Microsoft makes trial tenants available to potential customers with [Office 365 E3 or Office 365 E5 licenses](#) to try for up to 30 days without payment. This is a practical way to try out the functionality available with different products and decide which is the best fit for your company. It is also an excellent method to allow administrators to become accustomed to managing a tenant as they can make as many changes to settings as they like without running the risk of impacting users. You can transform a trial tenant into a production tenant if you want, but in most cases, it is best to start over and treat the trial as a sandbox that can be discarded when no longer needed. For this reason, never use a corporate domain name for a trial.

## Developer Tenants

The Microsoft 365 Developer Program supports the work of developers who need access to Microsoft 365 to build solutions for internal use or to sell to customers. By [signing up for the developer program](#), you can create a fully-functional test tenant with 25 Microsoft 365 Developer E5 licenses. The tenant is available as a general sandbox for development activity from testing code to validating performance. It's an excellent way to host application development, demos, proof of concepts, and deployment and testing of software builds. To facilitate testing, Microsoft has [sample data packs](#) to populate the tenant with test accounts and user data, but you can add your test data too.

The tenant comes with an initial 90-day subscription. If the tenant remains in use for development purposes (using tools like the SharePoint Foundation or Microsoft Graph), Microsoft automatically renews the license every 90 days to ensure that development activity can continue for as long as necessary.

# The Commercial Success of the Microsoft Cloud

Only Microsoft knows the details of how successful its cloud services are and how many paid subscribers each service has. This is commercially sensitive data, and it should come as no surprise that they reveal as little real data about the numbers as they possibly can. Commercial Cloud is an artificial mixture of revenues derived from different Microsoft cloud services including Office 365, Azure, Dynamics 365, and LinkedIn. It is not one of the formal business segments Microsoft uses to report results like Intelligent Cloud or Productivity and

Business Processes. Nevertheless, the number reported for commercial cloud revenues (now referred to as the Microsoft Cloud) is useful because Microsoft has reported comparable data for several years.

## Microsoft Cloud Revenues

In the 2011–2015 period, much of the initial growth for Office 365 came from small to medium business companies that moved from on-premises to the cloud. Today, even the largest enterprises are comfortable with the security, privacy, and operational aspects of cloud services. The growth rate for Office 365 has slowed over the years, but still reflects a continuing movement of work from on-premises systems to the cloud. Today, annualized revenue for the Microsoft Cloud, a large percentage of which comes from Office 365, is a major contributor to Microsoft's overall income. Some estimates put the revenue generated by Office 365 annually at over \$60 billion. Importantly, the revenue comes with a healthy gross margin (reported to be 70% for Microsoft Cloud in Microsoft's FY22 Q2 earnings call).

In their FY16 Q4 results, Microsoft reported that the annualized revenue run rate (ARR) for commercial cloud products [had reached \\$12.1 billion](#). Fifteen months later, when they announced their FY18 Q1 results (October 2017), Microsoft said that the run rate was \$20.4 billion, a remarkable jump of \$8.3 billion over five quarters which comfortably allowed CEO Satya Nadella to attain the goal of \$20 billion set in 2015. To calculate the annualized run rate, Microsoft takes the result achieved in the last month of a reporting period and multiplies it by 12. In April 2022, Microsoft reported revenue for the Microsoft Cloud to be \$23.4 billion, which equates to an \$93.6 billion annualized run rate. Because Microsoft books varying numbers of customer projects over a year, a difference always exists between the actual revenue achieved in a year and the ARR. Nevertheless, the headline figure used for ongoing comparison is the ARR.

Although Azure continues to grow faster than Office 365, the higher prices customers pay for Office 365 subscriptions (computed by Microsoft as the ARPU, average revenue per user) and the number of Office 365 users indicate that Office 365 represents the bulk of the Microsoft Cloud revenue mix. At the FY18 Q3 earnings call, a Goldman Sachs analyst estimated that Azure brought in roughly \$8 billion of the then \$24 billion annualized revenue. The figure for Dynamics 365 was probably another \$1 billion, which implies that Office 365 was then responsible for up to \$15 billion in annualized revenue. At the time, this equated to an ARPU of \$111, or \$9.26 monthly. Microsoft attempts to increase ARPU by moving subscribers to higher-priced plans or by selling add-ons to license additional features. To encourage customers to upgrade, Microsoft restricts access to many new features, especially in the areas of compliance and security, to tenants with high-end licenses. The importance of this tactic to overall cloud revenues can be seen in the [revelation by Microsoft's CFO](#) that in FY20 the E5 SKUs contributed \$7.5 billion in revenue, even though only 5% of the Office 365 user base had these licenses, an increase from the \$4.2 billion recorded in FY19. In July 2021, Microsoft said that 8% of the Office 365 installed base now had Microsoft 365 E5 licenses.

The latest quarterly revenue for the Microsoft Cloud business segment was \$23.4 billion (April 2022). See Table A-1 in the Appendix for a full list of quarterly revenue results since April 2015.

As of July 2022, Office 365 is available in [249 markets and 44 languages](#). It is now easier to say where Office 365 is unavailable (Cuba, Iran, Democratic People's Republic of Korea, Sudan, and Syria).

**Price Increases in 2022:** In August 2021, Microsoft announced [the first major price uplift for Office 365 and Microsoft 365 products](#). Both the Office 365 E3 and E5 plans increased by \$3/month on March 1, 2022. Microsoft justified the price increases because of the new Office 365 apps added since 2011, the volume of individual changes made in that period, the use of artificial intelligence and machine learning to automate processing for organizations and users, and the introduction of better security and compliance functionality. Some will be more convinced than others, with doubters pointing to Microsoft's need to



continue growing cloud revenues to please Wall Street as a more likely driving force. The price increase did not affect Office 365/Microsoft 365 academic and personal products.

## Growth in Usage

The increase in Office 365 users is in line with the growth in revenues. Microsoft uses a concept called *monthly active users* to report usage. This number is not the total of licenses sold to customers. Instead, it reflects those who have licenses who log on and use the service at least once a month (which is not much for anyone using any of the applications). A definition of what monthly active users are for a workload is:

*The maximum daily users performing an intentional action in the last 28-day period across the desktop client, mobile client, and web client.*

Intentional action is something that a user does to perform some measurable activity, like creating and sending a message or scheduling a meeting. For SharePoint Online, countable operations are things like uploading or editing a document, while Teams counts actions like attending an online meeting or starting a chat. The definition of an active user does not count actions like starting or exiting an app. Signals captured in the Microsoft Graph record user actions. The data is deduplicated based on user identifiers (the GUID for user accounts) to ensure that people are not counted twice.

The number of active users reported by Microsoft understates the total usage because it does not include accounts that Microsoft has provisioned for tenants which are not yet in active use, free accounts (such as Microsoft's use of the Service), and trial domains. The numbers are even higher if we speak about mailboxes instead of accounts. The number of cloud mailboxes is considerably higher than the number of cloud user accounts because of shared mailboxes, inactive mailboxes, resource mailboxes, and group mailboxes.

The current reported number of paid Office 365 customer seats is 345 million (April 2022). Given that some seats are paid for but not used, the likely number of active paid seats is probably around 321 million. You can add a few extra million free seats (used by Microsoft and their partners) to this number. Table A-2 in the Appendix details the growth in Office 365 users since November 2015.

## Workload-Specific Usage

When it reports overall numbers for Office 365, Microsoft doesn't break the data down into precise figures for each workload. In November 2019, Microsoft said that SharePoint Online has 100 million monthly active users; they updated the number for [SharePoint Online to 200 million in December 2020](#). Much of the increase in SharePoint usage comes from Teams. Every team has a SharePoint team site and Teams stores shared files and recordings of Teams meetings in SharePoint Online and OneDrive for Business. By itself, Microsoft claims that Teams has 270 million monthly active users but does not break out how many of these users are commercial, personal, or educational.

Given that a very high percentage of active Office 365 users have an Exchange Online mailbox, it's reasonable to assert that Exchange Online is the largest workload and that Exchange Online is hugely influential on the overall success of Office 365. In 2017, Microsoft said that Exchange Online alone handles about 60 billion requests per day from all connected clients (many people use more than one client, and each client handles multiple transactions). At the BUILD 2016 conference, Microsoft revealed that Exchange Online users create more than a trillion meetings monthly and that Exchange Online had processed over 4 trillion outbound messages to that point. The volume of work handled by Exchange Online was massive then and is probably three times larger now given the growth since 2016, even if some of the communications traffic previously handled by email is now in Teams.

Other products benefit from the consistent growth in Office 365. In their FY22 Q3 results briefing, Microsoft revealed that Enterprise Mobility and Security has 218 million users. Microsoft often sells EMS to enterprises

along with Office 365 (standalone or as part of Microsoft 365), which points to high penetration of EMS in Office 365 enterprise tenants. In the same briefing, they said that Azure AD had 550 million monthly active users (fifteen months prior, the reported number was 425 million monthly active users). Previously, Microsoft said that over 200,000 organizations pay to use Azure AD. It's worth remembering that among the 90 billion sign-in events handled daily by Azure, many come from non-Microsoft applications like WorkDay and ServiceNow which use Azure AD as their directory of record.

In September 2017, Microsoft Corporate VP Rajesh Jha said that he expected [more than 70% of the Exchange installed base to be in the cloud by Microsoft's 2019 fiscal year](#). Because it's hard to count on-premises licenses assigned to real people, no one knows the precise number for Exchange on-premises users at the height of its popularity. Some industry sources put the number at around 300 million, so 70% of that base is 210 million. Taking SharePoint into account, the base that Office 365 can grow off is larger. In October 2020, CFO Amy Hood said that Office 365 now accounts for "over 70 percent of our existing Office commercial paid installed base." Using the 258 million paid seats reported by Microsoft at the time, this figure indicates a total base of around 360 million, so there's still some room for Office 365 to grow.

## Service Level Agreements for Online Services

Two formal documents govern the provision of Microsoft Online Services to customers. Microsoft updates both documents quarterly. The [Volume Licensing Online Services Terms](#) lay out the terms under which Microsoft delivers the service to customers. The document is available in multiple languages. More information is in the [Service Level Agreement for Microsoft Online Services](#), also available in multiple languages. This document explains how Microsoft measures the SLA for its Online Services.

Each of the individual workloads has its own SLA definition. For Exchange Online, Microsoft defines downtime to be "Any period when end users are unable to send or receive emails with Outlook on the web." The actual calculation is a little more complex:

*"The "Monthly Uptime Percentage" for a Service is calculated by the following formula:*

$$\mathbf{((User\ Minutes - Downtime)/User\ Minutes) * 100}$$

*where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident."*

Incidents might only affect a few users, so the number of user minutes generated by an incident varies. Many incidents are transient, last just a few seconds, and are due to a temporary condition. Over a working day, micro-outages can happen five or six times without disrupting end users. These micro-outages can sometimes be detected as a failure of a page to load or a server to respond promptly. They usually resolve themselves without the need to intervene on the part of either Microsoft or the tenant administrator.

An incident that lasts several hours but only affects a small group of users (for example, in a single data center or region) will not have much effect on the overall SLA. To see any movement in the SLA, a global outage for a couple of hours needs to happen. Due to the distributed nature of the data center infrastructure, incidents usually stay constrained within a single region. High-profile but region-constrained outages like those that affected U.S. tenants in June 2014 and July 2015 seldom move the SLA needle very much. Indeed, a daily glance at the Service Health Dashboard will invariably show that at least one workload is currently experiencing some level of difficulty that might or might not affect your users. This is part of the joy and the pain of sharing in massive multi-tenant infrastructure.

## Financially Backed SLA

If the reported Monthly Uptime Percentage falls under 99.9%, Microsoft guarantees that they will refund customers with a service credit for a set amount ranging from 25% to 100% (for less than 95% availability).

An SLA of 99.9% (“three nines” in high availability parlance) allows a service provider to have total downtime of up to 43.8 minutes per month or 8.76 hours annually. As explained above, lost minutes for Exchange Online accrue when users can’t send or receive mail with OWA. For SharePoint Online, it is *“any period when users are unable to read or write any portion of a SharePoint site collection for which they have appropriate permissions.”* These conditions are in the Service Level Agreement for Microsoft Online Services.

Counting minutes for SLA calculation only happens when Microsoft support records an incident and accepts responsibility for the problem. In most cases, this means that the problem must happen within a Microsoft data center due to a component that is under Microsoft’s control. Anything that happens within the control of the customer, such as a misconfiguration of client software or a local network malfunction, does not count against the SLA. This is logical because Microsoft is responsible only for the parts of the service that it controls.

Microsoft is not the only cloud provider to limit SLA measurement at the boundary of its data centers. No one controls the Internet, and no one controls how data flows from your client to a data center. Many complex steps occur between a client connecting to a cloud service, perhaps using multi-factor authentication from a mobile device, to access and interact with data. It is therefore impossible for a provider to offer an SLA guarantee as measured at the client. Well, perhaps possible because anyone can offer such an SLA, but certainly foolish in terms of their chances of ever meeting the SLA.

Microsoft excludes scheduled downtime from any SLA calculation. This is downtime that Microsoft needs to maintain its service and is advised to customers at least five days before it happens. On the other hand, any sudden outage that comes without warning always counts against the SLA, if Microsoft accepts that the incident is caused by their software or infrastructure and did not happen because of any action taken by the tenant.

At any time, multiple incidents affecting Microsoft services are active across the world. Some incidents are very local and affect a single server. Some are more widespread and degrade the service delivered to large numbers of tenants or even stop an application working for those tenants. Some incidents are transient and go away on their own accord and some linger for days, albeit perhaps without affecting tenants. The point is that an infrastructure that is so large cannot run in a perfect state all the time. Because cloud software depends on hardware, software, and humans, you can be sure that something is always happening for the wrong reason. Of course, users will not care unless an incident stops them from sending emails, accessing documents, conducting a meeting, or some other operation.

In most cases, any issue that affects the SLA is soon obvious because users are unable to connect or do work. Microsoft has automated systems in place to interpret problems and map them against tenants so that the incidents highlighted in the Microsoft 365 admin center are those that might affect smooth operations for your tenant. Details of incidents appear in the Service Health Dashboard (SHD). We discuss how to understand the SHD, navigate within it, and file service incidents for your tenant in Chapter 4.

A summary of incidents for the last 30 days is also available in the Microsoft 365 admin center together with post-incident reports (PIRs) by going to **Service Health** and then selecting [History](#). However, it is easy to forget to check the portal from time to time to see the current status of all parts of the service.

**What is a PIR?** A PIR is a Post-Incident Report with the formal analysis of an *“unplanned customer-impacting service incident”* or outage where there was a *“broad and noticeable impact across a large*

*number of organizations.”* Few incidents that you see in the Service Health Dashboard affect enough tenants for Microsoft to write and publish a PIR. When these incidents do happen, Microsoft makes PIRs available to customers through the Service Health Dashboard. The support team publishes a preliminary PIR within 48 hours of the incident closure followed up by a full and detailed PIR within five business days. A PIR includes the following sections:

- **Incident ID:** Every incident has a unique identifier. For example, EX29054 is an Exchange Online incident.
- **User Experience:** A brief description of the impact on end-users. It might be that the incident affected only administrators if the component involved is something like PowerShell.
- **Customer Impact:** What business impact (if any) the incident caused.
- **Incident Start Date and Time:** In UTC format.
- **Incident End Date and Time:** Logically, the difference between the start and end date is the incident period used to calculate the time lost against SLA.
- **Root Cause:** Microsoft’s explanation of why the incident occurred.
- **Actions Taken:** A timeline of all the actions taken from the time when the incident occurred to its resolution. Sometimes several hours go by before engineers have correlated enough data from multiple tenants to be able to home in on the components that might lie at the root of the problem.
- **Next Steps:** What Microsoft proposes to do to ensure that the same problem will not recur.

As in all administrative documents, sometimes you must read between the lines to understand just what happened and why it happened. The PIR also exists in internal and “customer ready” forms, the former being much more detailed than the latter.

Social media is important to Microsoft when it comes to monitoring the service too. The sheer number of users connecting to Microsoft 365 means that any problem can quickly escalate to affect tens of millions of people. Microsoft monitors information flowing through social media to detect problems that users report (and complain about). This is just one of the ways that they try to figure out the quality of service delivered to end-users.

Seeing a tsunami of updates about a problem with an application like Teams or SharePoint Online is not a reason to panic because the problem might not affect you. Remember that Microsoft 365 uses data centers around the world. Many data centers offer local services to users in a single country. The way services run within a data center means that any problem which appears in a data center likely only affects people connected to that data center. Sometimes this isn’t the case when a component running in a single location fails. For example, this happened when the [Azure AD multi-factor authentication service suffered a couple of outages in late 2018](#). As time goes by, the number of single points of failure declines within the Microsoft 365 infrastructure, but issues can still happen, as occurred with [another authentication outage in March 2021](#).

The average administrator has no real chance of understanding the data used in SLA calculations because much of it is invisible to anyone outside Microsoft. The SLA equation excludes factors such as Internet outages, local network delays, or client misconfiguration. This means that while one tenant believes they receive an excellent service measured against SLA, another tenant who experienced some problems, perhaps some of their own making, has quite a different perception of the service quality received. Beauty is in the eye of the beholder, and it is obvious that the sheer size of Microsoft 365 creates a blurring effect across its regions. Service is excellent overall (as seen in the reported SLA figures) but is awful when experienced by individual tenants affected by different bugs or operational issues.

To be fair to Microsoft, they have met their financial commitment to refund customers when the root cause for obvious outages turned out to be mistakes that were under their control. It is important to recognize that most outages are of short duration and only affect a small percentage of the overall user base.

## Performance Against SLA

Microsoft reports Office 365 SLA performance [against its 99.9% target](#) every quarter. Microsoft aggregates the data for all regions to create a worldwide result and does not give SLA information for individual regions or applications. The latest SLA figure (Q1 2022) was 99.98%. Table A-3 in the Appendix lists all the quarterly SLA results since Q1 2013.

SLA results are available six weeks after the end of the quarter. Remember that the SLA data reported by Microsoft is a unified view taken from across the complete service. Also, failures in some dependent services affecting Microsoft 365 do not impact the SLA. For example, several problems experienced by tenants in 2020 were associated with failures in the Azure infrastructure. Some of these failures influenced the SLA; others did not.

## Leveraging the Breadth of Office 365

When Microsoft CEO Satya Nadella told the audience at the October 2014 Gartner conference that *"Office 365 is the new Exchange and one will cannibalize the other. The key is to ensure that current Exchange customers can transition on their own terms,"* he reflected the reality that many customers selected email as the first workload to move into the cloud. The reason why he called Office 365 "the new Exchange" is that the movement of email into the cloud mimicked the movement of email from expensive mainframe and minicomputers to PC-based Windows NT servers some twenty years previously. Exchange Server was the first Windows server application considered "enterprise-ready." In other words, Exchange could scale to offer the kind of email service needed by the world's largest corporations. Over the next twenty years, Exchange delivered on that promise.

Email is a big reason why many companies decide to move their workloads to the cloud, but what of the other cloud applications? After all, you pay for these applications alongside Exchange, so it is wasteful if you do not seek to take advantage of them. Any plan to migrate should consider what advantages a company can derive from the full spectrum of Microsoft 365 rather than applying a narrow email-centric focus. Here are the services that round things out:

- **SharePoint Online:** While organizations were initially slow to move on-premises SharePoint deployments to SharePoint Online, that point is long past and SharePoint Online now boasts over 200 million daily active users. The mission of SharePoint Online is different from the way organizations often deploy the on-premises variant. SharePoint Online is the document management service for Microsoft 365 applications and is closely integrated with many apps including Teams, Planner, and Yammer.
- **OneDrive for Business:** OneDrive for Business is core to enabling people to work offline or across low-bandwidth environments. Its synchronization client uses differential synchronization to deal with files of up to 250 GB stored in SharePoint Online or personal OneDrive for Business accounts. Microsoft provides "unlimited OneDrive storage" to Office 365 users with enterprise accounts to allow them to move files traditionally kept on local PCs to the cloud.
- **Yammer:** Microsoft bought Yammer in July 2012 to gain its enterprise social networking technology. Since then, Microsoft has integrated Yammer into Microsoft 365 in several ways such as using Azure AD for identity management, using SharePoint Online to store files posted to Yammer, and using the Groups service to manage the membership of Yammer communities. Some commentators observed that Microsoft also acquired a fast-paced web-centric development culture from Yammer which helped them to move to the release cadence expected from cloud services. Yammer is all about

sharing information on an enterprise-wide basis, an aspect that makes it more interesting to larger organizations.

- **Planner** and **Teams** are two applications available to enterprise tenants. Both use Groups to manage team members. Planner focuses on task-based planning for team activities while Teams offers a chat-based workspace for users to work together with audio and video calling, federated communications with Skype consumer users, an application platform, and a replacement for traditional phone systems.
- **Stream**: Stream provides video recording and consumption services to applications. See Chapter 10 for information about managing video.

Perhaps the most interesting thing that has happened since the introduction of Office 365 is how the barriers that existed between on-premises applications disappear when Microsoft has total control of deployment and operations. Relatively few on-premises environments successfully integrate SharePoint and Exchange, but the configuration and operational difficulties go away with the cloud versions. Microsoft takes advantage of this to build new applications that exploit Exchange Online for email and calendaring and SharePoint Online for document management, adding whatever extra software is necessary to complete the picture. Teams and Planner are examples of new applications that could not exist on-premises. Both are unique to the cloud.

If the right licenses are in place, all this functionality is available to tenants. Many companies plan to move to the cloud for a specific purpose and can benefit even more by exploiting the other technologies mentioned above. The best thing is that Microsoft takes care of all the work to deploy and manage the technology, meaning that the normal learning curve needed to master the details of capacity planning, server deployment, and application configuration and installation is not necessary. All efforts can focus on the question of how best to use the features available through SharePoint Online, Teams, Planner, and other applications to solve business problems.

Like any migration, the transition from on-premises servers to their cloud equivalents are projects that benefit greatly from experience. If you don't have the necessary expertise within the company, you should find some experienced consultants to help you plan for and execute the change.

# Chapter 2: Embracing the Microsoft 365 Cloud

**Paul Robichaux**

## Should You Move to the Cloud?

If you pick ten companies from a random selection in any country or industry, you'll find *at least* ten different lists of reasons for or against moving to the cloud. These reasons might include modernizing a creaky or unsupported IT infrastructure, the desire to treat applications like email and document management as utility functionality instead of running software on in-house servers, or fear of a loss of control due to turning over services to an outside provider.

Today's truth is that most enterprises that want the benefits of cloud services have already started taking advantage of them. Organizations that haven't yet moved at least *some* of their workload to the cloud probably either don't want to or cannot because the cloud can't support their business requirements. Industries with heavy regulation or compliance requirements, or unusual security needs such as the ability to process classified information, have historically lagged in cloud adoption. However, even in these unique corners of the market, some companies have decided that the benefits outweigh the risks, costs, or limitations.

This chapter will help you work through some key topics related to moving to the cloud, broken into three broad categories:

- Things you should know or consider *before* you start moving to the cloud. (Even if you've already moved to the cloud, you may find some useful nuggets here as you consider how and when to adopt additional cloud services.)
- Things you should know and consider *while* you prepare and execute your move to the Microsoft 365 cloud.
- Things that you should be prepared to do *after* your cloud migration is substantially complete.

Let's start with some basic terminology. When we refer to "the cloud," you can consider that shorthand for "Microsoft 365 and its related services." We won't cover Microsoft's other cloud services (including Azure or Dynamics 365) or non-Microsoft cloud services such as Amazon Web Services, SAP Hana, or Salesforce.

Next, we'll distinguish between the process of *migrating to the cloud* (meaning the process of moving data and services into Microsoft 365) and *adopting the cloud* (meaning convincing users and business stakeholders to *use* Microsoft 365 once the migration's underway). Migration must always start before adoption, although they will often run in parallel: the first users to migrate start adopting the new system while later waves of users wait to be migrated.

Just understanding the pro- and anti-cloud arguments isn't enough by itself. You must put these arguments in context for your company. Your knowledge about how the company works and its business goals will help you put the points made here into the right context for your business. Never take a consultant's opinion as definitive—or ours, for that matter!

# What to Do Before You Move

Different organizations may have radically different reasons for moving to the cloud. Some do it to replace an unmaintainable or expensive legacy infrastructure; others are chasing improved productivity or shiny features. Regardless of the basic reasons behind a specific migration, there are several important points to evaluate when you're considering a migration to Microsoft 365.

## Dispel Any Lingering Cloud FUD

The use of fear, uncertainty, and doubt (FUD) is a much-beloved tactic to delay or deflect a decision away from certain technology. Salespeople have used FUD (a term first used in 1975 by [Gene Amdahl](#), a former IBM employee whose company was the first real threat to IBM's mainframe business) for decades to try to discourage the adoption of PCs, mini-computers, client-server implementation, the web, social networking, mobile devices, and now the cloud. The tricky thing about FUD is that some of the fears expressed are valid concerns and deserve inclusion in your decision-making process as part of your decision; other fears are over-hyped.

Common FUD-related statements you might have encountered include:

- You can't depend on the cloud (or the Internet).
- Cloud services are immature.
- Cloud platforms are insecure.
- On-premises IT delivers better service to the company and end-users.
- You lose control when you move to the cloud.

These aren't the only specimens of FUD out there, but in every case, it is important to separate fact from fallacy. Let's look at each of the issues listed above.

### Is the Cloud Dependable?

*"You can't depend on the cloud"* can have several meanings. It could be that users or business leaders lack trust that they'll be able to connect to cloud services as easily as they can with on-premises services because they're worried about your internal network services, local Internet connectivity for users, or poor performance. It might mean they believe that your IT department can deliver a more robust service than a cloud provider would. It might mean that you're worried that a cloud provider will remove, change, or limit services once your business has come to depend on them.

You could short-cut the question of whether the cloud is dependable enough for business use by saying "of course it is; look how widely it's used!" That would be true but might not be enough to quell doubts. If we dig a little deeper, we can come up with a better answer, starting with the network.

The network is indisputably the essential link between cloud services and end-users. If you have a slow or unreliable network, your users will have a slow or unreliable cloud experience, period. Your network has three parts that work together to deliver end-to-end connectivity: the internal network used by your users (including its connection to the Internet); the Internet itself; and the network controlled by the cloud provider.

Your existing company network was probably built to connect users to on-premises systems, so it's reasonable to assume that you'll need to upgrade it to support your transition to the cloud. The best way to justify this upgrade is to keep in mind that you're shifting, *not* reducing, the traffic that would formerly have stayed completely within the corporate network. A user who has 50 GB of files stored on a file server and then moves them to OneDrive for Business is still going to generate roughly the same amount of network traffic when she copies, edits, or moves a file, but the traffic's endpoints will be different. One possible



counterargument is that the movement of so many people to partial or permanent telework has offloaded some traffic from corporate networks—but now the experience for a user will only be as good as the quality of her home network... which opens an entirely different set of performance, security, and connectivity problems!

Beyond this basic shift in traffic, you can expect a move to the cloud to generate some extra external network traffic for things such as hybrid connectivity, directory synchronization, and remote administration. Of course, any migration of data to the cloud platforms will generate additional traffic too. Some of this traffic will be short-lived (for example, moving mailboxes from on-premises servers to Exchange Online), while other traffic may recur (such as directory synchronization traffic).

Your move to Microsoft 365 may also cause some brand-new traffic. One example might be new network traffic generated by Teams voice or video calls that traditional PSTN networks used to carry.

Anyone who plans to move to Microsoft 365 needs to pay attention to both their internal network and their Internet connectivity to ensure that enough high-quality network resources are available to handle the expected load. For instance, do you have enough capacity available to handle outbound traffic to Microsoft 365? Are sufficient network egress points in place to handle the outbound traffic or how many access points exist to process internet traffic? Adapting to the cloud's network requirements is often more a project management and budget issue rather than a technical challenge. As we'll discuss later in the book, Microsoft has published some [very specific actionable recommendations](#) for how you should tailor your internal and client networks to work best with the service, and it's critical to understand and adapt those recommendations in your network early on. The most critical recommendation is to reduce the round-trip time required for a packet to travel from a client to the Microsoft network; the devil is in the details of how you do this. In efforts to lower this round-trip time, companies often change their topology by adding content distribution networks (CDNs) or various types of cloud application security (CAS) capabilities.

Before you move any significant amount of your business to the cloud, it's a good idea to gather as much data as you can about the reliability and service quality of the Internet providers you use. Depending on where you're located, you may have multiple providers available to you. Whilst it's probably not possible to know with certainty how a provider's network will work over the long term, a move to the cloud is the perfect opportunity for you to talk to other organizations in your area to understand how your local providers have fared. Most of us have home Internet connectivity and know that different providers differ in speed, service uptime, reliability, customer service, and so on. Does your current business Internet provider deliver the quality and performance you need? If not, the time to look for alternatives is *before* you start your move. Because network performance is so critical, don't accept any migration plan that calls for you to start moving to the cloud until you're happy with the quality and performance of your network. Don't try to change network providers and move to the cloud at the same time.

Microsoft understands Internet connectivity is never perfect, and they want to limit the impact of network failures and slowdowns, so over the years, they've added features to let users continue to work over degraded or unavailable connections. Examples include "[Project Nucleus](#)," the sync technology behind offline and disconnected access to Microsoft Lists, Teams survivable branch appliances (which allows users to continue to make and receive PSTN calls from a local office when the Teams cloud service isn't reachable), and the ability of OneDrive, Outlook, and Teams clients to allow some types of changes to be queued while disconnected. You should expect Microsoft to continue to improve offline support for the Microsoft 365 ecosystem, although this improvement won't be equal across applications and workloads.

After more than a decade of experience with Office 365, it's become clear that the root cause of most user connectivity problems is a problem with the device or its connection to the internet, *not* problems with Microsoft's internal network or the service itself. In particular, the same kinds of DNS problems that have

bedeviled enterprise administrators since the birth of the commercial Internet are still around and still cause problems for users. Microsoft experiences occasional glitches with its network but has built-in enough capacity and alternate network paths to handle most issues.

In summary: the cloud is only as dependable as your ability to connect to it. Taking the time to review your network design and capacity to improve user connectivity will be time and money well spent.

### Is the Cloud Mature?

The second issue often raised by cloud doubters is the claim that the platform is immature. In other words, technologists better understand on-premises deployments because they have managed these deployments for so many years and have met and overcome all their challenges. This argument used to be much more credible when cloud servers were still new and used the same production code that we could run in our on-premises servers. However, after more than a decade of Office 365, this argument is no longer remotely believable; Microsoft has invested incredible amounts of resources in developing, provisioning, securing, and operating the Microsoft 365 services at a scale that no on-premises deployment can even attempt to approach, and with uptime that only a handful of organizations worldwide could deliver. The services now available in Microsoft 365 were purpose-built for the cloud and are deployed, managed, and monitored using high-scale automated tools supported by a server fabric that just does not exist in the on-premises world. That doesn't even begin to mention the huge wealth of features that Microsoft has put into the service but not into the surviving on-premises editions of their products!

Of course, fairness dictates that we recognize that there *are* excellent on-premises deployments of Exchange, Active Directory, and SharePoint that exhibit superb execution, operational maturity, and dependable delivery of service to end-users. On the other hand, many on-premises deployments *don't*. In fact, many on-premises deployments will need some degree of rehab work before they can consider moving to any other platform, whether that is to use a new on-premises version of their existing products, a hybrid deployment, or moving everything to the cloud.

From a maturity standpoint, you can't argue the fact that cloud platforms deliver highly functional software in the form of highly accessible applications across the Internet and do so with a level of reliability and security that would be the envy of many CIOs responsible for internal IT systems.

### Is the Cloud Secure?

No organization will move to the cloud if they believe that their data will be less secure than it is on-premises. When you move to Microsoft 365, you're placing a large bet that Microsoft will do a better job securing your sensitive data than you can do yourself. In this case "securing" means maintaining all three aspects of the security triad: confidentiality, data integrity, and data availability. Add to that the requirement to ensure data privacy, and the additional need in many jurisdictions to maintain [data sovereignty](#), and you quickly find that the degree of security expertise and investment required to maintain a truly strong security posture is outside the budget and capability of the vast majority of on-premises customers.

Microsoft invests roughly US \$1billion each year in security across its entire product line. This investment covers the typical "gates, guards, and guns" measures that most people think of for physical security, but the investment goes much deeper. All Microsoft 365 workloads are designed and implemented per the [Microsoft Security Development Lifecycle](#). This is described as "*a software development process that helps developers build more secure software and address security compliance requirements while reducing development cost.*" On a practical level, this means that data is secured by being encrypted at rest, that auditing information is available, and that the unified audit API is supported to allow customers and third parties programmatic access to information protection and compliance events. For instance, BitLocker is used to protect Exchange Online databases while files stored in [SharePoint Online and OneDrive for Business libraries are individually](#)

[secured](#) with their own 256-bit Advanced Encryption Standard (AES) key. Strict adherence to this standard costs Microsoft money; for example, even the small 3-5% performance penalty imposed by BitLocker encryption adds up quickly when you're talking about millions of disk volumes across hundreds of thousands of servers. However, Microsoft believes that spending heavily on security is a wise investment, as it builds customer trust and credibility for the service. To further build credibility, they make extensive use of their products, including the various Azure and Defender security offerings.

Microsoft does not usually provide specific details about the methods used to protect information because that could potentially undermine security by revealing too much to potential attackers. However, a reasonable amount of information, including several white papers that address the security topic in-depth, is available in the [Microsoft 365 Trust Center](#) and details of the data encryption technologies used by the different workloads [are available online](#). The completeness of the security features developed for Exchange Online and Microsoft 365 in general was enough to warrant the inclusion of Microsoft in the "leaders" section of Gartner's 2015 Magic Quadrant for Secure Email Gateways, a status Microsoft has held since.

The methods used to protect data include:

- **Physical security:** Anyone who has ever visited a Microsoft data center can attest to the extremely high level of security processes and procedures that apply. Access is strictly controlled to the servers and other hardware and every action is verified and logged. Any faulty disk drives are removed and demagnetized before they are destroyed to ensure that customer data cannot be compromised.
- **Reducing the number of administrators:** Microsoft has steadily been reducing the number of employees that have persistent administrative access to applications or the servers running in the data centers. They don't discuss the specifics often, but their goal was to make all administrative access "just in time" (JIT) so that it can be monitored and regulated. Restricting access in this manner ensures that any interaction with customer data is controlled (and audited). The so-called "Lockbox" process ensures that administrators hold zero standing permissions, which removes the problem of how to control the growth of those holding elevated permissions over time. Administrators receive just-in-time permission to perform tasks and have the permissions removed when the task is complete. All administrator access is conducted from specially secured single-purpose workstations known as [privileged access devices](#) that layer additional physical and data security capabilities on top of the existing Windows platform.
- **Customer Lockbox**, a premium E5 feature giving more customer control over access to user data like mailbox content, documents, and chats. Any time an engineer needs access to tenant data, [a request appears in the Microsoft 365 Admin Center](#) for a tenant administrator to review and authorize (or deny). [Customer Lockbox](#) works for Exchange Online, SharePoint Online, OneDrive for Business, and Teams. You can also buy Customer Lockbox as an add-on for any enterprise plan.
- **Extensive auditing:** Microsoft audits all access to tenant data to provide a traceable record of who accessed what data, when, and from where.
- **Data isolation:** Microsoft takes great care to ensure that data cannot leak from one tenant to another (see this document for more information about [how tenant data is isolated](#)).
- **Transport encryption:** Transport Layer Security (TLS) encryption protects information sent between Microsoft 365 data centers and clients and messages sent between Exchange Online tenants. Exchange Online also uses opportunistic TLS to negotiate secure connections with external servers. You can also create transport connectors that use TLS to protect messages sent to specific domains, such as those belonging to trusted partners. Any client, including browsers, which connect to Microsoft 365 services should use TLS 1.2. Microsoft [began to deprecate TLS 1.0 and 1.1 in 2020](#) and formally ended support for those versions in Exchange Online in October 2020, but the protocols still work, at least for now. It will take some time before Microsoft has completely removed earlier versions

from all Exchange Online servers worldwide. When this happens, Exchange Online will no longer send or receive emails from servers that want to connect with the earlier versions. In the meantime, they have added a [new, opt-in endpoint](#) specifically for older clients and services that need the older versions. This page gives the technical details about the [current TLS cipher suites](#) used for encryption. Keep in mind that some devices, such as Teams Room Systems, may require additional updates from their vendors to fully support TLS 1.2.

- **Data at rest encryption:** servers and data in Microsoft 365 data centers are protected by [BitLocker](#). Protection covers both Exchange Online mailbox databases and SharePoint Online and OneDrive for Business document libraries. That means that rogue administrators cannot remove customer information and sell it to other parties or otherwise misuse the content. See [this page](#) for information about the encryption of SharePoint Online, OneDrive for Business, and Exchange Online data.
- **Customer-applied encryption:** Users can choose to protect or encrypt messages using sensitivity labels (Microsoft Information Protection) or S/MIME. To be complete, you might include features such as Data Loss Prevention (Chapter 19), Microsoft 365 Defender for Office 365 (Chapter 7), and information protection using sensitivity labels (Chapter 20) in the overall assessment of security capabilities. Microsoft also offers large enterprises the ability to provide their own security keys which can be used to encrypt data (known as “customer key” or “bring-your-own-key” encryption, discussed in Chapter 4).
- **Endpoint protection:** not only does Microsoft offer its Defender line of products as part of Microsoft 365, but they also run Defender on the servers hosting the workloads, as well as the workstations used by engineers, administrators, and support personnel. In addition, if you subscribe to Windows 365, those Microsoft desktop environments are protected automatically as well.
- **Azure Active Directory security:** Application sign-in and authentication depend on Azure Active Directory, which has its own set of protections. For more information, see [this whitepaper](#). In addition, features such as [Privileged Access Management](#) allow tenants to control administrative access to PowerShell cmdlets that create or edit data. Administrators must request access to specific functions and give a reason for the access. After review, the access might be approved, in which case the administrator can go ahead and execute the task.

Besides these designed-in security measures, there’s one other additional trick Microsoft uses: they automate server configuration and management using Desired State Configuration (DSC). Automation ensures that cloud servers run known configurations and that newly added servers are correctly baselined from the start. The automated and ongoing updating of servers with new versions of the operating system and applications means that a high degree of confidence exists that known bugs are not met in production.

With all this said, it is critical to keep in mind that Microsoft owns the security of the cloud resources and services that you use but *you* are still responsible for the security of your servers, services, and users. As the world saw with the [Exchange Server “HAFNIUM”](#) and ProxyShell attacks throughout 2021, it is *never* safe to assume that you’re immune to attacks; in the case of HAFNIUM, even organizations with no on-premises mailboxes were still vulnerable as long as they maintained at least one Exchange server on-premises. Don’t let the potential of a move to the cloud lead you to skimp on your security practices, procedures, or investments.

## Will I Get Good Service from the Cloud?

Quality can be measured in diverse ways. One way is by measuring how well a service meets its service level agreement (SLA). Microsoft has an excellent track record of meeting and exceeding its SLA for the Microsoft 365 services (although in fairness, they get to unilaterally define the metrics in the SLA). Some argue in favor of counting the number and impact of software flaws as a measurement. This is great for the organization creating the software, but it’s a terrible measure for outsiders (which means “all of us”) who don’t have full

visibility into the range, scope, and the number of bugs. It's a fun fantasy to imagine that any software vendor would eliminate all its bugs, but that's not going to happen for four reasons. These are:

- **Competitive pressures:** Microsoft must respond to technical and feature advances made by its competitors to ensure that its applications remain an attractive offering in the market. Every time you touch code to improve functionality or "make things better," the possibility exists that bugs creep in.
- **Customer demands for "more":** Customers also ask for many different forms of enhancements. Requests vary from a slight change to the way some feature works to an update designed to facilitate the onboarding of large enterprise customers. Microsoft monitors forums, Twitter, and support tickets closely to focus on areas that cause users to have problems and makes changes to improve matters. An example of this is how many changes were made in the Teams application to help customers cope with the upsurge of home working because of the Covid-19 pandemic.
- **A dynamic client base:** A staggering number of client platform combinations (browser, Windows, mobile) and protocols interact with Microsoft 365. Clients have been known to introduce problems (for example, several Apple iOS upgrades have caused issues for Exchange Online), and the support and testing matrix is very complex. Testing complexity creates more possibility for bugs to creep into software and stay undetected until a user meets a problem. It's also the case that sometimes Microsoft clients can cause problems —so the test burden is on the applications *and* the service.
- **Dependencies between services:** Although email is often the first workload moved to the cloud, it is important to understand that many interdependencies exist between what may appear to be separate workloads. A better and more productive view of Microsoft 365 incorporates Exchange Online, SharePoint Online, Teams, OneDrive for Business, Azure Active Directory, Microsoft Information Protection, the Microsoft Graph, and so on. Behind the scenes, a complex set of interconnections link these services. For example, Groups are based on features contributed by Exchange Online and SharePoint Online. eDiscovery depends on features contributed by Exchange Online and SharePoint Online (among others). SharePoint Online is the storage for applications such as Teams while Azure storage supports other applications like Planner and Teams. Each of these workloads and features has its own set of developers, testers, and plans. Maintaining the connections between the different applications requires substantial effort, which is one of the reasons why some of these features are not available on-premises.

It's hard to eradicate software bugs when software evolves and flexes all the time. We accept this situation because we like new features and functionality. Remember that one of the reasons why customers adopt cloud services is the promise of "evergreen technology", to move away from the more restrictive upgrade cycles used in on-premises deployments. With evergreen or ever-evolving technology comes the risk that things don't work quite as well as they should from time to time. That's when you enter the world of cloud support.

### Will I Have to Give Up Control When I Move to the Cloud?

Consumers of cloud services generally don't get to vote on when the service provider introduces, changes, or removes functionality. It just happens. Cloud providers like Microsoft make these changes to meet *their* goals and priorities rather than yours. This is part of the deal you're making when you move to the cloud. Moving to Microsoft 365 means that you're giving up a degree of control, but you're also able to shed the responsibility for installing the latest versions of Windows and Exchange on a server, bringing it up to the latest patch level, and making sure that any required additional products are installed.

Sometimes problems arise when the characteristics of the service change in a way that you prefer it did not. For instance, as part of the replacement of Skype for Business Online with Teams, Microsoft at one point took away the ability for smaller tenants to choose between Skype for Business Online and Teams and forced them

into Teams. You can argue that this was the right decision because Teams is newer and better technology, but some tenants didn't like the feeling of being forced. Microsoft also has an unfortunate habit of introducing features that admins may not want and enabling them by default, as they initially did with self-service license purchasing for Power Platform or the Bing search add-in for the Edge browser.

The positive way of looking at how things flex and change without warning or any ability to influence an outcome when you use cloud services is to not focus on any loss of control over the servers and other infrastructure that you no longer manage. Instead, focus on how best to use all the added hours that you gain and figure out how to use that time more productively.

## Answering Some Additional Questions about the Microsoft 365 Cloud

Besides the basic (but necessary and challenging) work of dealing with FUD questions, some additional considerations may come into play depending on your industry, background, and regulatory requirements.

### What Industry Security Standards Should I Consider?

None of the practices used by Microsoft are secret, and they aren't rocket science. What they *are* is implemented at massive scale and with huge attention to detail—but the basic practices they use to protect and secure data in Microsoft 365 applications are well-known and can be implemented by any company (and indeed, Microsoft [documents how they handle security incidents](#)). (Of course, some of the lessons Microsoft learns from their massive environment don't translate to on-premises environments!)

However, Microsoft realizes that customers want to see independent audits proving that they are following good security practices. Among the other audit-based attestations and certifications that Microsoft has earned, we have:

- [ISO 27001/27002](#): Information Management Security System and Best Practice for Security Controls.
- [SSAE 16 SOC1 Type II](#): Reporting on Controls at a Services Organization.
- [NIST 800-53](#): Security and Privacy Controls for (U.S.) Federal Information Systems and Organizations., Microsoft's audit controls for Microsoft 365 are based on the NIST 800-53A (Rev. 4) special publication.
- [HIPAA](#): U.S. Health Insurance Portability and Accountability Act.
- The [HITRUST](#) Common Security Framework.

A complete set of documents describing how Microsoft meets regulatory requirements and security standards can be found in the [Service Trust Portal](#), including the latest audit reports covering Microsoft 365 and other Microsoft cloud properties. Another [set of documents](#) helps customers perform a risk assessment and understand the compliance of Microsoft cloud services with industry standards and regulations. Microsoft also publishes information about how Microsoft 365 meets requirements that exist in specific geographies, such as the [European Union model clauses](#), which are *"standardized contractual clauses used in agreements between service providers and their customers to ensure that any personal data leaving the EEA will be transferred in compliance with EU data protection law and meet the requirements of the EU Data Protection Directive 95/46/EC."*

The European Union's General Data Protection Regulation (GDPR), in use since May 2018, affects how companies active in the EU gather and handle data. More information on GDPR and how it relates to Microsoft 365 [is available from Microsoft in the service trust portal](#) and Chapter 17. The [Microsoft Compliance Manager](#) gives customers a framework to help them organize the steps to achieve compliance with regulations like GDPR.

In addition, Microsoft sometimes seeks external proof that its applications meet the specific requirements of different industries. For instance, in the financial sector, they commissioned a [report](#) authored by [Cohasset Associates](#) to give an independent opinion as to why the compliance features built into Microsoft 365 meet the electronic records storage, retrieval, and management requirements of rule 17a-4 of the U.S. Securities and Exchange Commission (SEC), Rule 4511 (c) of the Financial Industry Regulatory Authority (FINRA), and regulation 17 CFR of the Commodity Futures Trading Commission (CTFC).

Microsoft also requires that its suppliers meet certification standards. They have a certification program, the [Microsoft Supplier Security and Privacy Assurance](#) program to require vendors who sell software to Microsoft to prove their adherence to privacy and security practices. Within and outside Microsoft's ecosystem, it's important to remember that no software exists in a vacuum; all the components that services like Exchange Online depend on should also receive external oversight and certification. For example, [Azure](#) is certified against the [ISO/IEC 27018](#) code of practice for the storage of personally identifiable information in public clouds.

**Security across the board:** Although Microsoft assigns talented personnel to the tasks of building security solutions and protecting customer data, there is no doubt that sometimes some extra help is necessary for organizations to achieve full protection. For example, tenants can deploy data loss prevention technology for many workloads, but this technology does not protect all the data that exists within companies. It might therefore be necessary to deploy some other solutions to ensure that sensitive data is not transmitted outside the organization. When a company moves workloads to the cloud, it is wise to consider whether some more protection is needed for the full spectrum of data that is in use.

Companies that need more security than the capabilities built into the Microsoft 365 applications have a wide range of options available to harden different components. [Microsoft Defender for Cloud Apps](#) is available to protect cloud applications at an extra monthly cost per user. Other companies that are active in the provision of additional security capabilities include [Mimecast](#), [ForcePoint](#), [Proofpoint](#), and [Skyhigh Networks](#).

## Which Cloud Am I Moving To?

Most of the time, we think of Office 365 as a single cloud service delivered worldwide by Microsoft. It's more than that, of course; there are the other services that Office 365 leverages, such as Azure AD, but there are also multiple independent instances of Office 365 offering what appear to be very similar services, but for specific regions or markets. The best-known example of this is probably the family of Office 365 solutions for government use. The basic Government Community Cloud, or GCC, is a separate instance of Office 365 that is dedicated for use for "United States Federal, State, Local, and Tribal governments, as well as contractors holding or processing data on behalf of the US Government," according to the [GCC service description](#). It's intended for unclassified information but is approved for use with various types of other sensitive information, including tax and financial data and certain types of sensitive law-enforcement-related information. Some of the approvals required to process certain data types (such as [the CJIS certification](#) required for law enforcement use) require special configuration, which of course Microsoft will assist with as part of a consulting engagement. There's a separate GCC instance (known as GCC High) that is approved for handling certain types of sensitive defense information, plus a [US Department of Defense \(DoD\) instance that has its own separate Office 365 cloud](#) implemented in physically separate data centers that have little or no infrastructure in common with the commercial and GCC versions. In late March 2022, Microsoft announced the forthcoming availability of a new cloud, Office 365 Government Secret, that's approved for processing certain types of classified information, joining the Azure Government Secret and [Top Secret](#) clouds.

Figure 2-1 shows the current set of publicly-acknowledged Office 365 clouds for government use.

# Office 365 Government:

Meeting the full spectrum of government data needs

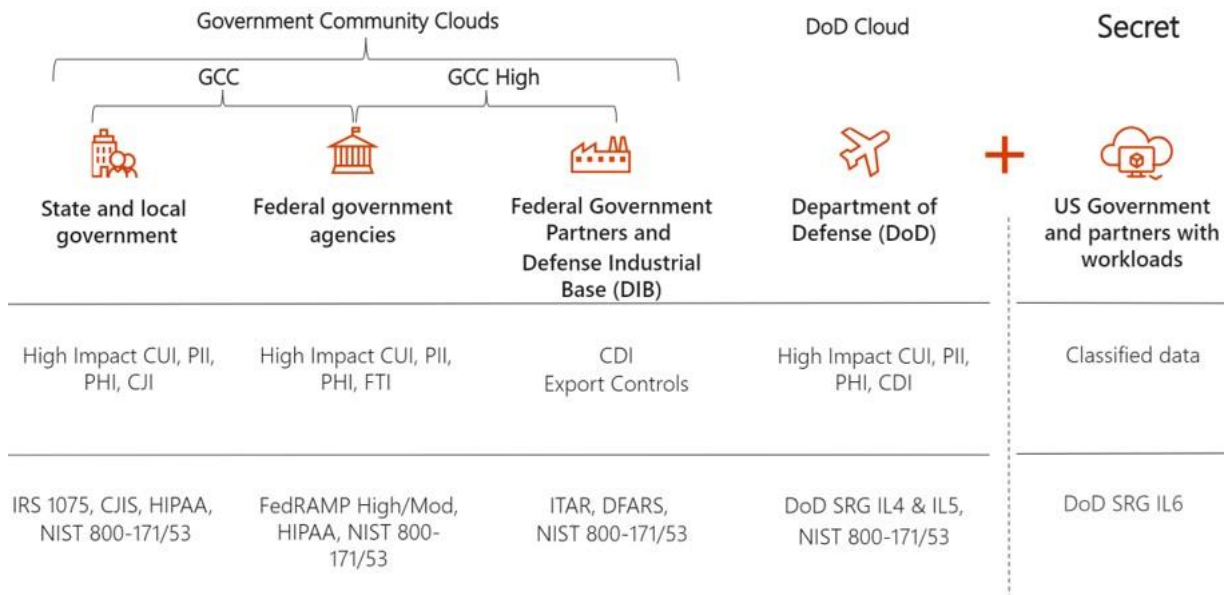


Figure 2-1: the 2022 version of Microsoft's government clouds for Office 365

Other government-specific Office 365 instances exist. To use any of these government-oriented services, you must be sponsored by a government agency and approved by Microsoft. This allows companies that contract to provide services to agencies that use GCC to have GCC tenants also, assuming that their customer agency is willing to sponsor them.

The Chinese government requires that most cloud services doing business in China have local instances that are logically separate from their overseas counterparts. These instances must be maintained by Chinese citizens in Chinese data centers. In the case of Office 365, this service is known as "[Office 365 operated by 21Vianet](#)" (21Vianet being the Chinese company that runs the service for Microsoft). Microsoft carefully emphasizes that they don't operate or manage this service. There used to be a similar service for German customers known as [Office 365 Germany](#), too, but as Microsoft added data centers in Germany, that variant went away. It's interesting to wonder about how much work Microsoft does to create a "containerized" version of the basic Office 365 services for deployment in a new country and to ask whether this will ever become a feature available to commercial customers.

Wherever they're located, each service instance has its own set of supported features. When Microsoft introduces a new feature into "the service," they're really introducing it into dozens of different environments. For example, as of July 2022, Office 365 in China supports Power Apps and Power Automate only for selected customers, and with some missing features (such as the Power Automate mobile app); various Teams features available to commercial customers in North America are still rolling out to tenants homed in other countries and Office 365 Education tenants; [Yammer isn't available in GCC High](#), and file sharing using SharePoint and OneDrive works differently in GCC High than in the lower-security GCC and commercial tenants.

Besides the obvious complications of having different sets of services and functionalities in different cloud environments, many third-party applications and services that work with the commercial Office 365 services may have incompatibilities, or not be supported, within some environments. There's a complex process to get approval to run third-party applications in GCC and GCC High, for example, so you may find that your preferred migration tool, management utility, or other application cannot be installed or run in your



government tenant, or that, when installed, it doesn't work properly. Be sure to check compatibility with your vendors before moving to one of these environments.

The bottom line for these differences is that if you are considering moving to anything other than the standard commercial service, you should take the time to research the differences in features and service capabilities between the well-known baseline services in Office 365 and the specific version that you'll have access to.

### Who's Responsible When the Cloud Breaks?

It's quite tempting to assume that, given all the work and money Microsoft puts into making Microsoft 365 secure and robust, you no longer need to worry. As we saw with the previously-noted HAFNIUM attacks, that's not the case. In fact, even if there isn't an active attack against your organization, you should be aware that Microsoft expects you to bear responsibility for a good part of the overall whole. To be specific, [Microsoft says](#) *"For all cloud deployment types, you own your data and identities. You are responsible for protecting the security of your data and identities, on-premises resources, and the cloud components you control (which varies by service type). Regardless of the type of deployment, the following responsibilities are always retained by you: Data, Endpoints, Account, and Access management."*

Microsoft's position is clear, although you may not have thought of it in those terms. Microsoft [expects that you are responsible](#) for data classification and marking, securing your devices and endpoints, and managing and securing your user identities and associated credentials. You may choose to use Microsoft's tools and services (whether free or paid) to help you with these things, and Microsoft may even assist in some ways (as they do by providing sign-in data and warnings of unusual sign-in activity to you), but they deny responsibility for those aspects of the end-to-end service. That means that you must stay aware of both security threats and potential problems that might interrupt your end users' work, then be ready to remediate any of them that fall under your control.

Given the increasing number of ransomware attacks against hybrid organizations, it's also worth mentioning that attacks (or failures) that knock your on-premises Active Directory infrastructure offline will have an outsize impact on your users' ability to work in the cloud unless you are 100% moved over to cloud-based identities. Part of your planning for a move to the cloud must cover recovery or resilience strategies to let people keep working if any important on-premises infrastructure is compromised or degraded.

### Can I Get Good Support for the Cloud?

Microsoft 365 tenants can get both free and paid support through multiple channels, including forums, telephone and email support, and service requests that flow to dedicated support personnel. Cloud observers often refer to support as the "Achilles' heel" of cloud services because it is so difficult to deliver support in a timely and effective manner to tens of millions of consumers. Many consumer-focused services (such as Netflix, Spotify, and Microsoft Xbox Live) force their users to seek support through self-service forums and community-based support. Now that Microsoft 365 counts more than 340 million active users, the temptation for Microsoft to do the same is strong.

By comparison, on-premises services are usually supported by in-house help desks and support teams (or the work is outsourced to a third party). A company has control over the quality of those services based on the investment made to staff up the help desk and train support personnel. The transition from a position where you have total control over all the moving parts needed to resolve a support ticket to depending on the intervention of cloud support when you do not control many of the moving parts is difficult for many organizations. IT professionals are frustrated with having to deal with phone support, users do not like how long it takes for even simple issues to be resolved, and managers worry about the control they have ceded to the cloud provider.

All of this is true. When you move from a customized in-house solution to use a shared utility infrastructure you must make some adjustments. Think of it like this – if you use a diesel-powered generator to power your house, you have total control over the infrastructure but you’re also responsible for everything that might go wrong with it. But when you sign up with a power company, you hand over responsibility for most of the operations to that provider. All you must do is make sure that the equipment is plugged in and turned on. The same is true of using a utility IT service. You don’t have to maintain the servers or make sure that their networking equipment is functioning, but you must ensure that your equipment is functional. (The interesting extension to this metaphor: some businesses have such specific needs for electricity that it makes sense for them to run dedicated generating plants even though utility-scale electricity is cheap, dependable, and ubiquitous throughout much of the world—this explains one reason why some companies cling so fiercely to their on-premises environments!)

When workloads move to the cloud, the job of the in-house support team changes. Instead of driving a resolution for every problem, the in-house team must manage interactions with Microsoft to work the problem as effectively as possible within the constraints placed on both you and the Microsoft support organization.

You control local equipment and settings and can make sure that you cover all the basics – that the user’s client is configured correctly, that they can connect to the Internet, and whether the problem is unique to a certain user or is shared by a group, and so on. Microsoft’s support, delivered remotely, will focus on making sure that the service is working properly and assist with troubleshooting when they can. This division will usually be invisible to the end-user unless they learn about it when they are asked to check some local setting or variable on their device. Microsoft owns the responsibility for recording the support ticket and following through until the issue is finally resolved. But you should remember that Microsoft support professionals cannot make changes within a data center. All they can do is record, probe, and diagnose problems. Sometimes that will be enough to get a resolution and sometimes they will need to escalate from first-level support through second-level support right up to the elevated heights of the few who can make changes that affect how an application works if that’s what is required to fix a problem.

Progressing a problem through many interactions with different levels of support takes time. Multiple phone calls and email interchanges are likely going to be necessary to keep an issue moving toward resolution. Don’t blame the support personnel – they follow a strict playbook to gather information, suggest fixes, and escalate when required. Think of how you’d function if you had to understand the myriad of different circumstances that customers who report problems are in. It sometimes takes a lot of time to simply understand what’s going on before you can move toward resolution.

Microsoft uses a mixture of subcontractors and Microsoft badged employees to provide first-level support ([access is described here](#)). First-level support handles incoming calls. Like every software-as-a-service vendor, Microsoft wants to resolve as many calls as possible on the first contact, and ideally, that first contact will be handled using automation. They have an elaborate process for call handling. This starts when calls are registered in Microsoft’s ticketing system and a structured approach is taken to recording details of the incident that can lead to a fast resolution if the problem and its solution are already known. Calls that cannot be resolved by first-level staff are escalated to second-level support. These individuals usually possess deeper product knowledge and are better able to diagnose and resolve complex support situations. Because Microsoft uses a [DevOps](#)-based model for its cloud systems, the most difficult calls eventually escalate to the product group responsible for the area where the fault lies. At all times, customer confidentiality and data privacy are respected.

When you report a problem to Microsoft, it is critical to keep careful notes of the date and time for each call or other contact and record to whom you spoke, what was discussed, what actions should occur, their

expected outcome, and when this should happen. Apart from simply being good sense to record information about support incidents, this information will be invaluable if you need to request an escalation or wish to review why a problem is not being resolved, perhaps by contacting the supervisor of the support professional with whom you are working.

Communication with end-users is also critical. They probably won't realize, or care about, the details of the complexity of cloud services, nor should they be expected to understand the links between the client software and an application. They care about being able to get their work done. Making sense of the complexity and managing support calls is the new role for the local IT department. When you communicate with them, try to do so frequently and in clear, plain language, including a timetable for when you will be able to give the next update. Keep in mind that you may need to establish more than one way to notify users of outages or problems since you can't automatically assume that email will be available.

**Limited responsibilities:** Microsoft makes it quite clear what they will and will not support when it comes to calls. Basically, they won't support anything having to do with your networking infrastructure, your hardware, any third-party software, or any Microsoft software that isn't part of Microsoft 365. They also specifically disclaim that they won't support your operational procedures or service management processes.

## Making the Decision to Embrace the Cloud

So far, we have concentrated on removing some of the FUD that surrounds cloud services. Now let's move on to consider a few scenarios—some where the decision to embrace the cloud and move to the cloud is easy, and some where complications might make things a little more "interesting".

### Reasons Why Moving Might Be Easy

As you might expect, the decision to move to the cloud is much easier for some companies than it is for others. Broadly speaking and accepting that these are high-level generalizations, the characteristics of companies in this category include:

**Start-up companies:** Why would any start-up want to create its own IT infrastructure when most, if not all, of the services they might need to use are available in the cloud? Apart from anything else, buying at a fixed cost per month is usually better for cash flow. The (rare) exception to the rule is companies that have been created to develop a product associated with an on-premises server product like Exchange. In this case, it is a good idea to eat your dog food and run some Exchange servers.

**Small to medium companies:** Anecdotal evidence shows that much of the initial migration to Office 365 came from companies with less than 500 employees, many of whom had run older generations of Microsoft servers such as Windows Small Business Server. Replacing old servers (some that resided in the proverbial cupboard or under someone's desk) with the promise of evergreen functionality offered by the cloud is very attractive, especially when the migration of mailbox data is relatively straightforward as the volume is small enough to accomplish a smooth switchover over a weekend. These companies tend to be "all in" the cloud and don't need to operate hybrid environments so they aren't interested in aspects such as federation and single sign-on. Microsoft provides a [FastTrack deployment service](#) for Azure, Microsoft 365, Office 365, and Dynamics to help companies like this that purchase 50 or more licenses (the exact number varies according to [what you buy](#)). The fact that Microsoft considers that many of the onboarding operations necessary to move these companies to the cloud can be executed on an almost factory-like production-line basis indicates the relative ease of the switchover, especially for small to medium companies whose needs do not extend past email. More complex email scenarios (for instance, those involving specific regulatory oversight) and SharePoint migration usually need the services of a specialized migration partner.

**Fortune 500 companies:** Including Fortune 500 companies after small to medium companies might seem strange, but the fact is that most if not all the Fortune 500 companies have some element of work running on Microsoft 365. Sometimes a tenant domain is taken to reserve it for future deployment. Sometimes it is used to test the waters for a planned move to the cloud. And sometimes it is because the company is in the process of moving some, and maybe all, of its workload to Microsoft 365. Hybrid deployments are the usual rule of thumb for very large companies in this category because this approach is most flexible and allows transparent co-existence across the two platforms, including a shared directory, integrated mail flow, and access to data.

**Moving from a non-Microsoft email system:** Microsoft supports the migration of IMAP4-based email systems (including Gmail) to Exchange Online. Older POP3 accounts are more of a challenge because these typically must go through an interim step where messages are downloaded to a PST with Outlook and then imported into Exchange Online, again using Outlook. It's a manual process that can be tiresome. Migration from other email systems such as Lotus Notes or Novell GroupWise can be done using separate tools bought from companies that specialize in the migration area. Table 2-1 lists some of the well-known companies in this space. Take the time to download trial versions of their software to establish the tool that is the right choice for you.

<b>Company</b>	<b>Migration Tool</b>
BitTitan	<a href="#">MigrationWiz</a>
Code Two	<a href="#">Office 365 Migration</a>
Quest	<a href="#">On Demand Migration Migration Manager for Exchange</a>
Skykick	<a href="#">Enterprise Migration Suite</a>

Table 2-1: Email Migration Software

**Strong needs for flexible scaling:** some organizations have strong requirements to scale up or down based on demand. One obvious example is schools and universities, which have long experience in moving from full operations during the academic year to a reduced mode (or even shutdown) over holidays and breaks. Other examples include businesses with high numbers of seasonal or temporary workers. Microsoft's Windows 365 offering is designed to play to this market, but the overall flexibility of Microsoft 365 serves it as well by greatly reducing the requirement to plan, buy, deploy, and manage on-premises infrastructure to provide collaboration tools to users.

**Educational institutions:** Many universities and colleges used earlier versions of Microsoft cloud-based email or have experience running a competitor's cloud-based email. These institutions find it easy to move from on-premises servers to the cloud.

Again, these are generalized comments and individual country-level markets differ in terms of the mix of accounts that have moved to Microsoft 365.

### Reasons Why Moving Might Be Difficult

The categories listed above are examples of organizations that are relatively easy to move. Now let's look at some of the circumstances that can complicate matters. They include:

**Multi-location or multi-national:** It is much easier to coordinate the movement to Microsoft 365 if everyone (users and IT staff) shares a common location. Planning becomes more complex as the number of locations grows and more complex still if multi-national locations are involved. The needs and legal requirements of each country, including security and privacy, should be factored into a migration project. Some applications (including Exchange Online, SharePoint Online, and Teams) have multi-geo capabilities, but others (like Planner and Yammer), do not.

**Tight integration between applications and business processes:** Many companies have integrated email with business processes in some way. For instance, when HR registers a new employee, a mailbox is created automatically, and a welcome message is sent to the mailbox. Exchange Online supports Exchange Web Services (EWS), the Graph API, and PowerShell, so it might be possible to move existing email-enabled processes to the cloud. The same issue might occur if SharePoint is integrated with some business process, but the APIs used on-premises can't be used or are more limited in the cloud. Microsoft's longer-term direction for extensibility is based on the Microsoft Graph covering all its cloud workloads. Alternatively, you could keep some on-premises servers to handle applications that cannot be moved easily into the cloud.

**Large tenants:** The more users are involved, the longer a migration will take. Migrations are never popular exercises as they involve lots of repetitive operations (moving mailboxes or documents) that are prone to problems (corrupt items and other issues) that take a long time to achieve, especially when moving mailbox data across the Internet. Given the right experience, good project management, and network connectivity, it is possible to move several thousand users per week – but it is not easy.

**Capable of running a private cloud:** Small companies have small IT departments and the people working in IT there have to be jacks-of-all-trades who are expected to be able to work with many different technologies. Larger companies can afford to have specialists and specialization tends to enable a more sophisticated IT environment. These companies might have the necessary expertise and operational maturity to be able to run an internal cloud that can host on-premises versions of products like Exchange at a price point comparable to the costs of using the cloud equivalent. On the downside, more functionality is available in the cloud; Exchange Online and SharePoint Online have many features that will never be available on-premises and there are of course no on-premises versions of Yammer, Teams, Stream, or Planner.

**Poor network links to the Internet:** Sometimes we forget that not every company enjoys high-speed and reliable Internet connections. It might be the case that offices are stuck at the end of an extended hub-and-spoke network that simply needs to be upgraded or it might be that poor connectivity is a fact of life in some or all locations. Microsoft 365 users depend on the Internet to connect to Microsoft data centers. The exact bandwidth and latency required are dependent on the client mix (Outlook creates more demand on the network than OWA does, for example), the number of users, and working patterns (always size for peak demand plus contingency). Remember that you also must move user data to the cloud, a process that won't be quick if the network lags.

**Old clients:** As we'll discuss in more detail later in the book, Microsoft defines specific requirements for clients to connect to the service. Exchange Online only supports certain versions of Outlook and browsers, and they are rapidly moving away from supporting legacy protocols such as basic authentication. Organizations with older versions of clients or devices in use should factor upgrades into the migration. Deploying a new version of the desktop Office applications, performing a complete desktop refresh, or replacing all your old multifunction printer/scanner devices is expensive and takes time and attention to detail to ensure no disruption of user productivity. The Microsoft 365 enterprise apps use streaming and virtualization technologies to simplify the installation process; this is an excellent way of addressing the issue of outdated desktop software as it allows workstations to download and install updates automatically.

**Recent on-premises upgrade:** A company that has recently (in the last two years) upgraded its IT infrastructure might not have the appetite to go through the additional step of moving some workload to the cloud. They want to get a return on their on-premises investment. However, this is becoming a less common issue as Microsoft slows down the rate at which they release new versions of on-premises software.

**Special regulatory or operational circumstances.** Organizations that operate in highly regulated sectors, including financial services, life sciences, and aerospace, may have specific requirements that make the cloud inhospitable. Microsoft has steadily been broadening its offering to help address regulatory needs for things

like compliance with US financial reporting laws and EU data privacy laws, but there still may be very specific scenarios where the cloud doesn't support a government or other requirement.

**Institutional distrust.** Some people, and organizations, distrust the fundamental concepts of cloud computing. For example, I've talked to multiple organizations in the European Union who are reluctant to move to Office 365 specifically because it is operated by a US-based company, and they distrust Microsoft's ability to resist demands for access from intelligence agencies either in the US or in their own countries. You can certainly argue whether this distrust is reasonable or rational, but it absolutely exists, and for some organizations, it may be the key factor that slows or blocks their movement to the cloud.

Of course, other complexities exist that are not on this list. For instance, any company that is already running on a hosted email service such as a managed service based on on-premises Exchange will have their own difficulties to overcome.

## The Unique Challenges of Tenant-to-Tenant Migrations

Most of the time, migrating to the cloud is conceptually straightforward: you have some on-premises data and you want it in the cloud. The details of how you accomplish that goal can be quite complex, but the underlying idea is simple. That's not the case for migrations where you want to move data from one Office 365 tenant to a different one; these "tenant-to-tenant" (or just T2T) migrations can be frighteningly complex for a variety of reasons.

You might wonder why anyone would do such a thing. The biggest driver of T2T migrations is merger, acquisition, and divestiture (MAD) activity. For example, it's common in the pharmaceutical industry for company A to buy company B (a merger) and then spin off some part of the company to form a new enterprise, company C (a divestiture). Tenant consolidation is also common; companies that grow through acquisition (as is common in the media and IT industries) don't want to have separate tenants for each acquisition, so part of the process generally involves consolidating data and users from the acquired companies into a single entity. The largest of these that I've personally been involved with was for a large media and advertising company that moved 135 separate Office 365 tenants, with more than 110,000 total users, into a single large tenant—a challenging process at the best of times.

Microsoft does not yet have a single consolidated toolset for performing these migrations, so you will need to choose a third-party tool that meets your needs (such as Quest's [On Demand Migration](#) tool). Some of the issues you should consider when planning a T2T migration include:

- **Workloads:** Migrating Exchange mailbox data and OneDrive files is pretty simple. Past that, workload migration rapidly gets more complex and difficult. For example, if you have extensive SharePoint customizations or lots of Power BI data, flows in Power Automate, or videos stored in Stream, you'll find that different workloads have wildly different levels of support. Part of this is because Microsoft has not provided complete import, export, or migration APIs for some workloads (notably Power Platform and, regrettably, Teams); part of it is because the more complicated the workload, the more individual components or repositories are involved (Teams again).
- **Planning:** as with any large-scale migration, understanding what you have is a critical step in planning and scoping your migration. As you'll see in Chapter 4, the built-in reports Microsoft provides are limited in flexibility and scope, and you'll probably find that they're not usable to plan your migration. Quite apart from that, the logistics of migrating data from one tenant to another in a way that allows users to continue their work are challenging themselves.
- **Dividing the data:** consider a divestiture, as when General Motors sold the assets of its Saab business unit to the Dutch carmaker Spyker. Someone had to decide, for every user, file, and mail message, whether to include that item in the divestiture. Of course, most of this work is automated and pro-forma—if a user named Joelle Smith is moving to the divested unit, the normal assumption is that all

her files and messages should go with her. However, determining what to do with data in SharePoint libraries, OneDrive libraries belonging to users who have left the company before the divestiture or will be laid off as part of the move, and shared repositories in Teams, on file shares, or in Exchange public folders can be a daunting problem of its own.

- **Soft issues:** a surprising number of T2T migration efforts become hung up on completely non-technical points. For example, the 135-to-1 migration I mentioned earlier featured some passionate fights about the naming of conference rooms in the Exchange Online room resource list; each faction had arguments and supporters for the naming scheme they used, and resolving the issue required executive attention. Now multiply that by the number of political, reputational, and control issues we've already discussed, and you'll begin to get a sense of the scope of problems that can arise, and for which migration tools cannot provide any help or even guidance.

## Following the Money

Now that we understand the various categories of companies that might consider moving to Microsoft 365, let's consider the money question.

On the surface, saving money seems like an excellent justification to move to the cloud. Many commentators focus on the fixed per-month subscription fee and compare this to the all-in cost (if it is known) to run an in-house system. Salespeople who want to push the virtues of the cloud love to talk about the fixed price per month, especially if they can get you to focus on the lower-cost plans. Once you start thinking that you can buy cloud services for \$6/month per user, that impression fixes itself in your mind, and a certain view that "the cloud is cheap" forms. It can be very difficult to change this impression, even when you realize that you need one of the more expensive plans and will pay \$23+/month per user.

Invariably the comparison is positive for the cloud because organizations often do not fully understand the complete cost picture and therefore do not factor all the cost elements into the equation. It's true too that the notion of paying a fixed monthly cost is attractive because it allows the company to manage cash flow more precisely than the sometimes-erratic spending patterns of IT. It is also clear that cloud systems are more flexible than in-house systems when the time comes to add or reduce capacity.

## Figuring Out Cloud Costs

You will never quite know what the actual cost base for a cloud deployment is until after you complete the deployment phase (migration of users and applications) and operations stabilize. Only then will you know exactly how much the cloud is costing and where those costs lie.

Before you decide to embrace the cloud, it is wise to figure out exactly what the current on-premises environment costs—assuming you have one. You obviously can't make a 1:1 comparison for workloads such as Teams that don't have an on-premises equivalent. You will need data about your on-premises environments to compare against cloud expenditure to show any savings. Let's discuss the typical cost buckets that you should consider, beginning with an understanding that most organizations don't have a clear understanding of exactly what they spend in any of these categories:

- **Hardware:** An on-premises deployment likely contains several different server roles, including Active Directory, Exchange servers, mail hygiene or bastion servers, SharePoint servers, Skype for Business servers, administration systems (workstations and servers), and servers for ancillary services like backup and restore or network monitoring. Some of these servers might run as virtual machines on a hypervisor (like VMware or Hyper-V). You need to capture details of all the servers involved in the workload to move to the cloud and figure out their annual running cost, including capital depreciation, warranty and support costs, finance charges for leased equipment, replacement cost, power consumption, cooling, and other data center expenses. You'll also need to consider that,

thanks to the normal evolution of things, you might be decommissioning some of those servers anyway (such as Skype for Business servers).

- **Software:** although you will probably save a good bit of money on server OS and application licenses, if you remain in hybrid mode your server count won't go to zero. You may still end up with on-premises servers for Active Directory, directory synchronization and federation, service monitoring, message hygiene, and archiving. Some of these servers might be replaced during the migration but some might not. Of course, you also must buy client access licenses (CALs) for many of Microsoft's services, which means you'll need to map the CALs you currently have for user services to the equivalent Microsoft 365 products. Remember to include all on-premises functionality in your calculations and then find what plans best match your needs. Standalone Exchange Online and SharePoint Online plans exist for those who do not need the full range of functionality available in a plan like Office 365 E3. Kiosk (web-only) plans exist for "frontline" workers, people who do not have workstations and only need intermittent access to applications like Outlook or OneDrive. You must decide whether you need functionality like Defender for Cloud Apps, Advanced Threat Management, or Advanced Compliance to justify buying Office 365 E5, and how many people need such a plan. Another thing to consider is whether it is cheaper to buy a lower-cost plan and then buy add-on components to make specific functionality available to selected users.
- **Legacy Microsoft Office applications:** You might be quite happy to run Office 2010 on your desktops. Unfortunately, Microsoft has a different opinion; they only support a subset of Office versions when connecting to the service (see Chapters 4 and 15 for more details). If you have unsupported versions of Office, you should consider upgrading to Microsoft 365 apps for enterprise to ensure that the Office desktop applications receive updates automatically over the Internet.
- **Complex line of business applications:** Most companies have applications that they have built in-house. These applications range from Visual Basic programs and Excel macros to full-blown database applications. Some of these applications might have dependencies on the workload that you plan to move to the cloud and might therefore need code changes so that they can continue to function afterward. That is, if you can find the code and have someone available who understands how to make the necessary updates to the code. One major factor to consider is that Microsoft 365 does not allow programmatic access to the data that it manages in the same way that is possible on-premises when you control all the moving parts, so programmers must use new APIs (the Microsoft Graph is the preferred approach) to interact with online services. In turn, this might involve some retraining or knowledge acquisition. However, you also get some significant benefits from Microsoft's support of the Power Platform for low-code application development and the possibilities of easier integration with customer- or accounting-oriented data stored in Dynamics.
- **Your network:** The cloud can't function unless you can access its services, and that means connecting across the internet to Microsoft's data centers. You need to understand just what kind of access your company will need to the public Internet to carry the anticipated traffic to Microsoft 365 and all other Internet traffic that your business needs, plus a reasonable uplift for anticipated growth over the next few years (as there is no point in upgrading to a connection that just about copes and runs into problems soon thereafter).
- **Your people:** Running an on-premises environment needs knowledge of and experience with Windows, Exchange, SharePoint, and all the other products that combine to form an enterprise environment. Over time, the people who run IT systems develop a certain insight into the company and how it conducts its business. For that reason, I hesitate to say that companies can achieve workforce savings by transitioning workloads to the cloud. Perhaps it is possible to save by letting some operations staff go because administrators no longer perform tasks such as server maintenance, but usually, you find that other work appears to fill the vacuum left by work now done in the cloud.



For example, in a hybrid environment, someone must manage tasks such as directory synchronization, single sign-on, certificates, and mail flow. Even though Microsoft may be doing the work of running the service in daily operations, someone still must administer those applications on behalf of your company. During the migration period, administrators will have plenty to occupy themselves as workloads move to the cloud. Afterward, they will have new tasks to perform (such as hybrid management) and can take on new responsibilities, such as setting up enterprise voice operations based on Teams or creating a document management strategy for the company that exploits SharePoint Online and replaces old Windows file servers with OneDrive for Business.

- **Change management:** Hopefully, some of your existing staff will have both the desire and the ability to act as change managers for the new deployment. In this role, they can investigate new capabilities that the company has never been able to exploit because of staff unavailability, such as figuring out how best to leverage the Microsoft Graph for business-specific solutions. And they can dedicate time to set up a monitoring framework that deals with both on-premises and cloud components and is available to administrators, support staff, and potentially even users. Change management is an ongoing process, and it's critical—don't be tempted to skimp on the cost! The time and money you will save through properly tracking, anticipating, and dealing with changes in Microsoft's platform dwarfs the cost of implementing an effective change management system even if you must hire additional people.
- **Training:** Moving to the cloud will cause considerable upheaval within your technical infrastructure. That upheaval affects users, administrators, support personnel, and managers alike. Users need training to work with new software (like how mastering the most effective technique to find documents with Microsoft Search or the proper etiquette for using Teams) and how to recognize and cope with common issues, such as a glitch with Internet connectivity that means they cannot connect. If you enable something like multi-factor authentication to augment security or use new features exposed in the service like sensitivity labels and message encryption, users need education on these points too. You'll need to address training for your administrators, your service desk and support staff, anyone involved in network or infrastructure monitoring, your security team, and even the executives or managers who oversee all these services—in fact, educating them may be most important of all as they cannot make good decisions without a thorough understanding of the differences between their old on-premises environment and Microsoft 365.
- **Everything else:** most organizations have third-party products that may or may not be useful in a Microsoft 365 deployment. Security, monitoring, reporting and auditing, backup, compliance, and automation are all areas that are often handled by non-Microsoft products, and you may find that the solutions Microsoft offers for these areas in the service are good enough for your needs... or the opposite. In all cases, before you commit to the extra cost of a third-party product, make sure that you do the research to check if some functionality already exists that satisfies your requirements. Given the speed at which cloud services evolve, it could be the case that an assessment made a year or two ago is no longer valid, so always check to understand the current situation.

## Updating Your Network to Support Cloud Services

In many respects, "the cloud" is shorthand for "the Internet." No cloud project can succeed without reliable high-speed, high-capacity connectivity to the Internet, so the first and most fundamental question that you must answer when considering a move from an on-premises infrastructure to use cloud services is this: can your network infrastructure cope? It is surprising just how many companies discover that they might have difficulty securing enough high-quality bandwidth given how much advertising for Internet providers floods the world, but it does happen. In these instances, all you can do is wait for the local network provider to upgrade their pipes.

Apart from Internet connectivity, organizations that move to the cloud often discover that their internal network infrastructure is not fit for purpose. This is quite normal. You cannot expect a network design created to serve the needs of an internal infrastructure where servers and clients are co-located on the same wide-area network (WAN) to suddenly transform itself to be able to cope with completely different traffic patterns and user demand. Thus, a move to embrace Microsoft 365 is an opportunity to review your network from several perspectives.

- Is your existing internal network configured for maximum performance? For example, customers sometimes find that they have ancient and incorrect network routes that send traffic along suboptimal paths stashed away from previous network upgrades, or that their TCP or IP network settings for things like auto-negotiation aren't properly set. Before your migration starts it would be a good idea to review the underlying quality of your network. You'll want to ensure that you have the most efficient routing possible from where your clients are located to Microsoft's service "front door" locations.
- Do you need to upgrade your internal network to handle a higher volume of client traffic to Microsoft cloud services? We already discussed the need for more bandwidth, but you also need to account for a possible increase in load caused by additional TCP/IP sessions consumed by users as they connect to Microsoft 365. A single user can consume up to 20 sessions to connect to Exchange Online, SharePoint Online, OneDrive for Business, Teams, Planner, and other services. To help, Microsoft publishes information about [network planning and performance tuning](#). Another page covers [network capacity and devices](#), including elements such as WAN accelerators. Another interesting source of information is the [Microsoft Cloud Networking for Enterprise Architects poster](#). This poster shows some of the changes in an on-premises network that tenants might need to make to support connectivity to Microsoft cloud services.
- What ports and IP ranges must be available through firewalls to allow users to connect to cloud services? To help in this task, Microsoft regularly publishes updates for the [URLs and IP address ranges](#) they use with a [web service](#) to inform customers about changes in their network, including updates to route traffic to Microsoft 365 and how to avoid latency and other network challenges.
- If you're upgrading applications as part of your move to the cloud, will those changes themselves require additional network resources? For instance, both the cloud and desktop Microsoft 365 applications have an AutoSave feature to capture changes made to documents during editing sessions so that users do not have to save files. Every AutoSave operation creates an update and syncs it to OneDrive or SharePoint Online. The save might then trigger OneDrive sync operations back to other devices.
- Do you need to make network security changes? For example, if you deploy MFA to protect Microsoft 365 applications, can you use Azure MFA to protect other applications? Is your cloud migration an appropriate time to deploy new services (such as deep packet inspection, WAN acceleration, or cloud application security improvements) or would you be better off doing those separately?

## Understanding Changes in Client Traffic Patterns

Client traffic patterns change with the cloud because the target servers are no longer within the internal network. Instead, clients connect via HTTPS across the Internet to Microsoft 365 and Azure service endpoints. The outbound links that connect your company to the Internet must be able to support the new volume of traffic. The exact type of link needed varies depending on the number of clients that you expect to be active at any time, the type of clients, and the cloud services they will access. Although Microsoft has created a large network of local front-end points of connection (or "edge nodes") to speed traffic to its data centers, unless the link from your company locations to Microsoft's network can carry the amount of traffic your users generate, they will not be able to work as well as they should. Once traffic enters Microsoft's network, it is

carried across fiber to its destination data center. Microsoft's cloud network, including other services like Azure, spans enough fiber to stretch to the moon and back three times. The combination of edge nodes and fiber networks means that it is usually possible to make a fast connection to Microsoft's network, wherever you are in the world. That is, if your local or last-mile connection can carry the traffic.

Microsoft used to post static maps of their network, but instead, they now offer [a flashy interactive model](#). The model shows data flows, renewable energy generators, Azure and Microsoft 365 data centers, and network points of presence, with each type of item clickable so that you can get more information on specific items.

If you run on-premises servers today, you probably have some data to show how much network bandwidth applications and users consume. One rule of thumb is to double (some say triple) this figure and use that amount as the basis for planning. Advice and guidance on this topic evolves over time as cloud services change and knowledge increases, so it is best to ask Microsoft or an experienced consulting company for their recommendation as to the capacity you need – and to then add 30% or so to cater for growth in user demand and to accommodate new applications and features which affect usage patterns and consumption. It is therefore important to keep an eye on changing demand over time.

If this seems like an unnecessarily high increase in bandwidth, consider the experience of companies that deployed Teams for voice and video calling before the COVID-19 pandemic—Microsoft has worked very hard to optimize the codecs used in the service, but users have been so eager to adopt these communication methods that average bandwidth consumption at every customer I've worked with since Teams' debut in 2017 has gone up significantly year-over-year, and above increases in employee count. Organizations that moved rapidly to telework because of COVID-19 often have legacy VPN systems, and people who are used to them; if your users VPN into the corporate network and then connect to Microsoft 365 services, both your network and your users will suffer poor performance—split tunneling (described in Chapter 15) will help with this. If you have a consolidated network where traffic is backhauled from remote offices to a central location, you should also investigate providing a separate link per location. This will deliver better performance and uptime.

Remember that moving to the cloud will not make a bad internal network any better. In fact, if users rapidly adopt new cloud applications (especially features that use more bandwidth, such as video conferencing), you may see a disproportionate impact on your overall network quality. If that network is barely able to cope with the demands of on-premises client-server connections, it will do no better and is likely to do worse when cloud services take over.

During the migration period, the volume of network traffic will be higher than normal because of the need to move information like mailbox data and documents to the cloud. After the migration period is over, network demand should settle into a more predictable load. There will be peaks and valleys in demand, but you should be soon well acquainted with the general shape of network demand for each of your company's locations and be able to make whatever changes are necessary.

### What About ExpressRoute?

Although an expensive option, it is possible to implement a dedicated network connection between your company's network and Microsoft's network, such as [Microsoft's Azure ExpressRoute for Office 365](#) service. ExpressRoute uses a dedicated MPLS connection between your network and Microsoft's network to ensure that traffic flows as if it were on your internal network. Microsoft doesn't recommend the use of ExpressRoute for Microsoft 365, and they won't let you use it unless you're able to persuade them that you should be granted permission. If you think you might need it, read their [service guide](#) first for more details.

### What about Informed Network Routing?

Software-defined network (SD-WAN) solutions have become increasingly common in large enterprises because of their flexibility and capability. Instead of using a fixed network configuration, SD-WAN solutions

can be reconfigured, often automatically, to adjust to changes in loading, network damage or outages, or changing requirements. If you deploy SD-WAN solutions, special devices known as SD-WAN [controllers](#) aggregate usage and quality data and use it to push configuration changes to the network.

Microsoft recently added limited support for SD-WAN integration through a feature they call [Informed Network Routing](#). If you have compatible equipment (limited as of January 2022 to Cisco IOS XE), you can enable performance and network quality data from Office 365 to be funneled to your SD-WAN controllers so that your network will dynamically adapt to changes or problems with the network. This is a promising technology, but in practice, it's still too limited to be useful to most of us. It requires equipment from a specific vendor, and you won't get any real benefit from the feature unless your network has multiple redundant paths to the Internet—if you only have one route to the service in the first place, the adaptive nature of SD-WAN can't do much for you. However, as with many other high-end enterprise technologies, as SD-WAN becomes less expensive and more widely deployed, it may add value to your network.

### Planning for Different Types of Network Traffic

When you plan for new network capacity, make sure that you accommodate four types of traffic that may be required *de novo* as part of your upgrade:

- **User connections to the new cloud environment plus existing work.** Remember that users often make use of company networks to access external social sites such as Facebook and Twitter during work hours. You can try to restrict this access but, in many cases, this is like pushing water uphill, so it is best to include a certain overhead for personal network activity. Different clients often have different network characteristics. For example, Outlook desktop clients consume more bandwidth than OWA. As explained in Chapter 14, Teams voice and video conversations are sensitive to network conditions and need certain basic capabilities to be able to function. Microsoft has [network calculators for Teams and Exchange](#) that help assess likely demand (further information is available [here](#)).
- **Migration of user data to the cloud.** You might be able to move 500 users in a single cutover operation during one weekend—but migrating 50,000 users will probably take a little longer. In either case, you'll be moving data from your internal network across the Internet to a Microsoft data center. Depending on how much data is involved and the available throughput from your network to Microsoft, that movement could take longer than you think. It's also worth considering the impact of where your Office 365 tenant is homed; it's entirely possible, based on its usage location, that you will be migrating data from a physical data center in one country or region to an Office 365 data center in a completely different part of the world. For example, consider a US company that creates a tenant (homed in the US) to consolidate its worldwide operations. As it migrates data from its regional data centers in the UK, EU, and Australia, much of that data will end up by default in US data centers, meaning it must transit from its original location across at least one ocean.
- **Client updates.** The way that Office desktop applications are updated in Office 365 may be different than the way you've upgraded them in the past. Because the default state is to let every device fetch updates whenever it wants to, you may find that your client OS and application updates consume more bandwidth (or at least consume it in different patterns) than they did when you used centralized updates.
- **Administration operations.** Intelligent applications like Outlook isolate users from the effect of network outages and allow them to work offline. However, all Microsoft 365 administrative operations occur online, and a reliable (and fast) network connection is necessary if you want to be sure that administration can be performed without difficulty.

You should also keep in mind that Microsoft has gone through a lot of effort to categorize its network endpoints according to the network priority they should receive; see Chapter 15 for more on how you can work with these categories to prioritize the most important client-facing network traffic to Microsoft's data centers.

## Getting Executive Buy-in and Communications

Few major projects succeed without executive buy-in and support. A good communication plan to explain why the project exists and the benefits expected to accrue to both the company and employees is also important to offset the upheaval that might happen as mailboxes are migrated, networks change, and clients are updated.

Executives should focus on simple messages like:

- We are upgrading our employees to the best-of-breed office productivity software from Microsoft.
- Using the cloud will help us be more flexible and adaptive to changing business circumstances.
- Employees who worked on our previous on-premises systems will now work on more important projects, such as figuring out the corporate strategy for using SharePoint Online, making better use of Teams meetings and the Microsoft Phone System (perhaps as a replacement for a traditional PBX), or determining whether Teams, Yammer, or Microsoft 365 Groups is the most suitable collaboration platform for the business.
- Moving to the cloud is a safe and secure choice.

The communication plan needs to cover:

- A summary of what's happening – the company is upgrading and improving its email and collaborative capabilities by moving to Microsoft 365.
- When the changeover will occur and when employee mailboxes, SharePoint sites, etc. will be migrated.
- What functionality changes and improvements that users will see and how these updates will make work easier. Some simple scenarios might be included, such as how to set up a group mailbox to support project teams.
- Tips and techniques for making a smooth changeover.
- Anything an employee must do to enable connectivity with Microsoft 365. For instance, those running older versions of Outlook or old browsers might have to upgrade their software.

Every company is different, and the suggestions outlined above are merely the start of the discussion about how to lead people through change.

## Working with Microsoft Partners

Not every company moving to the cloud will have the interest, skill, time, and budget to do all the work themselves. Many organizations choose to work with outside Microsoft partners in the hope that doing so will make their journey to the cloud simpler, faster, and/or less expensive.

Microsoft stokes this hope through their [FastTrack program](#), which promises free onboarding aid for customers who buy more than 50 seats and will provide some funding to accredited partners to help with migrations. FastTrack funding can be very helpful to offset the overall cost of a migration. However, FastTrack migrations are very structured and follow a strict playbook, so they might or might not be enough for your needs. A small company that wants to move from on-premises Exchange to Exchange Online will probably be fine, but FastTrack is unlikely to be enough for complicated migrations. This is when you might need to either run the migration yourself or seek the help of a Microsoft partner.

Of course, given enough time to learn and experiment, any experienced Exchange administrator could prepare a tenant and move mailboxes to it. The steps needed to configure hybrid connectivity are largely automated and well understood so that attention to detail and good preparation will get most administrators through the process. In general, the migration of a normal Exchange setup composed of a single organization based on a single Active Directory forest is usually straightforward.

On the other hand, the kinds of details that cause problems may include single sign-on (SSO), combining multiple Exchange organizations, providing identity federation, handling complex mail routing or message hygiene, and hybrid co-existence. As the complexity and/or size of your migration increases, so does the likelihood that you'll need outside help. Of course, during a typical enterprise migration, you'll need to keep your existing on-premises infrastructure running, you might need to engage some outside help to make the project happen in a reasonable period.

When you start considering how to find a reliable partner, you can of course work directly with Microsoft Consulting Services (MCS).

There are many other types of Microsoft partners, though, including partners who also sell migration tools and Microsoft Cloud Solution Providers (CSPs) that specialize in delivering managed services. Many small to medium tenants use CSPs to manage their tenants on an ongoing basis. However, note that some CSPs only resell licenses and don't provide any management or migration services.

Although expensive, MCS has the advantage of being part of Microsoft and offers the reassurance of having a single throat to choke if anything goes wrong. On the other hand, independent partners often offer more realistic opinions about how to run successful projects and the likely pitfalls that exist along the way. In addition, they often have deep knowledge of third-party solutions that can help solve some of the more complex situations that occur in migration projects.

When you consider whether a particular partner is a good fit, here are some things to think about:

- Who will be doing the work, and what are their credentials and experience? Apart from Microsoft certifications, non-Microsoft accreditations such as the CISSP security credential might also be useful, depending on the areas covered in the project.
- Besides accreditations and certifications, you should ask about the experience the partner and its staff have with cloud infrastructures and what projects with similar requirements they have successfully delivered in the past.
- Is the partner a member of the Microsoft Partner Network? A partner who is not a member misses out on many important resources and tools that Microsoft makes available through its network.
- Ask about the skills available in the service provider's Help Desk and how problems are handled. Good partners have substantial resources available to handle a high percentage of problems quickly and effectively without escalating everything to Microsoft. Bad partners take the easy route and act as a channel to Microsoft Support without adding any value.
- What areas of specialization does the partner cover? A partner who specializes in SharePoint or Yammer is probably not a good choice for companies who want to migrate Exchange 2013 to Exchange Online. On the other hand, if your focus is on moving workload from SharePoint on-premises to SharePoint Online, experience with Exchange will not be much help. The deployment of Teams-based voice over IP often needs experience with voice systems, especially when the need arises to replace traditional PBXes. These projects are often done alongside a migration from Skype for Business Online to Teams and might involve hardware such as replacements for room systems.
- If you want to deliver newer parts of the Microsoft 365 suite, or high-end features such as Microsoft Defender for Cloud Apps, information protection, or Advanced eDiscovery, you should find out whether the partner has experience with these lesser-known parts of the platform.

- Application development is different in the cloud. Many newer tools, such as Power Automate and Power Apps, do not have equivalents in the on-premises world and new APIs like the Microsoft Graph are available to expose more data to applications than ever before. If you need to write some code, look for people who know the new tools and APIs before you decide on your partner.
- Do not forget that Microsoft 365 is an ecosystem. Microsoft does not have solutions to every possible feature request or customer need, and it is good to find a partner with knowledge of solutions in areas such as service monitoring and reporting, security, external archiving, backup, and so on.

Apart from these points, you might also ask what contribution the partner makes to the local technical community, how useful their website or blog is in terms of the information that you find there, and whether any of their personnel are Microsoft MVPs as they have background channels to the development group that might be useful in resolving problems that arise during the project. Because Microsoft 365 covers such a wide space, it's also a good idea to understand the area of specialization that the MVPs have. Someone well-qualified in SharePoint Online development technologies might struggle with the finer details of Teams and voice communications.

## Maintaining a “Plan B”

One of the first things I learned during my instrument-flight training: when considering flying in cloudy weather, always have a plan to get out of it *before* you get into it. The Microsoft 365 cloud is no different; as part of your preparation to move into the cloud, you should have a plan to pause or reverse, your migration if you run into trouble or conditions change. Well-planned migrations usually proceed to plan and are successful. Most cloud migrations are, not least because of the expertise available to assist companies in the transition. But wisdom and prudence dictate that you should always have a Plan B – at least in outline – just in case your project experiences problems.

### You Own the Data... But So What?

Whether or not you start your use of Microsoft 365 by migrating existing data into it, you'll certainly use it to create new data. Microsoft makes it clear that tenants [own their data](#):

*"Your data is your business, and you can access, modify, or delete it at any time. Microsoft will not use your data without your agreement, and when we have your agreement, we use your data to provide only the services you have chosen."*

That's reassuring (and obvious!) but, on closer examination raises the issue that there's no obvious way to export the complete set of data contained within a tenant. That means that your ownership doesn't automatically guarantee the ability to pack up your data and take it somewhere else. For instance, you can copy documents from a SharePoint document library to your PC—but is this a realistic approach if you want to capture the entire corpus of information existing in all the SharePoint Online and OneDrive for Business sites within a tenant? That corpus includes lists, document libraries, OneDrive for Business sites, and eDiscovery cases. Even if you move the data items themselves (lists, documents, and so on), you may not be able to capture all the metadata (including audit data and revision histories) for those objects.

Likewise, videos stored in Stream pose a challenge, and if you use Planner, Yammer, Power Automate, and Forms, you need a way to recover their content too. Finally, the increasing role of the Microsoft 365 substrate in knitting different parts of the ecosystem means that apps are more connected and dependent on each other than ever before. Teams is the best example of a cloud-only app that relies on many different components drawn from across Microsoft 365 and Azure. How easy will it be to move content from Teams to another platform?

Third-party software vendors offer utilities to help move content to the cloud, but the same capabilities do not exist to move information back to on-premises servers or other cloud platforms. Given current trends to move away from on-premises servers, this situation is natural as it would take a bold decision to create a migration tool to move content from the cloud.

## Preparing an Exit Plan

No CIO wants to implement a new platform without knowing what will happen if things go wrong. All projects carry a certain element of risk and a project to move a workload to the cloud is no different. Good planning, solid project management, and technical expertise all mitigate the risk, but some cloud projects will end up at a point where management decides that on-premises is the better option. You must then figure out how to move content back from the cloud platform and resume operations on-premises.

Thanks to the engineering investment made by Microsoft to enable hybrid connectivity, Exchange is the easiest application to move back on-premises. Mailboxes can be transferred back to on-premises servers as easily as they can move to the cloud. The only downside is the amount of storage that might have to be deployed to accept inbound mailboxes because users might have become accustomed to the liberal mailbox quotas assigned by Exchange Online. Another complication is how to deal with content stored in inactive mailboxes or auto-expanding archives; it is easy to overlook these items because they do not show up in all administrative interfaces and reports.

Mailboxes and documents are usually easy to move but everything else is hard. You must make sure that you move any customization made to Exchange Online back on-premises so that things like PowerShell scripts, transport rules, Data Loss Prevention rules, retention policies and tags, and so on are moved across. Again, if you are in a hybrid environment these settings should already be duplicated to ensure that the cloud and on-premises platforms remain in tandem, but there's work to be done to make sure that everything is synchronized.

If you choose to use sensitivity labels to protect Microsoft 365 data, you might find that it is difficult to access protected content moved back on-premises because it is no longer possible to decrypt that content. A hybrid configuration will help if you keep a connection to the Azure Information Protection service.

Sadly, it's impossible to move some workloads back to on-premises servers. For instance, if you make a big commitment to Teams as a collaboration platform with or without voice integration, what do you do if things do not work out in the cloud and you decide that the best course is to revert to on-premises? You can move documents back to SharePoint Server, but good luck finding an exact replacement for the suite of capabilities baked into Teams.

## Canceling a Tenant Subscription

As long as you continue paying the monthly subscription fees, Microsoft is quite happy to keep your data online. If an administrator removes a license from a user account, its data stays accessible for 30 days and is then removed. The exception is when a mailbox is put on hold before the account is deleted. In this case, the mailbox is regarded as "inactive" and is retained while the hold remains in force (see Chapter 6 for more information).

What happens if you decide that the cloud is not for you and go ahead and cancel your tenant's subscription? (Hopefully, you will have carefully considered this decision in advance and made a plan to recover all the information from the tenant!) When you cancel a tenant subscription, a formal lifecycle process begins that eventually results in the complete and permanent removal of all tenant data. Table 2-2 lists the steps in the tenant removal process. For more information, see [Microsoft's online documentation](#).



<b>Period</b>	<b>Subscription Status</b>	<b>Effect on data</b>
1-30 days	Expired	Users will see warnings about losing access but can continue accessing their accounts and work with data.
31-120 days	Disabled	Only global or billing administrator accounts can access the tenant to either recover data or reactivate the tenant.
After 120 days	Deprovisioned	An automated process starts to systemically remove all tenant data, including inactive mailboxes and any others that were on hold.
Within 3 days of ending the subscription	Expedited deprovisioning	Upon customer request, the process to remove tenant data can be expedited to ensure that all data is removed within three days of the request.

Table 2-2: When tenant data is removed after a subscription is canceled

Note that the processes that remove tenant data after 120 days are constrained by resource availability. The data might therefore be removed after 120 days, 121 days, or soon thereafter. If you need the data to be removed sooner, you must ask Microsoft to start the expedited deprovisioning process by filing a support request. To confirm that expedited removal should go ahead, Microsoft support gives the tenant a lockout code that the tenant must input to the Microsoft 365 admin center. After entering the code, the services will remove the data within three days. It is critical to understand that once the data is removed it can no longer be recovered or reactivated.

Canceling a tenant is a big step. The responsibility for the decision and what happens afterward stays with the customer, who can decide to let the cancellation process play out or reactivate the subscription within the 120-day grace period.

## Understanding User Adoption Behaviors

You may have had the unhappy experience of taking the time to select a thoughtful gift for someone and then learning that they didn't enjoy it, or even use it (or, in extreme cases, even *unwrap* it)! There's a similar problem with Microsoft 365 migrations: suppose you went through the effort of planning and migrating your organization's users to the cloud, and then they didn't like or use the services in the new environment?

While that sounds painful and unpleasant, it happens with depressing frequency. The process of getting users to adopt new technologies is always tricky because users are human, and humans have a wide range of coping strategies to help them deal with change. Some of those strategies are positive, and others are not. You may see any or all the following behaviors in a typical Microsoft 365 migration:

- **Helplessness.** Pity the users who display this behavior; they are often so beaten down by a steady stream of technical and business changes that they can't even muster the energy to be reluctant or resistant. Working to educate and enable these users can be quite challenging.
- **Malicious compliance.** In the spirit of the original Luddites, who gleefully wrecked industrial machinery at the beginning of the Industrial Revolution, users who display this behavior will "work to rule" in an attempt to slow down or completely derail the pace of change. This behavior is often accompanied by passive aggression that shows up as users complaining that they were "just following your instructions" when something breaks.
- **Resistance.** Some users will forcefully resist change. Sometimes this resistance is rooted in personal fear of obsolescence or not being able to keep up; other times it comes about from concerns about the impact of the change on the business. The actual degree of resistance can vary from flat-out refusal to use a new process, tool, or technology to mild, low-level grumbling.

- **Reluctant acceptance.** Users who show this behavior start as resisters but eventually transition to acceptance of the inevitability of the change. That doesn't mean they have positive feelings about the change, merely that they accept that "it is what it is" and will use the tools on offer.
- **Enthusiasm.** Enthusiastic users are good. However, enthusiasm without knowledge is a bit of a dangerous combination because it encourages experimentation, which may not always be what you want.
- **Evangelism.** This is the best you can hope for when introducing a new technology into your organization. Evangelists not only embrace technical changes; they actively communicate its benefits to other people and encourage them to try it for themselves. Building an active core of evangelists (which I'll discuss shortly) will make your migration much easier by distributing the load of educating and enabling users.

## Enabling and Empowering Users

The keys to getting buy-in and adoption for your Microsoft 365 environment are simple to explain but harder to achieve. Three primary elements will unlock your users' ability and desire to use the new system you're building:

- **What they want and need.** Every human being has certain psychological and physical needs. People want to feel respected; they want to have autonomy in their work and to feel that their work makes a difference. Put squarely in the context of Microsoft 365 deployments, the basic tier of needs is that people want to feel listened to when changes are being proposed; they want to feel like their input is being given due consideration, and they want to be able to do their work without what to them may seem like needless hassle. Depending on your organization, its culture, the national cultures in countries where it operates, and the job roles of individuals, there may be many other wants and needs that you should be aware of. The IT needs of someone working on a factory floor manufacturing car tires will probably be very different from the IT needs of the factory manager, and those in turn will be different from the IT needs of the company CIO. The more you can do to investigate and understand these needs, the easier it will be for you to tailor adoption plans and messaging to meet those needs.
- **What they do and how they do it.** As an IT professional, you probably don't know the details of everyone else's job in your organization (although plenty of people probably think they know *your* job well enough to tell you how to do it)! A basic understanding of the major job *roles* in your organization, though, is key to developing good adoption plans. How many people do you have who would be considered classic information workers? How many front-line workers are there? Of the dozen or so most important or most numerous jobs, how much do you know about what their daily IT life is like and how they use your IT systems to do their job? There's an interplay between wants, needs, and job tasks. For example, if you interviewed a field sales agent, she might say that she *needs* access to data about her customers so she can contact them while she's out working. She needs that access because she *wants* to successfully reach and sell to her customers. To get that access, she will *do* something, such as keeping a printed list of contacts on hand to make sure she can still work when there are IT problems.
- **What they know.** A user who wants to do something but doesn't know how will be frustrated. A user who knows how to do something but doesn't *want* to do it will be inert. Both states are bad for adoption.

## Finding Your Champions

I mentioned the role of evangelists earlier. We all know people who are so eager to adopt new technologies that they wait in line for gadget releases, clamor for early access to beta software, and experiment (perhaps

even recklessly) with their computing tools and environments. You may even *be* one of those people. That adventurous and experimental spirit is a big part of what makes a successful evangelist—the other major part is an ability and willingness to share knowledge, information, and enthusiasm with other people who maybe are not quite as far up the learning curve.

The reason that evangelists are important in your migration is that they can multiply your leverage. By using them as guinea pigs to test your plans, communicators who help spread information throughout the user base, and feedback collectors who give you actionable data about how well the migration's being received by users, you can greatly improve the quality of your migration. A good evangelist is essentially a member of your migration team—with the exception that they will usually be delighted to participate in the process and don't have to be budgeted for and hired because they already work for your organization.

How do you find and connect with these potential evangelists? You probably already have a few names in mind just based on the description at the start of this section. You undoubtedly know people who enthusiastically adopt even the newest and most rickety technology, ask for their department or team to be part of critical pilots, or go out of their way to help other people solve their technical challenges. Approaching these people and asking them to help you plan and evangelize the transition to Microsoft 365 can pay significant dividends; they can give you early feedback, help you improve your communications to end-users, provide early warning of potential problems, and generally help keep the process flowing smoothly.

Microsoft has introduced a tool that they modestly call the "[Champion Management Platform](#)" to help you find and support adoption champions in your organization; it consists of a leaderboard, badges, and some enrollment management tools. Provided as open-source, you can modify and customize it as you see fit to meet your organization's needs.

## Improving Migration Planning and Adoption with Data

If you don't have data about the existing environment before you start your migration, then you are essentially just guessing at some of the most important aspects of the migration, especially how long it will take to complete the migration. Perhaps more importantly, if you have good data about what people are doing with the existing system, you'll be much better able to understand their needs and deliver a post-migration experience that leaves the users more satisfied and more productive than one where you turn them loose in a strange new environment without any guidance or planning.

### Understanding User Activity and Tasks

It's tempting to assume that, based on your understanding of the business, you already know what users do. However, the nuances of how users use the systems you're migrating away from, and how they *might* use the new workloads in Microsoft 365, can have a huge impact on whether users take advantage of the services available to them.

Start with email. Why? It's a familiar workload that, in most organizations, nearly every user will already be using. Using Microsoft or third-party reporting tools to familiarize yourself with which parts of the organization send and receive the most email, how large those messages are, and what types of devices or apps are most used is key.

If you're using an existing IM, conferencing, or telephony system, that should be your next step. Who's using that system, and how much? Is most usage between internal users, or with external parties? It can be very enlightening to compare the amount of email sent and received by a department or team with their usage of Skype for Business, Lync, Zoom, or other conferencing systems to see if you can discern activity patterns that point out where to focus your evangelism and communications first.

File storage and usage are important too. If you have on-premises SharePoint sites, the more you know upfront about how much data they store, how (or if!) it's classified and labeled, who creates and edits the data, and who consumes it, the better able you will be to plan the stages of your migration. For example, if your users depend daily on large on-premises SharePoint document libraries, you will probably prioritize moving those libraries above moving away from an on-premises Skype system as part of the same migration, but you can't make that decision unless you know who's using what services *and* how important those services are to the overall business.

## Deciphering Activity Patterns

Knowing when and how your users use various features of the service is useful. Microsoft makes much of the fact that Teams usage across the service skyrocketed starting in mid-March 2020, and it's been high ever since, but the more interesting aspects of this change are in the details.

Some activity patterns are predictable across organizations—you'll probably see the most activity during weekday mornings, especially on the first day of the work week. If your organization does a lot of seasonal or shift work, you'll probably see different patterns; for example, schools and universities see high activity at the beginning and end of academic terms, but much lower usage in between terms; hospitals and factories tend to see three peaks per day, one per work shift, and so on. However, different organizations may have very different baseline profiles despite these generalities—the number of frontline workers, knowledge workers, shared mailboxes, and so on will all play into what the overall profile looks like. An accounting firm might expect to see more usage at the end of each calendar quarter, with a sharp peak during tax season, while an identically-sized travel company in the same country might see usage spikes corresponding to when people plan their holiday travel.

Measuring the baseline of activity that you see *before* the migration starts is crucial to understanding what "normal" looks like so that you can understand what post-migration usage patterns are telling you about people's use of the service.

Individual workload activity may be steady or there may be bursts. This activity may, or may not, show correlations between workloads, too. If your organization's users send, on average, 10 messages per day, it's reasonable to assume that they'll still send about 10 messages a day after migrating to Exchange Online. On the other hand, if you deploy Teams at the same time, you may find that the average email volume drops because some email conversations have been moved to Teams. In fact, this is a very common usage pattern!

## Understanding User Knowledge

What do users know? That sounds like more of a philosophical question than an actual one. The truth is that knowledge in your user base won't be evenly distributed. Some users will know a great deal, some will know almost nothing, and the majority will fall somewhere in between these extremes. Remember that one key to adoption is equipping users with the knowledge necessary to make use of the new environment. This can be a difficult problem to solve because it isn't a technical issue for the most part. Microsoft has greatly improved the quality of its online help for the desktop and web applications in the Microsoft 365 suite, and there is an ocean of supplemental content from Microsoft, third parties, and the broader IT community covering every aspect of using the service, from basic tutorials about how to format Word documents up to complex and abstruse tutorials on Excel that I don't even pretend to understand. Identifying the training that your users need can be a challenging problem, and you would do well to get your organization's training department involved (if you're fortunate enough to have one).

At a minimum, users will want to know the basics of how to sign in and use the new environment. This doesn't have to be complicated or lengthy training. The simpler and shorter you can make it, the better, because that

makes it easier to update when things change, and shorter “snackable” training is easier for users to access, use, and retain.

You might find it valuable to look over the help desk tickets generated by users over the preceding 12 months and see if you can identify patterns that show you where users need training. For example, if you get a lot of tickets complaining about problems with audio devices in Skype for Business on-premises, it’s probably worth thinking about what users need to know about using devices in Teams, then find (or create) training to help address that specific topic.

## Understanding Wants and Needs to Arouse an Eager Want

The famous salesman and coach Dale Carnegie said that one of the key principles of sales was to “arouse in the other person an eager want,” and it’s this principle you need to adopt to successfully drive the adoption of Microsoft 365. To do this, first put yourself in the shoes of a typical user who will be using the new system. What does this person need to do her job? What are the technical or process obstacles she faces? Are any of those obstacles or needs solved by Microsoft 365?

Microsoft likes to use the phrase “IT hero” to cast a bright spotlight on people who enable business success by deploying their products, and certainly many people are driven by a desire to be recognized as an excellent salesperson, a top-notch developer, or whatever. Think about the different job roles you should have previously identified and what people in these roles need and want. Put these needs and wants in the unique context of the organization—a manufacturing company where everything is managed through rigid monitoring of key performance indicators will tend to produce a different set of needs and wants than a university, for example. Then think about these needs and wants through the lens of the changes you’re preparing to introduce. Your communications can then focus on explaining, as specifically and clearly as possible, how those changes will benefit users. For example, consider an organization where the users have an on-premises Exchange 2013 server that requires users to connect to a VPN when they’re out of the office. This may seem laughably ancient by current standards, but it’s not an uncommon configuration. Explaining to users that the move to the service will do away with the flakiness of VPN connections and improve their Outlook performance will get their attention. By contrast, the usual approach of focusing on features themselves is likely to do nothing for the users; if you’re used to waiting for a minute or more for Outlook to move a single message between folders, having your IT department pitch the move to Office 365 as a way for you to get dark mode in Outlook isn’t going to be very exciting.

### Getting Adoption Help from Microsoft

Microsoft’s continued success in the cloud depends heavily on getting end-users at their customers to *use* the services they’ve bought. To that end, they have a comprehensive site, <https://adoption.microsoft.com/>, that contains a curated set of resources to help you plan how you’ll get users to take full advantage of Office 365 after your migration. There are specific learning paths for each of the major workloads (including [Microsoft Search](#), the first time I’ve seen it treated as equal to SharePoint, Teams, and so on). Having a collection of learning, planning, and training material all in one place is very helpful, but it can seem a little overwhelming to see how much material is there. With that in mind, the earlier you can start browsing the content to get a feel for what’s there, the better; you can also take the list of topic areas and pick one at a time to become more knowledgeable about before moving on to the next one. Of course, Microsoft also offers consulting services, some included with large subscription purchases, and some sold separately, that may be of use to you in your adoption planning.

## Planning for Disaster Recovery and Business Continuity

It's tempting to think that moving your business operations into the cloud means that you can outsource all the worries and planning around disaster recovery (DR) and business continuity (BC) to your cloud provider. Of course, this is an idle fantasy; no cloud provider will ever care as much about your business operations than you do, nor will any provider have the specialized knowledge needed to ensure that your business keeps working smoothly in the event of a disaster. Organizations located in areas where natural disasters are more common (such as California or the Gulf Coast in the US or fire-prone regions of Australia) are already used to the need to plan for operational continuity even during disasters. For example, one US hospital on the Georgia coast proudly showed me their hurricane-recovery plan, which they thought would allow them to keep their office staff working after a hurricane. Unfortunately, their design featured only one Active Directory Federation Services (AD FS) server, located in the hospital data center; if it had failed, none of their users would have been able to log in to the cloud even if they otherwise had power and connectivity.

A complete discussion of DR and BC planning for the cloud is far outside the scope of this book, of course. A robust DR and BC plan will touch on every aspect of your business, from which employees are considered critical to how you communicate with them (e.g., what if email stops working?) to who has the authority and responsibility to declare a disaster. As part of any move to the cloud, you should ensure that you have a plan to deal with problems you might encounter that are unique to *your* cloud deployment. One example might be something like the [2001 Howard Street tunnel fire](#) or the Christmas 2020 bombing in Nashville, both of which affected Internet connectivity over an unexpectedly broad area. Another example might be planning to fail operations or critical communications over to another service or system in case your systems become compromised due to ransomware or an attack like [Solorigate](#).

## What To Do While You Move

Once you've decided to move and have handled all the issues described in the preceding section, you can begin dealing with the unique issues that may arise during the move itself. "During the move" is a pretty vague concept, as for many organizations the move can last months or even years. As one example, a large enterprise with 130,000 seats that I worked with took more than 18 months to fully migrate into a single Microsoft 365 tenant from their previous environment—and while this may seem slow, it was as fast as they were able to go considering all the business and organizational issues they had to resolve along the way.

## The Value of Hybrid Connectivity During a Migration

If you're moving to Microsoft 365 from a Microsoft-based on-premises environment, you will almost certainly end up in hybrid mode, Microsoft's term for an environment where some resources and data are on-premises while others reside in the cloud. In hybrid environments, the goal is to maximize integration between the two environments so that users can sign in with a single identity, use data and services wherever they're provisioned, and move easily between different services.

Any time you see the word "hybrid" used in the context of a Microsoft 365 environment, it means that there is a persistent connection designed to allow data to move between the cloud and on-premises servers. This is different from the typical "cutover" or "switchover" processes you may be familiar with from migrations involving older versions of Exchange—in a typical cutover migration, you'd move all the mailboxes from the source to the target, do some cleanup (such as redirecting mail flow by editing the DNS MX record), and immediately decommission the old system. Hybrid mode is in some ways like the process of raising children from birth: it's an ongoing process that takes a while and may proceed at different speeds and with different constraints in different families.

Various parts of the Microsoft 365 stack take different views of what “hybrid” means. Some services, such as Exchange Online and SharePoint Online, can function effectively in three ways: all on-premises, all in the cloud, or hybrid mode. For example, you can freely mix mail users who have mailboxes in Exchange 2019 servers on-premises with users in the same tenant who have mailboxes in Exchange Online. Of course, some workloads have no hybrid option: Teams, OneDrive for Business, Yammer, Planner, Whiteboard, and Stream have no hybrid capability.

Hybrid connectivity starts with identity management, the foundation of all the services in Microsoft 365. As described in Chapter 3, user identities can be maintained in a hybrid environment through a variety of different methods, including identity federation or the less complex and more easily managed directory synchronization process. Once you have provisioned identities in the cloud so that the service can see them, you can start exploring the cloud by using Microsoft 365 services on a trial basis, either with cloud-hosted identities or identities that come from your on-premises environment. This allows you to transfer some work to the cloud, observe the results, and decide on the best way to proceed. In some cases, the decision will be to move more workloads, and perhaps even to proceed along the line so that most mailboxes are transferred. In others, the decision will be to retain a substantial portion of work on-premises and use cloud applications for tactical purposes, such to support specific types of employees. The transfer of workload follows the pace set by the company and could extend over many years.

The existence of hybrid connectivity means that you can move gradually, as your business needs and user base will allow, instead of having one giant spasm of reconfiguration as in a cutover migration with an older version of Exchange (or a migration from Microsoft’s products to Google Workspace). Hybrid coexistence means you can plan and carry out your migration considering the factors that make it harder to plan for a fast switchover, including the distribution of users across multiple company locations, especially when international employees are involved.

Another major advantage to hybrid connectivity is that it gives customers a “plan B” if moving to the cloud does not deliver the advantages expected when deciding to use Microsoft 365. The same facilities that allow mailboxes to be moved from on-premises to Exchange Online can be used to move mailboxes back to Exchange on-premises. It’s possible that the same tools could be used to move documents from SharePoint Online back to SharePoint Server. Some gaps exist – like how to move public folders back because the public folder migration tools run in one direction – but overall, hybrid connections are two-way and can therefore be regarded as reducing dependency on Microsoft 365 should the need arise.

In early 2021, we saw a giant drawback to hybrid connectivity: the on-premises Exchange hybrid server represents an additional point of vulnerability if it is not properly managed and secured. The fast and broad spread of the HAFNIUM attacks which leveraged zero-day vulnerabilities in Exchange can be partially attributed to poor patching and information security practices at customers who didn’t realize or didn’t take seriously, that they had remaining Exchange servers even after their mailboxes had been moved. In mid-2022, Microsoft finally provided a supported way to remove the hybrid Exchange role from tenants that have completed their migration. However, it is very important that you *not remove the Exchange server*. You can shut it down, and even reformat it—but if you uninstall Exchange on that server, you will cause all sorts of spectacular damage, as described in the full [documentation for removing the hybrid role](#).

Until you’re ready to get rid of those on-premises servers, be sure that your Exchange hybrid servers are patched and protected with the same level of vigilance you provide to domain controllers and other critical server roles.

## Dipping Your Toes into The Cloud

Even after doing extensive research to verify whether Microsoft 365 is the right platform for your company, there is nothing quite like getting your hands dirty to make a technologist happier with a technology. Microsoft makes test drives easy by allowing you to sign up for a 30-day free trial. Essentially, you can create a fully functional test tenant and use it for 30 days to decide whether the cloud works for your company. And if you're still not sure after that period, you can simply discard the original trial, extend the trial, or start again with a new tenant.

Using a trial tenant as a proof of concept is a great way to get a sense of how cloud applications and management work in practice. The exercise will not expose some of the more complicated challenges, such as how to migrate public folder data or how hybrid connections or single sign-on work over an extended period. However, the information gained from a trial is certainly enough to gain a broad understanding of the different parts of the service and how the administrative experience differs from your on-premises environment. You'll be able to figure out how to move work to the cloud and what extra functionality exists – or where functionality gaps exist because something like a preferred third-party add-on is unavailable.

Some companies start with a trial tenant and then convert the trial into a paid-for service after the 30 days. This is OK if you plan for the eventuality and go ahead on that basis. However, you should be aware that Microsoft only provides a [limited capability to move data from one tenant to another](#). Third-party tools are available, but these come at added cost.

All-in-all, it is usually a better idea to regard a trial tenant as no more than a test and to accept that if the company decides to embrace the cloud, you will start over from scratch and use a new tenant.

## What To Do After You Move

If the key phrase for the pre- and trans-migration phases is “measure twice, cut once,” the useful cliché for the post-migration phase is “it’s a marathon and not a sprint.” Ideally, you will have conducted your migration in manageable phases to avoid swamping your service desk and IT team, but no matter how carefully you plan, on “day zero” after each phase you will probably see a spike in support requests as users encounter problems. That’s normal and shouldn’t discourage you—it’s important to deal with those requests quickly to keep users unblocked and to start building forward momentum.

Microsoft has historically tied incentives to adoption; they offer free consulting, license incentives, and other goodies to get their customers to *use* what they’ve already bought. Each year, Microsoft salespeople, and most partners, are given specific numeric targets for various features they’re supposed to focus on. Sales commissions, partner incentives, and even membership in Microsoft’s partner programs are tied to whether sellers can get customers to buy *and use* these targeted features.

The COVID-19 pandemic changed this mindset pretty quickly, since Teams usage in specific and adoption of Microsoft 365, both shot skywards... but we haven’t seen the last of the mindset at Microsoft that calls for continued growth in service usage from customers and partners. More to the point, to get the full value of the Microsoft 365 licenses you’ve purchased, you need to do two things: purchase the right mix of licenses for your needs and then properly assign and utilize them.

## Building Adoption Momentum

Much of the earlier discussion about adoption revolved around understanding what users *were* doing in the previous system. Armed with that understanding, you’ll be better able to assess how users are taking to the new system. Before they will (or can!) adopt the new system, you must work with users to help them:



- **Want the new system.** If users see Microsoft 365 as an unwelcome change forced on them by an uncaring IT department, its adoption will probably be reluctant, slow, and incomplete. On the other hand, if you've properly communicated the benefits to "arouse an eager want," users will look forward to, and embrace, the system.
- **Become capable.** Some of the features in Microsoft 365 will be new, and maybe even strange, to users. Teams is a great example; before the advent of Teams online meetings, most users never could easily record and share meetings and presentations, much less use live captions, machine translation, and the other nifty features that Teams offers. It falls on you to help them understand what's possible and to see the value these new features offer considering what you know about their needs, wants, and job responsibilities. The downside is that if users are to take full advantage of Teams, they may need headsets, webcams, and so on, and so providing these devices may turn out to be an important part of your strategy. License assignment is important here as well; users cannot use features for which they aren't licensed and provisioned.
- **Get the knowledge they need.** Users who want the new system and are enabled to use it, but don't know how, obviously can't adopt it. This is one area where the corps of evangelists you've built can make a huge difference in the success of your adoption planning by providing unofficial help, support, training, and cheerleading.

## The Cloud Is a Journey

If you haven't already started your move to the cloud, what you've just read may have worried you—it seems like there's so much to learn and consider before and during the move that it can be daunting. Good news, though—many, many organizations of all sizes have successfully moved to the cloud before you and there's lots of help available, starting with this book! Microsoft provides a wealth of training material online and in print, and many of your counterparts have published material that will help inform you. If you have already moved to the cloud, you should still be prepared for the fact that Microsoft's always growing, expanding, and changing its cloud services. You'll need to be ready to keep up with Redmond by learning new things and adjusting as things change.

It may help to keep in mind that you're providing a service to help the people you work with—it's easy to get caught up in the technical details of how the service works, what PowerShell cmdlets to use, and so on, but the reason so many customers have given Microsoft their money to subscribe to Microsoft 365 is that they believe that the cloud services will help their end-users work more productively and more securely. The service is just a means to that end.

# Chapter 3: Managing Identities

## **Brian Desmond**

Your online identity, and by extension the authentication process that establishes the identity, is the cornerstone of security in Microsoft 365. Once a user or workload identity is successfully authenticated by Azure Active Directory (Azure AD), access to any authorized resource is automatically granted. Throughout the book, we will distinguish between two separate processes:

- *Authentication* is when a system verifies the identity of a user. For example, the familiar Windows logon process you go through is how Windows authenticates your domain or local user account.
- *Authorization* is when a system decides whether a specific user or workload identity should have access to an object or service.

In its simplest form, the Microsoft 365 authentication process needs an identity, represented by the combination of a User Principal Name (UPN) and something by which their identity can be confirmed. Often the authentication method is a password, but it can also be a specific device such as a FIDO2 token, a biometric gesture, a certificate, or a combination of several methods (also referred to as multi-factor authentication [MFA]). Once Azure AD has confirmed a user's identity, it returns an artifact that an application or service can accept as proof of identity.

Although the default identity and authentication options suit most customers, not all organizations have equal security requirements. Nonetheless, protecting your (digital) identity is an important task; both inside or outside Microsoft 365 or one of its workloads. After all, once someone can gain access to your account, they hold "the key to your kingdom", granting them access to all the data accessible within the boundaries of your identity. To meet the need for stricter security policies dealing with authentication, Azure AD supports a plethora of options to customize and further secure the authentication process, as discussed here.

Because identities and authentication touch on so many aspects of Microsoft 365, it is important to carefully consider the implications of the various identity models and authentication options. The importance of factors such as complexity and cost is different for every organization and can greatly influence the decision as to which solution is best for you.

In this chapter, we explore the various authentication options that are available to you, spanning both cloud-only and hybrid deployments. As you will see, a robust identity infrastructure is fundamental to a successful deployment. Before exploring the available options, we should first look at the Azure AD infrastructure which supports both identities and authentication for Microsoft online services.

## The Role of Azure Active Directory

All the Microsoft 365 workloads depend on Azure AD for identity and directory information. Azure AD provides authentication and authorization for Microsoft 365 workloads as well as third-party applications. While similar in name, you will discover that Azure AD has few other similarities to traditional on-premises AD. However, we will explore how you can extend your AD to have a hybrid relationship with Azure AD.

Since Azure AD is a cloud service, there are significant differences between how on-premises workloads like Exchange Server or SharePoint Server use on-premises Active Directory (AD) versus Azure AD. Some

workloads implement workload-specific directories behind the scenes to bridge this gap. The Azure AD deployment supporting Microsoft 365 is, like most cloud implementations, designed to be a highly available multi-tier, multi-tenant service that is capable of handling load at an immense scale.

Every Microsoft 365 tenant has a corresponding instance within Azure AD, and each instance is isolated from others so that no tenant has access to data belonging to another tenant. When you create a new tenant, a new Azure AD instance is also created. The Azure AD instance is not licensed separately, and you do not pay anything extra for it. As you add users and groups, those objects go into this Azure AD instance.

Azure AD is deployed in multiple Microsoft data centers around the world and [operated with a service level agreement](#) (SLA) of 99.99%. In addition to Office 365, Other Microsoft cloud services like Microsoft Dynamics 365 and Microsoft Endpoint Manager consume Azure AD, as do numerous third-party applications. Microsoft has made significant investments in Azure AD to offer a 99.99% SLA. As with any cloud service, Office 365 has suffered from large-scale outages on occasion, and in some cases, an issue with Azure AD was the root cause of the outage.

To address this, Azure AD includes a [Backup Authentication Service](#) (BAS) that automatically steps into action to process authentication requests from supported applications if the primary service becomes unavailable. To do this, the BAS caches successful authentication requests for three days. The BAS relies on features such as continuous access evaluation and conditional access, discussed later in this chapter, to deliver resilience in a secure and configurable manner.

Like other cloud services, Microsoft updates Azure AD regularly to introduce new features and improve quality. See the [Azure Active Directory release notes](#) for details.

You can access the Azure AD admin center through the Azure Active Directory link in the **Admin centers** section of the Microsoft 365 admin center navigation bar or via the Azure portal.

## Workload-Specific Directories

Because of the way that some workloads work, Azure AD cannot offer the necessary functionality to support all the features required by those services. For example, the mailbox-specific attributes used by Exchange Online are not all stored in Azure AD. The same is true for many of the user profile properties used by SharePoint Online. To support workload-specific needs, and to isolate the configuration and other data owned by a single tenant inside the multi-directory Azure AD architecture, some services add another layer on top of Azure AD. This layer is a workload-specific directory store and is named after the workload:

- Exchange Online Directory Services (EXODS).
- SharePoint Online Directory Services (SPODS).
- Yammer Directory.

This workload-specific directory is essentially a cache of information held in Azure AD combined with workload-specific information. The workload-specific directory is designed to deliver a certain level of redundancy against network or other disruptions. To further explain the concept of workload-specific directory stores, we can use Exchange Online as an example.

### Exchange Online Directory Services

Exchange Online uses EXODS to hold its configuration data, including information about mail-enabled recipients. Exchange Online deploys its service across multiple forests with tenants divided between the forests. Each forest is unique to a data center region (multi-geo tenants spread mail-enabled objects across multiple forests) and uses an instance of EXODS together with a synchronization endpoint that is used to replicate information with the directories used by Microsoft Online Services (for user accounts and licensing), other applications (like SharePoint for mailbox information used in eDiscovery), and Azure AD. The introduction of support for spreading a single tenant's data across multiple data center regions means that

Microsoft has made additional changes to the way these forests are structured and synchronized so that all locations where a tenant's data may exist have complete and consistent data.

Synchronization occurs across the directories on an ongoing and continuous basis to ensure that the latest information is always available to all workloads. Occasionally, glitches can happen with the synchronization process, and you might have to wait for a new user account to be visible.

Exchange Online can create new user accounts during the mailbox creation process. In this case, Exchange Online pushes the information about the new account to the EXODS and Azure AD [in parallel](#). Exchange Online is unique in this capability, which exists because many customers have built extensive mailbox provisioning workflows with PowerShell that would break if forced to connect to Azure AD instead. When on-premises identities synchronize with Azure AD via the directory synchronization process in hybrid deployments, Azure AD is the target, and the resulting objects will synchronize to EXODS afterward.

## Licensing

As explained earlier, each tenant includes a free version of Azure AD that provides all the core functionality necessary to use Microsoft 365 services. For many organizations, the core capabilities of Azure AD are not enough to meet security requirements in the context of today's cyber threats. Microsoft bundles many of the security features that today's organizations require into Azure AD Premium. Azure AD Premium is available in two variants: Azure AD Premium P1 and Azure AD Premium P2. You can purchase Azure AD Premium on an a-la-carte basis, or get it bundled in the Enterprise Mobility + Security (EMS) or Microsoft 365 offerings.

For many organizations, Azure AD Premium P1 is enough, but Azure AD Premium P2 brings a set of governance tools and risk-based controls that can dynamically enforce security policies based on Microsoft's determination of the risk factor of each user and their sign-in attempt. A comparison of the different types of Azure AD licenses is available [here](#). Throughout this chapter, we call out whether a capability requires Azure AD Premium. Unless otherwise noted, you can assume that only Azure AD Premium P1 is necessary.

### Licensing Guest Accounts

Guest accounts need user rights to use Azure AD Premium functionality too. Microsoft's billing model for external identities (including guest accounts) uses telemetry to measure the number of unique identities that access premium Azure AD features monthly. Microsoft does not charge for the first 50,000 unique identities measured in a month and the charge thereafter is small. For instance, access to Azure AD Premium P1 features incurs a charge of \$0.00325 per month for all identities past the 50,000 threshold.

The Azure AD tenant must link to a valid Azure subscription through the External Identities section of the Azure AD admin center to support billing for external identities. Before doing this, the tenant must create the Azure subscription and create a suitable resource group to use.

Microsoft charges for each SMS-based multi-factor authentication request processed for an external identity. The charge levied pays for the telephony fees and is circa \$0.03 per attempt (successful or not). The charge does not apply when external identities use the Microsoft Authenticator app for MFA.

Although Microsoft does not currently enforce licensing restrictions for guest users, it is important to make sure that your licensing plans take this point into account if you deploy applications that support guest accounts, like Microsoft 365 Groups, Teams, and Planner. See [this guide](#) for more information.

## Identity Architectures

There are two options for how you create Azure AD user accounts (also referred to as "identities"):

- **Standalone identity:** In this model, user accounts exist only within the cloud environment and are not linked or related to any other AD forest. This means that a user may (or may not) have an account

in an on-premises AD, as well as an account in the Azure AD tenant that supports the organization's tenant. The accounts may happen to have the same username and password, as well as other attributes, but are separate and independent objects, and need to be managed independently as well. This duplicates administrative effort and introduces risks if attributes conflict or are not maintained correctly. Some organizations use standalone identities for users that do not require an account in their on-premises AD, such as third-party vendors.

- **Hybrid identity:** In this model, user accounts are created and managed in the on-premises AD environment and subsequently synchronized to Azure AD. Azure AD Connect or Azure AD Connect Cloud Sync performs the synchronization. In the hybrid identity model, the on-premises AD forest is the "source of authority", with objects and attributes replicating from on-premises to the cloud. Because the on-premises AD is now the source of authority, an administrator can only change limited aspects of the synchronized identity in Azure AD. More information on these limitations and how to configure Azure AD Connect is in Chapter 5 of the companion volume.

Each approach has its pros and cons. For example:

- Standalone identities are convenient because they do not need an on-premises AD. Many small organizations do not have or want to operate an on-premises AD. However, if an on-premises AD exists, standalone identities can be more time-consuming to manage because organizations must manage two separate accounts for each user. There is no arbitrary threshold for how many user objects you should have before synchronizing identities makes more sense. It typically comes down to how much time and effort an administrator must put into managing dual identities versus maintaining a synchronization infrastructure.
- Hybrid identities simplify administration because all changes occur to the on-premises AD and are then synchronized to the cloud by Azure AD Connect. The hybrid identity model doesn't prevent you from creating and using standalone identities.

You can optionally add identity federation to your hybrid identity design. Federation gives organizations much greater control over how to enforce security policies such as logon hours, third-party multi-factor authentication, and network locations from which users can access cloud resources. Most of the additional controls that federation offers are also available natively in Azure AD if you have Azure AD Premium. It is critical to remember that the identity federation infrastructure must scale to meet performance and workload requirements. It must also be highly available and resilient to failure because Azure AD may not be able to authenticate user logins if the federation service is unavailable.

You can deploy a mixture of these two models. Although it is more common to settle on a single identity model, organizations can choose to combine different identity solutions to fit their specific needs. For example, you may need to provide customers or contractors with access to certain resources. Using standalone identities for these users and hybrid identities for employees allows you to isolate the on-premises environment from customers and contractors. Ultimately, the decision about which identity model to use is driven by the business and technical requirements of the organization.

## Standalone Identity

The on-premises versions of Exchange Server and SharePoint Server use AD as the repository for user objects and to process authentication requests. In the case of Exchange, every mail-enabled object is represented by an AD object. The properties of the objects are managed by different tools, including the Active Directory Users and Computers console, the Exchange Admin Center (EAC), and PowerShell. Much the same division of responsibilities exists in the cloud versions of these services, with the exception that Azure AD is used instead of on-premises AD. For example, Exchange Online uses the same properties for mailboxes, public folders, groups, and other mail-enabled objects as the on-premises version, but these properties are divided across,

and managed through, both Azure AD and EXODS. Although a synchronization process keeps the two online directories in step, Azure AD is always the master.

Every user account has a unique user identifier, known as the User Principal Name (UPN), used to authenticate with Azure AD. For example, *britta.simon@office365itpros.com*.

By default, Microsoft 365 sets the primary SMTP address for a standalone identity to be its user principal name and things usually stay that way. However, this does not have to be the case as the two properties serve different purposes: the UPN identifies the object to Azure AD, and the primary SMTP address routes messages to its mailbox if it is mailbox-enabled.

Administrators often need to change the UPN, SMTP address, and display name for user accounts. In the following example, we use the Microsoft Graph PowerShell SDK to switch names for a cloud-based user from Jane Smith to Jane Jones and update the UPN and primary email address to reflect the same. It is good practice to retain old addresses assigned to accounts in the past so that Exchange can still route replies to those old addresses to the correct mailbox. To set the new values for the account, use the *Update-MgUser* cmdlet. The updated information will synchronize from Azure AD to the workload-specific directories such as EXODS. Finally, we use the *Set-Mailbox* cmdlet to update the mailbox's primary SMTP address to match the new User Principal Name.

```
[PS] C:\> Update-MgUser -UserId Lotte.Vettler@office365itpros.com -Surname Smith -UserPrincipalName Lotte.Smith@Office365itpros.com -DisplayName "Lotte Smith"
[PS] C:\> Set-Mailbox -Identity 'Lotte Smith' -WindowsEmailAddress Lotte.Smith@Office365ITPros.com
```

When you change the UPN for an account, the account owner must provide the new UPN the next time they sign into any Microsoft 365 service.

## Service Accounts

Not all the accounts you use belong to users. Some applications, such as those that monitor workload health or report on various aspects of Microsoft 365, might require the creation of special service accounts used solely for administrative purposes. These accounts usually do not need a mailbox or a license, but probably need the assignment of some form of administrative permissions to be able to view and access data on behalf of the tenant.

Most service accounts have passwords that never expire. This is done on the basis that managing password expiry for service accounts is often complex and time-consuming. While setting service account passwords to never expire is the path of least resistance, you should rotate service account passwords regularly in the same manner that you maintain on-premises service account passwords.

Configuring a password to never expire is easily done with PowerShell. The first example below sets the "password never expires" flag on a service account. The second example scans the set of user accounts and reports those that have this flag set. These examples only apply to cloud accounts. If you synchronize an account from the on-premises AD and password synchronization is enabled, the password of the corresponding cloud account is automatically set to not expire unless you configure Azure AD Connect to apply a password policy to synchronized users. In such cases, it suffices to also configure the on-premises account's password to not expire. Note that after successfully updating the password policy, Azure AD does not return a message. However, Azure AD signals if an error occurs.

```
[PS] C:\> Update-MgUser -UserId Service.Account@Office365ITPros.com -PasswordPolicies "DisablePasswordExpiration"
```

```
[PS] C:\> Get-MgUser -All -Filter "userType eq 'Member'" | Where {$_.PasswordPolicies -eq "DisablePasswordExpiration"} | Format-Table DisplayName, UserPrincipalName
```

While configuring a service account to have a password that never expires is the traditional approach, it also presents security challenges. Service accounts authenticating with a username and password use legacy

authentication methods which are particularly vulnerable to attack. Microsoft plans to block legacy username and password authentication in the future and has announced plans to block basic authentication for many connection protocols for Exchange Online as quickly as possible. The username and password must also be stored somewhere for the application to use it. This presents the risk that attackers could compromise the stored credentials.

Instead of using a username and password, consider using an Azure AD application registration (also known as a service principal). Application registrations can be configured to use a certificate for authentication. Application registrations can also be granted permission to use specific APIs via the Microsoft Graph. This approach requires that the API or PowerShell endpoint that is being called supports modern authentication methods. While this is not yet a possibility for every scenario, we recommend that you begin by investigating this method when you are designing a new script or unattended task.

## Hybrid Identity Authentication Infrastructure

Even after you have synchronized your identities to the cloud, your users will still need a way to authenticate. In a hybrid identity model, there are three ways to do this:

- **Password hash synchronization:** In this model, a cryptographic hash of the user's password, but not the password itself, is synchronized to the cloud. The hashing process is discussed in detail later. When a user requests access to a service component, the password they provide is hashed in Azure AD; if that hash matches the stored hash, the passwords are considered to match. This process relies on directory synchronization but does not require any other servers or components. Combined with Seamless Single Sign-On (discussed later), you can achieve single sign-on (SSO) with password hash synchronization.
- **Pass-through authentication (PTA):** In this model, authentication requests from the service are sent to a queue, where they are retrieved by a small agent that is installed by Azure AD Connect. The agent validates passwords against the on-premises domain controller (DC) and returns a status (success, failure, password expired, or user locked out) to Azure AD. Like password hash synchronization, PTA can be combined with Seamless SSO.
- **Federated authentication:** In this model, authentication requests from the service are passed to a federation server or service. This can be AD FS, or a third-party federation service such as Ping Identity or Okta. The federation server is responsible for authenticating the user by passing an authentication request to an on-premises DC and then returning an authentication token for the user to access cloud services. This process provides the end-user with an SSO experience that can be used to access various Microsoft 365 services.

You can combine password hash synchronization with PTA or federation. This provides a manual fallback mechanism for authentication if the on-premises PTA or federation infrastructure is unavailable. As of November 2019, 91% of Azure AD tenants globally enabled password hash synchronization.

If you are thinking about switching authentication methods, there is a feature that can make this much easier. Rather than switching from federated authentication to password hash synchronization, PTA, or seamless single sign-on for every user at once, you can control which authentication method is used on a per-group basis. To do this, access the Azure AD Connect blade in the Azure AD admin center. Click **Enable staged rollout for managed user sign-in** and you will be taken to the configuration area. Using staged rollout is a great way to test the user experience and better plan what amounts to a major change to your authentication infrastructure.

## Password Hash Synchronization

When directory synchronization is implemented with password hash synchronization, users can log on to services using the same password as their on-premises AD user account. This is referred to as “same sign-on”, which is not to be confused with “single sign-on” even though they both can be abbreviated to SSO.

The concept of synchronizing passwords tends to raise immediate concerns within an organization, as people inevitably assume that real passwords are being transmitted over the Internet and stored on Microsoft servers. The reality is that password hash synchronization is a secure process, and the passwords themselves are not transmitted or stored. Instead, a password hash is used. Despite this, some organizations may still object to the use of password hash synchronization, so it is important to understand what is being synchronized.

On-premises AD stores passwords as hashed values that are said to be *irreversible*. In other words, a password hash cannot be used to determine a user’s plain text password. Azure AD Connect extracts the password hash from AD, combines it with a user-specific salt value, and then hashes the combination 1,000 times before transmitting the hash over a secure HTTPS channel to Azure AD.

When password hash synchronization is enabled, the password complexity and expiration policies of on-premises AD override the policies set in Azure Active Directory. When a password is changed on-premises the new password is synchronized to Azure AD. This process usually occurs in under two minutes. You can [selectively synchronize](#) password hashes using a custom synchronization rule beginning with Azure AD Connect version 1.6.2.4. This adds complexity to your deployment and impacts password writeback so you should only use this capability if it is truly required.

If an on-premises password expires, the expired password will continue to work in Azure AD. This is because when password synchronization is enabled, Azure AD account passwords are set to not expire. For this reason, you should not rely on expiring or changing passwords to prevent a user from logging on. Instead, you should disable the on-premises user account and either force a directory synchronization, or block sign-in for their account in the Microsoft 365 admin center. If you want synchronized passwords to expire in Azure AD, you can configure Azure AD to apply the password policy in your tenant to synchronized users by following [these steps](#). You can use the `Set-MsolPasswordPolicy` cmdlet to configure your tenant’s password policy. Both these settings are global settings for your entire tenant.

Another important benefit of password hash synchronization becomes available if you have Azure AD Premium. With Azure AD Premium, Microsoft provides leaked credential risk event information for users if they discover the user’s UPN and password in a list of stolen/lost credentials. To perform the comparison, password hash synchronization must be enabled. Microsoft applies the same hashing process to the leaked password that is performed on-premises. They then attempt to match the hashed version of the leaked password to passwords in Azure AD. This works even if you use another sign-on methodology. It simply requires password hash synchronization to be enabled. In November 2019, Microsoft stated that they had processed over 5.5 billion leaked credentials and matched them to over 14.2 million Azure AD users. This feature is one of many reasons we strongly recommend that you enable password hash synchronization in your tenant.

## Pass-Through Authentication

PTA, if enabled, allows you to offload authentication from Azure AD to your on-premises AD without the need to deploy AD FS or a third-party federation solution. As such, it can greatly simplify your deployment if you are seeking to keep control of the actual authentication process.

For this to work, you must install and configure one (or more) on-premises PTA connectors to validate user authentication requests. The connectors can be installed as part of Azure AD Connect and are very similar to Azure App Proxy connectors. They share the same architecture; the PTA connector is a customized version of the Azure App Proxy connector.



From a high-level perspective, here are the steps an authentication request goes through when PTA is enabled:

1. The client connects to a service endpoint and is redirected to Azure AD for authentication.
2. The client connects to the Azure AD authentication endpoint and is prompted for credentials.
3. After the user submits their credentials and assuming PTA is enabled, Azure AD encrypts the credential data and holds the authentication request in a queue for validation.
4. The PTA connector makes an outbound connection to Azure and retrieves the authentication request from the queue.
5. The connector decrypts the data from the request and validates the decrypted credentials against the on-premises AD. The result of this verification process is then communicated back to Azure AD.
6. If the credentials were verified successfully, a token is issued, or the MFA flow is started.

## Seamless Single Sign-on

Seamless SSO enables Azure AD to accept a Kerberos ticket from the on-premises AD to authenticate the user if you are using password hash synchronization or PTA. The addition of Seamless SSO to either authentication model enables true SSO for users that are connecting from a domain-joined client computer. SSO is when users authenticate to both the on-premises organization and Azure AD using the same username and password without having to type their credentials again whenever they access a cloud resource. At a high level, this is what happens when Seamless SSO is enabled:

1. The client connects to a Microsoft 365 service and is redirected to Azure AD for authentication.
2. The client connects to the Azure AD authentication endpoint and is challenged for a Kerberos ticket.
3. The client turns to AD and requests a Kerberos ticket for the URL which is associated with Azure AD.
4. Active Directory locates this URL in the service principal name (SPN) of a computer account associated with Azure AD and encrypts a service ticket using that computer account's secret.
5. The client sends back the service ticket it received from AD.
6. Azure AD decrypts the Kerberos ticket using the computer account's secret which it received during the setup of the feature. If the ticket can successfully be decrypted, Azure AD will craft a token for the user.

For Seamless SSO to work, the following requirements must be met:

- Modern authentication must be enabled, and the client must support modern authentication.
- The client must be domain-joined and able to directly communicate with an AD domain controller. This is necessary for it to request a Kerberos ticket. Without direct access to a DC, SSO would fail and the process would fall back to the regular authentication option (username/password). That means that Seamless SSO cannot be used with Mac OS X (unless joined to a domain), mobile devices, or any other device that isn't joined to an AD domain.
- The Azure AD authentication endpoints must be added to the computer's browser Local Intranet zone settings; by default, browsers do not send Kerberos tickets to public endpoints.

Microsoft describes the SSO-feature to be "opportunistic". This means that SSO will be attempted if enabled. However, if anything goes wrong during the process, the authentication process will fall back to the default authentication option. We strongly recommend Seamless SSO combined with password hash synchronization if it meets your business and technical requirements. For more information about configuring Seamless SSO, refer to [this document](#). You should also explore whether hybrid Azure AD join can meet your SSO needs. We discuss hybrid Azure AD join later.

## Federated Authentication

Identity federation offloads credential validation, and optionally MFA, to on-premises federation infrastructures such as AD FS or Ping Identity, or a third-party cloud identity solution. A by-product of configuring identity federation is that it can provide an SSO experience.

**Note:** You can also achieve SSO much more easily with the Seamless SSO feature in Azure AD discussed earlier.

When you federate a domain, the responsibility for validating authentication requests is shifted towards the federation solution. The application that's asking the federation system to perform authentication is known as the *relying party* (because it's relying on the federation service). A federation standard named WS-Fed/WS-Trust is almost always used to federate with Azure AD. The standard specifies the format and content of data and metadata that the service and federation broker can use to negotiate and perform the authentication.

Federation depends on the concept of *claims*, which are statements made by one party about another. For example, the application might send a claim to the federation server that says "the user requesting authentication is coming from IP address 172.16.0.204" and the federation server might reply with a set of claims that includes "the email address associated with the account you are trying to authenticate is britta.simon@office365itpros.com." You can configure AD FS *claims rules* that allow or restrict authentication based on the contents of these claims.

Federation is enabled per domain and all users in Azure AD for whom the domain portion of the UPN (the UPN suffix) matches a federated domain are automatically considered to be federated identities (except if the user is included in a staged rollout group). Even with federated identities, the first authentication request is still received by Azure AD. The user's UPN suffix is examined and, if the domain name is federated, Azure AD will redirect or proxy the authentication request to the customer's federation solution. This process is called *home realm discovery* and needs to be performed for each initial authentication (such as the first time that someone logs on). You can see this process at work when you have federation set up and you visit a logon page; at first, you see the standard Microsoft logon dialog, which contains code that looks at the UPN and determines if it's for a federated user, redirecting to the organization's federation solution if so.

The details of the organization's federation solutions and all the relevant details such as the federation endpoint and certificate information are stored in Azure AD during the federation setup process. The initial request (home realm discovery), and the subsequent redirects all happen within a matter of seconds, often transparent to the user.

Most organizations choose to federate all their domains, but that is not a requirement. As such, you can have multiple authentication methods for different users in different domains of your organization. Information on how to configure identity federation with AD FS is described in Chapter 10 of the companion volume.

When authenticating with Azure AD, you will notice that most of the time the user is asked to enter their email address. However, that is not entirely true. Although the web page or application might ask for the "Email Address" on the screen, the user should enter their logon name (UPN) to complete the logon. Because of this, the general recommendation is to align the user's UPN and primary email address to remove any potential for confusion for the end-users. If your organization can't align the UPN and primary email address, you may be able to use Alternate Login ID, discussed later in this chapter.

### To Federate or Not?

Various elements influence the decision of whether identity federation is the right solution for you. Apart from the added infrastructure that is needed, complexity is often the reason why organizations avoid federation, especially when they must manage the federation solution themselves. Unsurprisingly, password hash synchronization is the most popular authentication method in Azure AD, as it is much easier to deploy and maintain, and it offers a similar login experience to end-users when combined with seamless single sign-on or

a device that is Azure AD or Hybrid Azure AD joined. In addition, users will still be able to authenticate, even when the on-premises infrastructure is down; the only thing that would (temporarily) stop working in such a scenario, is the actual synchronization of passwords.

On the other hand, identity federation (for example, through AD FS) can solve very specific problems, like whether to block authentication from outside the corporate network, limit access to Microsoft 365 services to members of a specific group, or use a third-party MFA solution. These types of problems can also be solved without AD FS if you have Azure AD Premium.

In general, we often recommend that you start with password hash synchronization and Seamless SSO for authentication with Microsoft 365 and Azure AD. If you have complex business or technical requirements that cannot be met without implementing federation, only then should you do so. Be sure to consider the tradeoff of complexity with the requirements you are trying to meet.

### Third-party Federation Solutions

AD FS is not the only federation solution to enable SSO. Several third-party solutions, such as those from Okta, Ping Identity, and OneLogin take a similar approach to federation. The decision to use a third-party solution depends on several factors such as integration with other cloud systems. Different solutions have different requirements and capabilities.

Azure AD also offers advanced SSO features like those supported by third-party vendors. In addition to these SSO features, Azure AD contains additional features such as enhanced multi-factor authentication and self-service password management. Many of these features require users to have a license for [Azure AD Premium](#) or EMS. While these capabilities come at an additional cost, they are worth evaluating along with other third-party solutions.

The pitfall you will run into if you elect to purchase a third-party federation solution is whether you can avoid the cost of Azure AD Premium. Many of the security controls that organizations require to secure their Office 365 implementations are native to Azure AD Premium. Without Azure AD Premium it may be difficult or impossible to implement these controls. Where it is possible to implement substitute controls, the user experience is often lacking. For these reasons, we find that many organizations find Azure AD Premium, or one of the bundles that includes it, is better value than third-party federation solutions.

## Alternate Login ID

Like many other cloud services, Microsoft 365 requires the logon ID to be *Internet routable* because ownership of non-routable domains like internal domains ending in ".local" cannot be verified. Users for which the on-premises UPN suffix does not match any registered domain are given a UPN suffix based on the default routing domain. For instance: *britta.simon@office365itpros.onmicrosoft.com*. Note that the tenant's default domain cannot be federated.

By default, when you configure AD FS, the UPN is used as the primary logon ID for Azure AD. Generally, it is recommended to ensure the UPN matches the user's email address so that they only need to remember their email address to sign in to Microsoft 365. Unfortunately, sometimes you might not be able to change the UPN to match the email address. If you are unable to change the UPN for your users, for example, because a legacy application needs a specific value, you can use a feature called "Alternate Login ID". This feature allows you to specify which attribute – other than the UPN – should be used to sign on to Microsoft 365. For instance, you can configure the actual Email Address attribute to be the new identifier. Although Microsoft supports Alternate Login ID for Hybrid deployments, there are some severe drawbacks from an end-user perspective (like extra authentication prompts). Unless you absolutely cannot reconfigure the UPN, we do not recommend choosing Alternate Login ID.

In [this](#) article, Microsoft describes the end-user experience for various applications and protocols if Alternate Logon ID is configured. Connections from within the corporate network are most likely to work just fine. External access, regardless of the client, can be problematic and result in extra authentication prompts. No matter what approach you use, the end-user experience suffers as soon as Alternate Logon ID is enabled.

You can [configure alternate login ID in AD FS](#), or you can configure it in Azure AD if you use password hash synchronization or PTA. When you [enable alternate login ID](#) in Azure AD (currently in public preview), users will be able to sign in to Azure AD using any email address configured for their account. These email addresses must be in their proxyAddresses attribute and belong to a domain namespace validated in the tenant. You can also use [staged rollout](#) to enable alternate login ID for a subset of your users.

## Passwordless Authentication

Passwordless authentication offers significant security benefits over passwords. By removing passwords, you can remove a significant attack vector. You can also improve user experience by removing the need for users to remember passwords, periodically change passwords, etc. If you have frontline workers that do not frequently use a computer as part of their job, passwordless authentication can reduce the friction involved in deploying IT services to this population. While it is not yet possible to entirely remove passwords, we strongly recommend beginning the journey of exploring how passwordless authentication can be integrated into your organization's IT strategy.

Whether you are in a hybrid or standalone identity model, Microsoft offers four passwordless authentication methods:

- **Text message sign-in:** In this model, users can authenticate to Azure AD with a one-time passcode sent via SMS/text message to their previously registered mobile device.
- **Microsoft Authenticator App:** In this model, users enter their UPN, and Azure AD sends a push notification to the Authenticator app. The user sees a random two-digit number on the sign-in screen and must enter the two-digit number into the Authenticator app. If the user provides the correct number in the Authenticator app, enters a PIN, or provides a biometric (e.g., Apple TouchID), they are signed in to Azure AD.
- **FIDO2 key sign-in:** [FIDO2](#) is an industry standard for authenticating to devices and online services using a standards-based hardware token.
- **Certificate-based authentication (preview):** this model allows users to sign in using a client certificate, typically stored on a smart card.

To use a passwordless authentication method, you must first enable the authentication method. To do this, log in to the Azure AD admin center and navigate to **Security** and then **Authentication methods**. From here, you can enable text message, Microsoft Authenticator, certificate-based authentication, or FIDO2 Security Key sign-in for a group of users, or you can enable it for all users. You can also use a preview feature to configure the Microsoft Authenticator app to show the estimated location of the authentication request and the application that generated the request when the user is prompted for approval. Once you enable a new authentication method or make configuration changes to an existing one, it might take a few minutes to start working.

Text message sign-in has several limitations. The most important limitation right now is that Teams is the only Office thick client application that text message sign-in works with. Otherwise, you can only use text message sign-in with web applications.

FIDO2 sign-in requires users to have a [supported](#) hardware token. Once the user has a hardware token, they can browse to <https://mysignins.microsoft.com>, click **Security Info**, click **Add method**, and choose **Security key**. The user will be walked through the process of pairing their security key with Azure AD. Once this is

complete, the user can click **Sign in with a security key** on future Azure AD sign-in prompts. You can also use FIDO2 keys to sign in to devices with [Windows Hello for Business](#) and to [sign in to AD joined](#) devices.

Certificate-based smartcard authentication is a critical scenario for a relatively small number of organizations, including the United States government. We do not expect that many organizations will use this authentication method unless they already have a mature smartcard deployment. Instead, we recommend that organizations invest in passwordless architectures based on Microsoft Authenticator and/or FIDO2 devices. For information on configuring certificate-based authentication, refer to the [Microsoft documentation](#).

## Temporary Access Pass

One of the challenges to deploying passwordless authentication is how to enroll a user with a FIDO2 token or another authentication method for the first time. If the user has a password, they can use their password (and MFA) to enroll. If they are a new employee, then you must give them a password for their initial sign-on, which defeats the concept of passwordless authentication. To address this challenge, Azure AD has a feature called Temporary Access Pass (TAP).

TAP lets you provide users with a short-duration unique code that the user can use to set up their passwordless credentials. When a TAP is created for a user, they attempt to sign in to Azure AD in the same manner as a normal user. Instead of prompting the user for a password, they will be prompted for their TAP. After entering a valid TAP, the user can enroll their passwordless authentication method such as a FIDO2 token or the Authenticator app. You can also use TAPs to restore access for a user that needs their passwordless credential replaced.

To enable TAP in your tenant, login to the Azure AD admin center and navigate to **Security** and then **Authentication methods**. Next, enable the Temporary Access Pass method. You can configure the lifetime of the TAP, whether it can be used more than once, and the length of the code. Once TAP is enabled, members of the global administrator, privileged authentication administrator, and authentication administrator Azure AD roles can create TAPs for users.

To create a TAP, login to the Azure AD admin center and navigate to **Users** the Azure and then find the user you want to create a TAP for. Click **Authentication methods** and then **Add authentication method**. On the screen that appears, you can create a TAP and optionally specify when it becomes valid. This capability may be useful if you want to create a TAP in advance for a new employee. You can make the TAP valid beginning on their first day of work, for example.

**Note:** If you do not see the Add authentication method button, you might need to click the **Switch to the new user authentication methods experience!** banner on the top of the screen first.

In addition to using the Azure AD admin center to configure TAPs, you can also use the [Graph API](#) to automate TAP management tasks.

# Understanding Azure AD Authentication

Establishing a standalone or hybrid identity to use with Microsoft 365 is only half of the work. Once you have an identity, you must first authenticate before you can access resources online. Leveraging the power of Azure AD, Microsoft 365 offers a variety of authentication options, ranging from the default username/password combination to more advanced authentication solutions.

Before diving into the advanced scenarios, let's first look at how Azure AD performs authentication. The simplest credential is the combination of a username and a password. This option is also the default in Microsoft 365 for all non-federated identities.

When a user attempts to access resources, the service prompts them for credentials. Depending on what client is used, the authentication prompt comes in various forms. Microsoft has worked hard to standardize the appearance of the Azure AD-based logon interface across clients and platforms, so the logon experience is remarkably consistent across devices and applications. No matter the interface, once the user provides a credential, various things happen behind the scenes depending on what client is used and what authentication method has been configured.

## Basic Authentication

When looking at the authentication process in Azure AD, one can distinguish two main approaches to authentication: basic authentication and modern authentication. The former is used when legacy clients (those that do not support modern authentication capabilities) connect to Microsoft 365. The number of these clients in use within Microsoft 365 tenants is steadily shrinking, but there are still examples of older clients and devices that should be updated.

When a basic authentication client connects, a basic authentication prompt is shown instead of the Azure AD-style prompt. Once the user enters their credentials, the credentials are forwarded to the service (e.g., Exchange Online) which, in turn, will connect to the underlying authentication solution (Azure AD or your federated authentication service) to validate the credentials. If the credentials are successfully verified, the service will authorize the client to connect to the resource it requested access to. This process is quite different from modern authentication where the client communicates directly with the authentication endpoint.

**Note:** Microsoft is pushing customers to stop using legacy authentication, or more precisely, to upgrade their clients to versions that support modern authentication. Legacy authentication methods, such as basic authentication, impose major limitations. The biggest limitation is that legacy authentication is subject to common attacks that do not affect modern authentication. In addition, you cannot perform MFA using basic authentication, nor can you exert the same levels of control (through policies) as with the modern authentication options. Because of these limitations, legacy authentication is also highly susceptible to attack. On October 1, 2022, Microsoft will disable basic authentication for seven email connectivity protocols in Exchange Online.

## Modern Authentication

Modern authentication enables improved authentication for a variety of clients, including Office 2013 (with the latest updates) and newer, Office for iOS and Android, the native iOS mail app, the Outlook mobile application, and Teams. Modern authentication is based on the use of the OAuth 2.0 authorization protocol. Modern authentication allows enabled clients to take advantage of the following:

- SAML-based sign-in with third-party MFA providers.
- Smart card-based authentication.
- Passwordless authentication.
- True multi-factor authentication.
- SSO across apps through token sharing (for example, the Office desktop apps).
- Application of Azure AD conditional access policies.

## Authentication Flows

As briefly discussed earlier, the biggest difference between modern authentication and legacy (basic) authentication methods lies in how the authentication happens and how communication flows between the client and Azure AD. Figure 3-1 depicts the authentication flow of a desktop client with modern authentication enabled:

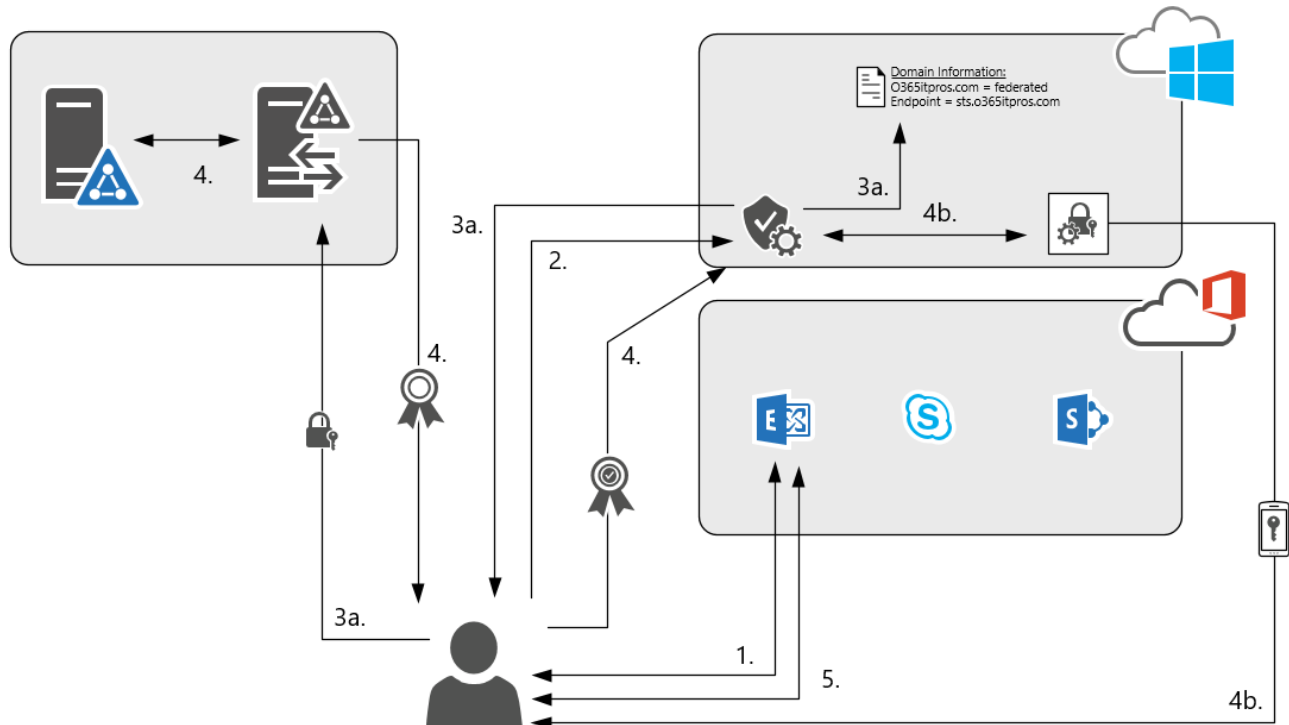


Figure 3-1: Modern authentication flow

1. The Outlook desktop client connects to Exchange Online and is redirected to Azure AD for authentication.
2. The client connects to the Azure AD authentication endpoint and is prompted for credentials. Because modern authentication is enabled, the client displays the forms-based modern authentication credential dialog.
3. After the user submits their credentials, Azure AD first verifies the identity type of the user. It does so by looking at the user's UPN and verifying what type of authentication the user is configured to use.
  - a. If the user is federated, the client is redirected to their identity provider. Depending on what client the user is connecting with and how the endpoint is configured, the user might need to enter additional credentials.
  - b. If Seamless SSO is enabled, the user's browser is challenged for a Kerberos service ticket to identify the account. If the browser returns a service ticket to Azure AD, the service ticket is validated, and the user is signed in.
  - c. If PTA is enabled, the user's request is passed on to the on-premises AD via PTA agents.
  - d. If the user's password is managed in Azure AD (standalone identity or using password hash synchronization), the username/password combination is verified by Azure AD itself.
4. When the credentials are successfully validated, one of the following happens:
  - a. In the case of federated authentication, the application receives a token (from the identity provider) and is then redirected to Azure AD where the user receives an artifact that identifies the user.
  - b. If MFA is required, and the user's federation service did not tell Azure AD that MFA was already performed, the user is first requested to enter additional verification after which (if successfully validated) the identifying artifact is created and returned.
5. The client is now redirected to the service it initially connected with (for example, Exchange Online). Using the artifact previously received from Azure AD, the client can now access the online resource.

Following successful authentication, Azure AD issues two tokens to the user: an access token and a refresh token. The access token is short-lived, meaning it has a limited lifespan (one hour by default). An app like OWA uses the access token to authenticate to Exchange Online. The refresh token, which can remain valid

indefinitely, is used to request new access tokens from Azure AD when they expire. With a valid refresh token, the user does not have to re-authenticate. Instead, the refresh token is used to obtain a new access token from Azure AD. The flow described above assumes the client has no valid refresh token, which would be the case when the refresh token expired or when the client has not yet connected to Azure AD.

Several things control the validity of the refresh token, such as how often it is used or whether the user changes their credentials. When the refresh token expires, the user must sign in again. Depending on the client, and how authentication is implemented, this might require a user to provide their credentials again. If true SSO is used, re-authentication should happen automatically and is fully transparent for the end-user. If no SSO is configured, the user will be presented with an authentication form.

## Revoking Access/Tokens

As mentioned in the previous sections, Azure AD authentication uses two types of tokens: access and refresh tokens. The access token ultimately grants you access to the online service: with a valid access token, you can access the requested service and do something like opening a mailbox. You can continue to work with the service until the access token expires or you interrupt or close the active session (like closing the browser window). When an access token expires, the client can acquire a new access token using a valid previously-received refresh token. If a valid refresh token exists, the user does not have to re-authenticate.

Because of how these tokens are used, organizations sometimes face a new challenge: what if you need to revoke someone's access immediately? The same problem also exists in an on-premises organization to a certain extent: when you disable an account in Active Directory, subsequent authentications are no longer allowed, but the user is not forced to close any open session/applications being used at that time.

Access tokens cannot be revoked. This means that once someone gets a valid access token, they can keep using that token for up to its maximum lifetime. You can, however, revoke the refresh tokens which prevents the user from acquiring a new access token when it expires, and thus force the user to re-authenticate (which will fail, if you disable the account). To revoke valid refresh tokens for a user, log in to the Azure AD admin center and navigate to **Users**, find the user you want to revoke the refresh tokens for, click **Profile**, and then click **Revoke sessions** on the toolbar. You can also revoke all the refresh tokens for an account by running the following PowerShell commands:

```
[PS] C:\> $UserId = (Get-MgUser -UserId Andy.Ruth@Office365itpros.com).Id
Invoke-MgInvalidateUserRefreshToken -UserId $UserId
```

Although revoking access tokens is not ideal, it will at least prevent the user from continuing to obtain valid access tokens. Some organizations need to customize the default sign-in frequency behaviors. If you have Azure AD Premium, you [can use a feature](#) of conditional access to control how frequently users must sign in to access applications.

Modifying lifetimes for tokens was a previously acceptable approach to mitigating risk. A better approach is for tokens to automatically get revoked when a risk event occurs. Microsoft has taken the first steps towards this with a feature called [Continuous Access Evaluation](#) (CAE). With CAE, when a risk event occurs, the workload (Exchange Online, SharePoint Online, OneDrive for Business, or Microsoft Teams) is notified so that it can take near real-time action to revoke access. When CAE is enabled, the one-hour access token lifetime discussed earlier is increased to 28 hours for CAE-capable applications. Since the workload is continuously evaluating the validity of the access token, the increased lifetime does not present a new or additional security risk.

The events below currently trigger a notification to the subscribing workloads:

- The user account is disabled/blocked.
- The user account is deleted.
- The user's password is changed (or reset).



- MFA becomes mandatory for the user.
- Refresh tokens are administratively revoked for the user.
- Azure AD Identity Protection detects elevated risk for the user.

CAE is enabled by default in all tenants. If you want to disable CAE or modify the behavior of CAE, you can do this with a conditional access policy (discussed later in this chapter). You can use conditional access to place CAE in strict enforcement mode for some or all users. In strict enforcement mode, CAE will revoke access if a user's IP address changes or when the client [lacks the capabilities to handle CAE processing](#).

## Modern Authentication in a Hybrid Environment

As of Exchange Server 2013 CU19, Exchange Server 2016 CU8, and Skype for Business Server 2015 CU5, modern authentication [is also supported](#) for on-premises Exchange servers, provided that a full hybrid configuration has been set up. This is good news, as it allows you to align the authentication capabilities (and experience) cross-premises.

The way modern authentication works on-premises is very similar to how it does in Exchange Online. However, instead of being a native capability, the on-premises Exchange or Skype for Business servers will pass authentication requests to Azure AD and use the OAuth tokens Azure AD returns instead of using the native on-premises authentication mechanisms. In effect, when a user attempts to connect to Exchange or Skype for Business, they are redirected to Azure AD to authenticate there. Once a user authenticates successfully, they will receive access and refresh tokens from Azure AD which can then be used to authenticate to Exchange or Skype for Business on-premises.

Note that enabling modern authentication for Exchange or Skype for Business on-premises is an all-or-nothing scenario; you cannot enable it on a per-user basis.

# Customizing the Tenant Sign-In Page

Tenants can customize the sign-in page presented to users when they sign in via Azure AD. Customizing the sign-in page lets you apply your organization's look and feel by modifying background colors and images as well as replacing the Microsoft logo. Applying your organization's brand has the added security benefit of giving your users a recognizable place to enter credentials. Customization is not hard, but it usually takes some homework to identify the right images and colors. Partnering with your marketing or internal communications teams is usually the best way to be successful.

When a user accesses resources from another organization using Azure AD B2B Collaboration, the branding experience changes slightly to provide context about the user's organization and the organization they are accessing. In this scenario, the background image (element 1) changes to the background of the organization that owns the resource. The remaining elements represent the user. For example, if a user from the Coho Vineyard is accessing a resource in Contoso's tenant, they will see Contoso's background image and the Coho Vineyard's banner logo, username hint, and sign-in page text.

The branding elements you choose apply to all your users, with one exception. You can define different branding elements based on the user's language. The user's language is determined dynamically by information provided by their web browser about language settings on their computer. To configure branding, login to the Azure AD admin center and navigate **Company branding**. If you edit the default settings, they will apply to all users. To create different branding settings based on user language, click **New language** on the toolbar.

If your organization uses AD FS or another identity provider, the branding settings in Azure AD will not be visible in some scenarios. You should also customize the sign-in page for your identity provider. More information about branding AD FS' sign-in pages is located [here](#). You will find that the customization

elements are very similar to Azure AD. If you're running AD FS on Windows Server 2019, don't forget to [enable paginated sign-in](#). This makes the AD FS sign-in process much more like Azure AD.

Figure 3-2 shows the sign-in page with all the available branding elements customized. The numbers in Figure 3-2 correspond to the branding elements in Table 3-1. Pay attention to the notes in Table 3-1 too. For the best user experience, you will need to select images that work best with how Azure AD displays the branding elements.

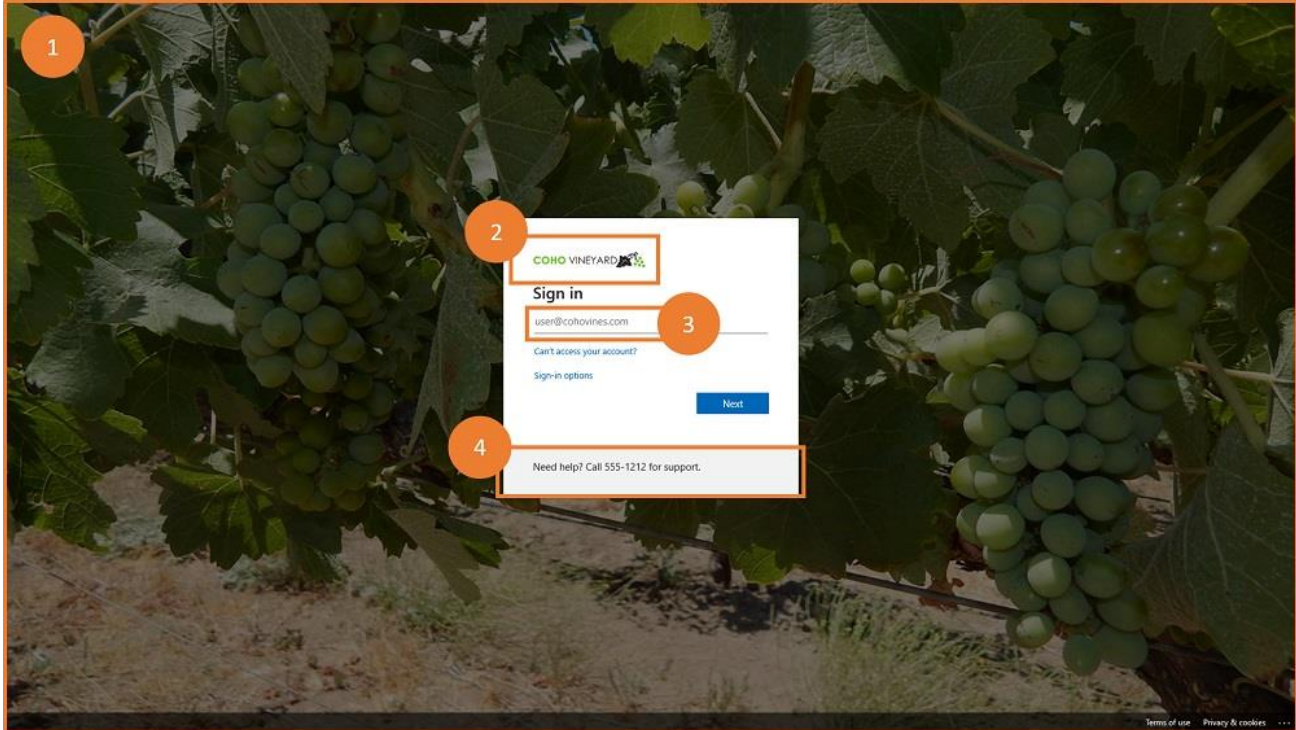


Figure 3-2: A Branded Azure AD sign-in page

#	Element	Requirements	Notes
1	Background image	<b>Dimensions:</b> 1920x1080px <b>Format:</b> PNG, JPG, or JPEG <b>Size:</b> 300KB or less	<ul style="list-style-type: none"> <li>This image is dynamically cropped based on screen size. Avoid using an image that includes text.</li> <li>Minimize the file size where possible to improve page load speeds.</li> </ul>
2	Banner logo	<b>Dimensions:</b> 280x60px <b>Format:</b> Transparent PNG, JPG, or JPEG <b>Size:</b> 10KB or less	<ul style="list-style-type: none"> <li>Use a transparent PNG file</li> <li>This image is shown in automatically generated email communications and other places like the MyApps portal.</li> </ul>
3	Username hint		<ul style="list-style-type: none"> <li>Text displayed as a reminder to the user about how to format their sign-in name.</li> </ul>
4	Sign-in page text		<ul style="list-style-type: none"> <li>Optional but useful for helpdesk contact info, legal disclaimers, etc.</li> </ul>
Not Shown	Sign-in page background color	RGB color code (hex format)	<ul style="list-style-type: none"> <li>Shown instead of the background image when the screen size is too small.</li> </ul>

#	Element	Requirements	Notes
Not Shown	Square logo image	<b>Dimensions:</b> 240x240px <b>Format:</b> PNG, JPG, or JPEG <b>Size:</b> 50KB or less	<ul style="list-style-type: none"> <li>Used on Windows devices during Azure AD join, Windows Autopilot, and other experiences.</li> <li>Provide two versions – one for light backgrounds, and another for dark backgrounds.</li> </ul>

Table 3-1: Branding elements

## Joining Computers to Azure AD

Windows 10 and newer versions (and to a more limited extent, earlier versions of Windows) support the concept of joining client computers to Azure AD. There are two different ways to join Azure AD: 1) Azure AD join only and 2) hybrid Azure AD join. In the first scenario, the client computer only has a relationship with Azure AD. Rather than signing into the device with a local account, users sign in with their Azure AD credentials. This can be a standalone identity or a hybrid identity. In the second scenario, the client computer remains joined to an AD domain (in the same manner it has for decades), but the device also maintains a parallel relationship with Azure AD. Users sign in with their AD credentials and the computer also acquires a primary refresh token (PRT) used to authenticate with Azure AD (and in turn dependent resources like Office 365 workloads).

Azure AD join can be useful for any organization, including smaller organizations that do not have an on-premises AD forest. If you have a workforce that primarily works remotely, you can achieve the benefits of domain join without the need to have network connectivity to reach a DC. Even if a client computer is solely joined to Azure AD, if there is an AD forest, that user can [seamlessly access](#) traditional resources and applications that use Kerberos authentication when they are on the network. Azure AD join requires no specific configuration to begin using the feature. You can control who can join a device to Azure AD by signing into the Azure AD admin center and navigating to **Devices** and then **Device settings**. Configure the **Users may join devices to Azure AD** setting. We also recommend that you require MFA to join a device. The best way to enforce this requirement is with a conditional access policy. We discuss conditional access and applying a policy to user actions later in this chapter.

Hybrid Azure AD join requires some initial configuration to use the feature. Windows 10 and newer devices locate the Azure AD tenant they should attempt to hybrid join by searching AD for a service connection point (SCP). The easiest way to create the SCP is to use Azure AD Connect. The configure device options task in Azure AD Connect will create the SCP record for you. If you use federated authentication, you will need to make sure that your identity provider supports Hybrid Azure AD join. AD FS supports hybrid Azure AD join, however, it must be configured to [issue certain claims](#) specific to computers. Azure AD Connect will also configure AD FS for you.

After configuring hybrid Azure AD join, you can begin using a client computer's identity as an input to conditional access, discussed later. You can also achieve SSO from Azure AD joined and hybrid Azure AD joined devices without configuring seamless SSO. Microsoft Edge and Mozilla Firefox (beginning with version 91) natively support this form of SSO, however, you can also deploy an [extension](#) for Chrome if your organization uses it. In both cases, the browser can submit your PRT to Azure AD at sign-in. The PRT identifies the user, whether they have performed MFA (or not), and the device. You can check on the PRT as well as the general health of a device's Azure AD relationship by running **dsregcmd /status** from a command prompt.

# Guest Access to Azure AD

Azure AD provides a rich and flexible fabric for bringing external people into your tenant in the form of guest accounts with Azure AD Business-to-Business (B2B). A guest is an external user whose account and credentials exist outside your tenant (often in another tenant). Azure AD B2B is the component of Azure AD that enables guest accounts to work. One of the most valuable benefits of Azure AD B2B is the ability for guest users to authenticate using their organization's credentials while accessing an application in another organization's tenant. For example, a user from contoso.com needs access to a file stored in SharePoint Online in the Office365ITPros.com tenant. The user from contoso.com will authenticate to Office365ITPros.com using their contoso.com credentials, rather than a separate username and password. This also means that when the user's contoso.com account is deactivated (such as when they leave the organization), their access to Office365ITPros.com will also be deactivated.

Azure AD assigns [guest accounts a more restrictive set of permissions](#) than it does to member accounts. However, you can change the **Guest users' permissions are limited** setting to No in the Azure AD admin center if you want guest accounts to have the same default permissions to see information in the directory (like group memberships) as member accounts do. Guest accounts can have administrative roles and execute functions associated with those roles.

Because guest accounts are Azure AD objects that share many of the characteristics of user accounts, applications can use guest accounts to control access to content. Azure AD uses email addresses to create user principal names to identify guest accounts. The email address can point to another tenant domain within Microsoft 365, an email domain for a company that doesn't have a tenant, or a consumer email service, like Gmail.com or Yahoo.com. You can view a list of guest accounts in a tenant through the Users section of the Microsoft 365 admin center, the Users section of the Azure AD admin center, the Contacts section of the Recipients tab in EAC, or by running the command:

```
[PS] C:\> Get-MgUser -All -Filter "userType eq 'Guest'" | Format-Table DisplayName, UserPrincipalName
```

You can edit properties of guest users through the Microsoft 365 admin center or Azure AD admin center. You can't make changes to guest accounts through EAC because the list of guests displayed there is for informational purposes only. It is common practice to update visible properties such as display names, telephone numbers, job titles, and addresses for guests. And as we'll see later, it is also possible to upload a photo for a guest account.

## Creating Guest Users

You can create guest users through [the Azure portal](#) by creating a **New guest user** rather than a **New user**, and the account creation process allows you to add the new guest user to groups. You can also create guest users with [PowerShell](#), but the most common method used is when applications issue invitations to external users, which is what happens when you add a guest to an Office 365 group or team. If you have regular user accounts created for external people in your tenant, these accounts can be [converted into guest accounts](#).

When Azure AD or an application add a new guest, it calls [the Invitation API to create and send an invitation](#) via email to the user's address to tell them that they have been invited to start using an application in the tenant. In addition to allowing applications to send customized invitations, the API supports the inclusion of a link for the user to access the application, like the access information needed for someone to join a group or team.

The *UserState* of a guest user created when an invitation is generated is *PendingAcceptance*. This is enough to allow the guest user to be added to a group. At this stage, the guest user is a placeholder. When the guest redeems the invitation, the *ExternalUserState* property changes to *Accepted* and Azure AD puts the account

into a state where it can access tenant resources. You can find out what accounts are in a specific state with the `Get-MgUser` cmdlet. In this example, we list the guest accounts in a pending acceptance state. The `ExternalUserStateChangeDateTime` property is the timestamp for the last account update.

```
[PS] C:\> Get-MgUser -All -Filter " userType eq 'Guest' and ExternalUserState eq 'PendingAcceptance' | Format-List -Property DisplayName, UserPrincipalName, ExternalUserState, @{e={Get-Date($_.ExternalUserStateChangeDateTime) -format g};n="Invitation Issued"}
```

It is important to realize that a guest user often uses their email address to sign in (ideally, the primary email address and the UPN are the same in their home tenant). The [New-MgInvitation](#) cmdlet has a parameter named `ResetRedemption` that indicates to Azure AD that the invitation is to change the email address used to sign in. All the other properties of the guest user are unaffected. After the user redeems the invitation sent to the new email address, the guest user switches to that address. This approach preserves access to all resources available to the guest user. The ability to change the email address for a guest account is [currently a preview feature](#).

## Guest Sources

The source for the guest is a property of the user object that tells Azure AD where the guest comes from and how it will authenticate. You can see the source by looking at the profile of guest accounts in the Azure portal, where accounts have one of [four sources](#):

- Users from **other tenants** have “External Azure Active Directory” as their source. These guests sign into Azure AD using the credentials from their home tenant.
- Users of **Microsoft consumer services** such as Outlook.com have a “Microsoft account” as their source. These guests sign into their Microsoft account (MSA).
- Users of **other services** such as Google.com also have “Microsoft account” as the source. In this case, guests sign into the Microsoft account created for them during the invitation redemption process and use those credentials to access resources.

The source for the guest account for anyone who has not redeemed an invitation is “**Invited User**.” When the redemption process is complete, the source changes to one of the sources listed above.

The host tenant extending invitations to guests can require multi-factor authentication to force guests to use a mixture of credentials: their password and another authentication method (like a one-time code sent to the guest’s mobile device).

All applications that use the Azure AD B2B collaboration framework generate and redeem invitations in the same manner. The difference is the link contained in the invitation – it can point to an individual team, group, or other application and the application referenced must be ready to allow access to that guest. SharePoint is an exception in that it creates a guest user account in your tenant when an external user from another tenant redeems a sharing invitation to a document or folder with the usual one-time code mechanism. SharePoint does this to allow these users to use their tenant credentials to open the shared documents with the Office desktop apps.

## Audit Records for Guest Additions

Azure AD captures details of the addition of a new guest account in an “*Add User*” audit record. Other audit records are captured for “*Add member to group*” events when a new member joins a group. Together, these audit records give enough information to allow administrators to track and report when new guests join groups in the tenant. To extract the information, use the Audit log search feature in the Microsoft Purview Compliance portal, or use the PowerShell `Search-UnifiedAuditLog` cmdlet. Several examples of how to use the cmdlet to extract, refine, and analyze audit log data are included in Chapter 21.

## Federated Identity Providers

Azure B2B Collaboration supports federated identity providers. Federation means that guest users can sign in to access tenant resources without needing to create a Microsoft account or account in the local tenant. The first [federated identity provider is Google](#), where support is available for sign-ins for guest accounts using the Gmail.com domain. [Federation with Facebook](#) is also supported. See [this article](#) for information about how to configure federation between Microsoft and Google. Teams also supports federated authentication for Gmail accounts.

## One-time Passcodes

One-time passcodes (OTP) for guest accounts allow users with any email address to authenticate using codes sent to their mailbox. The codes last for 24 hours. During this period, the user can access any resource they have been granted access to in the tenant. The [process to enable OTP](#) is simple and once the policy is active, any guest account whose email address is not part of an Azure AD tenant, a Microsoft account, or belongs to a federated identity provider, will receive a one-time code that they can use to redeem their invitation. The initial token lasts 24 hours and once it expires, the user must receive another code to reauthenticate.

## Creating Invitations

You can reproduce what an application does when it generates an invitation to a guest with a series of PowerShell commands. It is critical to understand that this code is an example not intended for production use. Instead, it illustrates some of the processing that occurs when a group owner invites a new guest to join a group.

First, we run the *New-MgInvitation* cmdlet to generate and send an invitation to the email address of the guest we want to add, in this case, [John.Doe@outlook.com](mailto:John.Doe@outlook.com). If the user accepts the invitation, the redirect URL brings them to the specified URL.

```
[PS] C:\> New-MgInvitation -InvitedUserDisplayName "John Doe" -InvitedUserEmailAddress
John.Doe@outlook.com -SendInvitationMessage:$True -InviteRedirectUrl
https://office365itpros.sharepoint.com/sites/GroupName
```

When you generate an invitation to a guest, Azure AD checks whether a guest users already exists in the host tenant. If one exists, it is used. If not, Azure AD creates a new guest user. You can check this by running the *Get-MgUser* cmdlet after sending the invitation. The account has a *UserType* of "Guest" and the User Principal Name is based on the email address used in the invitation. In our example, the UPN is `John.Doe_outlook.com#EXT#@office365itpros.onmicrosoft.com`. If you don't want Azure AD to send a standard invitation email, set the *SendInvitationMessage* switch to `$False`. You can later send your own customized message to the guest to invite them to join the tenant.

After they receive the message containing the invitation, the guest clicks the link to accept the invitation. If their account is already validated, they can go ahead and access the resource pointed to by the URL. If the guest account has never been used before, accessing a resource causes a redemption process to begin during which the account is validated. Following successful validation, the guest can go ahead and access the resources to which they have been given access. Azure AD updates the guest account to set its *ExternalUserState* property to `Accepted` and the *ExternalUserStateChangeDateTime* property with the current timestamp. We can therefore check for outstanding guest acceptances with some code to retrieve all guest accounts whose state is not accepted.

```
[PS] C:\> [array]$Guests = Get-MgUser -All | ? {$_.UserType -eq 'Guest'} | Sort DisplayName
ForEach ($Guest in $Guests) {
    If ($Guest.ExternalUserState -ne "Accepted") {
        Write-Host $Guest.DisplayName, $Guest.UserState, $Guest.Mail, $(Get-
Date($Guest.ExternalUserStateChangeDateTime) -Format g) }}
```

It all sounds simple, but the dependency on directory replication makes this approach to creating guests difficult to use in practice. For that reason, it is best to use the applications that you want to use guests with to invite those guests so that Azure AD creates their accounts through the redemption process.

## Restricting Guest Access

The [Azure AD Guest user access restrictions policy](#) in the [External collaboration settings blade](#) of the Azure AD admin center allows tenants to control the level of access guests have to Azure AD information. Three options are available:

- Guests have the same access as members (most inclusive) setting means guests have the same access to directory data as regular users in your directory.
- Guests have limited access to properties and membership of directory object settings. Guests don't have permissions for certain directory tasks, such as enumerating users, groups, or other directory resources. This is the default setting.
- Guests are restricted to properties and memberships of their directory objects (most restrictive). The ability to set this level of restricted access is currently in preview.

Most Microsoft 365 apps currently are unaffected if the most restrictive level is selected. This might change in the future as product groups work out how to exploit the feature. For instance, Teams might not allow guests to see details of team membership.

## Cleaning Up Old Guest Accounts

Many reasons exist to justify the creation of a guest account in a Microsoft 365 tenant. For example:

- An external expert is asked to review a document stored in SharePoint Online and receives a sharing link for this purpose. The link creates a guest account.
- Someone is asked to join a Microsoft 365 group or team. A guest account is created when the invitation is issued.

Over time, the reason why an external person has a guest account in a tenant might wane. For example:

- A guest account is used to review a shared document and is not needed thereafter.
- External people leave a team (or teams) and their guest account remains in Azure AD.
- People leave an employer and move on to new challenges. Their guest account is invalid because they can no longer authenticate using the Azure AD instance for the tenant of their old employer.
- Projects come to a natural end and the associated teams and/or private channels and their guest members are no longer needed.

You can, of course, leave guest accounts in place on the basis that they might be needed in the future, but usually, it is a better idea to clean things up and remove these accounts when they are no longer used, especially because these accounts have access to tenant resources. Several ways exist to control the longevity of guest accounts in a tenant:

- You can use the script covered in the PowerShell chapter (see the section “Finding Inactive Guest Accounts”) to find groups with external members and then check the individual members of each group.
- You can check guest accounts based on their activity level. The “Finding Inactive Guest Accounts Based on Activity” section in the PowerShell chapter describes how to use data from the audit log and email tracking logs to calculate the last activity for guest users. If a guest hasn't been active for months, perhaps it is time to remove their account.
- You can deploy [Azure AD Access Reviews](#) to force group owners to check the memberships of the groups they manage and attest that group members should keep or lose their status. Example reviews include all members of Microsoft 365 groups or Inactive guest accounts (accounts that

haven't been used for a set period). Access Reviews requires Azure AD Premium P2. A Graph API is available for programmatic access to access reviews (see this [example script](#)).

After target accounts are identified for removal, you can ask group owners to remove them or do so programmatically or through an administrative portal. Here's an example of using the *Remove-UnifiedGroupLinks* cmdlet to remove a member. If an account is a group owner, remember to remove it from the owner list first.

```
[PS] C:\> Remove-UnifiedGroupLinks -Identity "GDPR Planning Mark II" -LinkType Member -Links JackSmith_hotmail.com#EXT#
```

To remove a guest account completely, run the *Remove-MgUser* cmdlet:

```
[PS] C:\> Remove-MgUser -UserId JackSmith_hotmail.com#EXT#@office365itpros.onmicrosoft.com
```

Alternatively, you can leave it to individual guests to decide when it is time to remove their account from your tenant. A user can belong to up to 500 Azure AD tenants. They can choose to leave a tenant and remove their guest account at any time. To leave a tenant, the user goes to the Azure AD [Organizations page](#) via their Account details (they can also navigate to the page via the Teams desktop or browser clients). Azure AD lists the tenants where the user has a guest account. To leave a tenant, select **Leave organization**, and then confirm the decision with **Leave** (the user might have to sign into the tenant first). Azure AD then removes the guest account from the chosen tenant directory, signs the user out of the tenant, and sends an email to confirm that they have left the organization and can no longer access applications in that tenant. Leaving an organization only removes the guest account; it does nothing to remove any data which the user created in that tenant.

The guest account stays in a soft-deleted state for 30 days following removal. During this period, the tenant administrator can restore the guest account, or complete the process by removing the account permanently. When a guest account is removed, it loses all access to SharePoint and OneDrive documents, libraries, and lists that it has been granted access to on an individual basis or through a group. It also loses access to all Teams it was a member of in the tenant.

**Entitlement management:** [Azure AD entitlement management](#) is functionality to enable tenants to manage access to groups, applications, and sites at scale. You define access packages to describe the resources users can access and then assign the packages to users, making it easy to remove access to multiple resources at one time by removing an account's access to a package. Entitlement management requires Azure AD Premium P2.

## Cross-Tenant Settings

One of the challenges is the potential to allow guests from competitors access to tenant resources or to allow users from your organization to join teams and groups in a competitor's tenant. To mitigate this, Azure AD provides controls to limit guest collaboration. These controls come through external collaboration settings, as well as a preview capability called cross-tenant access settings.

Cross-tenant access settings are the primary control that you should plan to use where possible. The primary limitation of cross-tenant access settings is that they cannot control the invitation of consumer identities (e.g., Gmail) or other email domain names that are not verified in an Azure AD tenant. For these scenarios, you should use external collaboration settings.

You can find both groups of settings by signing into the Azure AD admin center and navigating to **External Identities**. External collaboration settings allow you to determine what email domain names can be invited to collaborate with your tenant as B2B guests. You can use the following settings under **Collaboration restrictions**:



- **Allow invitations to be sent to any domain (most inclusive)** – this allows B2B collaboration in your tenant with any guest user
- **Deny invitations to the specified domains** – this allows B2B collaboration in your tenant with any guest user, except guests from the domain names listed (e.g. a competitor)
- **Allow invitations only to the specified domains (most restrictive)** – in this scenario, you must whitelist individual domain names that you will allow guests into your tenant from

An important distinction between external collaboration settings and cross-tenant access settings, which we'll discuss next, is that external collaboration settings only control guests coming *into* your tenant. You must use cross-tenant access settings to control what tenants your users can be guests in. With cross-tenant access settings, you can apply much finer-grained control of guest behavior in your tenant ("inbound" access) and where your users can be guests ("outbound" access).

Rather than a domain name-based approach, cross-tenant access settings work on a tenant basis. You can control what guest users in a given tenant can access your tenant, as well as what applications they can access in your tenant. You might choose to allow all users in another tenant to become guests in your tenant, or you might choose to only allow specific users to access specific applications as a guest. The granularity of control that you choose will be a balance of managing risk and the operational overhead. You can apply similar controls in the opposite (outbound) direction to control what tenants your users can become guests in and what applications they can access in those tenants.

To configure controls for a specific tenant, go to **Cross-tenant access settings (Preview)**, and then click **Add organization** on the toolbar. Enter a domain name valid for the tenant you want to configure and then click **Add**. Note that these settings will apply to the entire tenant, not just that domain name. You will see that the *Inbound access* and *Outbound access* columns say *Inherited from default*. You can configure the defaults on the **Default settings** tab. These defaults apply to all tenants, even if they are not listed on the **Organizational settings** tab. Click **Inherited from default** and then choose **Customize settings** to configure tenant-specific settings:

- **External users and groups** – You can choose to **Allow access** or **Block access** from this tenant (inbound) or to this tenant (outbound). For inbound access, if you choose to **Select external users** and groups, you must provide the object identifiers for the individual users and groups that you want to allow access for. You can obtain the identifiers from an administrator in the external tenant. For outbound access, you can select the users and groups from your tenant allowed to become guests in the tenant.
- **Applications** – If you **Select applications**, you can choose the applications in your tenant that guests have inbound access to. For outbound access, if you **Select applications**, you must provide the application ID from the external tenant that users in your tenant can access. You can obtain the application ID from an administrator of the external tenant.

For inbound access settings, you can also configure trust settings for the external tenant. We discuss how these settings work in the Cross-Tenant Trust section in the discussion about Conditional Access. You might find that cross-tenant access settings do not work quite how you would expect, especially if you have used external collaboration settings before. External collaboration settings control what domain names can receive an invitation to become a guest. These invitations are often generated when a user tries to share content in SharePoint, OneDrive for Business, or Microsoft Teams. When a domain is not permitted to be invited, the user will receive an error message when they try to send an invitation. With cross-tenant access settings, users will still be permitted to send invitations to users in tenants that are not allowed to collaborate. When the guest tries to accept the invitation to collaborate, they will receive an error message that they are not permitted to access your tenant. We expect that this inconsistent experience will evolve as the preview of cross-tenant access settings matures.

# Protecting Your Identities

Microsoft has made significant investments in delivering features in Azure AD (and even on-premises) to combat the threats to identities. In today's world, every organization should have MFA deployed across the enterprise and enforced for all users. In October 2019, despite the widespread availability of MFA, Microsoft reported that MFA protects fewer than 10% of enterprise Azure AD accounts. Even with MFA, passwords are still a problem. Traditional password policies that require frequent password changes and complexity rules that are difficult to comply with lead to people picking even worse passwords. Azure AD Password Protection is Microsoft's approach to mitigating some of the risks of passwords.

End users are ultimately the most familiar with their behaviors. Azure AD makes sign-in logs available for users to review directly by visiting <http://mysignins.microsoft.com>. Information about each of the user's sign-in attempts including whether they succeeded and where the attempt was made are all presented in a searchable list. This same data is available tenant-wide to administrators in the Azure AD sign-in logs.

## Password Protection

Users often choose poor passwords. Sometimes they reuse passwords between services; other times they choose easily guessed passwords. Microsoft has long tried to educate both users and administrators about the rules of good password hygiene, and the on-premises Windows Server operating system has a long-standing [password filter feature](#) that allows you to install custom code to check proposed passwords when users try to change them on-premises. Two similar sets of Azure AD features are available to do this: [one for standalone identities](#) and the other for hybrid identities that are synchronized to the cloud.

First, for user accounts homed in the cloud, Azure AD applies a global banned password list when a password change is requested. All the old standards, such as "passw0rd" and "password123," are here, as are many others; the list is gleaned from password breach data that the Microsoft Security Intelligence center gathers, as well as from other factors. The suggested new password is first put through a set of normalization rules (similar in concept to the way phone numbers are normalized in Skype/Teams), then checked against the global bad password list, then checked against the organization's custom bad password list if you've defined one. The global banned password list is available to all tenants, but defining custom banned passwords requires an Azure AD Premium license for each user.

The on-premises equivalent of this feature is more complex. To use it, you install an agent on your DCs that checks password change requests against the Azure AD banned password list. There is also a proxy agent, which downloads the password policy from the service and makes it available to the DCs. When a user tries to change her password, the local DC agent (which is implemented using the password-filter mechanism that's long been part of on-premises AD) checks it against the most recent list and either allows or blocks the change. The good news: this approach means that your cloud and on-premises identities receive the same protection and use the same banned-password lists. The bad news is that the on-premises capability requires you to purchase Azure AD Premium licenses for each user.

**Hardening User Passwords:** Azure AD allows accounts to have passwords of up to 256 characters. Increasing the password length by just a single character dramatically increases the time to brute-force guess the password. However, as computers become increasingly more powerful, the processing needed to brute-force guess a password also decreases, and continually increasing the password length is not a long-term solution. If you want to use a long 200-character password, feel free, but there is a better solution now. After years of telling users that they needed long, complex passwords, the US National Institute of Standards and Technology (NIST) introduced a [new set of recommendations](#):

1. Use pass phrases instead of passwords. Phrases are easier for humans to remember.
2. Choose phrases with unique associations that only you will know.

3. Eliminate character-composition requirements.
4. Do not require passwords to expire.
5. Ban common passwords, to keep the most vulnerable passwords out of your system.
6. Educate users not to re-use their passwords for non-work-related purposes.
7. Enforce the use of multi-factor authentication wherever possible.
8. Enable risk-based multi-factor authentication challenges.

The [full NIST recommendation](#) makes for interesting reading, but if you implement the suggestions above, you'll greatly increase your security and, not incidentally, make your users happy. In addition to hardening user passwords, you should also monitor authentication attempts to track and report suspicious activities. Both options are discussed later.

Password Protection helps protect you from password spray attacks. Password spray attacks are a common technique used by adversaries to test common passwords and evade detection. In a password spray attack, an adversary will try the same password (for example, Winter2021) on hundreds or even thousands of accounts at a time. Because only one failed authentication occurs per-user account, typical detections for brute force attacks are bypassed. You can use Attack Simulation Training in Defender for Office 365, discussed in Chapter 9, to simulate a password spray attack against your tenant.

## Password Writeback

Password writeback allows users and administrators to change or reset Azure AD passwords and have that password propagated to the on-premises user account in real-time. The ability to write back passwords is integrated with Azure AD's [self-service password reset](#) (SSPR) capability and licensed through Azure AD Premium.

This is how password writeback works:

1. The user clicks a link to request a password change.
2. The user enters their old and new passwords (or completes one or more challenge gates in the case of a forgotten password).
3. The selected password is encrypted with a special key that was created during the setup of the password writeback feature for that specific tenant.
4. The encrypted password is sent over HTTPS to a tenant-specific service bus endpoint which is used to communicate with the on-premises password writeback service. Communications on this bus are protected by a shared credential that was created during the setup of the password writeback feature and is only known to your organization and Azure AD.
5. The writeback feature looks for the user account in the on-premises AD. To find a match, the *sourceAnchor* is used to look up the user in the Azure AD Connect connector space. From there, the object is traced back through the metaverse to AD.
6. If the user account is found, the password is reset. If the reset is successful, the user is notified.
7. If the password reset operation fails, an error is returned to the user. One of the reasons the password reset might fail is when it does not satisfy the on-premises password policy. If the password writeback service running inside of Azure AD Connect is unavailable, the operation will also fail.

The password writeback features can also be deployed without using password hash synchronization; you can also deploy password writeback alongside AD FS. You can also enable users to unlock their on-premises AD account using SSPR and password writeback without also changing their password.

**Real World:** You might expect Azure AD Connect to always connect to the nearest domain controller (DC). Azure AD Connect will prefer the PDC emulator (PDCE) role holder for its connections because the PDCE is considered the master authoritative source for password changes. If the connection is made to a DC in a

different site or over a slower connection, a variety of problems might occur like delayed password synchronization or slow synchronization performance.

## Configuring Password Writeback

Password writeback can be enabled as part of the Express and Custom installation modes for Azure AD Connect, or through PowerShell. The easiest way to enable password writeback is to select the password writeback option in the Optional Features section of the Azure AD Connect setup wizard. If you let the setup wizard configure the service account for Azure AD Connect, you do not need to make any further changes in AD.

If you used a custom service account, you must delegate the service account permissions to reset passwords and unlock accounts. The easiest way to delegate permissions is with the *Set-ADSyncPasswordWritebackPermissions* PowerShell cmdlet included in the *ADSyncConfig* PowerShell module.

## Configuring Self-Service Password Reset

The password writeback feature also works when an administrator changes the password for a user from the Azure AD admin center, but it is most useful when combined with the SSPR capability in Azure AD so that users can update their password without having to worry about the passwords getting out of sync.

Unlike most other features, password reset is not managed from the Microsoft 365 admin center. The link in the admin center redirects to the Azure AD admin center. After signing into the Azure AD admin center, open the **Password reset** blade. Once SSPR has been enabled, you can modify the policy to control various aspects of the password management features in Microsoft 365/Azure AD, such as whether users are required to provide multiple verification methods (to verify their identity), and if so, what verification options they should use.

We typically recommend that you enable only a few of the available verification methods for SSPR. Specifically, we suggest that you use the mobile app notification, mobile app code, and mobile phone options if they will work for your organization. Security questions are useful in situations where your users might need to reset their password in a location without access to their phone, or where you cannot require your users to provide a mobile phone to enroll. Be very careful when you pick your security questions. Security questions should capture details that only a user will know, and that cannot be easily researched by an adversary on social media.

Unfortunately, SSPR only lets you choose a single set of verification methods for all your users today. As a result, you will need to enable the methods that work best for everyone in your organization. You can also require your users to periodically confirm their SSPR registration details are still correct. We recommend that you configure this setting to be once or twice a year. After all, if the details are no longer correct, the user will be forced to call for support which defeats the point of SSPR.

## Using Self-Service Password Reset

Before users can use SSPR, they may need to configure additional verification options. If a user is not enrolled for MFA, they must register for SSPR through a process that is almost identical to MFA enrollment. In April 2020, Microsoft formally released a feature to “converge” the MFA and SSPR registrations, so users only need to register once. Once users have registered via the converged registration process, they are enrolled in MFA and SSPR. Microsoft enables the SSPR feature in tenants created after August 15, 2020. However, SSPR is disabled by default in pre-existing tenants to allow organizations time to plan for and adopt the converged registration experience. By September 30, 2022, Microsoft will complete enabling converged registration for tenants that do not already have it enabled.

To enable converged registration in your tenant, login to the Azure AD admin center. From there, select **Users** > **User Settings** and then click the **Manage user feature settings** link. Configure the **Users can use the**

**combined security information registration experience** setting. We recommend that you enable converged registration if you have not already.

Once the verification process is completed, there are several ways a user can change his password. One way is to go through the Office portal, click Settings (the cogwheel), and select **Password**. This will take you to the page where you can change the password. Alternatively, if the user forgot their password, they could click the **"Can't access your account"** link on the sign-in page after having entered the wrong password. This brings them to the online password reset portal where they are guided through a few steps to verify their identity based on the verification options they provided. If the user is successfully verified, the password can be reset.

## User Risk and Self-Service Password Writeback

Azure AD Premium P2 includes Microsoft's Identity Protection capabilities. One of the capabilities of Identity Protection is the ability to measure and report on risk scores for individual users. For example, if Microsoft finds a valid password for a user in a list of leaked credentials on the Internet, the user's risk score will be rated high.

You can configure the user risk policy in Azure AD to act when a user's risk score is rated by Microsoft as high. There are two actions you can take in the user risk policy. The first option is to block the user from further sign-in until manual action is taken. The other option enables self-remediation. The next time the user signs in, after completing MFA, they will be forced to change their password. The password change is written back to the on-premises AD using password writeback. The user risk policy is configured by accessing the Security blade inside the Azure AD admin center and then opening Identity Protection.

## Multi-Factor Authentication

Because of how people use them, passwords are inherently insecure. Some years ago, it would take a computer a very long time to crack or guess a password. As such, even the simplest passwords were secure enough to protect against most forms of attacks. However, as hardware became more powerful, the time needed for a computer to crack a password decreased exponentially. Today, with the right hardware, a simple six-character password can be cracked in a matter of minutes! To make it harder for passwords to be brute-force guessed, or at least to increase the time it takes, we are taught to use longer and more complex passwords. Unfortunately, research has shown that requiring longer, more complex passwords that periodically expire can lead to the selection of passwords that are sometimes easier to guess.

When you think about it, authentications are based on a combination of three things: something you know (such as a pass phrase), something you have (like a physical token), or something you are (a biometric marker of some kind). Traditional password-based authentication combines your username with something you know, but anyone else who knows the same thing can pretend to be you.

To explain why you should use MFA to secure access to user accounts, let's first take a somewhat simplistic look at how typical authentication against Azure AD is performed. Figure 3-3 illustrates the steps involved when a user tries to authenticate using a client like SharePoint Online. For the sake of simplicity, the finer details are omitted, as we only want to understand the principles behind standard authentication.

1. The user connects to the resource. For example, this might be SharePoint Online.
2. Because the client is unauthenticated, Azure AD will send the client a "challenge"; the user is asked to provide his credentials.
3. The user (client) enters his credentials (a username and password) and sends them to the authentication service which, in turn, will verify the credentials.
4. If the credentials are verified, the resource is "unlocked", and the user is granted access.




Figure 3-3: Traditional authentication flow

The last step is not entirely correct. Upon successful authentication, the client will get a token which it then can use to request access to a specific resource (like SharePoint Online). This is the difference between authentication (getting a token) and authorization (being allowed to access a resource). The exact contents of the four steps may vary according to how the application is implemented, too, but for the sake of brevity, let's just assume the user is granted access to the resource after the authentication happened.

An attacker who manages to steal a password, either by brute force guessing or through social engineering, might be able to leverage the stolen password to access multiple sites. Users who pick weak passwords, write passwords down, store them insecurely, or give them up to social engineering attempts aren't necessarily to blame; although many credential thefts are the result of user carelessness, it's difficult to blame non-technical users for making these mistakes when an average user must keep track of multiple credentials, each with its length and strength requirements. Although password management tools can help with remembering passwords, those tools are often protected by a single password themselves.

Since knowing a user's password allows an attacker to perfectly impersonate the user, poor password management means that users are vulnerable to impersonation. To help fix this, we have several options. One is to make it harder to guess or crack passwords; another is to harden systems, and train users, to reduce the chances that an attacker can steal the password. A third option is to require the user to provide more than one type of authentication factor, which is where the "multi-factor" part of MFA comes in.

The principle behind MFA is that you add a second (or even third) factor from the "something you have" or "something you are" categories. Typically, MFA is implemented as a combination of a password with a smartcard, one-time generated code, mobile app notification, or biometric information such as a fingerprint, facial recognition, or an iris scan. To gain access to a system, an attacker would have to have access to both the additional token and the user's password. Although adding an additional factor to the authentication is a huge leap forward in terms of security, it does not take away all risks that entail the use of passwords.

**Real-world:** MFA is a valuable security measure, but it is not a replacement for good password and account management, and it has vulnerabilities of its own. For example, there have been several high-profile targeted attacks where the attacker took control of the target's mobile phone number by convincing the mobile carrier to issue a new SIM so that the attacker received the MFA authentication code required to execute a password change ([here's one example](#)). Even with MFA deployed and enforced, you should still be monitoring your users' logon activity for anomalies, and where possible, you should discourage the use of SMS MFA. Recent NIST [guidance](#) also recommends against using SMS MFA.

## Office 365 MFA

Office 365 includes simple MFA capabilities through a feature called Office 365 MFA. Apart from Office 365 MFA, other solutions can provide multi-factor authentication, including Azure MFA and various third-party MFA solutions. Office 365 MFA leverages the Azure MFA platform. Although Office 365 MFA and Azure MFA are very similar in terms of functionality, the latter requires Azure AD Premium. The biggest difference

between Azure MFA and Office 365 MFA is that the latter can only be used to secure access to first-party Office 365 workloads (e.g., Exchange Online, SharePoint Online, etc.).

Azure MFA extends beyond these workloads and can also be used to protect other cloud and on-premises resources. Both Azure MFA and Office 365 MFA support the following additional authentication options:

- Mobile app ([Microsoft Authenticator](#)) available for Android and iOS:
  - Notification: the user receives a notification in the mobile app to ask them to confirm their identity by clicking the **Approve** button. Once the user confirms, they log in automatically.
  - Verification Code: the mobile app generates a new code every 30 seconds. After the user enters the current code (on-screen), the sign-in process continues.
  - Authenticator app sign-in: Also known as “passwordless sign-in”, in this mode the standard Microsoft sign-in dialog will display a two-digit number and generate a request to the Authenticator app asking you to tap or enter the matching number on-screen. This method lets users log in using their phone with no typing or password entry. See the [documentation](#) for more details on how to set this up.
- Phone call. The user receives a phone call, asking them to confirm that they are signing in by pressing the # key. After confirming, the user logs in automatically; no code is necessary.
- Text Message (SMS). A one-time passcode (6 digits) is sent to the user’s mobile device. This code must be entered before the sign-in process can continue.
- Oath Hardware Token. The user uses a physical token that generates a random one-time passcode every thirty or sixty seconds. These tokens are often useful in scenarios where the user cannot access their mobile phone when they need to complete MFA.
- Additional [third-party authentication methods](#), including those from RSA Data Security and Duo Security. If you want to use a third-party authenticator, you need Azure AD Premium licenses.

In February 2019, Microsoft noted that less than 9% of accounts with administrative access to Azure AD had MFA protection enabled. Although some excuses can be offered for not using MFA to protect accounts used to run PowerShell against Microsoft 365 endpoints, that excuse is less valid now that all the key PowerShell modules support MFA. You should secure all administrative accounts with MFA. Microsoft is making it easier to do so by introducing security defaults to all tenants that don’t already use MFA or conditional access policies. We cover the features available in security defaults later in this chapter.

## Configuring Accounts for MFA

Configuring accounts for MFA is straightforward. The simplest way to do so is through the Microsoft 365 admin center. Log in and navigate to **Users** and then **Active Users** and select **Multi-factor authentication**. From the MFA page, you can select users, individually or in groups, and enable them for MFA with a single click. Through the **service settings** tab, an administrator can control additional options, such as:

- Whether a user can generate app passwords.
- What additional verification options are available? By default, all options are enabled (mobile app verification/notification, text message (SMS), or phone call). Microsoft [strongly recommends](#) that you focus on using the mobile (authenticator) app for MFA.

You can configure the MFA registration policy for your tenant by accessing the Azure AD admin center, selecting **Security, Identity Protection**, and then **MFA registration policy**. You can also determine when a user can manage their MFA enrollment by using a conditional access policy. This is discussed in more detail later.

To set the MFA methods for a user account through the Azure AD admin center, find the user account, select **Authentication methods**, and then add the desired method. To reset the MFA enrollment for a user account, select **Require re-register multi-factor authentication** from the [...] menu.

**Note:** If your organization has Azure AD Premium licenses, you should use conditional access policies to require MFA rather than the script in this section. Conditional access policies offer significant flexibility compared to simply setting a user's MFA state to "Enforced".

If your organization has Azure AD Premium P2, you can configure an MFA registration policy to ensure that all your users are registered for MFA. If you wait for a user to be prompted to enroll for MFA, and you exempt users from MFA when they are coming from a trusted location or device, an adversary could enroll for MFA on a user's behalf and the user might not know. An MFA registration policy will ensure the user is prompted to enroll for MFA even if they would ordinarily be exempted from performing MFA during a sign-in.

## Scripting MFA Enablement

An administrator can also use PowerShell module to configure MFA or to enable/disable MFA for users. In the following example, we enable an account for Phone-based MFA necessary parameter values. You must be signed into the Microsoft Graph with the *UserAuthenticationMethod.ReadWrite.All* permission and select the beta profile for the command to work. This script shows how to use Microsoft Graph PowerShell SDK cmdlets to work with SMS-based MFA challenges. The script:

- Declares the value of the identifier used for mobile phone numbers (the value is always the same).
- Gets the identifier for the target user account.
- Populates a hash table with the phone number to use and its type. The phone number must have a space between the international code and the local number. It should be unique within a tenant but doesn't have to be.
- Runs the *Get-MgUserAuthenticationPhoneMethod* cmdlet to check if a mobile phone number is already present in the account.
- If a phone number does not exist, the script runs the *New-MgUserAuthenticationPhoneMethod* to add it as an authentication method to the account.
- Alternatively, if a phone number exists, the script runs the *Remove-MgUserAuthenticationPhoneMethod* cmdlet to remove the old number and then adds the new number.

```
[PS] C:\> $MobilePhoneId = "3179e48a-750b-4051-897c-87b9720928f7"
$UserId = (Get-MgUser -UserId Lotte.Smith@Office365itpros.com).Id
$Params = @{
    PhoneNumber = "+353 862267785"
    PhoneType = "mobile"}
$AuthPhone = Get-MgUserAuthenticationPhoneMethod -UserId $UserId -ErrorAction SilentlyContinue
If (!$AuthPhone) { # No authentication methods are there, so proceed
    New-MgUserAuthenticationPhoneMethod -UserId $UserId -BodyParameter $Params }
Else {
    Remove-MgUserAuthenticationPhoneMethod -UserId $UserId -PhoneAuthenticationMethodId
$MobilePhoneId -Erroraction Stop
    New-MgUserAuthenticationPhoneMethod -UserId $UserId -BodyParameter $Params }
```

The next time the account signs in, they will get an SMS MFA challenge (Figure 3-5).

Similar cmdlets are available to enable FIDO2 (*New-MgUserAuthenticationFido2Method*), the Microsoft Authenticator app (*New-MgUserAuthenticationMicrosoftAuthenticatorMethod*), temporary passwords (*New-MgUserAuthenticationTemporaryAccessPassMethod*), and so on.

Depending on the selected verification method, the user may need to complete some extra steps. For instance, if the Authenticator app is selected, the user must install the app on their phone and scan a QR code to link the application to Azure AD, and by extension to the tenant. Once the added verification method is configured, the user's enrollment is complete and the next time they sign in, they will use multi-factor authentication.



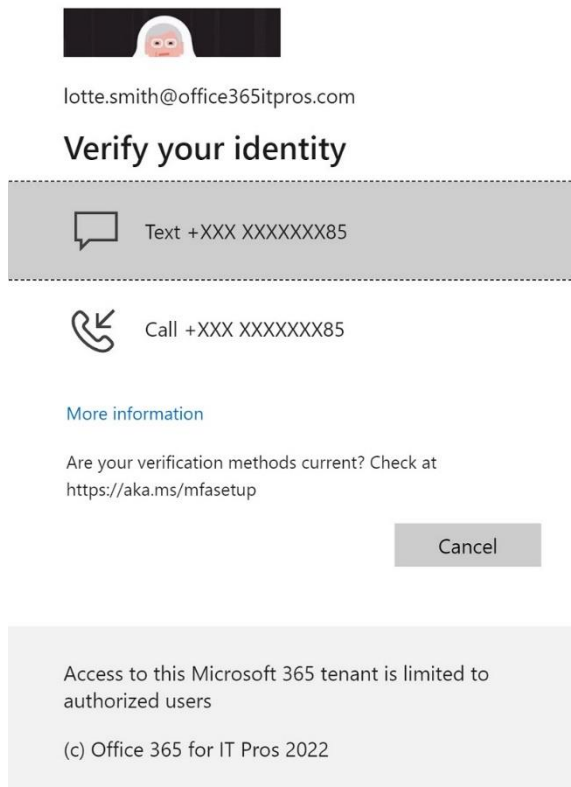


Figure 3-5: MFA sign-in using SMS

## Listing Office 365 MFA-Enabled Accounts

When managing Office 365 MFA for multiple users, it is useful to know which users are enabled for MFA, whether users have completed the onboarding process, or the verification options configured. You can access this information in the Azure AD admin center by accessing the **Usage & insights** blade and then selecting **Authentication methods activity**. From here you will be able to see (and export) information about MFA and SSPR enrollment. Note that accounts need an Azure AD Premium P1 or P2 license to access the Usage & insights blade.

It's also possible to retrieve the information about the authentication methods used by Azure AD accounts using cmdlets in the Microsoft Graph PowerShell SDK. See [this article](#) and example script.

## Tracking MFA Sign-ins

Like any other sign-in to Azure AD, all MFA-enabled sign-in appear in the [Azure AD sign-in log](#). You can view the log interactively or download events to a CSV file and examine them afterward in Excel or Power BI. The MFA Required column in the CSV file tracks whether an account is MFA enabled or not, while the status code captures the outcome of a sign-in attempt. Important codes include:

- **0**: Success.
- **50058**: An application tried to perform a silent sign-in (using an access token) to an MFA-enabled account but could not complete.
- **50076**: The user did not pass the MFA challenge because they did not respond in the allowed time.
- **50074**: The user did not pass the MFA challenge. For instance, they input an incorrect code. You can see the authentication method used in the MFA Auth Method column.

To look up the meaning of an Azure AD error code, use [this tool](#).

## Using MFA

Users sometimes need some time to become accustomed to dealing with MFA challenges, but since so many other cloud services and applications support it (including Google's applications, Facebook, Apple iCloud, and almost every online banking system), MFA quickly becomes routine, needing minimal effort or thought. We recommend receiving notifications to the mobile app for authentication requests as the default MFA method. Because the mobile app needs a connection to the Internet to receive notifications, you might not always be able to use the app, for instance, when roaming in a foreign country or while you are on an airplane. In these cases, you can revert to another method, such as generating a verification code in the mobile app.

By default, the user's preferred additional verification option is used. If the user registered for multiple verification options, they can select a different method when their preferred verification method is (temporarily) unavailable or when the user doesn't respond to a challenge within a period. For instance, when the mobile app is unavailable because of a lack of Internet connectivity, a phone call or text message can be requested instead. To register for added verification options, the users must access their **My Sign-Ins** page. On the My Sign-Ins page, click **Security info**. From the Security info page (Figure 3-6), users can add extra phone numbers, configure the authenticator app, manage their SSPR registration (e.g., security questions and alternate email addresses), create app passwords, or change the preferred authentication method.

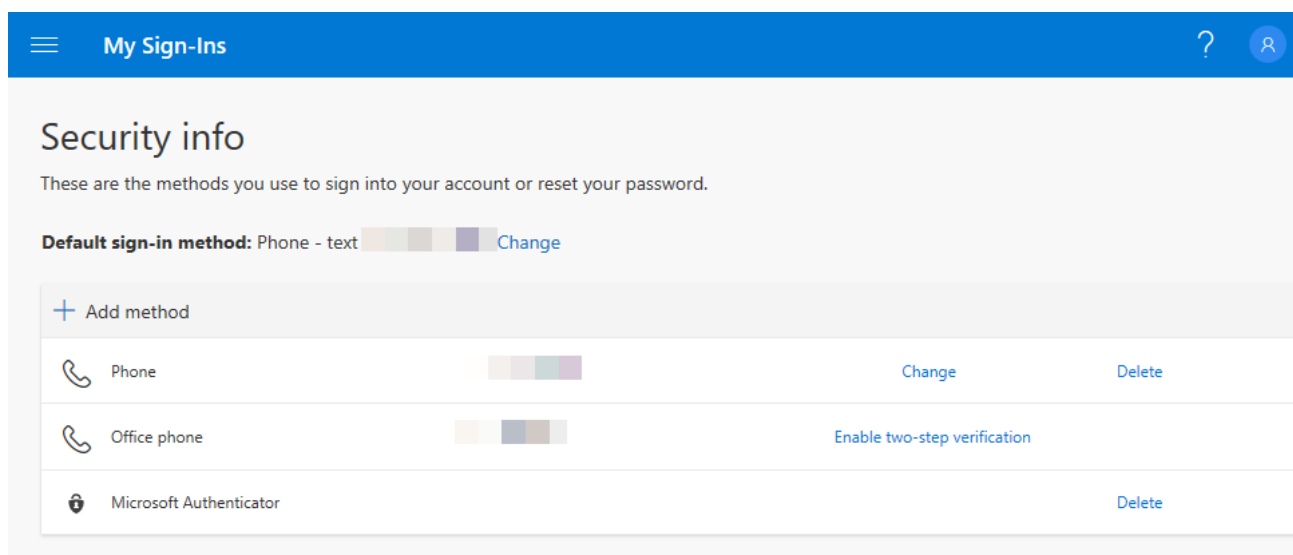


Figure 3-6: Changing MFA Verification Options for an account

Sometimes users complain about receiving too many authentication or MFA prompts from Azure AD. This can be a result of your tenant's policies; other times it is a result of devices or browsers missing important plugins or software updates. Microsoft has created an Azure Log Analytics Workbook to help diagnose authentication prompt frequency. This workbook is called [Authentication Prompts Analysis](#).

## Encouraging Use of the Authenticator App

If you have had MFA deployed for some time, chances are your users may be using a mixture of voice calls, SMS messages, and the Authenticator app to perform MFA. The Authenticator app provides the most secure verification, and it offers additional valuable capabilities such as one-time pass codes, passwordless sign-in, and authentication broker support on iOS devices. It is difficult to convert users from voice calls or SMS to the Authenticator app, though.

Azure AD can encourage this conversion through an in-line prompt to set up Authenticator. This prompt comes after the user completes their sign-in and MFA and only if they have not yet enrolled in the Authenticator app. To enable this prompt, log in to the Azure AD admin center and navigate through **Security > Authentication Methods > Registration Campaign**. From here you can enable the prompt, configure how

many days an end-user can ignore (snooze) the prompt, and optionally, target the prompt to a subset of your users.

In addition to this feature that prompts users to set up Authenticator, you can also enroll directly in the Authenticator app without scanning a QR code. Simply add a **Work or school account** in the Authenticator app and then click **Sign in** when prompted.

## Controlling Access

For most organizations, the thought of allowing access to all their data from anywhere, at any time, on any device is not palatable. Microsoft has recognized this concern and made the conditional access features a core capability of Azure AD Premium. With conditional access, you can exercise substantial granularity in controlling the who, what, when, and where of all connections.

### Conditional Access Policies

By default, all users with a valid license can access resources available to their account if they can reach the service endpoint and log in. Sometimes, an organization might want to limit who can access resources, or perhaps control from which locations users can access those resources. Most Microsoft 365 applications do not provide such functionality. Instead, they leverage the advanced capabilities of Azure AD Premium and Microsoft Endpoint Manager.

Name \*

Require additional authentication for ext... ✓

Assignments

Users or workload identities ⓘ

Specific users included

Cloud apps or actions ⓘ

1 app included

Conditions ⓘ

1 condition selected

Access controls

Grant ⓘ

2 controls selected

Session ⓘ

0 controls selected

Figure 3-7: Settings available when creating a new conditional access policy

Conditional access policies define rules that determine when and how a user or workload identity can access an application. Azure AD evaluates characteristics of the session such as the IP address, location, device, and sign-in risk score. If the device is enrolled with Microsoft Endpoint Manager, the device's compliance with organizational policy as well as security risk determined by Microsoft Defender for Endpoint (MDE) can also factor into the conditional access policy. Collectively, these factors allow you to create policies to apply granular decisions about what sessions can access applications. These factors can also apply to certain activities such as whether a user can view documents in SharePoint Online, or view *and* download documents.

Numerous ways exist to restrict access to resources using conditional access policies. You can restrict who is allowed access to a resource, define which devices can be used to access resources or control from what locations an app or service can be used. You can also add restrictions based on the characteristics of the user

login; for example, if a user logs on from an unknown location, you can require them to authenticate with MFA even if the device itself would normally be trusted.

Before diving into an example of a conditional access policy, it is useful to know that each policy consists of one or more *conditions*, each of which can have a set of linked controls as illustrated in Figure 3-7. Conditions may have exceptions.

Conditions you can apply to the policy include:

- **Users or workload identities** sets which users, groups, directory roles, or service principals the policy should apply to. You can choose to include or exclude individual users or groups for the policy, or you can include or exclude all guest users or holders of specific Azure AD directory roles. You can also apply the policy to built-in directory roles such as Global Administrators. Finally, you can apply the policy to service principals via the workload identities capability. This lets you restrict the use of service principals (also known as app registrations). Service principals often have highly privileged access to data in your tenant.
- **Cloud apps or actions** define which applications the policy applies to or what user behaviors the policy applies to. You can choose to enable the policy for all applications or select the applications for which it should be valid. Applications include the various Office 365 workloads or a third-party application that you have configured to rely on Azure AD for authentication, like Salesforce, Workday, Facebook for Work, etc.

Applying conditional access policies to individual Microsoft 365 workloads presents challenges to maintaining a consistent security posture. Microsoft has addressed this challenge by introducing a virtual application that represents all the workloads. This app can be selected in the **Cloud apps** section of a conditional access policy and is aptly named **Office 365**. Note that currently [some workloads](#), such as Microsoft Planner are not included in the Office 365 virtual application. Where possible, we recommend that you attempt to target your conditional access policies to the Office 365 virtual application rather than to individual workloads. Targeting individual workloads provides more flexibility but it also introduces hidden complexities due to the interdependencies of various workloads. A similar virtual application called **Microsoft Azure Management** is also available to target policy to Azure management tools like the Azure Portal, PowerShell, etc.

In addition to applications, you can also control when a user can register for MFA and SSPR (**Register security information**) and/or register or join a hybrid device (**Register or join devices**). This control is useful if your organization has a security policy that requires users to register for MFA/SSPR from a trusted device or location. To control what happens when the user registers or joins a hybrid device, you must also disable the corresponding MFA settings under **Devices > Device Settings** in the Azure AD admin center.

Finally, you can create conditional access policies that are triggered when users perform specific tasks in supported applications by using Authentication contexts. Authentication contexts modify a conditional access policy from applying to *all* access to an application to specific tasks (e.g., downloading sensitive information) within the application.

- **Conditions** define characteristics of the authentication attempt which lead to a policy to be applied or not. Examples of conditions include:
  - **Sign-in risk** evaluates the risk that the sign-in attempt is coming from someone other than the authorized user. Microsoft evaluates sign-in risk using a complex supervised [machine learning model](#). To use this condition, the target account must have Azure AD Premium P2.
  - **User risk** evaluates the risk of the individual user that is attempting to access the application. For example, a user whose password is in a list of leaked credentials is considered high risk. To use this condition, the target user must have Azure AD Premium P2.

- **Device platforms** determine from which OS platform the authentication must be made before an action is considered. Options include Windows, iOS, Android, macOS, and Linux.
- **Locations** specify locations that will trigger the policy when the user logs in; locations are determined based on the IP of the client making the request, the user's Global Positioning System (GPS) location reported by the Authenticator app, or the country Microsoft infers the user is coming from. You can define a set of known (trusted) corporate locations by IP address range.  
To define a location based on IP address, country, or GPS location, go to **Security > Conditional Access > Named Locations** in the Azure AD admin center and click **Countries location** (country or GPS location) or **IP ranges location** on the toolbar.
- **Client apps** determine what application type (browser, ActiveSync, mobile client, legacy authentication, etc.) should trigger the policy. Policies created before July 2020 only apply to the selected types of client apps (or "Other clients" for legacy authentication). All new policies apply to all client apps, including legacy authentication, by default.
- **Filters for devices** let you target the policy based on attributes of the device object in Azure AD. For example, you might want to have one conditional access policy apply to personal mobile devices, while another policy might apply to mobile devices that the organization owns. In combination with Intune, you can do this by creating a device filter rule. A filter rule of *device.deviceOwnership -eq "Personal"* will apply your policy only to devices marked as personally-owned in Microsoft Intune. A filter rule *of device.trustType -eq "ServerAD"* will apply your policy only to devices that are hybrid Azure AD joined.

On the controls side (which you access through the link labeled **Grant**) you can choose from the following controls. There is also a radio button to control requiring any one of the actions selected (an *OR* condition) or all of them (an *AND* condition):

- **Block Access:** this is self-explanatory.
- **Grant (Allow) access,** with or without the following optional requirements:
  - **Require multi-factor authentication.**
  - **Require device to be marked as compliant.** For this option to work, the device must be enrolled in Microsoft Intune and targeted with a compliance policy. The rules in the compliance policy determine if the device will be marked compliant (or not).
  - **Require Hybrid Azure AD joined device** verifies whether the computer is simultaneously joined to the on-premises AD and Azure AD tenant.
  - **Require approved client app** lets you restrict whether applications must be accessed through a [supported](#) client application instead of through a third-party application or a bare API.
  - **Require app protection policy** allows you to restrict connections to apps that are subject to a Microsoft Endpoint Manager app protection policy applied in your tenant.
  - **Terms of use** allows you to require users to view and accept text stored in a PDF document before they access an application. You must first create a terms of use document in Azure AD before you can use this control. To create a terms of use document, go to **Security > Conditional Access > Terms of use** in the Azure AD admin center and click **New terms** on the toolbar.
  - **Require password change** allows you to force a user to change their password through SSPR. This grant is usually combined with the user risk condition to lower the user's risk score.
- **Session:** this control lets you restrict what authenticated users can do within the context of a specific application session:
  - **Use app enforced restriction** lets you enforce what users can do in Exchange Online or SharePoint Online such as printing and downloading attachments/files.

- **Use Conditional Access App Control** integrates with Microsoft Defender for Cloud Apps (MDCA) to apply deep controls to actions users take within an application. This control proxies all the user's access to the application through MDCA.
- **Sign-in frequency** lets you control how often the user must re-authenticate.
- **Persistent browser session** controls how long the user can remain signed-in to browser-based applications.
- **Customize continuous access evaluation** allows you to disable CAE (discussed earlier in this chapter) or enable strict enforcement mode.
- **Disable resilience defaults** provides the option to disable the safeguards of the BAS during an Azure AD outage. The BAS still evaluates tokens using CAE to ensure they are valid (e.g., due to the deletion or disabling of a user account) but assumes that conditions in a conditional access policy that were previously satisfied when a cached token was issued (such as group membership) are still valid. If you are uncomfortable with these assumptions, you can enable this session control, but, in the event of an Azure AD outage, any users included in the conditional access policy cannot take advantage of the BAS.

As you begin building conditional access policies, you must be careful not to inadvertently impact user access to workloads and other applications. To help with this, Azure AD allows you to enable conditional access policies in report-only mode. When a conditional access policy is enabled in report-only mode, you will be able to see the expected effect of the policy when you review the Azure AD sign-in logs. Users will not be affected since the policy is not fully enabled.

One important caveat to this is the device compliance requirement discussed earlier. Even if a policy is enabled in report-only mode, if you require device compliance in that policy, users of iOS, Android, and macOS devices may be prompted to select a certificate if their device is not compliant. You can avoid this behavior by excluding iOS, Android, and macOS devices from your policy while it is in report-only mode.

Whether your CA policy is enforced or in report-only mode, you can troubleshoot CA behavior from the sign-ins log in the Azure AD admin center. You can access the sign-ins log in the context of the entire tenant, a specific application (e.g., Exchange Online), or a specific user. If you click on a sign-in and then click on a specific CA policy in the Conditional Access tab of the details pane, you can see what policies were applied to the sign-in. The Policies blade supports search, sort, and filtering, which makes it easier to manage many CA policies.

## Require MFA for External Connections

To illustrate the capabilities of the platform, consider the following scenario: a company employs sales representatives who travel often as part of their job. Sales representatives are permitted to use personal devices, so there is a mix of iOS, Android, and Windows devices in use. The company's security policy requires that all external connections to Microsoft 365 from sales reps must use an additional authentication factor for access from outside the corporate network.

We can break down the scenario into the following elements to create a new conditional access policy that meets the software company's requirements:

1. The policy should apply to all users who are part of the sales department. Because all sales representatives are part of a group called "Sales", the policy can easily be applied to the latter.
2. Even though a conditional access policy can apply to many cloud applications (including non-Microsoft applications), it must only be applied to Office 365. Be careful applying conditional access policies to only a subset of applications. You can inadvertently forget to require MFA in some scenarios which would present a security risk.
3. Only logon attempts from outside the corporate network should require an additional authentication factor.

4. Given that all sales representatives can choose personal devices, the policy should apply, regardless of the device used to access resources.

To create a policy meeting these four requirements, log in to the Azure AD admin center and navigate to **Security** and then **Conditional Access**.

1. Click **New policy** and name the policy "*Require additional authentication for external sales users*". Under **Assignments**:
  - a. Click **Users or workload identities**, select **Select users and groups**, and select the *Sales* group. Do not forget to click **Done** to save your progress before moving onto the **Cloud apps** blade.
  - b. Under the **Cloud apps** blade, select Office 365.
  - c. Under **Conditions**, click **Locations**, click **Yes** to enable this section and then click **Exclude**. There, select **All trusted locations**.
  - d. Under **Access Controls**, click **Grant**, and then select **Require multi-factor authentication** and **Require one of the selected controls**.
2. Under **Enable policy**, click **On**.

**Note:** Before you can create a new policy, you should configure your known locations. This can be done by going to **Security > Conditional Access > Named locations** and defining the IPv4 and IPv6 address ranges that are used for Internet egress in your organization. For most organizations, these ranges will be network address translation (NAT) addresses defined on your firewalls. If your organization uses public IP space on your network, this might also be the IP addresses of client computers and/or member servers. You should not include private IP addresses (e.g., 10.0.0.0/8, 172.16.0.0/12, or 192.168.0.0/16).

Without the information on known locations, the conditional access engine cannot determine if an authentication request stems from inside or outside of the corporate network, and the machine-learning systems that help rate sign-in risk miss an important signal. Named locations are also shown in the sign-in logs, making the logs much easier to interpret with context for an IP address.

When you configure named locations, you can optionally mark those locations as "trusted". If you mark a location as trusted, Azure AD allows you to target a conditional access policy to "all trusted locations" without individually listing them. Trusted location information is also used as a signal to Azure AD's security reporting functionality.

Now, if someone from the sales department tries to authenticate to Exchange Online, no matter whether they access it through the Outlook mobile app, desktop Outlook, OWA, or an Exchange ActiveSync client using modern authentication, they will be prompted for additional authentication factors if the authentication is attempted from outside the known corporate locations.

The above example illustrates one of many possible scenarios in which conditional access can be very useful. For more information on conditional access, have a look [here](#).

## Require MFA for All Privileged Roles

Privileged access to Microsoft 365 and Azure AD should always require MFA, with no exceptions. You might even consider requiring a hybrid Azure AD joined device or a device that is compliant with Intune as additional protection. Azure AD makes this easy to do by applying your conditional access policy to built-in directory roles. You should exclude "break glass" or "emergency access" accounts from this policy in case there is ever a situation where the MFA service is unavailable, or you do not have access to a device that will pass tests (if you require this)! Ideally, break-glass accounts should have as [few dependencies for authentication](#) as possible.

To achieve this, you will need to create a new conditional access policy. Configure the conditional access policy with the following settings:

- **Assignments > Users or workload identities** – Include **Directory roles** and select the built-in directory roles you want to include. We recommend that you select all the roles unless you have a valid business or technical reason not to. Exclude **Users and groups** and select your break glass accounts.
- **Cloud Apps or actions** – Include **All cloud apps**.
- **Access Controls > Grant** – Check the **Require multi-factor authentication** checkbox.

If you want to also require hybrid joined or Intune compliance devices, select the corresponding checkboxes on the Access Controls > Grant screen. Once you are ready, configure the **Enable policy** setting to **On**. Your changes will take effect immediately.

## Disable Basic Authentication for Exchange Online

Authentication policies allow tenants to disable basic authentication for email connection protocols such as POP3 and IMAP4. The [documentation on how to apply authentication policies](#) to enforce this change is mandatory reading before you make this change and the process is also explained in Chapter 7. If you have Azure AD Premium, you can (and should) also use conditional access to restrict the use of basic (legacy) authentication across all services that use Azure AD. Remember that many attacks against tenants try to exploit weaknesses in basic authentication, so it makes sense to deploy an authentication policy to disable basic authentication for as many protocols as possible and assign it to all accounts.

**Note:** Exchange 2019 Cumulative Update (CU) 1 and newer allows you to turn off legacy authentication for your on-premises Exchange servers. This change is strongly recommended. See this [blog post for more details](#).

## Restricting Legacy Access with Conditional Access

Microsoft recommends the use of conditional access policies to block legacy clients. This seems like a straightforward thing to do, but there are some nuances. For example, it is a bad idea to block these clients without first looking at your sign-in logs to see which applications your users are using with legacy authentication, and how many of them there are. Turning off every scan-to-email device in a large organization at once, for example, would not be good.

To block legacy clients, you will need to create a new conditional access policy. Configure the conditional access policy with the following settings:

- **Assignments > Users or workload identities** – Include **All users**.
- **Cloud apps or actions** – Select **All cloud apps**.
- **Conditions > Client apps** – Check the **Exchange ActiveSync clients** and **Other clients** checkboxes.
- **Access Controls** – Select **Grant** and then **Block access**.

Unlike other conditional access policies, this change may take up to 24 hours to become fully effective. You should test with a small set of pilot users before you target the policy to all the users in your tenant.

## Restricting Access to the Outlook Mobile Application

While blocking access from legacy clients that do not support MFA is an important first step to securing email access, you may want to go a step further and require the use of Outlook Mobile to access Exchange Online. You can simply require the use of Outlook Mobile, or, you can go a step further and require that Outlook Mobile be managed by a Microsoft Endpoint Manager App Protection policy.

To do this, you will need to configure a new conditional access policy with the following settings:

- Cloud apps or actions – Select **Office 365 Exchange Online**.
- Conditions > Client Apps – Select **Mobile apps and desktop clients**.
- Conditions > Device platforms – Select **Android, iOS, and macOS**.
- Grant – Select **Require approved client app**.



If you need to ensure that users are connecting with an instance of Outlook Mobile that is managed with an App Protection policy from your organization, also select **Require app protection policy** in the **Grant** menu and ensure **Require all the selected controls** are selected. For more information on App Protection policies, refer to Chapter 16.

**Note:** You must select macOS under device platforms due to a change Apple made in iOS 13. Beginning with iOS 13, iPads are identified as macOS devices. It is very important to note that this policy will also apply to connections from Office applications on Mac devices!

We recommend that you start by assigning this policy to a small number of users or groups to test functionality. Once you are satisfied with the results of the CA policy, you can target the CA policy to all users in your tenant.

**Note:** If you have an on-premises Exchange organization that is configured for Hybrid Modern Authentication (HMA), this CA policy will also affect client access to on-premises mailboxes.

## Controlling SharePoint Downloads

A common requirement is to restrict downloading content from SharePoint Online to corporate devices. At the same time, users should be allowed to view content through browsers on unmanaged devices. This can be accomplished with a combination of conditional access policies and configuration in SharePoint.

First, you will need to configure SharePoint Online. To do this, complete the following steps:

1. Access the SharePoint admin center.
2. Browse to **Policies > Access Control**.
3. Select **Unmanaged devices**.
4. Select **Allow limited, web-only access**.

Next, log in to the Azure AD admin center and navigate **Security** and then **Conditional Access**.

**Note:** You will find two new policies prefixed with [SharePoint admin center]. **Disable these policies immediately! They are automatically enabled by SharePoint Online and are very restrictive.**

SharePoint will also disable legacy authentication access (e.g. from Office 2010 and outdated Office 2013 clients) to SharePoint Online when you make this change.

Configure a new conditional access policy with the following settings:

- Cloud apps or actions – Select **Office 365 SharePoint Online**.
- Conditions > Device platforms – Choose **Select device platforms** and then select **Windows**.
- Conditions > Device state (Preview) – Select **Exclude** and then **Device Hybrid Azure AD joined** and **Device marked as compliant**.
- Session – Select **Use app enforced restrictions**.

This policy will enforce restrictions on all Windows devices. If the Windows device is hybrid Azure AD joined or Intune compliant, users will have complete access to SharePoint documents. If the user is not using a hybrid Azure AD joined (or Microsoft Endpoint Manager compliant) Windows device, they will only be able to view documents in the web browser. Downloads and printing will be disabled, and the user will see a notification across the top of the browser.

Mobile devices and Macs will not be affected by the policy. To apply restrictions on mobile devices and Macs, you need to add Microsoft Endpoint Manager to the solution and adjust your policy to target all device platforms. We discuss Microsoft Endpoint Manager device compliance and app protection policies in more detail in Chapter 16.

While the example in this section targets all of SharePoint, you can also enforce policies on a site-by-site basis. Microsoft documents this in [this document](#). Regardless of the path you choose, your changes might take ten-to-fifteen minutes to become effective.

## Require MFA and a Hybrid Joined Device When Accessing Labeled SharePoint Sites

While the previous example applies to SharePoint as a whole, chances are you have certain SharePoint sites that contain much more sensitive information. To protect this information, you may need to require more assurance that the data is being accessed in a manner that meets your security policies. In this example, we will combine multiple Microsoft 365 features with conditional access. Specifically, we will require users who access data in a Microsoft 365 group labeled with a Microsoft Information Protection (MIP) sensitivity label called *Highly Confidential* to connect with a hybrid Azure AD joined device and perform MFA every time they access the data. We will accomplish this using a feature of conditional access called authentication context.

Authentication context is currently in public preview but is available in all tenants. With authentication context, applications can request that Azure AD re-evaluate authentication and authorization when a user takes certain actions. You can define up to 25 authentication contexts in your tenant that applications can use in their request. These contexts are simply a name that you can choose from a list in the application and your conditional access policy. For example, you might create a context called *Require MFA*, and another called *Require Trusted Location*. Today, Microsoft lets you use authentication context in SharePoint Online as well as via Microsoft Defender for Cloud Apps (MDCA). The interface is [publicly available](#) and third-party applications and in-house developers are free to use authentication context as well. Applications retrieve the contexts you have configured via the Graph API.

To begin, we will first create an authentication context for our example. To do this, log in to the Azure AD admin center and navigate to **Security, Conditional Access**, and then **Authentication context (Preview)**. Click **New authentication context** on the toolbar and create a context with the following values:

- Name – Require MFA and Hybrid Joined Device
- Description – This authentication context requires the user to complete MFA and have a hybrid Azure AD joined device
- Public to apps – Checked

Next, we will create a conditional access policy to enforce our requirements. Configure a new conditional access policy with the following settings:

- Users or workload identities – Select **All users**.
- Cloud apps or actions – Select **Authentication context (preview)** and then check the **Require MFA and Hybrid Joined Device** context.
- Grant – Select **Require multi-factor authentication** and **Require Hybrid Azure AD joined device**. Select **Require all the selected controls**.
- Enable Policy – **On**.

We now have an authentication context and a conditional access policy that will enforce it, but, we have not yet created a trigger for the authentication context to be requested. To do this, we will configure our MIP label. For this step, we assume that a sensitivity label called *Highly Confidential* already exists in the tenant. If you have not worked with MIP labels before, refer to Chapter 23 for more information.

Browse to the [Microsoft Purview Compliance portal](#) and go to **Information protection**. Double click the Highly Confidential label and click **Edit label** to begin editing it. Configure the following settings:

- Scope – Check **Groups & sites**.
- Groups & sites – Check **External sharing and Conditional Access settings**.
- External sharing & device access
  - Check **Use Azure AD Conditional Access to protect labeled SharePoint sites**.

- Select **Choose an existing authentication context (preview)**.
- Select **Require MFA and Hybrid Joined Device** from the dropdown list.
- Session – **Sign-in frequency** configured to 1 hour.

Configuring the sign-in frequency session control to one hour requires users to complete MFA again if their session is longer than one hour. This achieves a variation of the requirement to ensure that users re-complete MFA any time they access content labeled *Highly Confidential*.

Once you save your changes, SharePoint sites labeled *Highly Confidential* will begin triggering the conditional access policy we created earlier. There are many ways to label a site. One method is to assign the label to the group in Teams. To do this, find the Team, right-click, and click **Edit Team**. Set the **Sensitivity to Highly Confidential**.

While the authentication contexts feature is in public preview, you will find that there are several limitations to which clients and applications support authentication context. Microsoft documents those limitations [here](#). We recommend that you review this list before trying authentication contexts out in your organization.

## Policy Templates

The conditional access policy examples we have illustrated so far are very commonly deployed. While creating policies by hand provides the most flexibility for customization, and enables you to learn the platform, it also introduces opportunities for inadvertent errors and gaps in coverage. To help mitigate this possibility, Microsoft also provides a set of templates for commonly used policies that you can start with. To access the templates, login to the Azure AD admin center and navigate **Security > Conditional Access** and select **New policy > Create new policy from templates (Preview)** on the toolbar. Microsoft documents the set of approximately fourteen templates [here](#).

## Security Defaults

While we recommend licensing Azure AD Premium P1 (at a minimum) for all the users in your tenant, this may not be possible. You can achieve some of the most critical conditional access controls discussed earlier by enabling security defaults in your tenant. When you enable security defaults, you will get the following protections for your entire tenant, free of charge:

- Mandatory MFA for every privileged sign-in. This applies to members of [thirteen privileged](#) Azure AD roles as well as access to Azure management APIs including the Azure portal, CLI, and PowerShell modules.
- Mandatory MFA registration for all users using the Microsoft Authenticator app.
- MFA on an as-necessary basis for regular users once they're registered for MFA.
- Block legacy authentication protocols for all users, except for Exchange ActiveSync.

To enable security defaults, access the Properties blade inside the Azure AD admin center and then click **Manage security defaults**. Once you enable security defaults, the protections listed above will take effect immediately for your entire tenant. If you have Azure AD Premium, you can only enable security defaults if there are no conditional access policies enabled in your tenant. Since security defaults cannot be scoped to specific users for testing, we recommend using conditional access policies instead if you have Azure AD Premium. Newly provisioned Azure AD tenants have security defaults enabled by default and Microsoft is in the process of enabling security defaults in existing tenants that do not have compensating security controls already enabled..

## Cross-Tenant Trust

When guest users from another organization access resources in your tenant via Azure AD B2B, conditional access policies also apply to the guest users. This can create end-user experience issues because the user may

be required to enroll for MFA in both their home tenant and your tenant. You can configure Azure AD to trust the MFA that was performed in a user's home tenant to satisfy MFA requirements in conditional access policies. You can also require a user's device to be trusted by their home tenant to satisfy device compliance and/or hybrid Azure AD join requirements in your conditional access policies. These capabilities are part of a preview feature called cross-tenant access settings.

To configure cross-tenant access settings, log in to the Azure AD admin center and navigate to **External Identities** > **Cross-tenant access settings (Preview)**. You can configure your tenant to trust MFA and device identity from a guest's tenant globally by selecting **Default settings** and clicking **Edit inbound defaults**. On the **Trust settings** tab, you can check the following boxes:

- **Trust multi-factor authentication from Azure AD tenants** – if the user completes MFA in their home tenant, that MFA will also satisfy the require multi-factor authentication grant control in your conditional access policies.
- **Trust compliant devices** – if the user's device is marked as compliant by Intune or a partner solution in their home tenant, the device will also satisfy the device state filters in your conditional access policies.
- **Trust hybrid Azure AD joined devices** – if the user's device is hybrid Azure AD joined to their home tenant, the device will also satisfy the device state filters in your conditional access policies.

You may find that globally trusting these factors from any guest's tenant does not meet your security and risk management requirements. Instead, you can trust individual tenants. To do this, click **Add organization** on the **Organizational settings** tab. Enter a domain name registered in the guest's tenant and click **Add**. You can configure the same trust factors discussed above just for this tenant. Note that while you only need to specify one domain name in the guest's tenant, the trust settings will apply to the entire tenant, even if multiple domains are verified.

You can also use cross-tenant access settings to control a feature called Azure AD B2B Direct Connect. Teams is the first application to use B2B Direct Connect in its shared channels feature. Shared channels allow users to join a Teams channel without having a B2B guest account in the channel's home tenant. Users can also see the channel in their Teams client without changing tenants. The default cross-tenant access settings block B2B Direct Connect for all users in your tenant and inbound connections from external tenants.

There are two sets of settings for B2B Direct Connect. Inbound access settings allow you to control which external tenants can access your tenant. Outbound access settings control the users from your tenant can access resources in external tenants. You can enable Inbound or Outbound access globally, or you can enable it only for select users (or groups).

## Restricting Access to a Single Tenant

Throughout this chapter, we discuss several ways you can restrict user access using a wide variety of built-in features like custom claim rules or conditional access. However, none of these options address the problem of a user authenticating against another tenant. You might wonder why this is important or how this is different from disallowing someone to authenticate by disabling their account. In the latter scenario, the user is not able to log in to the corporate tenant, but that does not prevent them from accessing other tenants, using another identity. From a data leakage standpoint, this is a potentially dangerous situation as the user can log on to another tenant and copy corporate data to a repository on that tenant. The reason why a user can log in to various tenants is that most of the endpoints are the same for all tenants, worldwide. Blocking the endpoint would not be a smart move unless you want to block access to Microsoft 365 entirely. Here is an example to further illustrate the problem.

A user, Erica, works for a company that uses Exchange Online for corporate email, and she also has a personal tenant for her personal email. When no tenant restrictions exist, Erica can log in to either or both tenants from

her desk at work, but there is nothing to prevent her from signing into her personal tenant while at work. To lower the risk of intentional or accidental data leakage, the security team at Erica's employer wants to prevent Erica and her co-workers from signing into tenants other than the corporate tenant while at work. As mentioned earlier, blocking the sign-in endpoint (for example, through a proxy server) would also prevent Erica from signing into her company's tenant, and thus prevent her from doing her job.

Of course, other ways exist to restrict access. For instance, an administrator can define a set of internal IP addresses from which a user can authenticate into the corporate tenant. Although this effectively prevents someone from accessing the tenant outside the corporate network, it does not solve the problem described above.

To help organizations control access, Microsoft supports a feature known as "[tenant restriction](#)." The way this feature works is quite simple. When a user authenticates with Azure AD, the authentication platform also checks for an (HTTP) header called *Restrict-Access-To-Tenants*. The value of this header holds the names of all the tenants the user can access. You also need to add the *Restrict-Access-Context* header, which specifies the GUID of your tenant, so the service knows which tenant to apply the restrictions to. If these headers hold the name of the tenant the user is trying to access, the sign-in proceeds. If not, they receive an error message and are blocked from signing in. The message informs the user that *Your network administrator has restricted what organizations can be accessed. Contact your IT department to unblock access.*

For this feature to work, there must be a proxy server between the user and the internet that performs SSL inspection. Anything that can add an HTTP header will suffice. If you want to make sure that control is always exerted over user sign-ins, including outside the corporate network, users must always access the internet through a proxy server that can inject the header. To ensure this happens, other countermeasures (such as ensuring the user cannot modify proxy settings, etc.) must be present. Otherwise, users can bypass the restriction themselves. In addition, the proxy server should be accessible from outside the corporate network if you want to restrict access to a single tenant, regardless of the user's location. One way to achieve this is to use a "proxy-as-a-service" such as ZScaler.

The tenant restriction capability is not something you configure in Azure. Instead, it relies solely on your network infrastructure to inject the header. Currently, the feature only works with modern authentication, which is fine for most of Microsoft's applications, but if you use other applications or built-in ActiveSync apps, you must block access for those "legacy" apps if you want to maintain the restriction. You only need to apply the restriction to the authentication process (only for the Azure AD endpoints). There is no need to force users to connect through a proxy server to access applications such as Exchange Online.

Given the dependency on modern authentication, and because you cannot control how other tenants operate, if those tenants still allow legacy authentication (basic authentication), users can bypass the restriction. Thus, the feature in its current form only really is useful to control the potential for data loss if you fully control the endpoint, and thus prevent the use of applications that do not use modern authentication. With that said, it is a clunky way to prevent data loss; it is much less flexible or capable than Azure Information Protection (AIP), although (unlike AIP) it requires minimal client-side configuration.

You can use the feature for free with Office 365, but you must buy an Azure AD Premium license if you want to restrict access to other applications that rely on Azure AD for authentication.

Restricting access to a single tenant might or might not be useful. If preventing data leakage is a priority, anything you can do to make it harder for people to share information in an unauthorized manner can be helpful. Although the feature will most likely not prevent malicious users from leaking data, it can stop most regular users from (accidentally) leaking sensitive data to another tenant. Of course, it is only a small cog in a much bigger picture as many other features like DLP policies (Chapter 21) and Microsoft Information Protection (Chapter 23) help to safeguard your data.

# App Registrations and Permissions

Thus far we have focused on how users and their devices work with Azure AD. Azure AD also plays a very important role in supporting the integration of different Microsoft 365 workloads (for example, how Microsoft Teams gains access to SharePoint and Exchange Online), as well as third-party and in-house developed applications. Microsoft controls the permissions held by first-party applications like Teams, but the permissions that you grant third-party and custom-developed applications are controlled by you as the tenant administrator. There are two types of permissions that can be granted: delegated and application. Delegated permissions are used when an application (such as a website) signs the user in and then calls an API on behalf of the user. Application permissions are used when there is no user sign-in, such as in a background process or service.

Permissions are granted by consenting to different API permission scopes that are defined on various applications and resources. In Figure 3-8, a user is allowing an application to take various actions in Azure AD, Exchange Online, and SharePoint through a process known as consent. Azure AD manages and tracks consent and brokers the issuance of tokens that enable applications to communicate and use their permissions.

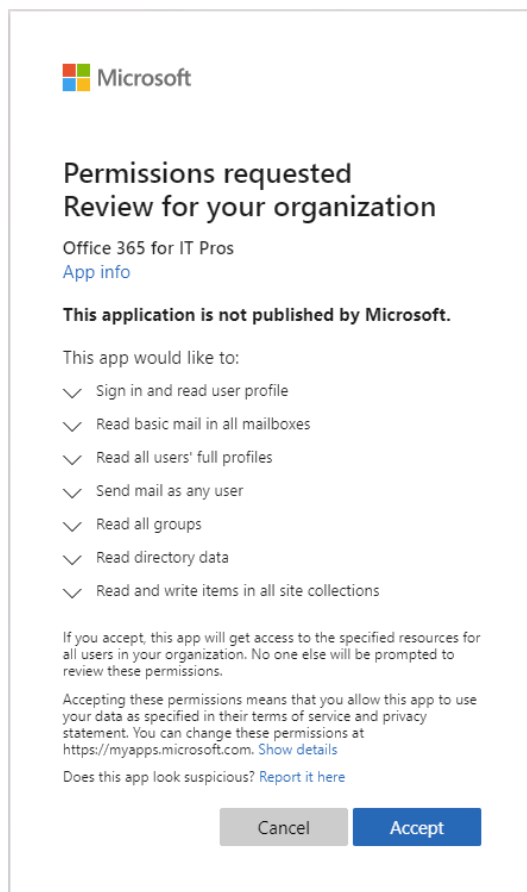


Figure 3-8: Granting Graph API permission consent to an app

Some of the permissions granted in Figure 3-8 are extremely broad and in the hands of an attacker or a poorly written application could present a significant security risk. For example, this application can read information about every user and group in the tenant, read email messages, send email as a user, and read and write any file in SharePoint. Historically, many organizations did not pay much attention to these consent requests, and they even allowed users to independently consent to certain application permissions.

Attackers have taken advantage of this historically overlooked vulnerability through avenues such as phishing attacks that induce users to grant [illicit consent](#). Once granted consent, the attacker might use their new

access to read a user's email or send a message impersonating them, for example. This gives the attacker a foothold and higher powers to expand their dominance.

## Controlling Consent

To combat the risks of granting illicit consent and over-privileged applications, Azure AD lets you control what types of permissions end users can consent to and provides an approval process for users to request consent for higher levels of permission. To configure this, browse to **Enterprise Applications > Consent and permissions** in the Azure AD admin center. Under **User consent settings**, you can configure what types of permission users can independently consent to. You can also decide whether owners of Microsoft 365 groups can grant an application access to the data in their groups.

For user consent for applications, we recommend that you choose the Microsoft recommended setting, **Allow user consent for apps from verified publishers, for selected permissions**, or the more conservative choice of **Do not allow user consent**. In a small organization, the conservative approach is likely practical, but, in a larger organization, the number of requests you will receive for relatively low-risk permissions is probably not manageable. The middle ground attempts to address this. Applications that are from verified publishers have completed an identity verification process with Microsoft and the publisher has agreed to certain terms and conditions. You can choose which permissions you consider low risk, including a curated list that Microsoft recommends.

Whether or not to allow users to grant access to group data is a question of risk tolerance. Your groups may contain very sensitive data that you would not want to risk granting access to a third-party application. This may mean that you require an administrator to grant consent on a per-request basis. You must however be prepared for the volume of approval requests that this may generate.

If you choose a restrictive setting for either user consent or group data consent, which we recommend you do, you will need a process to manage requests. Azure AD offers a solution to this with an approval process governed through a feature called admin consent workflow. To enable the approval process, browse to **Enterprise Applications > User settings** in the Azure AD admin center and configure the following settings:

- Users can request admin consent to apps they are unable to consent to – Yes.
- Who can review admin consent requests – Specify a list of users, one or more groups of users, or a built-in directory role.
- Selected users will receive email notifications for requests – Yes.
- Selected users will receive request expiration reminders – Yes.
- Consent request expires after (days) – 14.

Once you have enabled the admin consent workflow and configured the user consent settings discussed earlier, end-users will receive a modified version of the prompt in Figure 3-8 when they try to access an application requiring some permissions they cannot grant consent for. Instead, the user will be prompted to provide a business justification and submit a request for approval. The reviewers specified earlier will receive an email notifying them of the request that they can approve or deny under **Enterprise Applications > Admin consent requests** in the Azure AD admin center.

## Reviewing Existing Permissions

The admin approval workflow and restrictions on user consent add significant security for new requests, but they do not do anything for existing permission consents. In some organizations, the number of applications granted access to various permissions may be substantial. Some of these permissions may also be risky or even granted to an illicit application. You will need to create a process to evaluate previously-granted permissions and determine if they should remain.

This [blog post](#) explains how to use a script to collect an inventory of granted permissions. The script even attempts to determine the relative risk of different permission grants so you can focus your evaluation on the best place. While the script is helpful, the number of entries in a large organization can number in the thousands.

The App Governance add-on for Microsoft Defender for Cloud Apps (MDCA) aims to simplify the evaluation of permissions and provide continuous insight on the potential risks that your organization is exposed to through overly-permissioned apps.

## Restrict Workload Access to Approved IP Addresses

When you have app registrations that must hold sensitive or privileged access to support the functionality of an application or business process, you should take steps to control when that access can be used.

Conditional access can be used to do this using the workload identities preview. With this feature, you can create a conditional access policy, like the examples we discussed earlier in the chapter, which applies to app registrations. This feature requires Azure AD Premium P2.

In this example, we will restrict an app registration called *Office 365 for IT Pros* from being used except when running from approved IP addresses. You can use the egress IP address ranges of the data center(s) that host the application using this app registration. To configure the approved IP address ranges, log in to the Azure AD admin center and navigate to **Security > Conditional Access > Named locations**. If you have not configured a named location before, we discuss this in detail earlier in the Conditional Access section of this chapter.

Next, select **Security** and then **Conditional Access**.

1. Click **New policy** and name the policy "*Restrict Workload Access to Approved IP Addresses*". Under **Assignments**:
  - a. click **Users or workload identities** and select **Workload identities (Preview)** under **What does this policy apply to?**. Next, click **Select service principals** and select the Office 365 for IT Pros app registration from your tenant. Do not forget to click **Select** to save your progress.
  - b. Under **Conditions**, click **Locations** followed by **Exclude**. Select **All trusted locations**.
  - c. Under **Access Controls**, click **Grant**, and then select **Block access**.
2. Under **Enable policy**, click **On**.

Once this policy is enabled, the selected service principal can authenticate only from a trusted location that you configure. Note that the feature does not currently support service principals configured for multi-tenant use.

## Connecting to LinkedIn

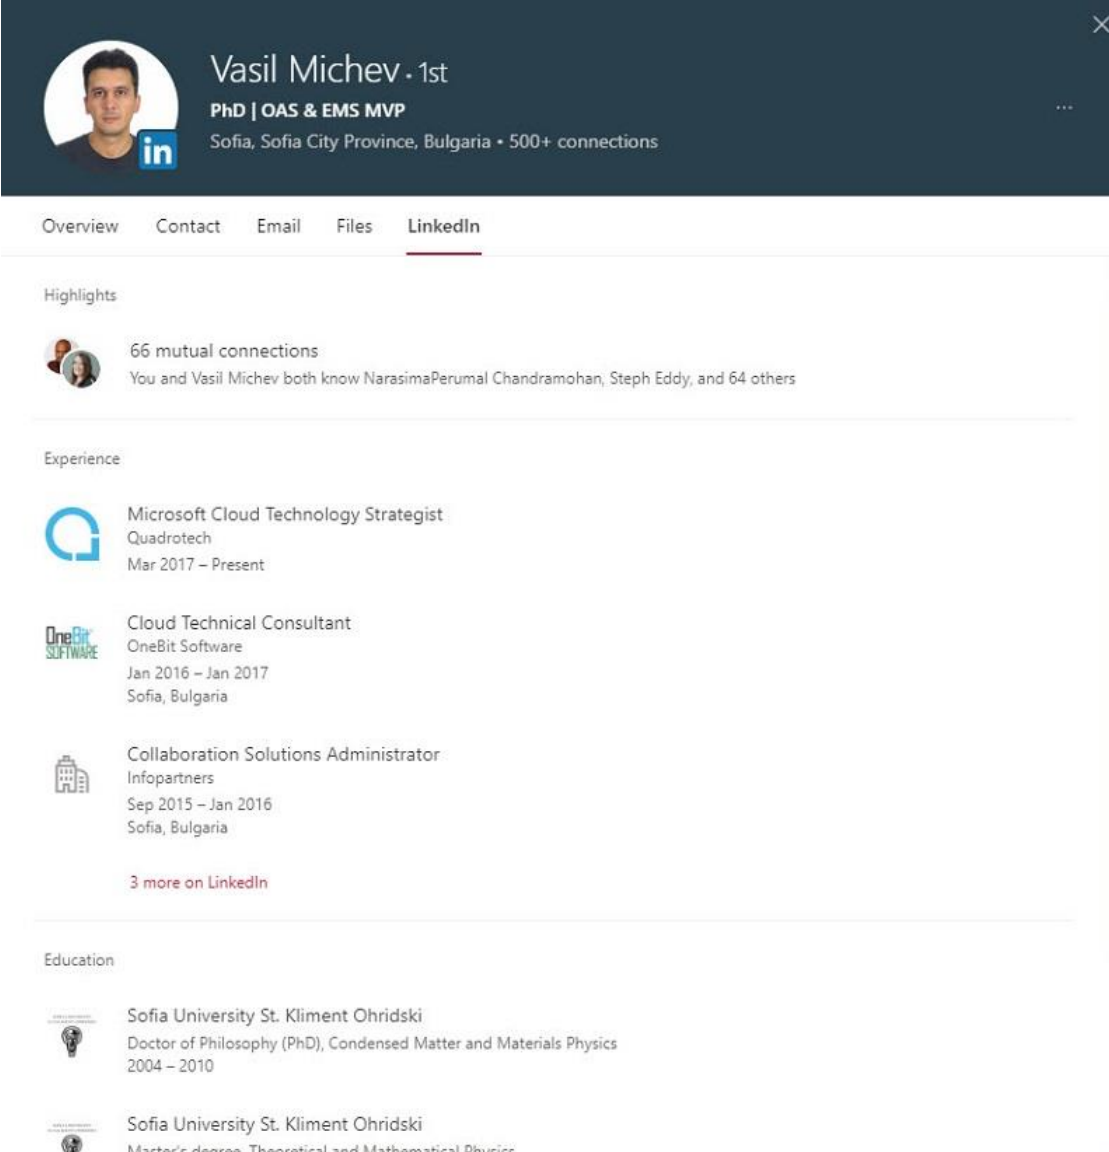
According [to LinkedIn](#), 830 million people are members of their professional network. Ever since Microsoft bought LinkedIn in June 2016, the two companies have looked for ways to co-operate, including the ability to connect LinkedIn to Azure AD. By default, all tenants except those in France and Germany have a LinkedIn connection enabled. Tenants of the sovereign clouds do not support the LinkedIn connection, so you cannot use the feature if your tenant is in the China or U.S. Government clouds. Enabling the connection only allows users to decide if they want to connect their Azure AD account to their LinkedIn account. No data is ever exchanged unless authorized by the user to whom the data belongs.

To control the connection, go to the **User Settings** section of the Azure AD admin center. You can then set the connection to allow all users to connect, some users to connect, or no users to connect. If you decide to restrict the connection, you select a security group containing the set of users allowed to connect their accounts to LinkedIn.



If the tenant allows users to access LinkedIn, a LinkedIn logo appears in the people cards displayed by applications such as OWA, Teams, and SharePoint Online. Apps can display publicly available information from matching profiles. Before LinkedIn reveals more detailed information such as job history or details about the user's connections, the user must connect to their LinkedIn account. The search can happen for both internal and external recipients and result in three outcomes:

- The person is a LinkedIn contact: You see the private profile for the person as shared with their LinkedIn contacts (Figure 3-9).
- The person has a LinkedIn account but is not a contact: You see the public profile of the person and can send them a LinkedIn invitation to connect.
- LinkedIn cannot find a match: Either the person doesn't have a LinkedIn account or several matching LinkedIn accounts are found. In this case, you can choose the best match.



The screenshot displays a LinkedIn profile for Vasil Michev - 1st. The profile header includes a circular profile picture, the name "Vasil Michev - 1st", and the text "PhD | OAS & EMS MVP" and "Sofia, Sofia City Province, Bulgaria • 500+ connections". Below the header is a navigation bar with tabs for Overview, Contact, Email, Files, and LinkedIn. The main content area is divided into sections: Highlights, Experience, and Education. The Highlights section shows "66 mutual connections" and lists "You and Vasil Michev both know NarasimaPerumal Chandramohan, Steph Eddy, and 64 others". The Experience section lists three roles: "Microsoft Cloud Technology Strategist" at Quadrotech (Mar 2017 - Present), "Cloud Technical Consultant" at OneBit Software (Jan 2016 - Jan 2017, Sofia, Bulgaria), and "Collaboration Solutions Administrator" at Infopartners (Sep 2015 - Jan 2016, Sofia, Bulgaria). The Education section lists two degrees from Sofia University St. Kliment Ohridski: "Doctor of Philosophy (PhD), Condensed Matter and Materials Physics" (2004 - 2010) and "Master's degree: Theoretical and Mathematical Physics".

Figure 3-9: Viewing LinkedIn profile information for a contact

The important thing to remember about the LinkedIn connection is that users control it. Although the tenant decides if the connection can be used, users decide if they want to connect their Azure AD account with their LinkedIn account.

## Fetching Email Addresses from LinkedIn

Another feature gained when someone connects their LinkedIn account to their account is that email addresses for first-degree LinkedIn contacts are included in the people suggestion list used by Microsoft 365 browser applications when the user addresses an email or shares documents. In addition to LinkedIn contacts, the suggested people list includes tenant and guest users and Outlook's auto-complete list. As an example of use, when you create a message with OWA and begin typing an address into the TO: or CC: field, OWA checks what you type against the suggested people list and shows any matches that it finds.

## PowerShell and Azure AD

Microsoft is [transitioning away](#) from the AzureAD PowerShell module. The Azure AD module uses the Azure Graph API, which Microsoft plans to [deprecate](#). As a replacement, Microsoft recommends that you use the [Microsoft Graph PowerShell SDK](#) module. If you are familiar with the AzureAD module, you will find the same tasks are possible with the new module, albeit with different syntax. As the name implies, the foundation for the Microsoft Graph PowerShell SDK is the set of Graph APIs, which cover much more than Azure AD accounts and is constantly evolving. New features first come to the "beta" endpoint before transitioning to the stable "v1.0" endpoint. All the tasks you can accomplish in the Azure AD admin center are not yet available in the Graph API.

More information about the Microsoft Graph PowerShell SDK is in the chapter covering managing your tenant with PowerShell.

# Chapter 4: Tenant Management

***Paul Robichaux***

There are many ways to manage tenants. As the service has evolved to add new features, capabilities, and apps, different products have adopted different administrative workflows. Depending on the task you're trying to perform, the functionality of the tools Microsoft provides may lead you towards one approach or another. Sometimes there's more than one way to perform the same task; in those cases, each approach will have its pros and cons. For example, the web-based admin interfaces described in this chapter are simple to use and make it easy to perform one-time tasks such as adding a domain name to the tenant. However, for bulk administrative tasks such as changing the department attribute for multiple users, PowerShell is more efficient, which is why we've devoted a complete chapter (Chapter 23) to this topic.

It's important to remember that in a hybrid environment where you're using directory synchronization, changes to users, groups, and other objects stored in on-premises Active Directory must occur in the on-premises Active Directory because it's the source of authority. As you learned in the preceding chapter, changes you make to cloud copies of on-premises objects won't necessarily replicate back. For example, to change a user's surname you must make the change using the on-premises Active Directory management tools and then allow directory synchronization to synchronize the change into Azure AD. The Microsoft 365 admin center and PowerShell will warn you when you try to change an attribute in the cloud synchronized from on-premises AD.

Because there are so many individual services in the Microsoft 365 portfolio, Microsoft hasn't succeeded in unifying administration into a single portal or a single PowerShell module, although they are making some moves to do so across Microsoft 365 and by using the Microsoft Graph as a single point of data access. It wouldn't make sense to do so because the workloads are so different; a single unified toolset would be complex, slow, and confusing to use and would limit the ability of the product groups that develop each service to optimize and improve their admin experiences.

Part of learning to manage the service is becoming comfortable with a sometimes-confusing variety of web-based admin portals, PowerShell modules, and endpoints that we can use. In the next section, we'll start exploring some of this variety to help you understand what tools to use and when.

## Cloud versus On-Premises Management

Any seasoned professional experienced in managing on-premises infrastructures should already be aware that some level of control and visibility is lost when a company migrates workload to the cloud. Microsoft's platform is sold as "Software as a Service" (SaaS), which transfers the responsibility for running the service to Microsoft. You give up access to anything relating to the underlying hardware, networking, and software. You don't get to select the make and model of servers used to run the service, nor do you get to choose the specifications and sizing of those servers.

Microsoft manages the storage for Microsoft 365 applications and services, both in terms of quotas assigned to individual applications and users and overall storage performance. At most you may need to help your end-users with managing their mailboxes and sites within the quota limits that the service imposes, through a combination of user education and policy controls (e.g., retention policies). Microsoft also manages the network and Internet connectivity for the service. You only need to ensure that adequate bandwidth is available to reliably connect to Microsoft 365 across the internet. Similarly, you must manage any firewalls or other network devices such as web proxies that affect the connectivity of your users to the service.

The underlying operating system and all the software running on Microsoft's servers are also out of your hands. Microsoft surfaces a lot of configuration options for you that apply to the users in your tenant, as well as some reporting tools such as message tracing. However, you simply do not get access to Windows event log data, performance log data, or any of the other log types that Windows and the applications generate. Similarly, you do not get access to stop or restart services on the servers or perform server reboots. On the positive side, you also don't need to perform any Windows patching or upgrades for Exchange, SharePoint, and other applications.

Because many common administrative operations you have to perform to manage on-premises servers become unnecessary, some IT professionals view the movement to the cloud as a loss of "control" and fear the impact that the cloud will have on their careers. Another way to look at it is that you are allowing the same people who develop the product to run it for you within a defined set of parameters. All the time that you previously spent monitoring, managing, and fixing server hardware, operating systems, and application installations is now available for you to spend on helping to deploy features to your end-users, align the configuration of different workloads with your organization's policies, and act as a broker between your company and Microsoft as a service provider. On the bright side, when something goes wrong, a Microsoft engineer wakes at 3 am to answer a pager instead of you.

## Break Glass Administration

An obvious difference between cloud and on-premises management is that Microsoft 365 won't allow you to sign into the console of a physical computer when all else fails and you need to access a server. Other factors which can limit the ability to sign into an administrator account include a working internet connection to Azure AD, blocks imposed by conditional access policies, the ability to satisfy multi-factor authentication challenges, and so on. During the setup of a new tenant, you nominate an account to be the global tenant administrator. This account is all-powerful and should be protected with multi-factor authentication and a complex password. In general, you should not use the global administrator account for day-to-day operations. As discussed later, it is better to assign limited administrative roles to accounts to allow them to perform specific tasks.

Because an outage might interfere with the ability of administrators to sign into their regular accounts, you should create one or more "break glass" accounts. These are highly-privileged accounts (perhaps holding the global administrator role) intended for use in emergencies with the following characteristics:

- Cloud account only (to remove any dependency on account synchronization with an on-premises directory). The user principal name for the account uses the tenant service domain (*tenant.onmicrosoft.com*). Consider giving break-glass accounts obscure names to draw attention away from their true purpose.
- Multiple layers of authentication such as MFA protect the account. However, you should take care to [minimize the number of dependencies used by authentication](#) to ensure that the account is available when needed. For instance, you should exclude break glass accounts from conditional access policies to ensure that a policy doesn't block a sign-in attempt for the account.

The account password for break glass accounts should be complex and deliberately obscure. Because these accounts have access to the entire Microsoft 365 tenant, it's important to store the password for break glass

accounts securely. The details of the storage location for the passwords and how administrators can access the passwords will vary from organization to organization. The important thing is that the process to retrieve passwords and use the break glass accounts is proven and effective. After each use of a break glass account, you should change its password and update the new password in the secure locations.

## Administrative Interfaces

The [Microsoft 365 admin center](#) is used to configure items shared across the applications that make up the platform, such as domain names, billing details, and license assignments. Some frequently used administrative functions are also in the Microsoft 365 admin center. For example, you can create a new cloud-only user account complete with a mailbox without opening the Exchange admin center. Users with Global administrator, Service administrator, Billing administrator, or application-specific administrator rights will have access to the Microsoft 365 admin center.

### Using the Microsoft 365 Admin Center

Figure 4-1 shows the current version of the Microsoft 365 admin center. Expect to see more changes as Microsoft continues to evolve its functionality and publish a [detailed change log](#) to help you keep track. The organization name shown in the top left-hand corner of the *Home* page of the Microsoft 365 admin center is a link. Clicking it takes you to the organizational settings page. More interestingly, if there are multiple tenants associated with a single organization (as determined by the existence of a partner-of-record relationship) there will be a small icon next to the organization name to indicate that you can switch between organizations by clicking the name.

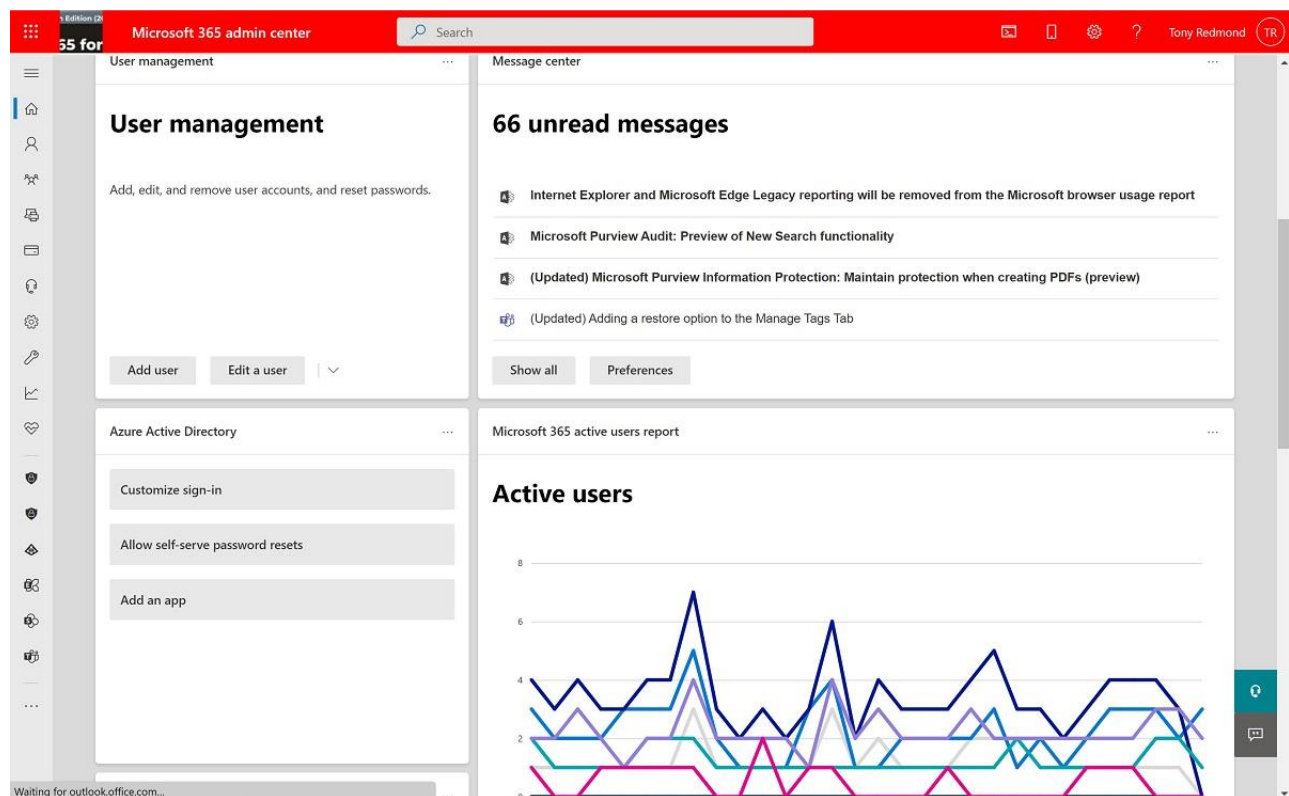


Figure 4-1: The Microsoft 365 admin center

You can customize the home page of the admin center to display specific tiles and the order and layout for the tiles (which Microsoft calls “cards”). The default set of cards you see include cards for message center notifications, reviewing your billing, and working with support incidents. You can customize the set of cards shown in the portal by reordering them, removing them, or adding new ones with the **Add cards** link. As a

customization example, tenants that synchronize Azure AD with an on-premises Active Directory can add a card showing the directory sync status to your home page. Clicking the card reveals other information about your directory sync status, including the version of the directory synchronization client (Azure AD Connect) installed in the on-premises environment.

The navigation pane on the left side of the admin center has expandable menus for different categories of tasks. The tasks that are available throughout the menu are those that Microsoft thinks are commonly performed by administrators. In some cases, a complex administrative task is provided with a user-friendly wizard in the admin center, while the workload-specific admin portal for that task will have a more complicated process involved. In other cases, admin tasks initiated from the Microsoft 365 admin center simply launch a wizard from one of the workload-specific admin portals. To manage the service, you'll have to use a mix of the native admin center functionality and the workload-specific admin centers. Microsoft frequently adds new functionality in the Microsoft 365 admin center, and otherwise rearranges things, to keep us all on our toes. You may see a small bullet icon next to new or renamed menu items when Microsoft wants to highlight changes.

For smaller tenants, you may notice a slightly different, and less detailed, "essentials view" targeted at tenants with 10 or fewer users. This view provides a stripped-down, card-based interface to a few common tasks that smaller tenants will be most likely to use.

## Microsoft 365 Admin Center Functionality

Assuming your account has the necessary permission, and you view the full version of the Microsoft 365 admin center, here's what you'll see:

- In the **Users** section, you can manage user accounts, contacts, guest users, and deleted users, change user license assignments, and change multi-factor authentication settings.
- The **Teams & groups** section allows you to create and manage Microsoft 365 Groups, Exchange distribution lists, shared mailboxes, and security groups. You can see the properties for security and distribution lists that are homed on-premises and synchronized to the cloud, and you can edit properties for cloud-homed objects. See Chapter 8 for more detail about Microsoft 365 Groups.
- The **Roles** section shows you a list of the roles available for assignment to accounts. From this page, you can see the roles available in your tenant and, manage role assignments. By selecting a role, you can also see which users are assigned to it and what specific permissions it has. See Chapter 5 for more on role assignments.
- The **Resources** section is where you can manage room and equipment mailboxes and SharePoint Online sites (although some SharePoint links in this section just take you to the SharePoint Online portal).
- In the **Billing** section, you can manage subscriptions, buy, and allocate licenses, and update payment details. This is also where you can buy added services, either separately or as part of a bundle. For example, if your organization uses Office 365 Enterprise E3 you can buy Microsoft Defender for Cloud Apps as a standalone license or buy the full Office 365 Enterprise E5 license which includes an Office 365-specific subset of MDCA, as well as other features.
- In the **Support** section, you can submit and monitor the progress of service requests for problems that you have encountered with your tenant. Any open Customer Lockbox requests appear here.
- The **Settings** section contains several tools for controlling the configuration of different parts of the service. For example, you use this section to manage the DNS domains associated with your tenant, what third-party applications and add-ins are available, how Microsoft Search works in the tenant, what individual services such as Cortana or Teams are enabled in the tenant, and what integrated applications (formerly known as "Office add-ins") are available. You can also apply restrictions to services such as preventing guest user access to Microsoft 365 Groups or Teams. Security and privacy settings for the organization are also in this section, including the password policy for cloud

identities, Customer Lockbox configuration, and sharing controls. The organization profile is also configured in this section, which includes company and contact details, and the options for configuring Targeted Release (previously known as First Release) for the advanced deployment of new features. This section also contains a link that allows you to see and manage any digital partner of record (DPoR) assignments in your tenant.

- In the **Setup** section, Microsoft presents a list of specific actions you can take to set up your tenant. For example, there are quick-access buttons for configuring self-service user password reset, applying a basic set of data loss prevention rules, and starting the process of migrating user data from other services.
- The **Reports** section presents a dashboard view of the usage of individual services such as Exchange Online, OneDrive, and Teams. You'll also find detailed activity reports, plus the Productivity Score reports. More information about reporting is in Chapter 21.
- The **Health** section is where you can access the Service Health Dashboard (SHD) to see the current health status of workloads in addition to details of outages, service history for the previous 30 days, and planned maintenance. The Health section also includes access to Message Center, where notifications for new and changed features in your tenant are published by Microsoft, a summary of your directory synchronization status identical to the one available in the dirsync status card, the network connectivity toolset, and controls for gathering product feedback from your users and sending it to Microsoft. The SHD, network connectivity tools, feedback tools, and Message Center are discussed later in this chapter.
- The **Admin centers** section has shortcuts to other administration consoles for services that you are licensed to use, such as Azure AD, Exchange Online, the Compliance portal, Teams, and Endpoint Manager.

By default, the Microsoft admin center collapses some of these items, meaning that they don't appear at first sight; you can use the **Show all** link at the bottom of the left navbar to reveal all available options. By hovering over any top-level item, you'll see a pin icon appear; clicking it will pin that item to the left navbar so that it's visible in the condensed view.

Let's take a quick look at the other workload-specific admin portals.

## Managing Exchange

The Exchange admin center (EAC) is a web-based console used to manage elements of Exchange Online including mail contacts, mail users, role assignment policies, OWA mailbox policies, public folders, and hybrid connectivity. Some elements, such as distribution lists and mail flow rules, are also available through other consoles. Two versions of EAC are available:

- The ["classic" older form of the EAC](#) is like the version available for Exchange on-premises servers. In the legacy EAC, you see a switch in the top right corner labeled "Try the New Exchange Admin Center." In September 2021, Microsoft [announced that they will retire this version of EAC](#) on September 1, 2022.
- Microsoft built the [modern EAC](#) for Exchange Online. The functionality available in the new EAC reflects that Microsoft 365 replaces many workload-specific features (like compliance and auditing) with service-wide equivalents managed through other consoles like the Microsoft Purview Compliance portal.

You only need to go to EAC when a Microsoft 365 console does not support an action. For example, the new EAC is the only GUI for an administrator to recover mailbox items for a user.

## Managing Microsoft Teams

When Microsoft first introduced Teams, there were separate admin portals and PowerShell interfaces for Teams and Skype for Business Online. The Teams admin center has replaced the older Skype admin center, and all management for Teams is performed in the [new portal](#). You can also navigate there by going to the **Admin centers** section and then clicking **Teams**. More information about using the Teams admin center is in Chapter 13.

## Managing SharePoint Online and OneDrive for Business

SharePoint Online of course has its own set of management tools; the URL to access them follows the SharePoint pattern of using the tenant name plus “sharepoint.com” instead of having a *Microsoft.com* address. For example, a tenant named Contoso will have a SharePoint admin center URL of <https://contoso-admin.sharepoint.com>. The SharePoint admin center looks a great deal like the Microsoft 365 admin center, with a similar card-based display mechanic and left-hand navigation bar. OneDrive for Business administrative settings are now integrated into the SharePoint Online admin center. For more information on using the SharePoint admin center, see Chapter 8.

## Using the Microsoft Purview Compliance Portal and Microsoft 365 Defender

The Microsoft 365 portfolio includes a broad set of security and compliance features available to tenants depending on the licenses purchased. Microsoft’s original 2016-era vision was to provide a single portal to manage them all, so the original Security & Compliance Center brought together compliance features from across the entire service with a specific focus on those that apply to multiple workloads rather than being specific to an application. Since then, Microsoft has added a large amount of new security- and compliance-related functionality. This growth led Microsoft to first merge the security and compliance functionality into Security & Compliance Center and then to break it back out again. As of July 2022, the current state of play is that we have two independent portals for this functionality:

- [Microsoft 365 Defender](#) (formerly called “Microsoft 365 Security Center”)
- [Microsoft Purview compliance portal](#). This used to be called the Compliance center but, as part of the ongoing global rebranding of all Microsoft 365 compliance features to use the Purview name, it got a new name.

The old Security & Compliance center at <https://protection.office.com> still exists but its links all point to the new pages; if at any point you see a page titled “Security & Compliance”, then you’re not looking at the new versions. The old page will probably disappear in the future. When you visit the Microsoft 365 Defender portal, it will offer to enable the portal-redirectation setting (**Settings** > **Email & collaboration**) that, when enabled, will redirect you from *protection.office.com* to *security.microsoft.com*.

As part of the rollout of the new centers, Microsoft has added a “getting started” interface that offers to lead you through a few of the major changes, as well as a way to provide feedback. Once you’ve dismissed this welcome section, you’ll see the familiar card-based interface like the one in the Microsoft 365 admin center. You can customize the appearance of the security and compliance centers by moving the widgets around to suit the needs of your tenant; the **Add cards** button at the top of the pages will allow you to make changes.

### Using the Purview Compliance Portal

The in-depth material covering classification, compliance, and related topics later in this book will be useful, but a summary may help as a quick reference. Figure 4-2 shows the layout of the Microsoft Purview



Compliance portal. The following sections form the Compliance portal; all are accessible through the left navigation bar.

- **Compliance Manager** is a new toolset that combines a scoring system (known as Compliance Score) with a set of recommendations that you can apply to improve your score. Microsoft also includes assessment templates for specific compliance regimes (such as GDPR or US state-level data breach notification laws) that you can run against your tenant to see how you do. For more on this, see Chapter 17.
- **Data classification** is where an organization manages the application of retention and sensitivity labels to control the protection and retention of user content through tools like the Content explorer and Activity Explorer. Chapter 18 covers retention labels and the policies used to distribute labels to workloads. Chapter 20 covers sensitivity labels.
- **Data connectors** create and manage connections to outside data sources that you want to include in your compliance coverage. For example, you can purchase a license to a connector that allows you to ingest data from AT&T's SMS network to archive and apply compliance policies to employee text messages.
- **Alerts** is where you view security alerts for the tenant. These alerts are created by alert policies, which allow you to set conditions that match the activity for which you would like to be alerted. There are dozens of alert conditions available covering a wide range of possible security concerns such as malware detection, file and folder sharing, DLP policy matches, permissions changes, and many more. Chapter 17 covers activity alerts and alert policies in more detail.

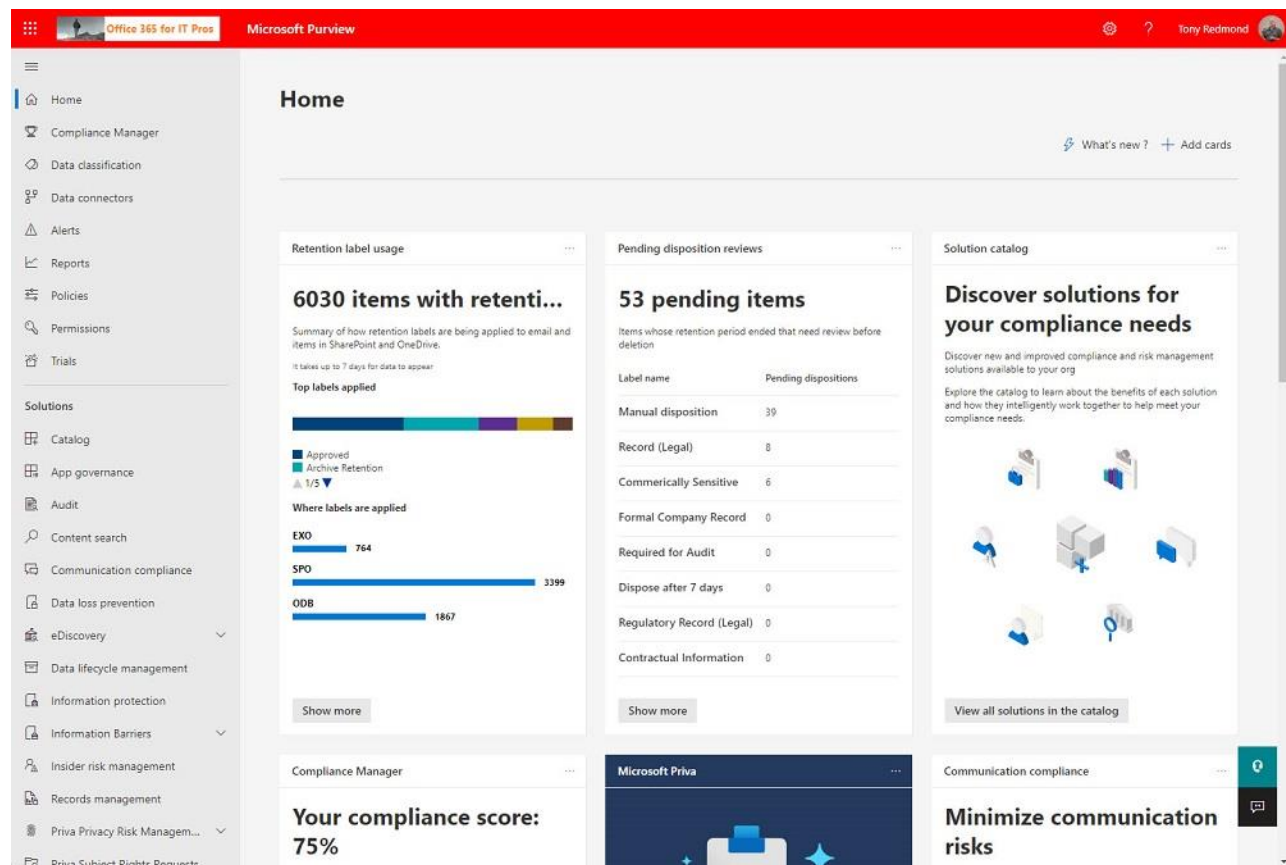


Figure 4-2: The Microsoft Purview Compliance Portal

- **Reports** gives you tools to view and download security and compliance-related reports from a dashboard customized for the tenant. The exact set of reports revealed in the dashboard depends on the functionality licensed by the tenant and can include, DLP, Defender for Office 365, and anti-

malware reports. You can schedule some of these reports for distribution via email to whoever you need to get the information. Chapter 21 covers reports.

- **Policies** – This is intended to eventually be the primary management location for all your compliance and governance policies, including policies for data loss prevention and alerting. For example, this section gives you links to manage DLP policies for content stored or transmitted through Exchange Online, Teams, OneDrive for Business, and SharePoint Online (Chapter 19).
- **Permissions** - View role assignments that allow users access to compliance functionality and data. You can assign compliance roles (but not Azure AD roles) from this section. Users who hold the Global administrator right will have access to the Security and Compliance centers. You can use the roles listed in the Permissions tab to grant selective access to other users.
- **Trials** is essentially like the old SkyMall catalogs that used to grace every seatback pocket on commercial flights in the US; it's where Microsoft advertises trial versions of services that they hope you'll buy.
- **Solutions** – This section collects the tools you can use to take various compliance-related actions. For example, the audit logs, content searches, data loss prevention, and information protection toolsets are all linked from here. The best place to start is probably the **Catalog** section, which provides a simple grid view of different compliance-related tools in Microsoft 365 along with short explanations of what they do.

## Managing Microsoft 365 Defender

[Microsoft 365 Defender](#) is a complex beast. Its unified portal has replaced the previous Security Center as well as the Office 365 Advanced Threat Protection portal. The sections and capabilities you see in your tenant may vary according to the region hosting your tenant, whether it's a commercial, academic, or government tenant, and the licenses you've purchased.

The Defender left navigation rail divides into at least three sections. You can customize which items are shown with the **Customize navigation** link at the very bottom.

The top section covers the broad functional categories that are available in the Defender suite. These options appear for every tenant:

- **Home** returns to the main Microsoft 365 Defender page.
- **Incidents & alerts** shows you data about security incidents that Microsoft has detected in your tenant.
- **Actions & submissions** collects data about messages that users have reported as spam, phishing, or malware. By default, entries from the last 7 days are shown, but submitted messages are kept for up to 30 days. For any user-submitted report, you can forward the report to Microsoft and tag it as clean, malware, phishing, or spam, or you can use the message to start an automated investigative process.
- **Secure score** provides access to the secure score dashboard described later in this chapter.
- **Trials** contains links to other Microsoft security and compliance services that you might want to try.

If you have assigned Defender for Office 365 licenses to accounts (either by buying them or starting a 90-day trial), you'll see additional categories:

- **Hunting** lets you perform security-related queries to look for suspicious activity, as well as configure custom detection rules to specify things that you want to be considered suspicious.
- **Actions & submissions** adds an Action center that lets you review messages or attachments that look suspicious (or that were reported by users) and then optionally send them to Microsoft for further analysis.
- **Learning hub** contains an extensive set of documentation, videos, and other learning materials to help you figure out how and when to use the Defender suite.

One of the interesting features of the new Microsoft 365 Defender portal is that it unifies data formerly stored across modules; for example, if you had both Defender for Office 365 and Microsoft 365 Defender for Endpoint, you'd have to look in two separate places to get a comprehensive view of incident data. Now all the incident data (or actions, or hunting rules, etc.) across all Defender workloads are collected into a single interface.

The **Email & collaboration** section is next. Every Office 365 customer will see some of the options here; some specific options (such as the Campaigns item) require the Defender for Office 365 P2 SKU, which you may or may not have purchased.

- **Investigations** lets you view the results of [automated investigation](#) processes triggered either by the service itself or by you, in response to items shown in the **Explorer** pivot or
- **Explorer** lets you query messages to look for specific patterns; you can filter and query by date or content, and there are predefined pivots that will show you what the service has characterized as phishing messages, malware, or campaigns. One interesting view in Explorer will show you all the URLs in the selected set of messages, which is interesting in a creepy sort of way when you realize how many clickable objects show up in users' inboxes each day.
- **Review** shows a simple card-based interface that gives you easy access to the Action Center, quarantine, and user blocking portions of the portal.
- **Campaigns** shows how the service has categorized phishing messages that appear to be part of structured campaigns. Microsoft [says](#) that a campaign "is a coordinated email attack against one or many organizations" and having a view that shows which campaigns may currently be targeting your users is very useful in responding quickly.
- **Threat tracker** shows the results of queries you've defined in Explorer—after you've defined a query, you'll see the query here and can use this view to review what threats may be associated with it.
- **Exchange message trace** redirects to the message tracing tools in the EAC, described in Chapter 7.
- [Attack simulation training](#) lets you run "benign cyberattack simulations" in Microsoft's words. The goal is to let you simulate various kinds of phishing attacks through SMS, web pages, and emails to see how your users respond and tailor your user education, policies, and response policies to improve your security. This feature requires licensing for Microsoft Defender for Office 365
- **Policies & rules** contains links that give you access to assorted security policies from throughout the Microsoft 365 platform. For example, there are links to alerting policies, device configuration and compliance policies for Intune (see Chapter 16), and anti-malware and anti-phishing policies from Defender ATP.

If you've purchased Microsoft 365 Defender for Endpoints, you'll see its capabilities, including device inventory, vulnerability management, and device configuration management, under a section in the left navbar labeled **Endpoints**. Microsoft has more details on how Defender for Endpoint is integrated with the new Microsoft 365 Defender interface [in this article](#).

Below the **Email & collaboration** section, you may see other sections depending on what other Microsoft products and services you've purchased. Below all those, if any, you'll see a standardized set of additional options:

- **Reports** accesses security-related reports including summaries of mail flow, user-reported phish or spam messages, and potentially compromised users. You can't customize these reports, but you can schedule them to be mailed, or download them.
- **Audit** lets you search the unified audit log; you can also see and change any audit retention policies applied to your tenant.
- **Health** contains links to the Microsoft 365 message center and service health dashboard.

- **Permissions** and **Settings** do what you'd expect—note that these sections are both very bare-bones since most of the other interesting permissions and settings controls are in the primary Admin center.

## Managing Other Microsoft 365 Services

There are many other services in Microsoft 365 besides the “big four”: Exchange, Teams, SharePoint, and OneDrive for Business. Microsoft has been adding new services at a rapid pace since it first launched the service.

Yammer is an enterprise social network for enterprises and teams. The Yammer admin portal includes options to activate your Yammer network, customize the appearance, manage users and policies, and perform community management tasks. The URL for the Yammer admin portal will vary based on your tenant domain, but you can use the general admin URL to access the portal, which then redirects to your specific tenant URL automatically.

Power Automate and Power Apps are two of the main components of what Microsoft calls the Power Platform. Microsoft bundles the administrative elements for these components together into a [single admin center](#). See Chapter 23 for more information.

Dynamics 365, the latest version of Microsoft's venerable customer relationship management (CRM) system that competes with Salesforce, has a link in the **Admin centers** section. However, the link takes you to an error page if clicked while logged in to an account that doesn't have admin permissions on the organization's Dynamics 365 instance.

## Managing Azure Services

The Azure management portal administers other Microsoft cloud services for your tenant, or for separate Azure subscriptions that your organization also has. If The link in the Microsoft 365 admin center takes you to the [Azure AD portal](#). You can navigate from there to the entry point for the [Azure portal](#) to manage the other Azure services that your tenant users and manage Intune mobile device and application policies.

Among the most common management operations are tasks such as creating new user accounts, updating their properties, and managing their permissions. As you learned in Chapter 3, Microsoft 365 supports standalone and hybrid identities in a variety of configurations. The thing that all identity types have in common is that the service relies on Azure AD, so it's helpful to have some familiarity with the Azure AD management tools.

When you open the Azure AD portal, you'll see that it has categories for managing users, groups, roles, devices, Azure AD Connect, and password reset (along with several other categories). It also has links to controls for conditional access policies, reports for risky logins and suspicious sign-ins, and troubleshooting tools. A complete exploration of this section is far outside the scope of this book, but we've covered selected parts of Azure AD in Chapters 3 and 16. It is well worth your time to explore the settings here, as many features and settings exist to improve the security and functionality of your environment that you may not be familiar with.

## Service Administration Apps for Mobile Devices

Microsoft ships a tenant management app for Android and iOS. The management apps can perform common administrative tasks from a PC or mobile device including editing or removing users, resetting passwords, turning email forwarding on or off, and viewing service health information. You can review billing alerts, items posted to the Message Center, and service health notifications. In fact, you can have the app post push notifications when specific service health notifications are issued. However, the management apps do not support managing any aspect of a hybrid configuration.

## ServiceNow Integration

As much as Microsoft might wish it otherwise, ServiceNow is the most popular IT service management (ITSM) solution in the world. It's common for Microsoft 365 administrators and support desk staff to have ServiceNow more or less permanently open in one tab with various Microsoft admin portals open in other tabs. In a welcome move, Microsoft published an [app in the ServiceNow app store](#) that provides some integration between the two. If you install the app in your ServiceNow instance, you can review SHD items and open new tickets with Microsoft. This is a useful capability and it'll be interesting to see what Microsoft does to enhance it in the future.

## Managing Licenses, Plans, and Billing

Everyone who wants to use an application besides the administration tools needs a license. Because Microsoft controls all aspects of the service, they can be more insistent on this point than they are with on-premises Client Access Licenses (CALs). You cannot use an application like Exchange Online without a license and an account that has its license removed will end up losing all its data. License management is therefore an important part of administration, with the goal being to ensure that you have enough licenses to allow people to work while not paying for unused licenses. The **Billing** section in the admin portal allows you to buy licenses, see and manage the licenses you have, and get various types of billing notifications.

The **Licenses** option in the Billing section in the admin portal provides an overview of the licenses available to, and used by, a tenant. As we can see in Figure 4-3, some Office 365 E5 licenses are unassigned. This might become an issue if the situation persists because Microsoft charges for every license monthly even if no one uses the license.

**Note:** There are several ways to buy licenses. Depending on what subscriptions you have, and whether you have a licensing agreement with Microsoft, you may buy licenses directly through the web portal, through a reseller or partner, or directly from Microsoft. For partner and Microsoft purchases, the license term and cost are whatever you negotiate; for purchases directly through the portal, the licenses you buy cover a 12-month term, but you may pay for them monthly or annually. In either case, you pay up front. Microsoft doesn't give refunds for unused licenses. As you add licenses, you will see them "roll in" at renewal time. For example, suppose that on January 1 you create a new tenant and buy 100 Office 365 E5 licenses. Then in March, you buy another 100 licenses. Come January 1 next year, all 200 licenses will renew at the same time.

The usage reports in the Microsoft 365 admin center provide you with a basic view of who uses the services and what licenses they have, allowing you to decide about reallocating licenses or reducing the number of licenses that you're paying for.

The **Requests** and **Auto-claim policy** pivots in this view are covered in Chapter 5, as they deal with automating (at least partly) the process of assigning licenses to specific users.

You can also perform some license management tasks in the **Your products** section under the **Billing** section, as shown in Figure 4-3. Functionally the two areas of the portal are very similar, allowing you to assign licenses or buy more if you have no unused licenses. The "Settings & actions" section in each license card allows you to take appropriate actions, including installing the Office desktop apps (if licensed) or assigning licenses to users. The **Your products** section also includes pivots for Azure subscriptions, purchased applications, and usage benefits associated with enterprise agreements (EAs) or other purchase contracts.

## Your products

These are products owned by your organization that were bought from Microsoft or 3rd-party providers. Select a product to manage product and billing settings or assign licenses.

Products Benefits

Microsoft 365 Domestic Calling Plan (120 min) will be deleted on 7/4/2021. [Reactivate now](#)

Search 4 filters selected

### Microsoft products (9)

Product name	Assigned licenses	Purchased quantity	Subscription status	Paid with	Purchase channel	Billing account	Choose columns
Azure Active Directory Premium P2	1	1	Active: Renews on 4/15/2022	Invoice	Commercial direct	Quadrotech Solution	
Exchange Online (Plan 1)	10	10	Active: Renews on 3/5/2022	Invoice	Commercial direct	Quadrotech Solution	
Microsoft 365 Audio Conferencing	117	125	Expired: 5/15/2021	Prepaid	Commercial direct	Quadrotech Solution	
Microsoft 365 Domestic Calling Plan (120 min)	20	22	Disabled	Invoice	Commercial direct	Quadrotech Solution	
Microsoft 365 Phone System - Virtual User	10	50	Active: Renews on 3/22/2022	Invoice	Commercial direct	Quadrotech Solution	
Office 365 E5 without Audio Conferencing	150	150	Active: Expires on 11/30/2021	Prepaid	Commercial direct	Quadrotech Solution	
Power BI (free)	3	200	Active: Renews on 6/30/2021	Invoice	Commercial direct	Quadrotech Solution	
Visio Plan 2	10	9	Active: Expires on 7/2/2021	Prepaid	Commercial direct	Quadrotech Solution	
Visio Plan 2	10	2	Active: Renews on 5/11/2022	Invoice	Commercial direct	Quadrotech Solution	

### Pay as you go (1)

Product name	Purchased quantity	Subscription status	Paid with	Billing account	Choose columns
Communications Credits	CHF 0.00 credits	Active	Not available	Quadrotech Solution	

### Azure (1)

Product name	Purchased quantity	Subscription status	Billing account	Choose columns
SendGrid - Free	1	Active: Renews on 7/3/2021	anton.sutjak@quadr	

Figure 4-3: Managing licenses in the Products and Services section of the admin portal

The other items in the **Billing** node allow you to select and set up payment methods for your bills, choose who gets billing notifications and in what language and format, and buy new products—but all of these options are simple enough that we won't discuss them further here.

For more information on how to assign licenses to users, see Chapter 5.

## Managing Integrated Apps

Microsoft has worked for literally decades to make Office an extensible platform, both because it enables them to extend and improve the platform and also because it enables their partners to do so, which in turn helps tie customers to it. This strategy has continued into the present day, leading to the new Integrated Apps page under the settings tab in the Microsoft 365 admin center. The controls here allow you to select and deploy add-ins that run inside various Office clients. Figure 4-4 shows the view from a sample tenant. You may notice that the caption at the top refers to applications developed by Microsoft partners, yet some of the applications (such as FindTime) shown in the list are from Microsoft. If you use the **Get apps** button to jump to the AppSource marketplace for Microsoft 365 Apps add-ins, you'll find several dozen add-ins, many of which come from Microsoft, so don't let the "partner" wording fool you.

Any add-in written to use the [correct set of APIs](#) can be deployed and managed here. The beauty of this approach is that a well-written add-in for, say, Outlook will work in any recent version of Outlook, including the desktop, mobile, and web versions. The advantage of using the **Integrated apps** page is that you can select an add-in, optionally create a test deployment, and then specify which users can access the app (just you, the entire organization, or a list of users or groups you specify). Because these apps are add-ins downloaded and run by the client application, you don't have to do anything further to deploy these capabilities to users.

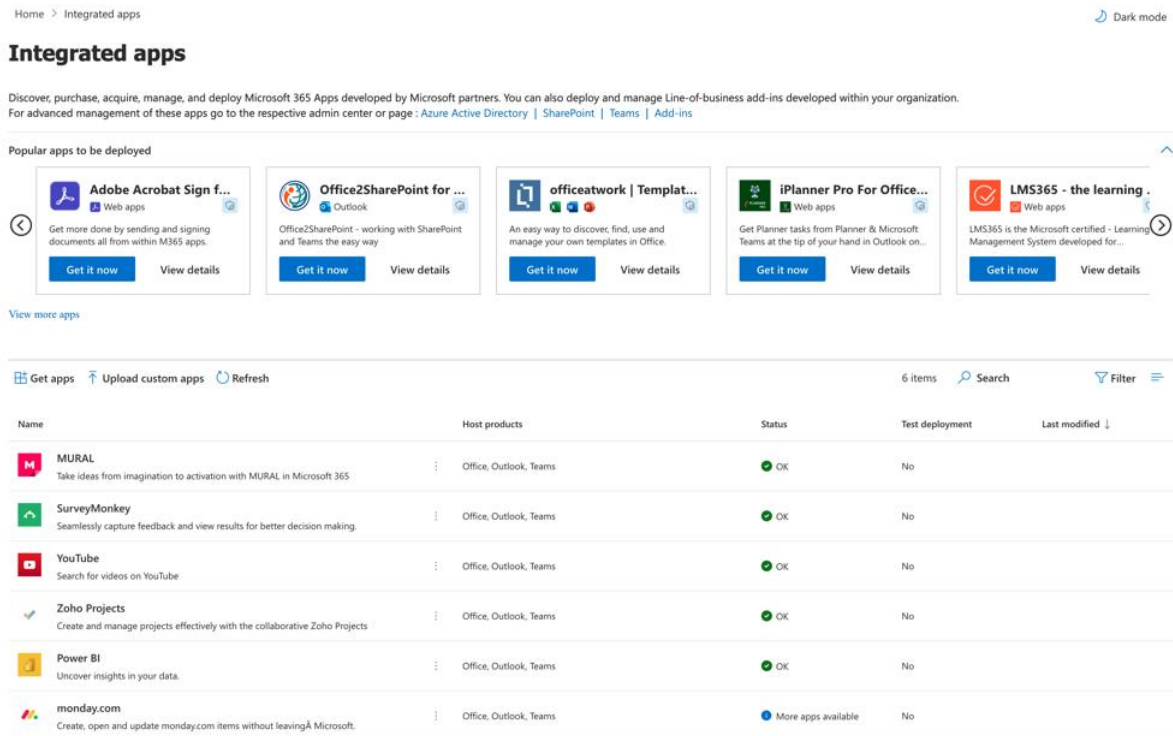


Figure 4-4: The Integrated apps view lets you centrally install and publish Office add-ins to users

**Protecting against app-based attacks:** The **Integrated Apps** node in the Settings section of the Microsoft 365 admin center controls whether users can allow applications to access their data, such as their calendar. By default, the setting is on. Because some malware exploits work by gaining access to user data ([here is an example](#) that encrypts user data for ransom), you might prefer to turn the setting off. It is worth emphasizing again that unless users are aware of threats and sensitive to granting access to their data, the possibility always exists that they will succumb to an attack if they are targeted. Microsoft refers to these as “[illicit consent grants](#)” and they have documentation and tools to help you review what applications users have consented to, including [two PowerShell scripts](#) that create a set of CSV files holding details of access granted by users to applications and their publishers.

## Managing Feature Releases

Microsoft uses a series of “rings” to control the release of new features. In general, the first ring is the development team, the next is Microsoft, the third is composed of tenants who have nominated themselves by signing up for “Targeted Release,” and the last is general availability (also known as “standard release”). This approach varies by workload; individual applications, such as Teams or Planner, may use more rings in their deployment model.

“Standard release” is the term Microsoft uses when they make a feature available to all tenants licensed for the functionality. Tenants who opt for Targeted Release see new features (or updates to existing features) a few weeks ahead of general availability. However, for some new applications, the period covered by Targeted Release can extend up to several months.

The ability to control how new features become available to tenant users is through **Release preferences** in the **Organization profile** pivot, available through the **Settings > Org settings** section of the Microsoft 365 admin center. Three options are available:

- Standard release for everyone.
- Targeted release for everyone.
- Targeted release for selected users (Figure 4-5).

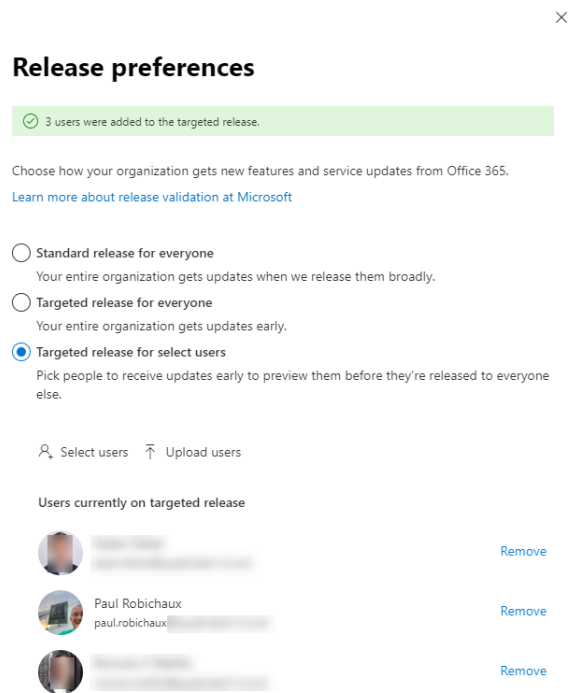


Figure 4-5: Enabling Targeted Release for a tenant

The last option allows the administrator to nominate specific individuals or groups (referred to as a staged rollout) while the remainder of the users continues to access the functionality available in the standard release. For example, you might decide that only the IT department should be exposed to a new application to allow some experience to be gathered about its functionality and so drive a decision about how it might be used within the organization.

Switching an account from Targeted Release to Standard Release or vice versa can take up to 24 hours before the switch is completed. Keep in mind that putting a user in Targeted Release just enables their ability to see or use a feature; client-side features implemented in the Microsoft 365 Apps software won't be visible to the user until she downloads the appropriate version of the software. Depending on your organizational policies, this may add an extra delay before the new features become available.

Although Microsoft's goal is to move features from Targeted Release into general availability reasonably quickly, the exact period between the two phases can vary from feature to feature. Timings can be affected by user feedback, bug reports, or the discovery of flaws such as performance or scalability issues that must be addressed before a feature can be made available to all. The shortest period is a week or so. The longest (to date) is six months. The rate of change also varies within applications. Exchange Online tends to see more new features over a certain period than SharePoint Online while the rate of change for new features can also be high. Yammer and Teams do not use Targeted Release because they use different mechanisms to give early access to new features to a select group of tenants.

Even when a feature is released to general availability, it doesn't mean it shows up everywhere at once. It can take a few weeks before every tenant in every data center sees a new feature. The exact speed of deployment depends on the results of Targeted Release deployments, plus other factors such as local market support. Microsoft also often staggers feature release dates by geography. In general, features tend to appear first in North America, the UK, Western Europe, and Australia, then in other areas. For example, more than two years after the first release of Teams phone calling features, the feature isn't available worldwide yet.

Another factor in feature availability is the type of tenant you have. In general, enterprise tenants get features first, and education and government tenants (including those in the Government Community Cloud,



or GCC) come later, if at all. The small-business versions of Microsoft 365 fall somewhere in the middle. Keep in mind that Microsoft appears to believe that they will make the most money by delivering all their features to all tenants everywhere, so this is their aim; when they don't make a feature available in a region or to a type of tenant, it's either because engineering, operational, or legal constraints exist to prevent it or because the feature requires enough investment that it may not be profitable for them to deploy it.

When new features arrive in your tenant, they will almost always be enabled by default for all users with the necessary licenses, regardless of your Organization Profile settings. That includes features such as Microsoft Teams and Planner that Microsoft added as service plans to current products; in other words, Microsoft added the features to existing licenses such as Enterprise E1, Enterprise E5, and so on, rather than creating separately licensed services like Intune or the various Dynamics 365 services. Applications that don't require a separate license are also usually enabled by default. Note that over time, some applications may change their licensing status; as one example, Whiteboard, which originally had no license, is now a licensed service plan within the Office 365 SKUs, which means that tenants can control if users have access to Whiteboard.

Features released in Public Preview before General Availability are not enabled by default; however, they will become enabled once the feature reaches General Availability. The approach of enabling features by default makes adoption easier but is a concern for organizations that have strict change management procedures, or those who want to control the rollout of new features so that they can provide appropriate training to their IT support staff and end-users. It is important to maintain awareness of the changes that are rolling out to customers by monitoring your Message Center notifications and keeping an eye on other sources of early information such as the [Microsoft 365 roadmap](#).

**Mastering Targeted Release:** Targeted Release allows tenants to gain faster access to new functionality. However, there is a downside to the value gained by seeing new features earlier than the norm. Microsoft tests new code before making it available through Targeted Release, but it is well known that frayed edges (bugs) can appear in new software when exposed to the stresses of production workloads. Other common challenges include new features showing up in user interfaces without warning or any sign as to how they should be used with little or no available documentation. Selective Targeted Release is available to allow some users in a tenant to access new features earlier, but this option is sometimes not supported by an application.

To mitigate these downsides of Targeted Release, some companies run two tenants: a "production" tenant for most users and another trial tenant configured as Targeted Release for a small subset of users responsible for testing new software as it appears. This is analogous to the situation often found in the on-premises world where customers keep a test environment to deploy new builds of Exchange as Microsoft makes them available.

It is important to realize that the appearance of a new feature to tenants who have enabled Targeted Release does not mean that the feature will be made available for all tenants soon afterward. Microsoft can and does use time to gain knowledge of how customers use a feature and can adapt and change the feature over an extended period before deciding that the software is suitable for general consumption and ready to be rolled out globally. Targeted Release is a large-scale beta program, and no assumption should be made when or if any feature will be available just because Microsoft makes the functionality available to tenants configured for Targeted Release.

## Managing Connectivity

Each of the workloads in Microsoft 365 has its unique requirements for connectivity. There are two basic invariants: every service needs reliable and correct DNS information, and every service needs to be able to pass traffic on TCP port 443 to allow TLS-protected HTTP communications. Individual services may have additional connectivity requirements for various features. This seems easy to deal with in theory... right up until users start complaining that some feature or another in their clients isn't working properly.

Some years ago, Microsoft introduced a tool called the Exchange Remote Connectivity Analyzer (ExRCA) for troubleshooting on-premises connectivity for MAPI and Exchange ActiveSync. Over the years that tool has grown in scope and power and is now known as the [Microsoft Remote Connectivity Analyzer](#) (Figure 4-6). It currently supports more than two dozen tests for Exchange, Teams, on-premises Skype for Business and Lync, and Office 365.

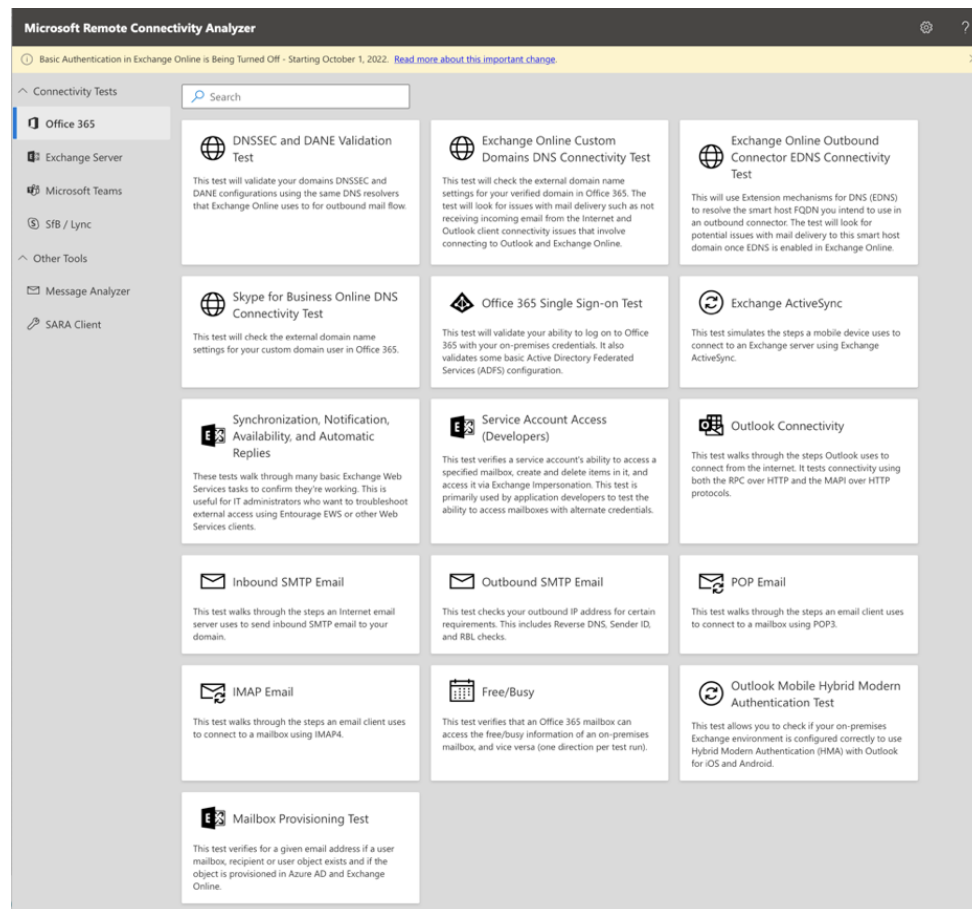


Figure 4-6: The Microsoft Remote Connectivity Analyzer main page

Currently, the analyzer supports the following tests on the Office 365 tab. Some tests require you to provide valid credentials for an active Office 365 account that's licensed for the workload you're testing because the test requires logging into a mailbox or Teams account; others (such as the Outlook Mobile modern authentication test) require an email address but not a password:

- **DNSSEC/DANE:** this test checks the DNS records you have defined using the same DNS resolution path that Exchange Online users. If any of these tests fail, so will DANE resolution for your domain.
- **Exchange Online custom domain test:** this test validates whether the specified custom domain is correctly set up and has valid MX records for mail flow into the service.
- **Exchange Online outbound EDNS test:** this test validates whether your DNS infrastructure can handle the [DNS extensions](#) required for using DNSSEC.
- **Skype for Business DNS connectivity:** this test checks for the presence and correctness of the DNS records needed for Skype/Lync hybrid connectivity.
- **Office 365 single-sign on test:** this test verifies whether single sign-on to the service works correctly.
- **Exchange ActiveSync test:** this set of tests emulates a mobile device and makes Exchange ActiveSync requests for synchronization to ensure that EAS network traffic works properly:
- **Synchronization, notification, availability, and automatic replies test:** this set of tests checks various Exchange Web Services (EWS) network flows that are used by Outlook clients.

- Service account access test: this test verifies whether the specified account is correctly configured to use Exchange impersonation for service account-level access to mailboxes.
- Outlook connectivity test: as its name suggests, this test suite does an end-to-end test of all the steps performed by the Windows Outlook client for it to connect over RPC or MAPI over HTTP.
- Inbound SMTP email test: this test checks whether the DNS and network configuration of your tenant is correct so that outside servers can send mail to it.
- Outbound SMTP email test: the inverse of the inbound SMTP test, this test checks whether your outbound IP address has correct reverse DNS, Sender ID, and real-time block list (RBL) settings.
- POP email test and IMAP email test: these tests perform the same steps that a POP or IMAP client takes to verify connectivity.
- Free/busy connectivity test: this test verifies that a cloud mailbox can access on-premises free/busy data, and vice versa.
- Outlook Mobile hybrid modern authentication test checks to see whether you've correctly set up hybrid modern authentication between Exchange Online and your on-premises Exchange environment.
- Mailbox provisioning test: this test checks to ensure that the specified mailbox is correctly provisioned in the service and that all the required directory attributes have valid values.

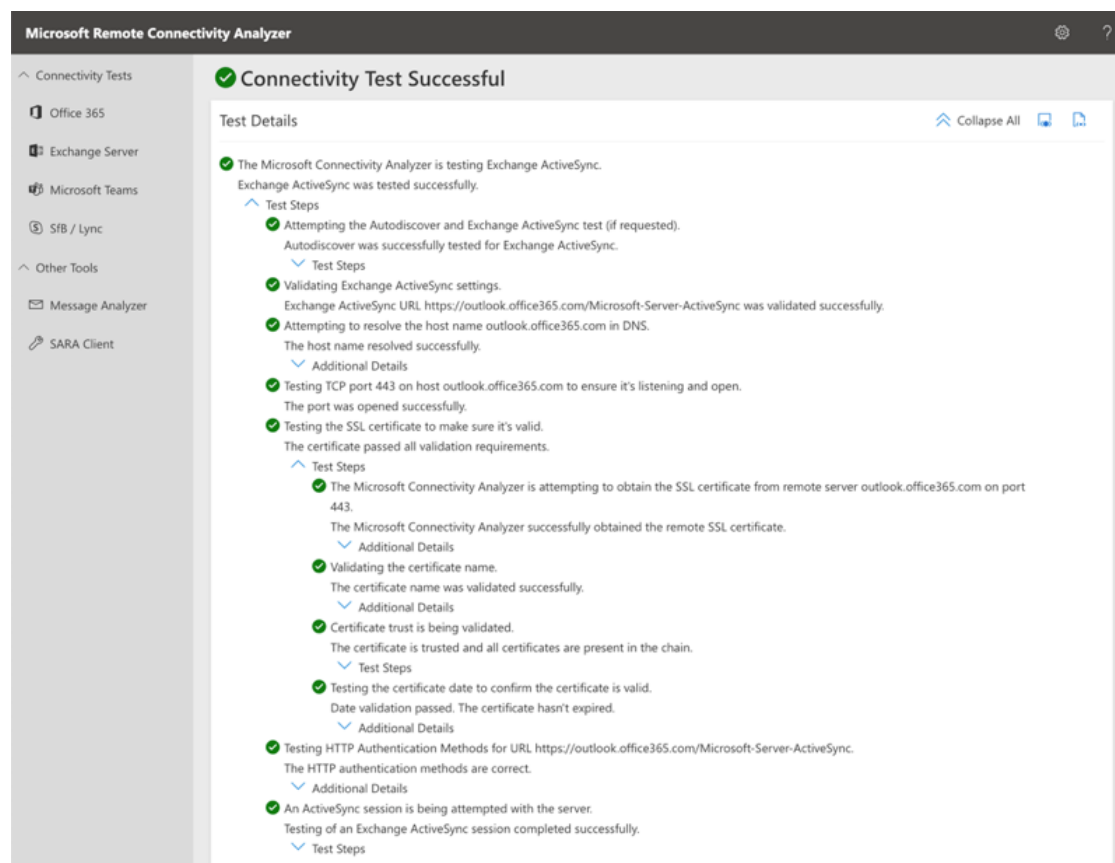


Figure 4-7: Running a test with the Microsoft Remote Connectivity Analyzer

On the Microsoft Teams tab, there's a separate set of tests:

- Teams Calendar Tab: this test verifies that the Teams service can connect to an Exchange mailbox.
- Teams Meeting Delegation: this test checks whether a specified account has the right permissions to schedule Teams meetings on behalf of a delegate.
- Teams Meeting Recording: this test verifies whether the specified account has permission to record a meeting and store the resulting recording.

- Teams Exchange Integration: this test checks whether the Teams services can connect to Exchange mailboxes; you can run this against either on-premises or cloud mailboxes.
- Teams Channel Meeting: this test checks to see whether the specified account can schedule a Teams channel meeting.
- Teams Voicemail: use this test to verify that the selected account can retrieve and store voicemail messages.
- Teams Presence Based on Calendar Events: this test will check to see if Teams can read calendar items and use them to update Teams presence.

The analyzer's tests are very simple to use. In general, you'll supply an email address or DNS domain, provide a password if required, and run the test. RCA presents the results of each test in a clear, easy-to-understand format (Figure 4-7). In most cases, these tests alone will give you enough information to figure out the specific element of your network, DNS, or on-premises environment that's misconfigured and lead you to get it fixed.

## Protecting Data with Encryption

Encryption is a complicated subject. The legal, technical, operational, and business implications of where, when, and how data are encrypted could easily fill a series of books. As described in Chapter 2, Microsoft uses two primary strategies for encrypting data: some data are encrypted in transit, and some are encrypted while at rest.

In general, encryption works best when it is automatic and ubiquitous, and Microsoft has done pretty well at building ubiquitous transport and at-rest encryption into the service (as [detailed here](#)). For example, by default, all communications to and from the service are protected with TLS, and servers running the Microsoft 365 applications all use BitLocker to protect their physical disks. However, Microsoft offers several additional [encryption services](#) that may be of use to you. For example, you can require the use of mobile-device encryption on Android and iOS or iPadOS devices through Intune, and you can create conditional access policies that only allow encrypted devices to connect to various services. If you have the infrastructure for it, you can enable or require the use of S/MIME for messages, and so on. Keep in mind that not every encryption feature is globally available; Microsoft 365 tenants in particular regions or government-associated clouds may not have the same features as commercial tenants.

### Service Encryption

The basic level of data-at-rest encryption provided in Microsoft 365 is called [service encryption](#). Along with the Windows-standard BitLocker tool, Microsoft has built a tool called [Distributed Key Manager](#) to ensure that the BitLocker recovery keys are protected and secured. Microsoft [describes DKM's function](#) by saying

*"Only members of a specific security group in Active Directory Domain Services can access those keys to decrypt the data that is encrypted by DKM. In Exchange Online, only certain service accounts under which the Exchange processes run are part of that security group. As part of standard operating procedure in the data center, no human is given credentials that are part of this security group and therefore no human has access to the keys that can decrypt these secrets."*

The good news is that, even if you never change any default settings, every file and message stored by the service will be encrypted by Microsoft using the BitLocker and DKM combination.

### Customer Key

One of the longest-running arguments in the encryption world is "who holds the key?" As you might imagine, some organizations want to retain complete control over the use of encryption in their enterprise.

This is usually for one (or both) of two reasons: they either don't want encryption used by an attacker to lock up data or they don't trust a third-party service provider (which in this case would include Microsoft) not to peek at their data. One solution to this is to encrypt data using a key that only the customer has access to. Microsoft offers a capability known as [Customer Key](#) (CK) that allows you to hold the root keys used by the service encryption system. You create keys and upload them to Azure Key Vault, then create data encryption policies (DEPs) that specify which keys to use to encrypt which data items. Microsoft manages encryption keys for any services that don't have DEPs defined.

You can currently define three types of DEPs:

- **Exchange Online mailboxes:** these DEPs are applied per mailbox and only encrypt the mailboxes they're applied to; when you apply a DEP to a mailbox, the mailbox is moved and then encrypted as part of the move.
- **SharePoint Online and OneDrive for Business:** this DEP type applies to content stored in SharePoint or OneDrive, including Teams files. You can create a single DEP per tenant unless you're using multi-geo, in which case you can create one per geo.
- **Multi-workload DEPs** apply to multiple workloads and components for all users in the tenant. You can encrypt additional Teams data (including 1:1 chats, group chats, meeting chats, conversations in channels, chat notifications, and images and videos). When you create a DEP at the tenant level it is applied in addition to, not as a replacement for, any DEPs that are currently in force on Exchange, OneDrive, or SharePoint data. The tenant-level DEPs will encrypt all data held by Exchange Online, meaning you no longer have to assign DEPs to individual mailboxes; for Teams, you can only encrypt data created after the time the policy is applied (meaning you can't encrypt historical data).

The [setup process](#) for CK is quite complex, so you should carefully consider whether you have the needed skill set and operational maturity to make use of it; it's easy to create a situation where your data items are encrypted in ways that might render them unusable in the future if you lose access to a particular key.

## Double-Key Encryption

Microsoft also supports another encryption mechanism that uses *two* keys: one that Microsoft has, and one that only your organization has. Data protected with this mechanism (known as [Double-Key Encryption](#), or DKE) is encrypted twice: first with your key and then with Microsoft's. Microsoft never has access to your key, so they can never read your protected data. For this reason, though, data protected with DKE can't be processed by several useful service features, including transport rules, content search, and eDiscovery. Sensitivity labels (see Chapter 23) can apply protection to files using DKE.

DKE requires you to set up a key-management system of your own, build and install a connector, and deploy sensitivity labels that end users can use to tag specific documents as needing DKE protection. Microsoft makes the point very clearly that DKE is only intended for the most critical items that need the heaviest protection and isn't intended to replace the use of CK or other protection mechanisms for everyday use.

## Protecting Items with Information Protection

Microsoft 365 already includes a broad set of features for protecting messages, documents, and containers. As described in Chapter 23, you can create sensitivity labels to automatically mark specific items as requiring protection; you can use [Purview Message Encryption](#) (OME) to enforce encryption for email messages under various conditions, and you can apply rights-management restrictions to control which users can do what with specific items based on their origin, location, or sensitivity. Microsoft broadly lumps these capabilities together under the heading of "information protection"; the encryption features enabled in their information protection solution are intended to reduce the risk of accidental or purposeful disclosure of sensitive data to unauthorized people, but they aren't necessarily intended to protect against other threats.

## TLS Certificate Updates

The certificates that Microsoft uses to provide Transport Layer Security (TLS) protection are all issued by root certificate authorities. Between January 2022 and October 2022, Microsoft is moving its certificate issuance to a new root CA, replacing Baltimore CyberTrust with DigiCert. This change should be transparent to most tenants and administrators. However, if you have applications that are hard-coded (or “pinned”) to use a specific root CA for certificate validation, this change may break your applications. Microsoft maintains a [page explaining the change](#).

## Handling the TLS 1.x Deprecation

As you read in Chapter 2, Microsoft has been planning to stop supporting older versions of the TLS protocol for some time. To prevent any unpleasant surprises for your users and applications, it’s important to know which endpoints still use the older versions. One way to do this is to use [Azure Log Analytics](#) to query your Microsoft 365 log data to reveal the details. Here’s an example query:

```
SigninLogs
| extend details = parse_json(AuthenticationProcessingDetails)
| mv-expand details
| where details.key == "Legacy TLS (TLS 1.0, 1.2, 3DES)"
| extend LegacyTLSEnabled = details.value
| project-away details
| where LegacyTLSEnabled == true
| project TimeGenerated, AppDisplayName, AppId, UserPrincipalName, Identity, ClientAppUsed, IPAddress
```

## Monitoring

Although Microsoft 365 removes much of the responsibility for ongoing support and maintenance for its applications away from administrators, the need still exists for tenants to know how well the services function and if any technical issues exist that might affect the availability or quality delivered by a service at any time. The Microsoft 365 admin center includes mechanisms to understand the current health of applications while third-party applications are also available to increase and improve the quality of monitoring. Service health and other important notifications are available in the left pane of the Microsoft 365 admin center, in the **Health** section.

### Service Health Dashboard

The **Service health** menu item under the **Health** section of the Microsoft 365 admin center takes you to the Service Health Dashboard (SHD). Microsoft uses the SHD to communicate the current health status for each of the services consumed by your tenant. Each service is displayed with an icon showing whether it is currently healthy, has an active incident that is causing an impact on customers (in other words, something is broken), or has a non-incident advisory. Advisories represent conditions Microsoft wants you to know about but aren’t confirmed as active incidents, such as planned maintenance scheduled for different regions. Figure 4-8 shows how the SHD displays the current service state and any known problems.

Using the **Customize** button, you can customize the SHD to only tell you about the state of services you care about. You can control this in two ways: what you see in the SHD, and what items Microsoft sends you email notifications about. For instance, if you don’t use Intune, it’s unlikely that you will care about knowing when Intune has an outage.

Clicking on an incident displays extra information, including ongoing details of the investigation for an incident, which may span multiple messages posted by Microsoft over a period. After Microsoft resolves an incident, they often inform customers what happened in more detail through the publication of a post-incident review (PIR). Microsoft uses Azure-based machine translation to translate advisories; if you enable

this feature and then use the admin center in a language other than English, the service will try to translate the incident title, user impact, and incident history.

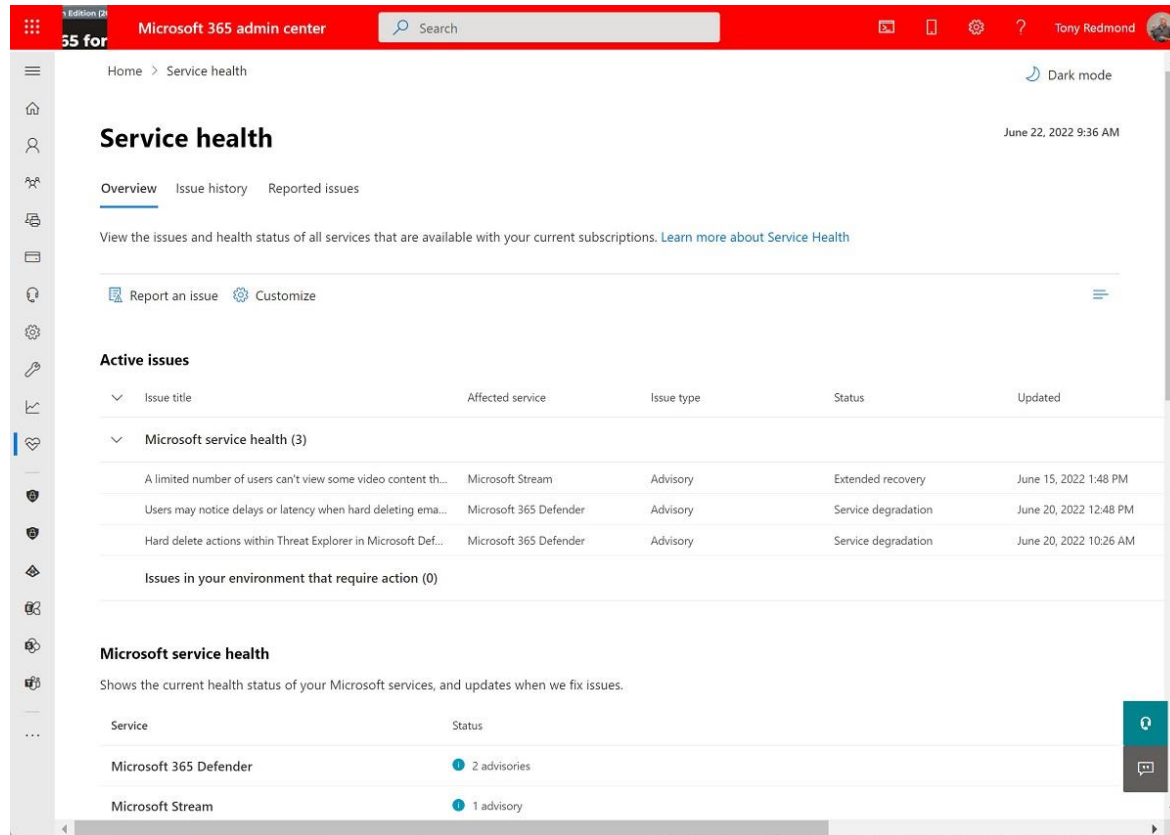


Figure 4-8: Viewing the list of current incidents in the Service Health Dashboard

Sometimes the SHD may alert you to issues that require you to take some sort of action. An example might be that you have a certificate that's about to expire. Issues that require administrative intervention are in the "Issues in your environment that require action" section and, as the name implies, aren't problems with the service... but they will turn into problems if you don't take the required action to resolve them. Microsoft will probably continue to roll out new features to detect environmental issues over time, so you should make checking this item a regular part of your work.

Because of the sheer scale of the service, you can expect that an outage is occurring somewhere in the service at any given point in time. The SHD is a view of problems determined by Microsoft as potentially affecting your tenant, rather than every single issue impacting any tenant across the service. It is also possible that your tenant experiences an issue that the SHD does not show. With that said, it is common to experience problems with the service that affect your users or your entire tenant without seeing it in the SHD.

Microsoft's automated monitoring systems are sometimes slow to detect a problem. In some cases, multiple customers need to report issues and escalate the issues beyond level 1 support teams before Microsoft accepts the issues as confirmed incidents. At that point, Microsoft might add details of the incident to the SHD. Furthermore, you might be experiencing an issue with your tenant that is due to a misconfiguration on your part, not a service fault. For example, if you misconfigure your DNS records and your Exchange Online recipients stop receiving emails, that is not something that the SHD will alert you to because it is not a fault with the service itself. Regardless of any limitations or issues with the timeliness of information, the SHD is a useful resource for administrators.

Microsoft continues to develop the capabilities and intelligence behind the SHD. For example, when logging a new service request the portal tries to proactively inform you of known issues, as well as notifying you

when Microsoft detects actions by your users that may lead to a support request such as clients connecting with out-of-date versions of Outlook. Microsoft also uses telemetry data and machine learning to detect issues, to reduce the amount of time between an incident occurring and the first notification posted to the SHD. There is also a feedback mechanism built into the SHD so that customers can rate the accuracy and usefulness of the information that Microsoft communicates to them.

## Reporting an Issue

You can use the button labeled “Report an issue” at the top of the SHD to register a problem with services such as Exchange Online, SharePoint Online, and Teams. Microsoft says that these reports are correlated with other telemetry data and used to trigger internal investigation (and presumably verification) of potential problems so they can decide which ones are serious enough to pass on to the engineering team and treat as “incidents,” the formal name for service problems that require investigation and repair. It remains to be seen how useful this feature will be either for tenant admins or Microsoft; the risk of false alarms seems high, so the prudent course (which Microsoft also recommends) is to continue to file support tickets in addition to using “Report an Issue.”

## Highlighting Incident Notifications in Outlook

The idea behind incident notifications is simple: when you have Outlook for Windows open for an account that has Microsoft 365 global administrative privileges, the service will post notifications in a new admin notifications pane to tell you about potential problems in your tenant. You might be excused for not having known about this feature before for a simple reason: good security practice dictates that you don’t use your normal work account—you know, the one you’re probably running Outlook against—for administrative tasks. If you’ve followed the normal best practice of assigning separate accounts for your global administrator role holders to use for day-to-day work, you may never see these notifications. Now that you know this feature exists, it might be worth testing it with your global admin accounts to see if you find it worthwhile.

## Getting Incident Notifications via Email

Tenants can opt to receive email notifications of incidents. Even though this capability lags far behind the monitoring notifications provided by third-party products, it’s still good to have. If you want to receive email notifications, enable them by going to the SHD and clicking the **Customize** icon, then choosing Email, then checking the “Send me email notifications for service health incidents” checkbox. Settings changes may take up to 8 hours to take effect, and of course, when you don’t receive an email announcing an outage, it may mean that there *is* no outage, or that one is happening but that it affects mail flow to your specified recipients in some way.

## Receiving Updates for Active Incidents

When you open an individual incident notification, you may notice a link labeled “Manage notifications for this issue.” You can use it to specify up to two email addresses that Microsoft will use to send updates on that specific incident, including any status changes or resolution information. This isn’t quite the same as the “Send me email notification for service health incidents” checkbox mentioned above—the “Manage notifications” link will send you updates only for the specific incidents where you’ve enabled it.

## Programmable Access to Service Incident Information

Microsoft offers the [Microsoft Graph-Service Health and Communications API](#) to help administrators get programmatic access to:

- Notification messages for service updates posted to the Message Center.
- Incident and advisory messages posted to the Service Health dashboard.
- Overview of current workload status.



- Historical incident information.

If you're curious, this [PowerShell script](#) is an example of how to fetch the same set of messages displayed in the Message Center from the API and process them into a form that can be consumed in different ways, such as posting to a Teams channel, export to Power BI, or formatted in an HTML page.

## Other Sources of Service Information

Microsoft maintains a global status page at <https://status.office.com>, but it normally says nothing more than "This site is updated when service issues are preventing tenant administrators from accessing Service health in the Microsoft 365 admin center. Alternatively, customers can reference <https://www.twitter.com/MSFT365Status> for additional insights into widespread, active incidents." Let that sink in: it's a dashboard to check the health of the *real* dashboard.

Users have access to a simple end-user-facing status page at <https://portal.office.com/servicestatus>, although it's not much use for identifying issues.

Finally, given the importance of Azure AD and other Azure-based services to all the Microsoft 365 components, [the Azure status page](#) can also be a useful source of information when things aren't working quite as well as they should be.

## Message Center

The Microsoft 365 admin center dashboard also includes a view of the most recent messages from the Message center. The various product and service operations teams use the Message center to notify customers of changes such as updated features and user interfaces, as well as changes that may interrupt service such as IP address range changes. Items shown in the Message Center are supposed to be specific to a tenant rather than the more general view of what is changing found in the roadmap.

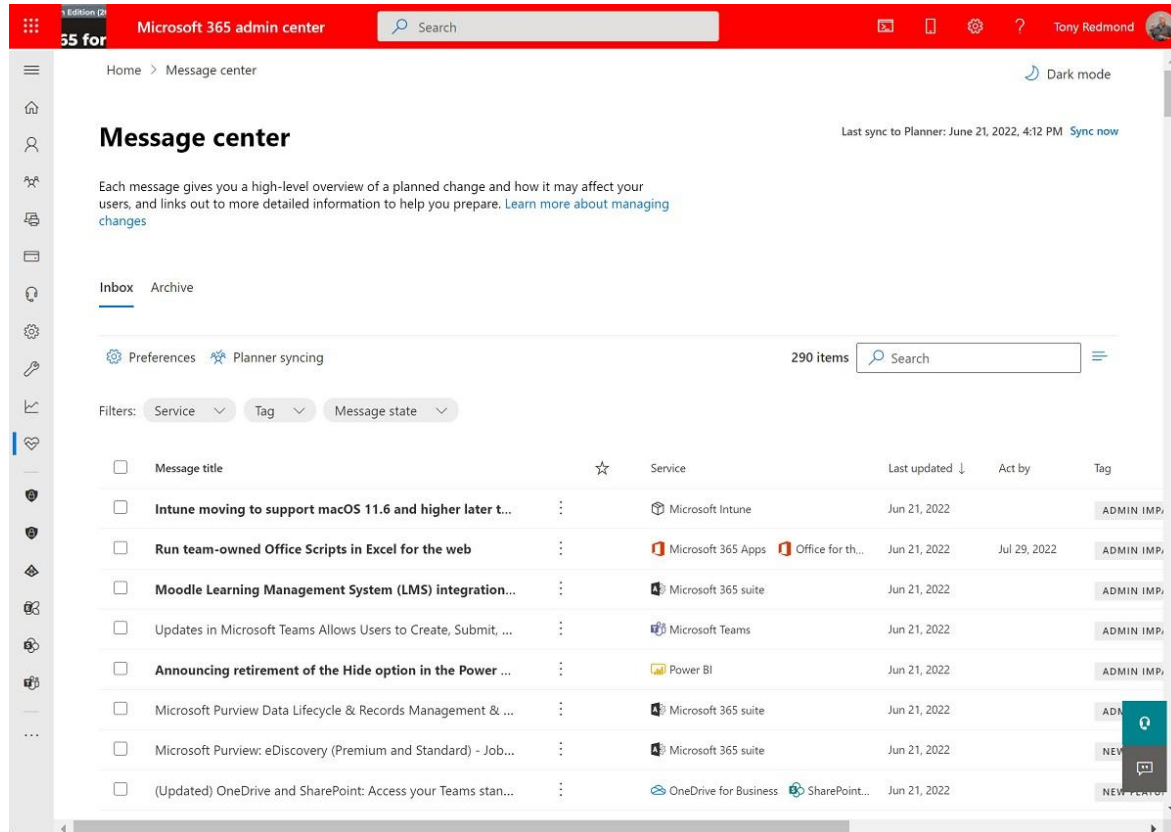


Figure 4-9: Message Center notifications about functionality updates

To make sure that you don't miss anything important, you should access the Microsoft 365 **Message Center** from time to time to review the notifications posted there (Figure 4-9). If you find something interesting in the notices, you can select **Share** to send the notice to someone else via email and include a personal message to add some context about the change. The email includes the complete text of the update notice. Microsoft added tags to notifications, along with controls to let you filter by tags, the services affected, or the state of the notifications. Tags include "Major Update," "Admin Impact," "New Feature," "Data Privacy," and "User Impact." These tag names are self-explanatory, but they are nonetheless useful to help you filter and sort through the steady river of notifications issued by the service.

In early 2022, Microsoft started showing the affected workloads, and the number of affected users per workload, in some Message Center notifications. This is meant to help you prioritize your handling of those changes.

As described in Chapter 9, you can also have Message Center notifications synchronized to Planner, so you can track every change and make sure that any change which might impact your organization is assigned to the appropriate people for action.

## Major Updates

The **Major updates** tag in Message Center features major changes to a service. Microsoft posts these updates at least 30 days before they come into effect. To decide whether a change falls into this category, Microsoft considers how the change might affect users and the tenant with questions such as:

- Does the change affect **the way people work daily**? For example, the introduction of the Focused Inbox changed the way that people deal with new emails. Changes to meetings, delegations, and sharing and access control are also in the "daily productivity" bucket.
- Does the change affect **how the tenant customizes the service**? For example, if Microsoft changes a theme or web part, it might affect how pages render and overwrite a change made by the tenant.
- Does the change **increase or decrease capacity**? A good example is when Microsoft changed the default storage quota for mailboxes (for enterprise plans) to 100 GB. Another example is the change in how OneDrive for Business manages storage quota.
- Does the change affect **what users see**? Any change that might affect the ability of a help desk to support users or change how users access functionality is a major update. A minor branding change such as the replacement of "Office 365" with "Microsoft 365" for the Admin Center is not in this category but forcing people to use content searches instead of workload-specific searches is. Changing a URL used by a feature also falls into this category.
- Does the change introduce **a new service or application**? Microsoft needs to tell tenants when they launch a new application like To-Do or Teams.
- Does the change **require administrative action**?
- Does the change involve **the storage location for data**? Often, applications launch with storage in the U.S. As the roll-out proceeds across the world, storage moves into other data center regions. Customers need to be told where their data is stored, such as when Teams moved its data services into the U.K. or India data center regions.

Once you see a notice about a major forthcoming change, it's a good idea to search for more information such as the blogs written up by people who have tested the change or already have the change in production in their tenant. Independent information added to Microsoft formal documentation creates a more complete picture of what you can expect to see happen in your tenant.

## Message Center Preferences

Despite Microsoft's efforts to improve the relevance of Message center items, some of the messages appearing in your tenant's Message center might not be directly relevant to you. For instance, if you do not use Office 365 Mobile Device Management, you are probably uninterested in any of those messages. You

can filter the messages you see by clicking the **Preferences** link and selecting the services you want to know about.

Message center preferences include the choice to have Microsoft send you a weekly email digest of new messages. Even if it is not a substitute for keeping your eyes open and looking out for information about developments inside the service, the weekly digest is a practical way to stay updated with Microsoft announcements. By default, Microsoft has historically sent the weekly digest to the global administrator for the tenant and addresses the note to both the primary and backup email addresses for that account. A third address is also available for use as the tenant wishes. You can use this address to send copies of the mail digest to a distribution list, a Microsoft 365 Group, or to a team channel for other people to learn about changes. If you do not wish to receive the mail digest, uncheck the boxes for one or more of the addresses and save your preferences.

Unfortunately, Microsoft changed this default so that admins who customized the digest settings before April 2022 will keep their defaults, but users who are granted an admin role after that date will get the digest for four weeks as a trial.

Users holding specific administrative roles can also receive update notifications. Originally, only global admins could see Message Center messages, but as of now users who hold roles such as Teams administrator or Exchange administrator receive messages (which they can opt-out of) for the workloads they manage.

## Automatic Translation

Update messages are composed in English. Administrators who use non-English languages can have the Microsoft 365 admin center translate messages to another language by selecting their preferred language through the admin center settings (cogwheel). If some difficulty occurs in understanding the translated version of a message, you can choose to view the message in English selecting English from the drop-down list of available languages. This is a dynamic choice that reverts to the language selected in settings when you navigate away from the Message Center.

## Managing User Feedback

Getting useful feedback from users is one of the key challenges of making software. Just like every other software company, Microsoft is eager to collect feedback from its users; it uses that data (which might include direct feedback, crash logs, survey results, or metadata about application usage) to see what people think about its products and services and how they can be improved. It may seem sometimes like the Microsoft 365 product teams aren't listening to the feedback their users generate, but nothing could be further from the truth—it's just that getting good-quality feedback is difficult, and at the scale of the service, it's also difficult to parse and interpret it.

There are several interlocking sets of settings that control what happens to this feedback and who can see it:

- The native Windows, macOS, Android, or iOS crash-dump systems all have controls over whether crash logs from an application failure are automatically shared with the application vendor; in general, I encourage users to enable these settings so that the vendors can identify patterns of crashes caused by OS updates, product bugs, and so on.
- Most of the Microsoft 365 apps have an "[in-product feedback](#)" feature (usually in **Help > Feedback**) that lets users submit feedback directly to Microsoft while using Excel, Outlook, or other apps. This feature set requires a relatively modern client version (see [this link's "Before you begin" section](#) for a full version-requirement table).
- Teams has a "Give feedback" feature in its desktop and web clients, and a similar feature in the mobile versions, that you can control with the `Set-CsTeamsFeedbackPolicy` cmdlet.

The admin center includes a **Product feedback** item under the **Health** section in the left navbar; here you'll see all of the feedback that users have submitted. However, before you see anything, users must submit some feedback... and before they can do so, you must enable that by creating a feedback policy with the Office cloud policy service, as described in Chapter 6. You can specify whether users can submit feedback, whether Microsoft can follow up with users who do so, and whether the feedback may contain log files, content snippets, or other organization-specific metadata that might potentially be sensitive.

Starting in February 2022, Microsoft added a new feedback mechanism that you can use to gather net promoter score (NPS) data from your users. It uses the same policy mechanism as the product feedback system described above but collates the data and uses it to present NPS data that you may find helpful as a guide for your adoption and rollout planning. They've made various other tweaks to this feature (including labeling feedback as positive, neutral, or negative) and there is probably more to come here.

Somewhat confusingly, there's a *completely different* mechanism for anyone to submit general feedback for a range of Microsoft 365 services. The [homegrown feedback portal](#) uses the Dynamics 365 Customer Service module, and the portal's design makes it easy to submit and share ideas with engineering groups.

## Network Connectivity Health Monitoring

It makes sense that Microsoft would know a great deal about network connectivity health at the service level, and they're making a subset of what they know visible in the admin center. The **Network connectivity** link under the **Health** section in the left navigation bar shows you data from client-side telemetry to illustrate the performance and quality of your sites' connectivity to the Microsoft network. These data are collected in three ways:

1. Enable Windows Location Services (WLS) on at least two machines in each location that are running Windows OneDrive clients (provided you have version 19.232 or later). When you do this, all the OneDrive clients in your domain will aggregate to the same metro area—that is, if you have 50 computers in Huntsville, they will appear as a single location at the city level. Microsoft also rounds location information obtained from WLS to the nearest 300-square-meter grid square.
2. Use the Locations tab to define a list of locations and corresponding public IP addresses; Microsoft will group clients whose service traffic enters their network from this public IP as “belonging” to that location.
3. Have users run the [Microsoft 365 network connectivity test](#) manually.

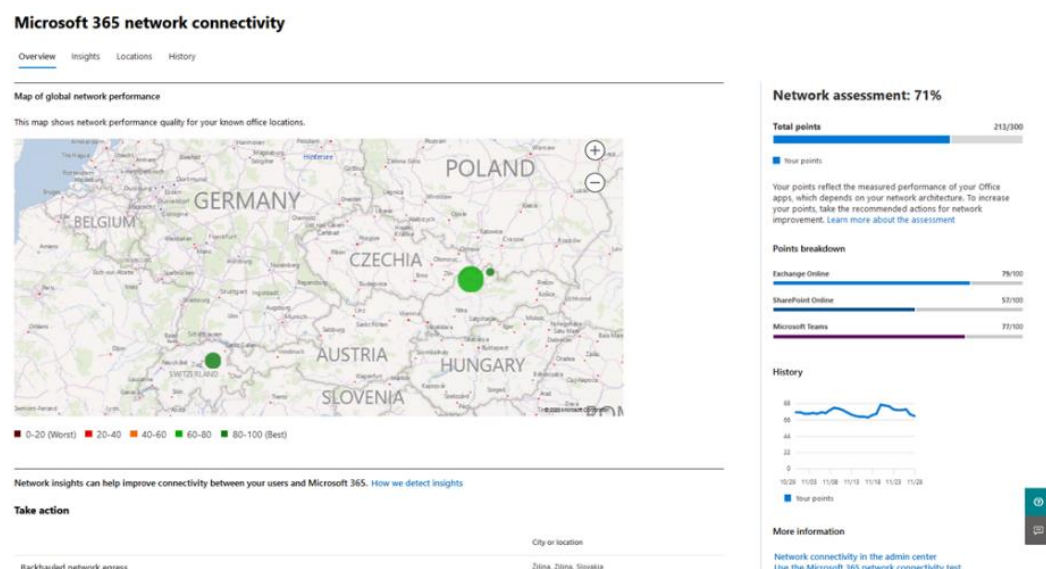


Figure 4-10: The network connectivity view shows you what Microsoft knows about your network quality

The resulting data is presented as a performance map and a series of scores, along with a list of recommendations, as shown in Figure 4-10. Microsoft has plans to continue enhancing this connectivity map by ingesting more telemetry signals from clients and refining the machine learning models that are used to produce insights and recommendations. This is exactly the type of application for which “big data” gathered at a very large scale can provide interesting results; the larger the number of clients that report for any given area, the more resolution Microsoft has for identifying potential network problems.

## Exchange Online Health Monitoring

Exchange Online and Google Gmail together delivered the practical reality of email as a utility service that “just works.” With that said, sometimes Exchange Online *doesn't* “just work,” so demand thrives for solutions to help administrators get early warning, and better insight into, problems that affect email routing and delivery. The huge wealth of telemetry that Microsoft gathers throughout the Exchange Online service is mostly reserved for internal use; it wouldn't do you any good to know the temperature of an individual disk sitting somewhere in an Exchange server pod anyway. However, Microsoft is slowly moving to provide first-party tools to help answer the most common concerns service administrators have: is there a problem? If so, is it *my* problem or *Microsoft's* problem? And what do I do about it?

In February 2021, Microsoft added a public preview of [Exchange Online health monitoring](#) to the service, and that feature has improved over time to [cover other Microsoft 365 services](#). However, you probably can't use it, because you must have at least 5,000 licenses of E3 or E5 licenses from either Office 365 or Microsoft 365 before you can enable the feature. When you do, you'll see health data shown on the **Health > Service health** page. The data collected by the service includes a list of advisories or incidents, plus links to show specific scenario-based data that may help identify problems in one of three categories: problems with the service itself (labeled as “Infrastructure” by Microsoft), problems with third-party network connectivity, and problems with your infrastructure (labeled as “Your org”).

Right now, Exchange Online has the most health monitoring data available. You can see rolling 30-minute counts of the number of users who have logged in (using either basic or Modern Authentication), the number of messages delivered without delay (the time count starts from when the service receives the message), and the number of active users. This last metric deserves a bit more explanation: Microsoft counts an active user as one who has read at least one email using a supported client (all versions of Outlook, plus the native iOS and Android mail clients). A user who sends mail, syncs mobile device folders without reading anything, creates a task item, etc. won't be counted as active. This leads to the unfortunate possibility that at times when you'd normally have low message reading activity, you might find that the health monitoring system doesn't indicate a problem.

For now, there are no real alerting or notification capabilities built into this dashboard. The dashboard will show you incident notifications when selected [mailboxes are nearing their quota](#), when a [Mailbox Replication Service \(MRS\) source mailbox is running slowly](#), or when [outbound messages queues start to build](#). As a place to check when you suspect that there might be a problem, this is a useful tool, but as a means of getting early warning of a newly emerging problem, it will require some refinement to be useful since you won't get a proactive notification

## Health Monitoring for Priority Users

In addition to the overall tenant health monitoring features, you can see health information for [priority accounts](#). You can tag users as VIPs by adding them through the **Users > Active users** view, or by adding them in the **Setup > Organizational knowledge** view. If you'd rather use PowerShell, you can call the `Set-User` cmdlet with the `-Vip` flag, like this:

```
[PS] C:\> Set-User Tony.Redmond -Vip:$true
```

Once you designate an account as a priority user, you'll see some additional health monitoring data, plus the ability to generate alerts when a priority user's account has mail flow problems. You can designate any licensed user as a priority user.

If you have assigned an account a Microsoft Defender for Office 365 Plan 2 license, you can also benefit from a feature Microsoft calls "priority account protection." This is essentially a set of additional quarantine reports and filtering options that let you more easily sift through and deal with quarantined messages sent to your priority users. For example, alerts on accounts that are marked as priority accounts display a unique tag in the alert list.

## Better Network Connectivity with Informed Routing

Microsoft gathers a cornucopia of network performance data across the service that it uses to tune its internal networks and to present the data described in the preceding section. They've been working on an additional use for this data, too: using it as a feedback mechanism to help you tune *your* network. If you're using a supported software-defined wide area network (SD-WAN) solution, a new feature called [informed network routing](#) allows your on-premises SD-WAN equipment to read network data coming *from* the service. The basic idea behind informed network routing is that the service (especially the front doors you connect to) can provide real-time feedback on the performance and reliability of your connections to the service, and your SD-WAN can reconfigure itself as needed to optimize traffic flow to the service. It's a fascinating idea, but as of June 2022 the feature is still in an opt-in preview, and it only works with a single vendor's SD-WAN solution. As Microsoft expands the scope of what this feature can do, it will be interesting to see how much practical value it provides for organizations that are comfortable allowing performance metrics from a single service to drive their network configuration with minimal human intervention.

## Monitoring Systems

As with many Microsoft products, there is a [System Center Management Pack for Microsoft 365](#). Organizations that are running System Center Operations Manager (SCOM) can create alerts and monitoring dashboards for the health of their tenant. Tenants who do not use SCOM implementation can consider a solution based on the various Microsoft Graph APIs available for monitoring and management. You can develop your scripts or software solutions based on the information available through the API or invest in a third-party monitoring solution that can do it for you. Examples of third-party monitoring products include Quest Nova, MessageOps Monitor, Analyzer for Office 365, and 365 Command by Kaseya.

## Service Requests

Service requests are problem reports filed with Microsoft, although you are almost certainly going to be working with third-party support engineers under contract with Microsoft rather than blue-badged Microsoft employees. [Telephone support](#) is also available if urgent problems occur. Because a lot of configuration and other information will need to be given to support engineers, it is logical that calls should be made only by users with administrative access to the tenant.

When you submit a service request, it is routed to a Microsoft support center covering the Microsoft 365 region for the tenant. The support request is examined by a support engineer who will probably call the tenant administrator to discuss the reported problem and to look for added information to help diagnose its cause. Because of the scale of the worldwide service, Microsoft follows a carefully structured approach to gathering and checking facts about service requests. At times, it can be frustrating for a tenant administrator when a support agent asks them what appears to be much the same question several times or to be asked to test components that are known to be failing, but it is all part of the process. Eventually, the problem will either be solved or escalated to second-level support (this can take several days) and, if necessary, to the

product group that supports the workload where the problem exists. Remember, support engineers cannot make changes to code or to how the service works as this kind of intervention can only happen through a process controlled by the product group. Microsoft is rightly cautious about making changes because a small change made for the benefit of a single tenant might have a ripple effect elsewhere within the service that impacts multiple tenants. In general, support engineers can only access tenant data when you allow them to do so. The need to preserve customer confidentiality and privacy is paramount, even if this slows things down at times. You should mentally prepare yourself to be asked to look up and provide data that you might think the service engineers should already have access to.

Naturally, before you add a new service request, it is a good idea to do some troubleshooting of your own as you might discover the answer much faster than through the Microsoft support process. Use your favorite search engine to check for obvious solutions and, if nothing turns up, try asking a question in a community forum such as [Microsoft Q&A](#) (the allegedly-official support site) or the [Microsoft Technical Community](#). Don't be lazy and expect someone else to do the work to solve your problem. Always take the time to investigate and share information in your requests to prove that you have done some troubleshooting for the problem and the results of your investigation. People in support forums generally want to help, but only if you are prepared to help yourself first. Even if the support forums do not help, the information you gather in your investigation will give invaluable background to the Microsoft support engineers who work on the service request.

## Creating a New Service Request

To add a new service request, click **Support** and then **New service request** from the Microsoft 365 admin center. You'll then go through a three-step process.

In the first step, you enter a few search terms, and the admin center will suggest some potentially helpful articles. Microsoft's intent here is to get you to solve your problem, whenever possible, using the search results instead of opening a support case. Microsoft uses a mixture of machine learning and telemetry that they have for your tenant, and searches against their support databases to find potential solutions based on the text you enter. As is typical of Microsoft search technologies, the quality of results here may vary widely. If you don't find any relevant or useful results, you can click the **Contact support** link to go to the second step.

In step two, you give some details about the problem you need help with. You'll need to specify who should be contacted, their email address and phone number, and whether you prefer an email or phone response. You can also specify a preferred language.

When you click the **Contact me** button, your request will be queued for action. The wait time for a response is usually quite short, but it will depend on the number of cases that the support representatives are dealing with. IT professionals love to complain when they have been let down by a poor support experience, so you will always find stories of very long waits for a call back from Microsoft. If the delay for a call back is too great a risk for your organization then you should consider a Premier Support contract to get faster support. You can request a specific callback time, which will help reduce the frustration inherent in playing phone tag with support engineers, and if you specify your time zone, they might even honor it (hopefully avoiding late-night phone calls). However, they have removed the option to request contact via email, which is not a good tradeoff.

Note that when you send a support request, if the phone number or email address you supply isn't listed in your organization profile, Microsoft will send a unique PIN to the registered email addresses that *are* in the support profile. This helps reduce the risk of spoofing but means that, if you're not watching for the PIN, your support case may languish until you've supplied it. Microsoft recommends making sure that customers update their profile email and phone details to avoid this problem.

When you create a support request, you can add attachments such as screenshots or network traces; these, along with the contents of the problem description field, are the only real ways that your assigned support engineer will have to learn about the problem before you talk to them. Be sure to be clear and descriptive.

When Microsoft receives a service request, they assign it a service request identifier such as 10312033. You should give this identifier to Microsoft whenever you communicate with the support team, as it allows them to track the progress of a service request through Microsoft's support and escalation systems. You should also receive a message from the support engineer assigned to the service request. The subject of this message has some tracking information to capture the interaction between support staff and tenant administrators in Microsoft's support databases. To keep a record of the case as it unfolds, reply to the message whenever you have something to add or want to request an update. It is a good idea for you to keep your record of interactions with Microsoft just in case this data is necessary to prove that a problem hasn't received enough or timely attention.

The details that you add to a service request can be extraordinarily useful to the support engineer assigned to handle the case. You should give information such as:

- Detailed steps to reproduce the problem. If the problem surfaces in different ways, include steps for each way that you know to reproduce the issue.
- Is the problem still evident after you sign out and sign back into Microsoft 365? An expired token can cause a failure to connect to a service, so it's important to check that the account with the problem is authenticated.
- Scoping information for the problem's impact. Does it affect one user, all users, or something in between? Are all the affected users in the same geographic region?
- If the problem occurs only in one location, is there something special about how users connect to the service or the Internet from that location?
- Has anything changed recently? For example, have you made any configuration changes that might be related to the problem?
- Did this functionality ever work, or did it stop working at some point?
- Screenshots showing any error messages that are visible to users.
- PowerShell commands and the output from those commands to help support engineers to understand the problem.
- Background information such as the version number of clients (including browsers) that are affected by the issue. Make sure that you use a browser supported by Microsoft 365 and try to replicate the problem on different browsers to narrow the conditions under which the problem appears. For instance, the Edge and Chrome browsers can behave differently from the Brave Browser. Does the problem occur with all browsers or just a specific browser? Can it be reproduced on multiple workstations or just one that is running a certain version/build of an operating system? Does the problem happen if you open a private or incognito browsing session?
- If the problem affects a hybrid component, details of the hybrid connection and other associated components such as how directory synchronization is performed and whether single sign-on is used.

Just like in any support system, the various kinds of problems that can occur need different periods to resolve. Some issues might never be resolved because they need substantial engineering investment that Microsoft considers unjustified or unnecessary. Microsoft should resolve straightforward problems in a day or so, but to be brutally honest, as the service grows, any problem that can't be addressed by a simple troubleshooting script is likely to linger. It's very difficult to attract and retain good support engineers, and Microsoft has struggled with doing so over the last few years. Problems that need engineering intervention, the acquisition of more detailed diagnostic data, or are simply complex will need more time. Microsoft is usually good at keeping tenant administrators up to date with the progress of service requests through



email. Updates are posted to the service request and are visible through the Microsoft 365 admin center. And if things go quiet, you can always email the assigned engineer to ask for an update.

## Gathering Tenant Information

While you work through a service request with Microsoft, you might be asked to give some information about your tenant, usually to allow the support engineer to understand what version of the software the tenant is running. Remember that a tenant can choose to use software released to the Targeted Release or Standard Release rings or a mixture of both. Even within these rings, Microsoft deploys software at different intervals to reach every tenant in all regions. It is impossible to deploy new software to everyone at the same time, so it is important to know what configuration is active within a tenant when you meet a problem.

The Microsoft 365 admin center doesn't provide a built-in way to gather and report tenant configuration, so you must use PowerShell for this purpose. Two cmdlets are especially important. The *Get-OrganizationConfig* cmdlet returns information about Exchange Online and some generic tenant data while *Get-SPOTenant* returns information about SharePoint Online. You will find examples of *Get-OrganizationConfig* in use for different purposes in other chapters, but we will concentrate on its use to report tenant data here. The information reported by the cmdlet changes over time and is difficult to review on-screen. It is usually best to dump the output to a text file and review it with an editor, which will also make it possible to extract information to share with Microsoft support. In the following example, the first command lists several important settings that might be of interest when troubleshooting a support case while the second redirects the output to a text file.

```
[PS] C:\> Get-OrganizationConfig

Name                : office365itpros.onmicrosoft.com
ObjectVersion       : 16500
ReleaseTrack        : FirstRelease
SharePointUrl       : https://office365itpros.sharepoint.com/
MapiHttpEnabled     : False
IsLicensingEnforced : True
IsTenantAccessBlocked : False
IsTenantInGracePeriod : False
RBACConfigurationVersion : 0.1 (15.20.218.12)
AdminDisplayVersion : 0.20 (15.20.218.12)
ServicePlan         : BPOS_S_E15_0

[PS] C:\> Get-OrganizationConfig > Config.txt
```

For instance, settings such as the *RBACConfigurationVersion* and *AdminDisplayVersion* will tell engineers what version of the software runs inside the tenant. The *ReleaseTrack* setting shows if the tenant uses Targeted Release for all or some users, while the *ServicePlan* setting shows the basic plan configured for the tenant.

SharePoint Online does not report as much configuration data for a tenant as Exchange Online does, but the devil might be in the detail when engineers are debugging a problem. Here is what the *Get-SPOTenant* cmdlet reveals:

```
[PS] C:\> Get-SPOTenant

StorageQuota                : 44032
StorageQuotaAllocated       : 20174
ResourceQuota               : 10900
ResourceQuotaAllocated      : 1900
CompatibilityRange          : 15,15
ExternalServicesEnabled     : True
NoAccessRedirectUrl         :
SharingCapability           : ExternalUserSharingOnly
DisplayStartASiteOption     : True
StartASiteFormUrl           :
```

```
ShowEveryoneClaim           : True
ShowAllUsersClaim          : True
OfficeClientADALDisabled   : False
ShowEveryoneExceptExternalUsersClaim : True
SearchResolveExactEmailOrUPN : False
RequireAcceptingAccountMatchInvitedAccount : False
ProvisionSharedWithEveryoneFolder : False
SignInAccelerationDomain    :
```

Sometimes, you might be asked for the tenant identifier. This is a unique value for the tenant used in different places by Microsoft 365 and Azure AD. One way to retrieve the identifier is to use the *Get-MgOrganization* cmdlets. This example uses *Get-MgOrganization*:

```
[PS] C:\> Get-MgOrganization | Format-List Id, DisplayName

ObjectId                DisplayName
-----
a662313f-14fc-43a2-9b7a-d2e27f4f3478 Office 365 for IT Pros
```

The *ObjectId* reported by the cmdlet is the tenant identifier. This is important information to have because it's the unique value used by Microsoft 365 to identify the tenant. You'll need to know this information if you ever want to sign up for a Microsoft beta program.

Alternatively, input your tenant domain name into this [website](#) to retrieve the identifier. The site uses information published on the internet to allow OAuth 2.0 sign-ins to function to report the tenant identifier.

## Customer Lockbox

When Microsoft support personnel are working on an issue for you, they may need access to some user data to resolve the problem. These situations are rare, as most of the support operations performed by Microsoft are automated. Where possible, any support tasks performed by support engineers are isolated from customer data. However, there will always be some cases, such as when Microsoft support is trying to help a customer with problems with mailbox contents, that access to the user data is necessary.

Customer Lockbox is an extension of Microsoft's Lockbox system for managing "just in time" access to the service infrastructure by support staff. Multiple levels of authorization are required before support staff can gain access. Upon approval, the support engineers receive access limited in scope and duration to the minimum needed for the task. Lockbox also includes comprehensive audit logging of any activities that support staff perform.

Under normal circumstances, support engineers receive approval from a Microsoft manager to access customer data. When a tenant enables the customer lockbox feature, an added approval by the customer is necessary before Microsoft support can access user data (this requirement does not cover access to system data such as logs). Microsoft sends the Customer Lockbox request as an email notification which administrators or users with the customer lockbox approver role can approve or deny in the **Support** section of the Microsoft 365 admin center. Access requests have a time limit (12 hours by default), and a scheduled cleanup task removes the support personnel's access automatically upon resolution of the issue or when the time expires.

Actions performed by Microsoft support or by the automated systems used in the service are captured in the audit log for the tenant. Like any other audit log records, they are accessible using PowerShell or APIs to allow third-party security monitoring systems to extract and include data in reports and dashboards.

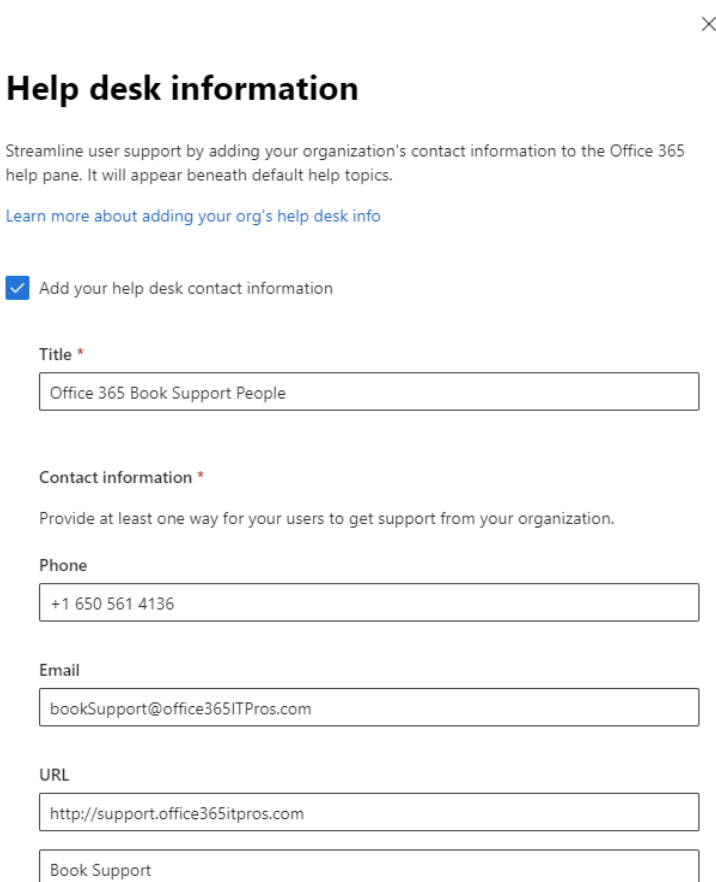
To enable Lockbox for a user, the user will have to have an Enterprise E5 license *or* the Advanced Compliance license.

# Customizing the Microsoft 365 Interface

One of the big selling points behind Microsoft 365 is that it has a consistent and familiar interface, both for users and administrators. Microsoft doesn't want any of us going overboard with interface customizations; they have included a few features to help tailor the Microsoft 365 client interfaces to the needs of specific organizations though.

## Creating Custom Help

Tenants can create a custom help card to be displayed alongside the standard help text revealed when a user clicks the question mark (?) icon in a browser while accessing an application. The idea is that you can provide users with tenant-specific information to allow them to seek help if they have a problem. To create a custom help desk card, open the Microsoft 365 admin center and go to **Settings** -> **Org settings** and then click the **Organization profile** pivot. Click the edit icon beside **Help desk information**.



**Help desk information** ×

Streamline user support by adding your organization's contact information to the Office 365 help pane. It will appear beneath default help topics.

[Learn more about adding your org's help desk info](#)

Add your help desk contact information

**Title \***

Office 365 Book Support People

**Contact information \***

Provide at least one way for your users to get support from your organization.

**Phone**

+1 650 561 4136

**Email**

bookSupport@office365ITPros.com

**URL**

http://support.office365itpros.com

Book Support

Figure 4-11: Populating custom help desk information

Office 365 displays the input screen shown in Figure 4-11 and you can enter:

- A title.
- Phone number details.
- Email address.
- Web page.

The next time a user signs in using a web browser, they see the custom help desk card when they access help.

Office 365 supports a single custom help card per tenant. You can't, for instance, have location-specific help information presented to users. For this reason, if you manage a tenant that has users distributed in multiple locations or countries, you should make sure that the information provided in the custom help card is valuable for all users. For instance, it does not make sense to direct a user in Japan to a help desk in the U.S. if that help desk is unable to speak Japanese or is only available during U.S. working hours.

## Custom Tiles

When users sign into the Microsoft 365 portal, or when they click the waffle menu (the App Launcher) from a browser application such as OWA, a list of tiles for the applications and services to which they have access appears. Users can select apps to appear in the launcher from their My Apps page, which lists all the apps known to the tenant. The apps that appear in My Apps come from several sources:

- Standard apps licensed from Microsoft as part of the plan assigned to the user account. For example, OneDrive for Business and Stream.
- Apps that are developed for or by the tenant. Administrators then configure the app to assign it to individual users. These might be line-of-business applications that are registered in Azure AD or applications written by ISVs and stored in the Azure application marketplace, such as DocuSign, DropBox, or Citrix GoToMeeting.
- Apps created with the Office 365 API Tools for Visual Studio that allow users to sign in using their Microsoft 365 credentials. These apps automatically show up on the My Apps page.
- Web apps from AppSource that support single sign-on, such as the [Starbucks app](#).

Custom tiles are available to accounts that have an Exchange Online mailbox and are the simplest and easiest way to customize the App Launcher. Essentially, a custom tile serves as a pointer to a web page. The URL defined as part of a custom tile can bring a user to a SharePoint site, Microsoft 365 Group document library, external website, or any other page accessible through a URL. The custom tiles that are defined for a tenant are listed in the My Apps page displayed to users and can be pinned to the App Launcher from there.

To add a custom tile, login to the Admin Center and navigate to **Settings** -> **Org settings**, and then click the **Organization profile** tab, then choose **Custom app launcher tiles**. Click + and you'll be able to add a custom tile by specifying the name of the tile, the URL that the tile should invoke when clicked, a description (purely for administrative purposes), and a URL for an image for the tile icon. The image must be in JPEG format and should be sized at 50 x 50 pixels. It is easiest if the image file is stored in a SharePoint library as this means that it is accessible to users. In addition, make sure that all users have at least read-only access to the tile image. If users don't have access to the image file, a blank space is shown for the tile in the My Apps library.

**Hiding launcher tiles:** Many administrators want to hide specific applications from users by removing the application tiles in the app launcher for all (or some) users in the tenant. Unfortunately, there's no way to do this. The list of tiles the user sees is dynamically built based on the licenses you've granted them. Users may add and remove tiles themselves but there's no way for you as a tenant administrator to do this in bulk. If you don't want users to use an application, instead of hiding it in the launcher and hoping they don't stumble across it, you'll have to remove the license (or deactivate the application) for those users. You can also hide [all custom tiles](#) or set up [application collections](#) that have predefined sets of tiles for users to see.

## Managing Themes

Microsoft has long offered the ability to apply themes or skins to various applications. This ability started with Outlook Web Access in Exchange 2000 or so, although creating themes for on-premises OWA was always sort of a hit-or-miss proposition and wasn't generally supported. With the move towards cloud-based services, Microsoft has provided a more robust theming mechanism that allows you, or your users, to select

themes that apply to the Microsoft 365 admin center and various applications. There's a similar facility in the [Microsoft 365 Apps](#) as well. The drawback? You can't customize these themes very much.

If you open the **Organization profile** settings for your tenant in the Microsoft 365 admin center (under Org settings), you can select **Custom themes** to edit items like:

- The corporate logo for display on Microsoft 365 web pages. The selected logo file must be sized precisely at 200 x 30 pixels and be less than 10 KB. Files in SVG format are best because these are supported by mobile apps. You can also associate a link that users go to when they click the logo.
- The navigation bar color. This should match the color of the background image.
- Text and icon color. This should highlight text and icons when overlaid on the navigation bar.

Of the other available settings, the most important one is probably the **Prevent users from overriding their theme** checkbox. If for some incomprehensible reason you need to exert a kindergarten level of control over your users' choices of theme, this is where you'd do it.

Speaking of themes: Microsoft occasionally releases new ones for users. Currently, there are more than 50 themes available to users under the gear settings icon. Some themes are whimsical, while others are merely solid-color sets.

You can create up to five separate themes of your own and then [assign them to members of specific Microsoft 365 Groups](#). Microsoft originally referred to this as "conglomerate branding" but now seems to prefer the term "group branding." Each of the themes can have a unique experience, including separate company logos (and a logo variant for dark mode).

## Reporting

Microsoft 365 includes a few different sets of reports that may prove useful. The Microsoft 365 Admin center contains a reporting section with an activity dashboard and a handful of reports that give organizations a view into the adoption level for the different services they buy with their subscriptions. For example, the email activity report might show a healthy amount of usage, justifying paying for Exchange Online mailboxes, while the SharePoint or OneDrive for Business reports might show a different level of usage that prompts the organization to either reconsider the licensing of that feature or embark on an adoption project to encourage more use. Similarly, the Activations report lets an organization analyze the consumption of licenses bought for their end-users. For example, if you assign international calling licenses to users who do not use them, then you can reassign the licenses to other users who need them rather than buying more licenses.

In addition to these canned reports, there are two more detailed and more flexible reporting tools that you should know about. The first is the Productivity Score toolset, which tells you what users are doing with the service. The second is the Secure Score toolset, which is part of Microsoft 365 Defender and gives you a synoptic view of your tenant's security.

See Chapter 20 for more information about reporting and auditing, including the Microsoft 365 Usage Analytics pack for Power BI and some third-party reporting alternatives.

## Microsoft Productivity Score

The Productivity Score section of the Admin center purports to solve a long-running problem: how do you know how effective people are at their jobs? It pulls data from individual users' activity in the service, data from their calendars, and telemetry from various applications, and synthesizes that data into a series of measures that allegedly show how productive they're being. Figure 4-12 shows an example Productivity Score dashboard; a glance will show that there's an aggregate score but that the aggregate is a composite of several individual measures.

## Productivity Score

Productivity Score provides insights into your organization's digital transformation journey through its use of Microsoft 365 and the technology experiences that support it. Your organization's score reflects people and technology experience measurements and can be compared to benchmarks from organizations similar to yours.

### People experiences

#### Communication

Organizations that use a variety of ways to communicate support different work styles, needs, and preferences.

0% of the people in your org use two or more modes to communicate.



#### Meetings

When people use online meeting tools effectively, they can save up to 104 minutes a week.

0% of meetings in your org follow one or more meeting best practices.



#### Content collaboration

When people collaborate with online files, they can save up to 100 minutes a week.

15% of the people in your org collaborate with online Microsoft 365 files.



#### Teamwork

When people share information and collaborate in a shared workspace, they can save up to 4 hours a week.


10% of people in your org are contributing to shared workspaces.



#### Mobility

Access to email and files, and communication with teammates on any device help people get work done on their schedule.

9% of the people in your org use apps across multiple platforms.



### Your organization's score: 31%

Total score: 250/800 points

■ Your org ▼ Peer benchmark

Your organization's Productivity Score is the total of its people experiences and technology experiences scores, which are each comprised of several categories of data. Scores are not provided at the individual user level.

Score components 250/800 points

People experiences: 24/500

Technology experiences: 226/300

▼ Peer benchmark

[Learn about how your org's score is calculated](#)

### Your organization's score history



### More information

[Learn more about Productivity Score](#)

[Learn more about privacy in Productivity Score](#)

Figure 4-12: Viewing the Productivity Score for a tenant

Clicking on one of the scorecards in the overview report shows a more detailed analysis of that specific scorecard. For example, Figure 4-13 shows the details used to calculate the content collaboration scores shown in the overview.

While the idea of an overall score measuring the productivity of an organization seems tantalizing at first thought, there are some problems with this approach. First, it immediately calls to mind the mostly-obsolete discipline of [scientific management](#), with its rigid emphasis on measuring every activity of every worker and continuously iterating to remove all inefficiencies of process or execution. There are good reasons why most organizations don't ruthlessly try to optimize every minute of workers' work times (not least of which is the importance of good labor relations in industries or countries where workers' unions are common).

Second, it's fair to question whether the data Microsoft presents tells you anything useful or actionable. Is it valuable to know, for example, what percentage of meetings have at least one user with video turned on? Perhaps. There doesn't seem to be any meaningful research showing that enabling video makes meetings X% more productive, so it's hard to quantify the impact of knowing e.g., that 20% of an organization's meetings have one person with video turned off. And if having a meeting where one person turns their video feed on is good, what about meetings where two people enable video? Are those meetings twice as productive?

Third, Microsoft can necessarily only gather metrics from things the service can see. If your users use non-Microsoft tools (e.g., Zoom, Dropbox, Google Meet, etc.), then Microsoft won't know what they're doing and cannot, therefore, try to estimate their productivity. And of course, the service knows nothing about meeting participants who are physically present in hybrid meetings (something that will become increasingly common as Microsoft Teams Room systems become more widely used post-pandemic).

## Content collaboration

We measure the number of people who create, read, and collaborate (edit and share) online for this part of your score. When people collaborate with online files, each person saves an average of 100 minutes, or almost 2 hours, per week. [See the evidence](#)

On April 22, 2021, we changed how the collaborators metric is calculated. As a result, you may see a reduction in your score. [Learn more about the scoring change](#)

### 15% of people in your org collaborate with online Office files


When people create and read files online, they are more likely to collaborate online as well. We define content collaboration as one person creating and sharing an Office file, and then at least one other person reading it. This data contributes to your overall productivity score. [How we calculate your organization's score](#)



Peer benchmark

[View content collaboration resources](#)

### Number of readers, creators, and collaborators over time



### Explore how your org collaborates

#### 30% of people who use Office create files in OneDrive or SharePoint

Creating files in OneDrive or SharePoint means they're backed up, available from other devices, and set up for real-time collaboration.

##### People creating files, by location



[View related content](#)

#### 100% of people share files as an email attachment

Sharing a link to a file in the cloud instead of attaching a copy in email makes sharing more secure and allows users to collaborate in real time.

##### People sharing files in email, by type



[View related content](#)

#### 0% of people share content externally

Customize SharePoint's external sharing settings to help people collaborate with external partners or people in your organization who have different licenses.

##### People sharing content




[View related content](#)

#### 2% of people collaborate on 4 or more Office files

Invite people to learn about saving and sharing files in the cloud, co-authoring in real time, and collaborating with @mentions.

##### People collaborating, by number of shared files



[View related content](#)

Figure 4-13: Details of the data used to calculate the content collaboration score in the dashboard

Fourth, the individual metrics that Microsoft gathers in the “People in your organization” section are questionable at best. For example, on the meetings detail page, for each user you can see the date of their last meeting, how many meetings they attended, how many hours those meetings took, how many meetings had video, and how many meetings had screen sharing enabled. (There are a few additional columns as well). The first two users listed in my organization have, respectively, “meetings attended” values of 39 and 81. Does that mean the second user is twice as productive as the first one? Or maybe half as productive? The idea that an individual or organization’s productivity can be measured by such a low-resolution metric brings to mind physicist Wolfgang Pauli’s lament that a theory “[is not even wrong](#).”

Fifth, Microsoft sometimes moves the goalposts—for example, as detailed in MC251864, they changed the way that the “Content collaboration” metrics are calculated in late April 2021, and they are free to change any or all the metric calculations in the future. On one hand, it’s good that they evolve the metrics as they come to understand whether a specific metric can lead to a measurable insight. The August 2021 retirement of the “embed link in email content” insight is a good example of removing something that didn’t provide any useful insight. On the other hand, changing the metrics and insights available means your ability to track trends in the productivity score over time will be affected.

Besides all of these points, there are legitimate privacy concerns in some organizations, and almost immediately after Microsoft released the Productivity Score feature, they were broadly roasted in the tech media for it, so they [backtracked and modified the feature](#) so that it no longer shows user-specific data. This doesn’t change the fact that most of the data about what users are doing remains available to administrators; you can still see how many meetings people have gone to and so on.

With all that said, there are certainly some useful data items highlighted in the Productivity Score report. The “Technology experiences” section, for example, summarizes data about device startup time, network connectivity, and Microsoft 365 application health that are useful in understanding what tools people are using and how well they work. This data is both useful and actionable in a way that much of the more

collaboration-focused data isn't. As Microsoft continues to refine the utility of the metrics and the associated analysis and recommendations, we'll have to see if this feature becomes more valuable.

## Microsoft Secure Score

Microsoft acknowledges that it can be difficult for an administrator to understand how to best secure a tenant. Many places exist in administrative consoles where you can tweak settings that affect how things work. In addition, a multitude of data exists within applications that administrators should check on an ongoing basis. Therefore, it makes sense to measure a tenant against a set of predetermined standards and score the tenant based on the actions taken to increase security. At the same time, Microsoft 365 can flag outstanding actions to the administrator, who then decides whether to implement the action and so increase their tenant score. This feature was called Microsoft Secure Score, and, after a few false starts, that seems to be the name used throughout Microsoft 365 now for this functionality. The Secure Score dashboard allows administrators to:

- Track the progress of their score over time.
- Understand the actions that contribute to the current tenant score.
- Understand how they can improve their score by completing various actions.

Secure Score now lives in the Microsoft 365 Defender portal. Here's how it works.

### Microsoft Secure Score

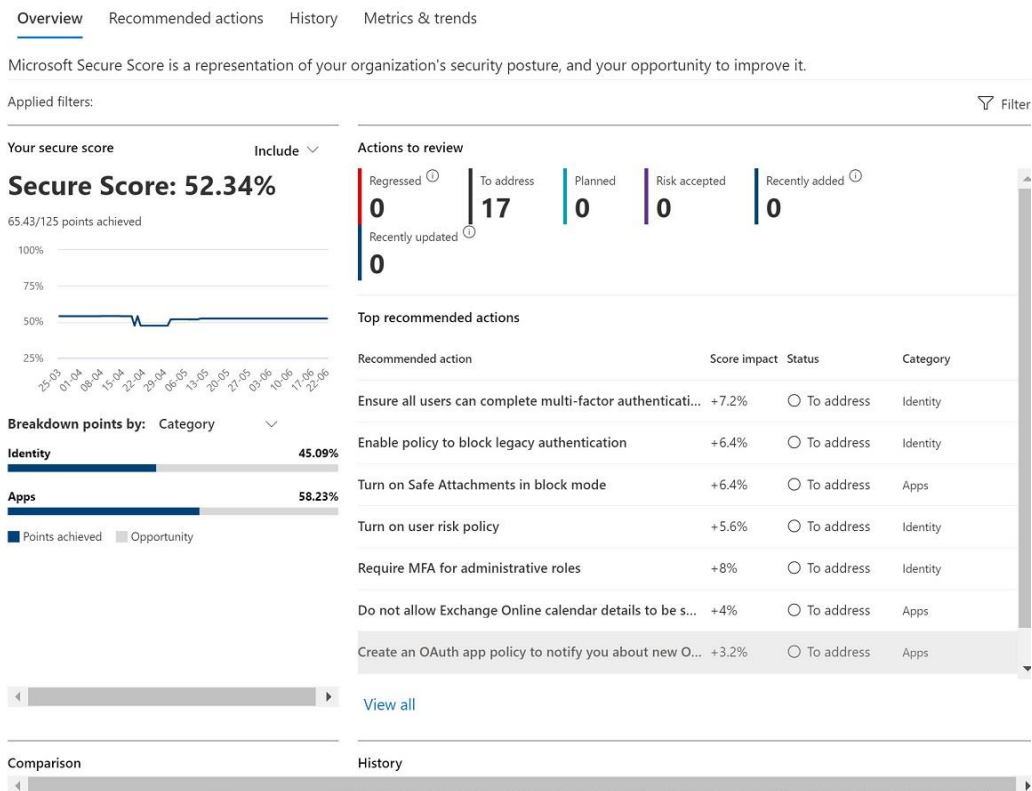


Figure 4-14: Viewing the Secure Score for a tenant

When you go to <https://security.microsoft.com/securescore> or click the **Secure score** link in the left navbar, you'll see a dashboard like the one shown in Figure 4-14. This version represents a significant change since its introduction in that the score is now shown as a percentage instead of a raw number of points, and it is gradually incorporating data from more settings and workloads across the Microsoft 365 platform. This dashboard summarizes your overall security score, calculated based on inputs drawn from multiple sources in the tenant including Azure AD, Enterprise Management + Security (EMS), and various Defender services, in



whichever combination you have them. As you enable more security features, your score goes up. The dashboard includes a list of recommended actions that, if taken, will increase your score. Keep in mind that some security features will increase your score but will irritate users or make it harder for them to do their jobs.

Over time, Microsoft gradually adds both new measurements and new recommendations to Secure Score, so your score will fluctuate even if you don't change anything.

The "Include" dropdown allows you to choose different metrics to measure your score against, and the filter icon lets you choose categories of score items that will, or will not, be included when the score's calculated, while the other areas of the dashboard show actions that you may want or need to take to improve your overall score. The pivot at the top of the dashboard lets you see specific actions that Microsoft recommends taking to boost your score, as well as a history of what your score has looked like over the last 90 days.

Suppose you configure Microsoft Information Protection to allow tenant users to protect confidential content; that adds five percent to your score. Even better, if users store documents in OneDrive for Business, adding AIP is worth ten percent. Although you can argue that OneDrive for Business is a more secure location for documents than a local hard drive or a network file share, assigning ten points to this measurement seems like more of an encouragement to do better. Other controls are easier to understand and more fundamental in terms of security. For instance, administrator accounts should use multi-factor authentication and there should be less than five global administrators within a tenant.

Secure Score combines percentages awarded for the measured aspects into a tenant score. The "Comparison trend" chart shown under the **Metrics & trends** pivot widget in the dashboard shows your organization's score compared to other organizations with a similar number of licensed seats and other organizations in the same industry you're in.

In addition to the base Secure Score mechanism, Microsoft also calculates a [separate identity security score](#) for Azure AD. Here's what Microsoft's [Chris Hallum said](#) about the apparent disparity between the identity score and the "real" Secure Score:

*"The vision for Microsoft Secure Score is that it will be the centralized user experience for all security-related points and Improvement Actions across Microsoft 365 and Azure workloads. Individual products can include a secure score experience scoped to their workload however they must align to the Microsoft Secure Score design patterns and branding. They must also forward their score and improvement action data to Microsoft Secure Score so that it can provide the end-to-end super set the view for an organization's security posture."*

Reviewing your identity score data, along with the list of [best practices recommended for Azure AD administrators](#), and then making appropriate changes can quickly boost your Secure Score overall *and* for identity management. At Ignite 2018, Microsoft said that organizations that follow the Secure Score recommendations are 30 times less likely to be breached, although they have been suspiciously quiet about updates to this metric since then.

When you review individual actions, you'll see a comprehensive page showing lots of information about the action including a description of the proposed change, what state the service thinks you're in now, what the user impact of the proposed change is, and any historical information about changes to this setting in the past. You can mark individual actions with an action plan that tells the system what you plan to do about the risk; Microsoft also helpfully includes a list of implementation steps, although they have some room to improve the display of HTML content there.

It's also the case that the Secure Score mechanism may fail to notice and score an action you're already taking or may recommend actions whose true impact can't be scored. For example, the "Review permissions & block risky OAuth applications connected to your corporate environment" score item promises 15 percent

if you use Defender for Cloud Apps to “block access to a risky OAuth app,” but the precise definition of “risky” is missing, so presumably you can earn 15 percent by blocking *any* OAuth app.

Interestingly, Microsoft 365 keeps the score data for only 3 months, so if you want to see longer-term trends, you’ll have to track the scores yourself. You can [access Secure Score data through Graph](#) to facilitate this.

The engineering team responsible for Secure Score constantly reviews threats and operational processes to ensure accuracy and relevance. Microsoft assesses feedback from the tenants as they analyze tenant scores to understand any gaps and weaknesses that might exist and fix the issues. The score for a tenant is likely to change over time as Microsoft adjusts its scoring scheme and measurements. Administrators should review their tenant’s Secure Score regularly to ensure that they leave no gaping holes for attackers to exploit.

## Backing Up Office 365

Do you need to back up your Office 365 data? This is a complex question. Interestingly, no one questions the value of backup for on-premises data. Why are things different in the cloud?

On one hand, many experienced administrators believe that it isn’t strictly necessary to take backups of cloud-based data. In fact, Microsoft only takes limited backups for SharePoint Online. They don’t take backups of Exchange Online, Teams, Planner, Yammer, or other applications. Microsoft doesn’t provide APIs specifically designed for backups for their cloud applications. Backup vendors must resort to using protocols or APIs not intended to stream large quantities of data to copy data across the Internet to another data center (whether in Azure or another cloud provider). In addition, cloud applications are often interconnected in ways that don’t exist on-premises. The result of this interconnection is that some workloads are much easier to back up and restore. Backing up the documents from a SharePoint site is relatively straightforward, but restoring them in such a way that Microsoft 365 Groups and Teams work properly is a different challenge.

On the other hand, many organizations believe that backups are a good thing because they want to have a method of restoring user and configuration data should the need arise. The need might arise from what I’ve started to call “the four Ms”:

- Malicious changes, such as those made by ransomware or a disgruntled employee.
- Mistakes, such as the infamous mass deletion of Teams chat messages caused by an administrator who misconfigured KPMG’s retention policies.
- Mishaps at the cloud service provider—a more polite way to say “data loss caused by Microsoft”
- Migrations often benefit from having a comprehensive backup of the “before” state in case problems occur during the migration.

As further ammunition, some companies cite [Microsoft’s Shared Responsibility model for cloud services](#) and point to its assertion that customers are always responsible for their data. In other words, Microsoft takes no responsibility for protecting data and it’s up to customers to ensure that they can recover. This is true in general terms but ignores the fact that many Microsoft 365 apps like Exchange Online and SharePoint Online include data protection and retention capabilities to mitigate against data loss.

The writing team for this book is not against backups. Our perspective is that the decision to use backups is one that each company must take after careful consideration of how Microsoft 365 works, the likelihood that problems will arise, the cost and extra complexity of a backup solution, and how quickly data can be retrieved from backup and the operational steps necessary to retrieve data, and so on. But the most important thing is to understand the full spectrum of technology available to a tenant and how to use it to mitigate some or all the reasons why people think they need backups. Following that review, if you still consider that backups are important, it’s a good and rational decision.

## What If You Do Nothing?

Before deciding to use a third-party backup solution, we should understand how out-of-the-box features can reduce the risk of data loss and where gaps might exist. You can think of this as the “do nothing” question: what level of data protection do you get if you don’t buy or configure anything and just rely on the data protection measures in the service?

To answer this question, it’s often useful to look at what Microsoft itself does. When Exchange first introduced the concept of “native data protection,” with multiple copies of each mailbox database substituting for keeping backups, many administrators were horrified, but in practice, this approach has worked very, very well when implemented properly—so taking guidance from the team that builds and runs the service is a valuable way to start.

As an example of a potential gap, many backup vendors point to the fact that a disaffected administrator or user might remove data in the period before they leave the company; they claim that backups offer a robust solution to this potential problem. This is true, but if you have the right licenses (see the [Information Governance section](#) of this page), retention policies can mitigate the problem by automating the process of keeping data for extended periods. Retention policies cover information in:

- Exchange mailboxes and public folders.
- SharePoint Online sites and OneDrive for Business accounts, including video recordings of Teams meetings and all Stream content.
- Teams channel (regular, private, and shared) and chats (personal and group).

For greater protection, preservation locks can secure retention policies and stop administrators from being able to change retention settings. Putting all mailboxes on litigation hold will stop their removal if a rogue administrator deletes some user accounts. And you can lock Exchange Online down further by using [Privileged Access Management](#) to limit what administrators can do to specific time-limited operations.

In fairness, though, these measures protect you against data loss but they don’t deliver some of the other benefits of backup—for example, maintaining a physically separate copy of your data isn’t something that retention policies can help with.

It’s also fair to point out that some Microsoft data protection features such as auto-label policies that find and apply retention labels to sensitive data need extra investment because of their licensing requirements. This element should be kept in mind as you consider the cost of a third-party backup because it might be cheaper to buy the backup instead of upgrading your user licenses, especially if the tenant supports many users with low-cost licenses (like Office 365 E1 or A1, or Microsoft 365 F3). On the other hand, if you need some extra features that are only available in a higher-priced plan, you might decide to use the money that would otherwise be spent on a backup service to take advantage of the high-end protection features bundled in that plan.

## Backup Considerations

When we discuss the service from a backup perspective, the following aspects should be considered to decide what data should and can be backed up. Knowing what data to backup and why the backup is necessary will drive the choice of the backup technology to use. You should consider:

- Backup of the **base storage workloads**: SharePoint Online, Exchange Online, OneDrive for Business, and the configuration and user information held in Azure AD.
- Backup of **applications that use Azure** for all or part of their storage: For example, how do you handle the text and graphics for Teams messages stored in Azure Cosmos DB or the tasks and plans used by Planner?

- Backup of **applications that use multiple components** such as Teams, Planner, and Groups. The lack of backup solutions capable of dealing with workloads outside Exchange and SharePoint is currently the biggest challenge facing those who consider using third-party backups.

In addition, you should consider:

- **How are backups processed?** A backup product might be able to process the volume of data generated by a small tenant and struggle to process the volume created by a large tenant. The backup application must move data from Microsoft's data centers to the backup location, which might mean that the data must travel via the internet, so network considerations and the ability of the backup vendor to process inbound data come into play. On the other hand, if the backup location uses Azure, the data might stay within the Microsoft data center network and the transfer is faster and easier, but you may have to pay ingress and egress fees to move the data.
- **What APIs are used for backups?** Microsoft 365 includes many APIs that can interact with mail, documents, tasks, groups, and so on. However, not every API can stream large volumes of data to a backup destination, so some testing should occur to ensure that a backup application can handle the quantity of data produced by a tenant, especially at peak load.
- **What data are backed up?** Some backup vendors have ported their on-premises products to work in the cloud. Applications like Teams and Planner don't have on-premises equivalents, so on-premises backup products can't deal with their data. Thus, you might select a product that can copy Exchange Online mailboxes and SharePoint Online documents but ignores everything else. You'll be much better off looking for cloud-native products that are explicitly designed to work on Microsoft 365.
- **How often is the data backed up?** A daily backup might be enough to deal with small tenants, but constant and ongoing backups (trickle mode) might be necessary to process the quantity of data produced by large tenants.
- **Where is the backup data stored?** Most backup vendors propose using their own cloud data center to hold backup data. This approach is perfectly acceptable if it meets the customer's data at rest and data sovereignty requirements. Unsurprisingly, few backup vendors can aspire to the same widespread distribution of data centers that Microsoft has, but a growing number of backup vendors use Azure or Amazon Web Services bulk storage as a backup target. Some vendors claim that having backups in Azure is preferable (for speed and security) because data travels across the Microsoft data center network from the service to the backup location while transfer to Amazon involves an internet connection.
- **How accessible is the backup?** If the backup is in the cloud, what SLA does the vendor give about its availability to restore? Is the SLA limited in terms of the amount of data, number of mailboxes or sites, or any other factor?
- **How easily and rapidly can the data be restored?** Having a backup copy of data is one thing; being able to restore it is another. How quickly can data be brought online within an application from the backup copy? Can the data be merged with live data or will it overwrite what's there (for example, does a complete document library need to be restored). Another factor to consider is how much effort is needed on the part of the backup vendor and tenant administrators to restore data. The ideal situation is to have a restore process that automatically connects to the backup source and inserts the data into the target repository in such a way that it is immediately usable without further administrator intervention.
- **Can data be restored in context?** Restoring a single mailbox or single document library is straightforward. It is much more complex to rebuild a compound entity like a Group, plan, or Team as it was at a point in time. This is because those entities depend on multiple components, each of which must be restored in such a way that the links between the different components are preserved and accurate. Data that cannot be restored in context might still be valuable, but it will be raw and

need manipulation before it is fully usable again. Another interesting question is posed by protected content (encrypted email, documents, and other items) as the backup solution must be able to backup and restore this content too. It's also wise to ask about the granularity of restored data to ensure that it is possible to restore something like a single document into a SharePoint Online document library.

- **How much does the backup cost per user per month?** What basis is used for this calculation (per site, mailbox, licensed user, etc.). How does the cost vary as additional applications are included in the mix?
- **Does the backup vendor comply with any industry or regulatory standards that might apply to your tenant?** Such regulations include the European Union General Data Protection Regulation (GDPR), the U.S. Health Insurance Portability and Accountability Act (HIPAA), Service Organization Control (SOC), and the U.S. Federal Risk and Authorization Management Program (FedRAMP).

In short, backing up cloud application data is not simply a matter of stretching the application-centric technology and techniques used to process on-premises data. The complexity and interconnectivity of Microsoft 365 make it a radically different environment that demands a different approach to backup.

## Backup for Exchange Online

Microsoft is deadly serious when it says that native data protection is the best approach for Exchange. This means that four copies of each database exist, including a lagged copy, and that data is spread across at least two data centers to ensure resilience. Users must recover deleted items if they make a mistake and if they don't do this before the default 14-day retention period (extendible to 30 days), they won't be able to retrieve the data.

Given that hundreds of thousands of mailbox servers run inside Exchange Online, all deployed in Database Availability Groups, the use of native data protection is understandable when you consider how many pieces of backup media might be used to copy mailbox data (according to Microsoft, some 1.1 exabyte of mailbox data was in use in September 2018), not to mention the enormous cost of managing and validating the backups and handling and storing the backup media.

Third-party software companies offer online cloud-based backups that can extract data from Exchange Online and send the data to their repositories from where the data can be retrieved if necessary. However, the cost of these services might not be justified for the benefit gained, especially if data is protected by features such as litigation holds. Other issues to consider include if the backup product can deal with expandable archives and encrypted content. Being able to stream data out of Exchange Online is easy (usually, backups use Exchange Web Services to connect to mailboxes and read their contents); it can be much more challenging to restore data in such a way that it is usable. For instance, if the mailbox holds email encrypted with Microsoft Information Protection sensitivity labels or S/MIME, how will the user be able to access this content in a restored mailbox?

Don't use backup utilities that extract mailbox data and store it in PSTs. This is a terrible idea from a performance, logistical, data protection, and compliance standpoint.

## Backup for SharePoint Online

Unlike Exchange Online, SharePoint Online doesn't have native data protection built into the product. Accordingly, Microsoft takes backups for SharePoint Online site collections (these backups cover OneDrive for Business too). Here is what happens:

- Microsoft backs up the content in site collections every 12 hours (approximately). The backups are kept for 14 days.

- Tenant administrators have no control over backups or restores. If a restore is necessary, it can only be started by contacting Microsoft support. Microsoft will ask you to say the best time for the restore (i.e. a time when the required information was known to exist in the site collection). Determining that time can be quite a challenge!
- Microsoft cannot restore a single item, folder, document, list, or library. A full site restore is the only option. The restore is targeted at the URL for the site collection and will therefore overwrite whatever data currently exists in the site collection. If you need to preserve data in a site collection that is to be restored, you must copy or move all the information that exists in the collection before the restore is done and then readjust the content afterward as needed.

Before running to ask Microsoft support to restore a site collection, it's important to understand that SharePoint Online allows information to be recovered in three other ways:

1. The **Restore this library** option (in the settings menu for document libraries) allows site owners and administrators to recover files within a document library to a point in time within the last 30 days (OneDrive for Business has a similar feature). Restore this library is intended to deal with scenarios such as a mass deletion of files by a disgruntled employee or the infection of files through malware. If you know when an incident happened, you can restore the library to a point just before the incident occurred. Restore this library depends on [versioning](#), the ability to store different versions of files created over time. The versions can be created through the auto-save functionality built into the Microsoft 365 desktop and online apps or the files can be saved explicitly. If you don't enable versioning for a document library, you might not be able to roll back to previous versions of files. If only one version of a file exists and it's the one that is compromised, SharePoint won't be able to recover it.
2. Deleted documents and other items are held in the [SharePoint recycle bin for up to 93 days](#). SharePoint's recycle bin is arranged in two parts. The site recycle bin is available to users and is sometimes called the first phase recycle bin. If items are removed from the site recycle bin, SharePoint moves them into the site collection recycle bin (second phase). Administrators can retrieve items from the site collection recycle bin. After the 93-day period elapses, background jobs remove files from the site collection recycle bin. If you can't restore a file because it's outside the 30-day window for a restore from the document library, you can look in the site recycle bin and retrieve the file from there.
3. If a site is under the control of a retention policy, SharePoint keeps items removed by users or automatically by background jobs in a special hidden library called the Preservation Hold library. The items stay in the Preservation Hold library until their retention period expires, which could be several years. Administrators can retrieve files if needed from the Preservation Hold library. See Chapter 18 for more information about data governance for SharePoint Online.

Although SharePoint gives administrators and users the ability to retrieve files deleted in error or corrupted in some way, two factors deserve consideration when deciding if external backups are necessary. First, some user education is needed to make people aware of how to use the **Restore this library** feature and how to recover deleted files from the recycle bin. Second, the methods of recovery listed above require manual intervention to restore content to sites. Better automation, the ability to deal with multiple sites at one time, and more granular restores are reasons advanced by third-party backup vendors to justify the purchase of their products for SharePoint Online. The basic idea behind these solutions is to use remote agents that make scheduled connections to applications like SharePoint Online to grab information and copy it out to some other data center or cloud service.

Remember that SharePoint Online includes many other elements than just documents and lists such as customizations made to the search schema or term store. It's a good idea to ask your potential backup

vendors what other SharePoint Online elements they can back up—losing e.g. the term store could cause a major problem.

The decision whether to invest in a third-party backup solution for SharePoint Online largely depends on the faith a tenant puts into the way that Microsoft manages SharePoint data and the perceived risks that exist. The increasing prevalence of ransomware raises the question of how you might recover from an attack, short of paying the ransom. The usual proposed solution is to restore to a point in time before the infection occurred. This might be possible with the functionality built into SharePoint Online (if you catch the problem early) but broad-scale recoveries from malice or mistakes might be easier with a third-party backup. As in the case of any ISV solution, you should test products thoroughly to ensure that they meet your needs, including aspects such as security, privacy, and data at rest.

## Backup for Teams, Planner, and Other Apps

Traditional backup products address the need to copy information belonging to an individual workload. For example, you take backups to protect data in on-premises Exchange databases or SharePoint sites. The different nature of the service creates problems for this approach because new functionality is built by combining features taken from different workloads. Take Planner as an example. Data is held in group mailboxes managed by Exchange Online, group document libraries managed by SharePoint Online, and the Tasks service running in Azure.

Teams is similar in the way that it brings data and functionality together from multiple sources to deliver to users and introduces still more complications. For instance, Teams can use Planner as one of its resources. Currently, there is no comprehensive public API available meant for backing up Teams or Planner data; some structural information (for example, the channels in a team) can be extracted using the Graph API, but no API exists to allow the totality of user data to be backed up from Teams and its associated apps.

This makes Teams the [most challenging of all Office 365 applications](#) from a backup perspective. This is slowly changing; for example, Microsoft has [released a Teams export API](#), but it only handles 1:1 and group chat messages, meaning that there are still several other data items (and many metadata items) that aren't available through the API and thus cannot be reliably backed up or restored. Worse, this API carries a charge to the customer—that is, if you use a backup product that calls the export API, you will incur a charge for every message read through the API.

Worst of all, there is no corresponding *restore* API to match the export API. It is critical to understand that your ability to restore Teams data will probably fall far short of your expectations, as Microsoft has *no* supported APIs for ingesting Teams channel content with original metadata. Ask any vendor you're thinking about working with to demonstrate a full end-to-end backup and restore cycle on a Team that you choose before you pull out your credit card.

**Note:** Backing up the compliance records captured for Teams chats and channel conversations in Exchange Online is not enough – the items are incomplete and cannot be restored into Teams chats or channel conversations in a usable fashion.

Likewise, no backup APIs exist for Planner, Stream, or Yammer, so care should be taken when considering how to extract information from these applications for backup purposes.

## Protection Against Ransomware

Despite the cautious attitude to third-party backup technology for Office 365 expressed here, it is possible to make a case to use backups to protect against the consequences of a ransomware attack. The most likely targets of ransomware are documents and email, both of which are adequately covered by backup solutions. It's also possible to restore documents and emails to other locations if Office is unavailable. For these

reasons, it is worthwhile considering using backups to have copies of data available should a ransomware attack succeed against your tenant.

Of course, it is much better to avoid being attacked, or at least to avoid being compromised by a successful attack. Before seeking solace in backups, organizations should make sure to secure their tenants by eliminating basic authentication (as far as possible), using multi-factor authentication for all accounts, and educating users about techniques such as phishing and business compromise email.

## Service Changes Affecting Backup Technology

Backup technology changes over time to improve functionality and make it easier to backup and restore Microsoft 365 data. It's also true that the service's feature set changes over time, which can make it harder for an organization to copy all the data in use. For instance, the lack of a suitable API to extract Teams data for backup is a problem for tenants who use Teams. Likewise, if you enable auto-expanding archives for Exchange Online mailboxes, you might find that your selected backup product is unable to process this type of archive mailbox.

Technology changes at a fast cadence, especially in the cloud. The need for external backups might not be the same today as it was a couple of years ago. For this reason, it's wise to review your backup strategy (or rather, the need for backups) on an ongoing basis to make sure that your organization only uses what is necessary instead of following the dogma that often comes from on-premises systems.



# Chapter 5: Managing User Accounts

**Paul Robichaux**

In the on-premises world, a user's account typically contained both identifying information (such as a security ID, or SID), plus all the settings that applications and devices need to decide what the user can do. As part of their effort to make Microsoft 365 fully hybrid, Microsoft has effectively split apart this traditional binding. In Microsoft 365, user identities (described in Chapter 3) control authentication to the service; other metadata associated with user accounts, such as the Office 365 licenses assigned and the management roles (if any) assigned to the account, control how the user may interact with the service. These two sets of data no longer must be stored in the same place. For example, Exchange and Teams have their own separate shadow copies of some user attributes that they use for its own purposes. In this chapter, we'll discuss how to manage the accounts themselves and the most important metadata associated with them.

## Managing User Accounts in the Microsoft 365 Admin Center

Many of the day-to-day admin tasks you'll deal with involve user accounts. The **Users** item in the left navigation bar of the Microsoft 365 admin center allows you to manage users, contacts, guest users, and deleted users separately, apply different filters to view only the set of users you need to work with, and quickly take common actions on those accounts. For example, the view in Figure 5-1 uses the licensed users filter to exclude guest users and other unlicensed accounts; the drop-down menu shows actions that you can take on the current selection.

The controls in these pages give you quick access to the users and contacts in your tenant, whether those accounts are synchronized from an on-premises Active Directory or natively homed in the cloud. When you look at the details for an individual user, you can, with one click, delete the user or block them from signing in or changing their password; with a few more clicks, you can easily assign roles, change their multi-factor authentication settings, and perform other common individual tasks.

Keep in mind that, following the guidance in Chapter 3, where you initially create user accounts matters a lot. Some properties are cloud-only (such as the set of assigned Office 365 licenses), whereas others (such as the account's password hash) may either be cloud-native or synced from the on-premises Active Directory, depending on where the account lives. The admin center tools are written so that they hide much of the complexity of setting properties for user accounts. You can view and change most properties without worrying about where the account is homed. However, you can't use the admin center to edit some key properties of on-premises accounts synchronized to the cloud using Azure AD Connect. For example, you can't edit a synchronized user's first or last name or change their email addresses, as those attributes are authoritative from the on-prem directory.

When you add a user from the admin center, the account will always be created in the cloud. You'll be prompted to specify some basic information about the user; you can have the service generate a password for you or you may specify one, and you can assign any of the licenses you've already purchased for the tenant. You may also save the user configuration as a template to speed up the configuration of future users

with similar settings. Templates allow you to specify profile information such as location and department, the service domain, and the licenses to assign.

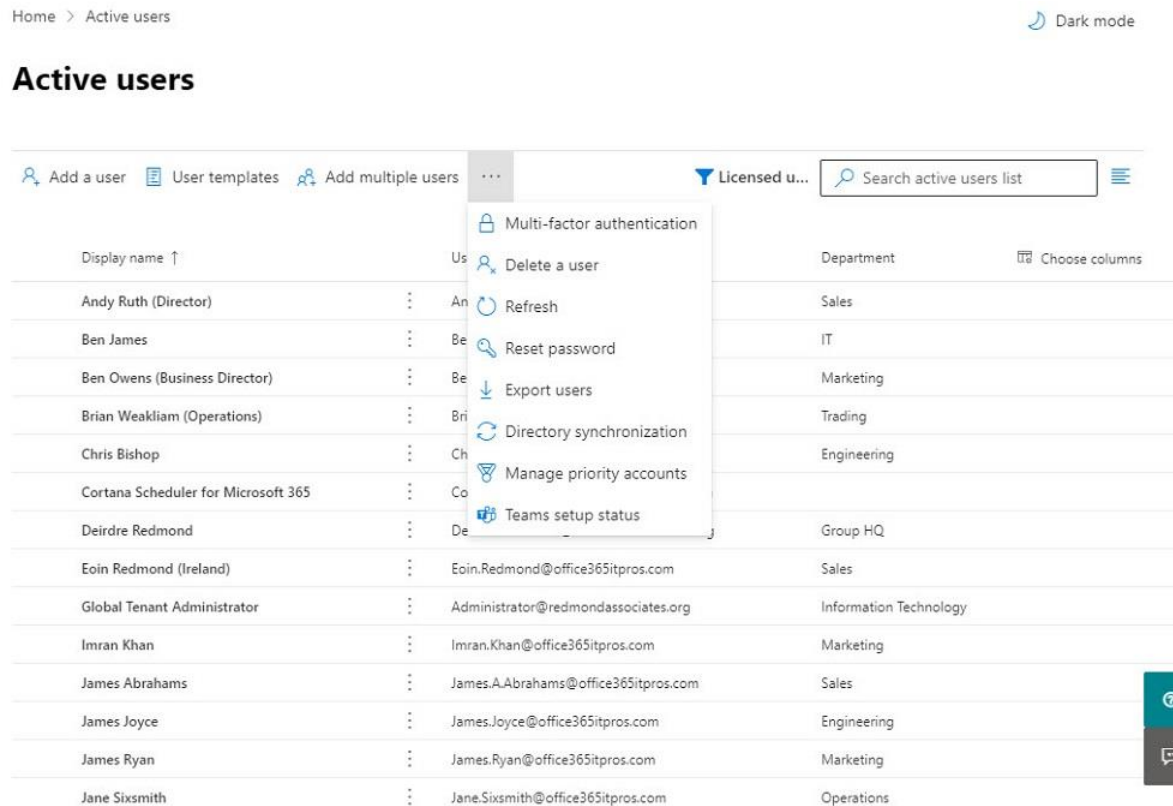


Figure 5-1: Viewing licensed users in the Active Users section of the admin center

The user detail view features tabs that allow you to see and manage devices associated with the user, review and change license assignments, and change email- and OneDrive-related settings. Because these controls are straightforward, there's no need to go into them in detail except for a few notes of interest.

Microsoft continues to gradually roll out better integration between different user-related actions. For example, when you delete a user, you can optionally choose to also remove that user's licenses and/or mailbox delegate permissions or grant another user access to the deleted user's OneDrive files or email, using a simple set of checkboxes. You should expect Microsoft to continue adding support for these kinds of multi-step operations under the banner of "improving efficiency" throughout all the admin center experiences.

## Enabling Self-Service User Management

Any Microsoft 365 user can log into [portal.office.com](https://portal.office.com); this is often how users get access to desktop application software. Most users don't know this, though, and fewer still understand the difference between the Azure AD and Microsoft 365 portals, the *My account* link, and what data and actions are available from each. When a user logs in and clicks on their profile icon in the upper-right corner, then chooses the [View account link](#), they'll see a view similar to Figure 5-2, from which they can see and change many of the settings associated with their account. For example, they can see what devices they have installed Microsoft apps on and what licenses or subscriptions are assigned to them (including ones they've purchased). Since the introduction of the view, Microsoft has added sign-in [information](#) from the Azure AD portal, and it's reasonable to expect that they will continue to add more self-service data and actions in this view so that users can directly manage their accounts to the extent that it's reasonable.

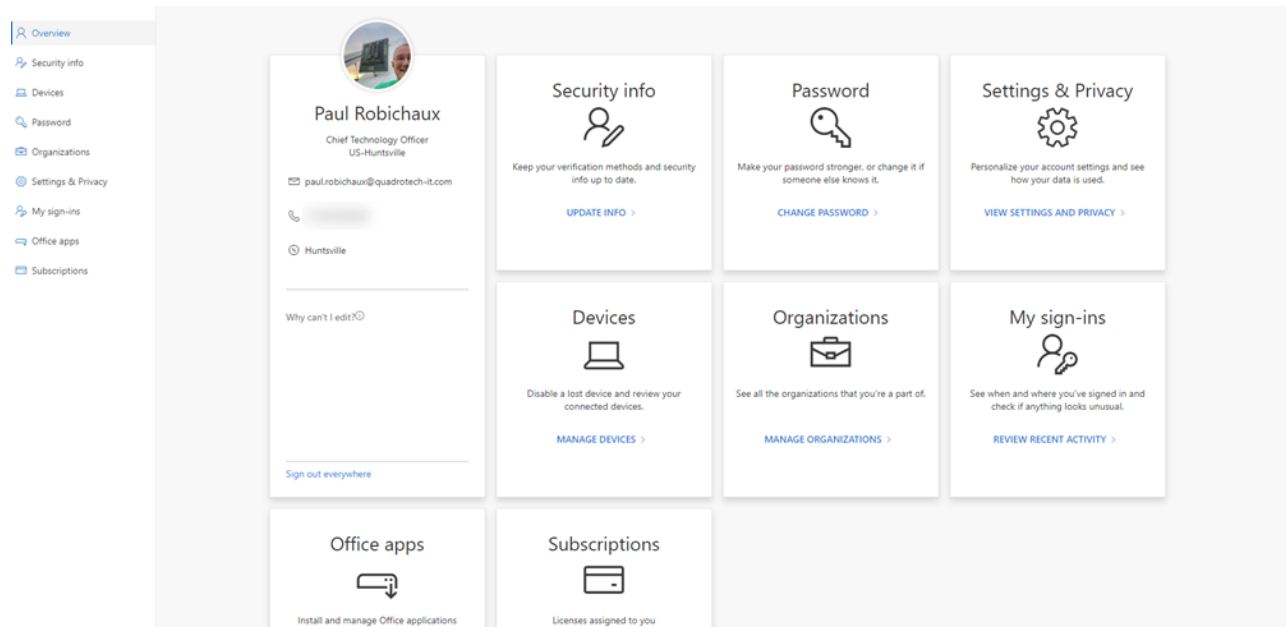


Figure 5-2: The integrated account view is available to all users

## Managing User License Assignments

Every user needs a license to use the service. While there are a few workloads or capabilities that Microsoft doesn't charge for (e.g., the free trial-version licenses formerly available for Teams and Power BI), the major workloads are all tied to licenses, so users will need one to use them. Normally you'll assign licenses to users when you bring their accounts into the service, either as part of a migration or when you onboard a new user. You can do this manually with the admin center, via a custom provisioning script, or using third-party products such as Quest's On Demand License Management tool. By the same token, when a user leaves, you may wish to reclaim the license, bearing in mind the limitations covered later.

### Azure AD Group-Based License Management

Managing licenses for a large user population has been a pain point for many customers. Automation is essential, as manual methods are far too time-consuming for any environment with a high rate of change, such as dealing with new and departed users, or licensing sub-features and extra applications.

In the past, some customers have invested quite a lot of time into scripting solutions based on AD group membership. These scripts typically require that the user be added to a group whose members should receive an Enterprise E5 license, and the custom script runs the necessary PowerShell commands to assign that license. This approach becomes challenging with more complex licensing scenarios such as sub-SKU features and assigning multiple products to an individual user (e.g., Enterprise E5 plus Project and Visio).

To solve these challenges Microsoft made group-based license management available through the Azure portal. The feature originally required a paid or trial subscription for Azure AD Premium P1 or greater. This subscription may be purchased on its own, but it's already included with Microsoft 365 Enterprise E3 or A3 and higher, and all the E plans of Office 365 Enterprise can natively use group-based licensing, so you probably won't have to buy anything. However, at the end of August 2021, Microsoft announced that a new licensing platform, originally promised in the first quarter of 2022 but still not released as of July 2022, will allow you to use group-based licensing [without an additional Azure AD Premium license](#)—a very welcome change.

The groups that you can assign licenses to can either be created in Azure AD or synchronized from on-premises Active Directory. The license assignments can either be static (i.e., to the members of a group) or

dynamic (e.g., based on user attributes such as *ExtensionCustomAttribute1*). For some organizations, a department-based model will be the preferred approach, with licenses assigned to groups representing the different departments within the organization. For others, a product-based approach will be more appropriate, with each type of license being assigned to a single group, and the users being added to that group regardless of which department they are in. If you're using static groups instead of dynamic groups then you will need to directly add the user accounts to the group, as nested groups will not be supported for group-based licensing until the 2022 licensing platform is delivered.

Transitioning from direct license management to group-based license management is simple. Once you put the group-based license assignments in place, users can have both a direct and group license assignment listed for their account for the same type of license. However, they will only consume a single paid license during this time, allowing you to remove the direct license assignments slowly to ensure there are no unexpected results. License assignments will fail if there are no available licenses to assign. The solution is to purchase more licenses, then force the group-based license assignment to reprocess.

As Microsoft rolls out new features to your tenant they will usually be enabled by default. This will apply to group-based license assignments as well, with new sub-SKU features being enabled by default. As new features arrive you should review your group-based licenses to ensure that any new sub-SKU features are enabled or disabled to meet your organization's requirements.

To configure the group-based licensing assignments, you use the Azure portal. After logging in to the Azure portal choose **Azure Active Directory** from the list of services in the portal, and then select **Licenses**. The **All products** view shows you what licenses your tenant currently holds.

Select a product license and click on the **Assign** button. From **Users and Groups**, choose the group that you want to assign licenses to, and then click on **Select**. In the **Assignment options**, you can select the sub-features for the license that you've chosen to assign to the group. In the example shown in Figure 5-4, all the EM+S features are enabled for the newly licensed users.

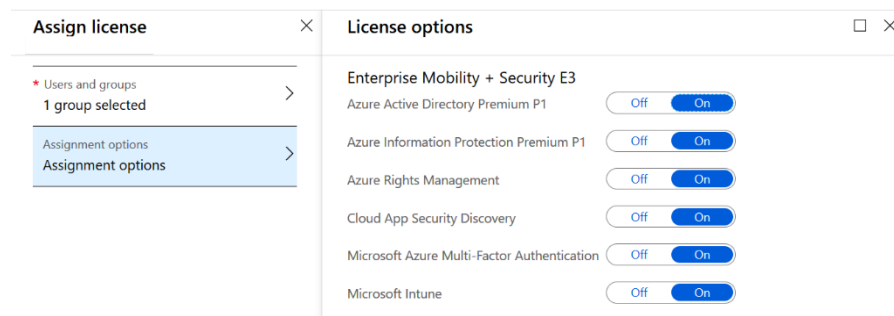


Figure 5-4: Configuring sub-SKU features for a license assignment

Click **Review + Assign** when you're happy with your selections to create the license assignment. If there are any errors at this stage, you'll receive a notification in your Azure portal. License assignments are updated almost instantaneously, but for on-premises Active Directory groups, you should expect the normal directory synchronization delay before group membership changes flow through to license assignment changes.

If you've already been using direct license assignments for your users, you will see each user in the Azure portal listed with an **Assignment Path** of both *Direct* and *Inherited*. To complete the transition from direct license assignment to group-based license assignment the direct assignment must be removed. You can select one or multiple users, and then click the **Remove** button. You should approach this with some caution, removing the direct license assignments for users in small batches so that you can be sure that you do not inadvertently leave your users unlicensed for an application that they are still using.

To reduce the risk of licenses being accidentally removed from users, if a user is removed from a group that is managing their license assignments the associated Microsoft 365 services will initially go into a suspended state, instead of a deleted state. The user will not be able to access and use those licensed services, but their

data will be safe while their IT administrators correct the licensing mistake. If the license removal was intentional, the suspended services will eventually age out to a disabled state and will begin their normal deletion and final purge processes.

User accounts can be members of multiple groups for license assignment. For example, if a user is a member of a group that is assigned the Enterprise E3 license, and a member of a group that assigns the Enterprise Mobility and Security E3 license, the cumulative effect will be that both licenses are assigned to that user. This allows you to approach group-based licensing in a modular fashion, instead of needing to create separate groups for all possible combinations of license assignments in your organization.

## Allowing Users to Buy Licenses

It's obviously to Microsoft's benefit to sell as many licenses as possible to the service. To help in this goal, in January 2020, Microsoft introduced a self-service purchase mechanism to allow people with a valid Azure AD account in a tenant to effectively bypass their IT departments and buy licenses for the Power Platform products (Power Automate, Power Apps, and Power BI). They have since added other parts of Microsoft 365 (Visio, Project, Power BI Premium, and Power Automate so far); in August 2021 they added the ability for users to buy licenses for the new Windows 365 Enterprise and Windows 365 Business offerings. Starting in January 2022 users can sign up for 30-day trials of Project and Visio using the self-service mechanism. Self-service purchases are not available to Government, Nonprofit, and Education tenants. All data management and access policies enabled by the company apply to self-service purchased services. Also, data created from products with these licenses remains owned and controlled by the organization.

By default, this capability is turned on. Turning it off requires you to use the *MSCommerce* PowerShell module, as described in Chapter 23.

## Allowing Users to Request Licenses

In an ideal world, we'd always have a big enough budget to buy enough of the right licenses to assign every user exactly what she needs. In practice, this doesn't happen much, so sometimes users don't have the licenses they need. Besides permanent assignments, it is sometimes the case that a user needs a license for a specific assignment or task. To help solve this problem, in the same way that you might check out a book from the library, Microsoft allows users to request licenses on their own, but only if you block self-service purchases first. The way this currently works is a little weird: users click the "buy now" link from the product page on Microsoft.com, then enter their email address. If the domain matches a domain that's enrolled in the service, the user logs in with her credentials, then is presented with a form. The contents of the form may vary depending on what you've done. Let's say that Alice wants to buy Visio licenses for herself, Bob, and Carlos:

- If the organization enables self-service purchases, Alice goes to the Microsoft.com page for Visio, supplies her payment information (which must be a valid credit card), logs in to the tenant, and specifies the email addresses for Bob and Carlos. The purchase proceeds and the licenses are assigned to those users.
- If you want to force Alice to use your organization's existing license process, you can go to the **Billing > Licenses** section of the admin center and click on the **Requests** pivot, then click the **Use your existing request process instead** link. This page allows you to set a message that users see when they try to buy a self-service license (e.g. "To request a product, file a service ticket at <https://help.internal/>"). Alice will see this message but isn't prompted for anything else; her request won't appear in the admin center.
- If the organization doesn't have a specific process, then when Alice tries to make a self-service purchase, she'll see the same form as described earlier, asking for payment details and so on, but when she submits the request, it will appear on the **Requests** pivot. You can approve or reject

individual requests, and for requests you approve, you can choose which users in the request are approved (e.g., you can accept the requests for Alice and Carlos but reject Bob). Remember that accepting the request allows it to be completed for the original purchaser, just as if you had enabled self-service purchases directly. The purchased licenses aren't considered as part of your tenant subscriptions.

## Assigning Licenses with an Auto-Claim License Policy

In addition to having a license assigned automatically through group membership, manually by admin action, or manually by the user herself, you can control license assignments through an *auto-claim policy*. This policy allows a user to automatically receive a license when they try to use a workload for which they are currently unlicensed—think of it as a just-in-time automated license assignment. The auto-claim policy is triggered the first time an unlicensed user signs into the workload, and the resulting license stays assigned to the user until you do something to remove it. Teams is currently the only Microsoft 365 app to support auto-claim.

Each tenant can have a single auto-claim license policy, configured in the Licenses section of the Microsoft 365 admin center. The policy specifies:

- The apps covered by the policy.
- The license to assign when an unlicensed user signs into the specified apps for the first time. Obviously, the plan should license the use of the app.
- One of more backup licenses to assign should no available licenses exist for the primary assignment. The backup licenses are in priority order and the claim is made for the first available license in that order.

For more information, read [this article](#).

## Managing User Role Assignments

Microsoft 365 (and some of its key workloads, notably Exchange Online and Azure AD) offers customers a variety of administrative roles that you can assign to users who need to perform management tasks. Much like on-premises environments, there is no single “admin” level permission that someone needs to perform some management tasks for your tenant should have. Instead, a set of pre-configured administrative roles and groups is available to allow the assignment of limited but necessary permissions to administrators to do their job. This model is called role-based access control (RBAC). Keep in mind that the service and workload-specific roles are different from the Azure roles that you may wish to assign; for more information on Azure AD role management, see [Microsoft's documentation](#).

### Assigning Administrative Roles to Users

There are several ways to assign a privileged role to a user. The easiest way is probably to select and edit their account in the Microsoft 365 admin center. You can then assign an administrative role to the user by clicking the **Manage Roles** link on the user settings page. You can also use the Roles section in the Azure AD portal to assign some roles, or even break out the [Microsoft Graph PowerShell](#) module and do it from the command line.

By default, no user account receives an administrative role. The exception is for the account that is created when a company signs up for a new tenant, which is assigned the **Global admin** role. This default account is based in the cloud and will remain permanently available unless you delete it. You should keep at least one cloud-based Global admin account available so that you can log in even if AD FS or Azure AD Connect is broken for your tenant.

Global admin holders get complete administrative access to all features within Microsoft 365, including individual services such as Exchange Online and Teams. Naturally, every organization will have at least one Global admin. While having too many of these accounts is a security problem, it's a good idea to keep multiple accounts to give you redundancy in case the one person with administrator permissions isn't available when you need them.

## Assigning Roles with Azure AD Groups

If you have an Azure AD Premium P1 or P2 license, you can [assign Azure AD roles to groups](#). This is a little more complex than the way that you'd accomplish the same task with on-premises AD. You must create the group in Azure AD, then give it the **Azure AD roles can be assigned to the group** option (which changes the *isAssignableToRole* property on the group object). Then you can assign roles to the group, then add users to the group. *isAssignableToRole* is an immutable property, so you cannot role-enable existing groups, and you cannot disable role assignment from a group once you've set this flag.

By default, only Global Administrators and Privileged Role Administrators can create groups with this flag set or manage the membership of such groups. Although you can delegate this ability, you should be careful when doing so to prevent accidentally giving people excess privileges. For the same reason, you cannot use dynamic groups with role assignments; all assignments to role-enabled groups must be done manually. These groups can be managed through the Microsoft Graph, provided that the caller has the *RoleManagement.ReadWrite.All* permission; the ordinary *Groups.ReadWrite.All* permission won't work.

## Assigning Roles Directly to Users

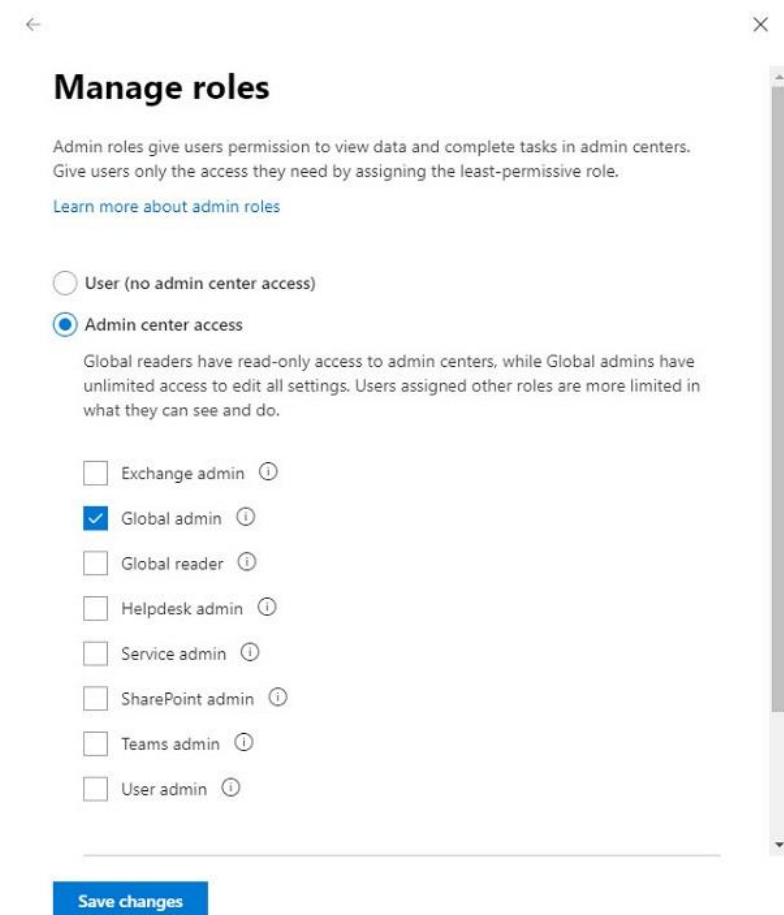


Figure 5-5: Assigning administrative roles to a user account

In the user management interface, you can assign a user to the **Global admin** role or assign them to one or more customized administrator roles (Figure 5-5). When you assign a user to one of these roles you must

give an alternative email address for them to use when they need to recover a lost password for their admin account.

## Understanding the Built-In Roles

The familiar security principle of “least privilege” dictates that you should only grant users whatever limited admin access is necessary for them to do their job. With more than 60 roles scattered throughout the Azure AD and Microsoft 365 portals, it can be difficult to track which roles do what, or even what all the roles *are*. Microsoft groups these roles into two basic groups: the first group (labeled with “Admin center access”) grants access to specific admin centers in the Microsoft 365 ecosystem:

- **Global Administrator:** you already know what this is—the all-powerful master admin account that can do anything, anywhere, within the tenant.
- **Exchange Administrator:** Users with this role use the Exchange Online admin center to manage mailboxes, groups, anti-spam policies, and access activity reports in the Microsoft 365 admin center. An Exchange Admin becomes an Organization Management role group member in Exchange Online, which is a high privilege role group. Exchange Online has a more granular permissions model known as Role-Based Access Control (RBAC), which you can use to assign least privilege permissions instead. RBAC is discussed later.
- **Global reader:** this is like the old “Exchange view-only administrator” role in on-premises Exchange; users who hold this role can read, but not change, most non-security-related settings in both Azure AD and Microsoft 365. Note that at its initial launch, the Global reader role worked with Microsoft 365 admin center, Exchange admin center, Teams admin center, Security center, Compliance portal, Azure AD admin center, and Device Management admin center, but not with SharePoint or some Teams management features; you may find that it doesn’t completely work with all workloads.
- **Helpdesk admin:** Delegating responsibility for password resets can relieve some of the support burden from IT teams. Many processes created to reset user passwords involve the need for an authorized person to check the identity of the user before any change is made to their password, for example by having them walk up to their desk and show a company identification badge. Once the user is recognized, the authorized person sends the request to the help desk. Giving that authorized person the right to reset the password themselves can save time. However, there is no granularity when it comes to assigning this admin role through the Microsoft 365 admin center. Helpdesk admins can reset any other user or helpdesk administrator’s password, but not other types of admin users. They could reset the password for someone in a completely different location, or someone whose identity is unverified. As such it is important to be careful to only assign this role to trusted individuals.
- **Service support admin:** Formerly known as the Service admin role, this role grants the ability to manage service requests and to access the Service Health Dashboard. You should assign the Service support admin role to users who are members of the Exchange administrator, Teams administrator, or SharePoint administrator groups to allow them to raise support tickets for those services. The role also gives the holder read-only access to information about groups, users, and licenses, which means that it is a good role to assign to people who want to know about license usage within a tenant. Interestingly, Microsoft changed the name of this role to match the role name in the Microsoft Graph API, rather than leaving the UI alone and changing the role name in the API definition.
- **SharePoint Administrator:** This role grants the ability to manage SharePoint Online, which also affects OneDrive for Business, as well as creating and managing sites, and managing user profiles. SharePoint admins can assign other users with administrative permissions within SharePoint Online for sites and term stores. They can also access activity reports in the Microsoft 365 admin portal.



- **Teams Administrator:** Users with this role can manage all aspects of Teams (including policies and calling) but cannot assign or remove user licenses. They can also manage Microsoft 365 Groups, which makes sense given how Teams uses them.
- **User Administrator:** This role can manage users, groups, passwords, service requests, and see the Service Health Dashboard. This admin role is ideal for a help desk or low-level support person. Although a User admin role holder can remove other user accounts, they cannot remove the account belonging to a Global admin nor can they reset passwords for Billing admins, Global admins, or Service admins.

Other admin roles are grouped according to service category (e.g., “Collaboration”, “Devices,” and “Identity”). Microsoft is continually moving and adding roles as they expand the set of supported role features. There are currently more than 80 individual roles, the details of which you can find [documented here](#).

## Creating Custom Roles

In addition to the built-in roles, you can create your own custom roles in Azure AD. These custom roles draw from the same set of permissions as the built-in roles—think of the built-in permissions as a deck of cards from which you can draw the specific cards you want to. You can’t make up a new “13 of diamonds” card, and you can’t make up your own new permission. However, you can combine these existing permissions into custom roles that grant exactly the access you want to grant to role holders. To create custom roles, you must have Azure AD Premium P1 or P2 licenses. The mechanics of setting up custom roles will be familiar to anyone who’s used custom roles in Exchange on-premises or Exchange Online: first you create a new role, then you add permissions to the role, then you assign the role to users. You can create these roles using the Azure AD portal, PowerShell via the *New-MgDirectoryRole* cmdlet, or through Microsoft Graph. All 3 approaches are [documented here](#).

## Using the Roles Section of the Admin center

The Roles section in the Microsoft 365 admin center may be easier to use for assigning roles, because it gives you a single list of all the supported roles and allows you to filter, search, and sort them. You can export a CSV file showing which users have been granted roles in the tenant, or you can see and change assignments for any individual role. Clicking on a role will give you a short list of capabilities the role has, plus tabs for viewing and changing user assignments and viewing a comprehensive list of the role’s permissions (which, sadly, you cannot change).

The Roles page uses multiple pivots (or tabs) at the page top. All the Azure AD roles are on the Azure AD pivot, while Exchange roles are on the Exchange pivot. Microsoft will occasionally add or remove pivots here as they make changes to the admin centers (for example, Intune used to have its own pivot here, but it’s now gone again.)

## Listing Who Holds Administrative Roles

The list of roles above may not be exhaustive because Microsoft adds new roles over time. To find the full set of currently-known roles that are active in your tenant, and the people who hold these roles, you can run this PowerShell command:

```
[PS] C:\> Get-MgDirectoryRole | % {$Role = $_.DisplayName; Get-MgDirectoryRoleMember -ObjectId $_.ObjectId} | Format-Table @{"Name"="Role"; Expression = {$Role}}, DisplayName, UserPrincipalName -AutoSize
```

In the list, you will likely see several service accounts alongside user accounts. For example, the Microsoft 365 admin center holds the Adhoc License Administrator role while a service like Secure Score holds both the Directory Readers and Directory Writers roles. In addition, apps like Secure Score use [service principals](#) as a

way to hold roles that allow them to access information such as the tenant directory. This code tells you what service principals have been assigned Azure AD roles in a tenant:

```
[PS] C:\> $Roles = Get-MgDirectoryRole
Foreach ($Role in $Roles) {
    $RoleMembers = $Null
    $RoleMembers = (Get-MgDirectoryRoleMember -ObjectId $Role.ObjectId)
    If ($RoleMembers) {
        ForEach ($RoleMember in $RoleMembers) {
            If ($RoleMember.ObjectType -eq "ServicePrincipal") {Write-Host $Role.DisplayName "service
principal:" $RoleMember.DisplayName }}
    }}
}}
```

You can assign administrative roles to licensed or unlicensed users. An unlicensed user does not have access to any licensed feature, such as an Exchange Online mailbox or being able to install the Microsoft 365 apps for enterprise on their computer. However, they can log in to admin portals and use PowerShell to perform management tasks. Administrators who perform content searches and want to preview search results need an Exchange Online mailbox to be able to see the results. In general, third-party services and tools that require the use of privileged or service accounts are likely to send mail to those accounts, so be sure you're paying attention to their mailboxes.

**Note:** Microsoft 365 uses the alternative email address for account recovery if the password for the account is lost. In some cases, the alternative email address is needed when the administrator is unable to access their account or mailbox at all. Therefore, it is best to use an email address that is independent of the tenant. This raises some obvious security concerns; someone could use the alternative email address to reset the password for the administrator account. You should ensure that the alternative email address is well protected by a strong password and multi-factor authentication. Services such as Outlook.com or Gmail can offer the necessary level of security.

## Exchange Online Administrative Roles

Exchange Online has its own administrative roles or management role groups (Figure 5-6), some of which link to Microsoft 365 admin roles and others that are independent. You can see the Exchange Online admin roles by logging in to the Exchange admin center and navigating to the **Roles > Admin roles** section.

Global admins automatically have Organization Management rights in Exchange Online. Users assigned the Global admin role join an Exchange Online role group called **TenantAdmins\_XXXXX**, where the last five characters in the group name are unique to your tenant. The TenantAdmins group is nested in the Organization Management role group and is displayed as **Company Administrators**. Organization Management is a powerful admin role that has access to manage all the features of Exchange Online, so you can see why the Global admin role should only be assigned to selected administrators in your organization.

Users who get the Exchange administrator role also automatically have Organization Management rights in Exchange Online. From an Exchange Online perspective, this is the same level of admin rights that a Global admin receives, but an Exchange service administrator is not granted any other admin rights within Microsoft 365 itself, such as managing billing, subscriptions, or domain names.

Helpdesk Administrator holders are automatically granted View-Only Organization Management rights in Exchange Online. Users assigned the Password admin role join an Exchange Online role group called **HelpdeskAdmins\_XXXXX**. The HelpdeskAdmins group is nested in the View-Only Organization Management role group and is displayed as **Helpdesk Administrator**. This role group can view the configuration and recipient details within Exchange Online but can't make any modifications, other than resetting user passwords.

## Admin roles

Admin role groups give users permission to view data and complete tasks in the Exchange admin center and use the Powershell cmdlets. Give users only the access they need by assigning the least-permissive role. [Learn more](#)

Role group ↑	Description
CLBAccessApprovers_-444955902	
Compliance Management	Membership in this role group is synchronized across services and managed centrally. This role group is not manageable through Microsoft Exchange. Members of this role group may include cross-service helpdesk or password administrators, as well as external partner groups and Microsoft Support. By default, this group is not assigned any roles. However, it will be a member of the View-Only Organization Management role group and will inherit the rights of that group.
ComplianceAdmins_1744690546	
Discovery Management	Members of this management role group can perform searches of mailboxes in the Exchange organization for data that meets specific criteria.
<input type="radio"/> ExchangeServiceAdmins_53add	Membership in this role group is synchronized across services and managed centrally. This role group is not manageable through Microsoft Exchange. Members of this role group include Exchange-Online service administrators only. By default, this group may not be assigned any roles. However, it will be a member of the Organization Management role group and will inherit the capabilities of that role group.
GlobalReaders_1331601875	Membership in this role group is synchronized across services and managed centrally. This role group is not manageable through Microsoft Exchange. Members of this role group may include auditor. By default, this group is not assigned any roles. However, it will be a member of the View-Only Organization Management role group and will inherit the rights of that group.
Help Desk	Members of this management role group can view and manage the configuration for individual recipients and view recipients in an Exchange organization. Members of this role group can only manage the configuration each user can manage on his or her own mailbox. Additional permissions can be added by assigning additional management roles to this role group.
HelpdeskAdmins_7ae4b	Membership in this role group is synchronized across services and managed centrally. This role group is not manageable through Microsoft Exchange. Members of this role group include Exchange-Online service administrators only. By default, this group may not be assigned any roles. However, it will be a member of the Organization Management role group and will inherit the capabilities of that role group.

Figure 5-6: Exchange Online RBAC role groups

**Note:** You might notice in the Exchange admin center that a View-Only Organization Management role group member can create distribution lists and manage distribution lists that they have created. This is due to the default role assignment policy for users in Exchange Online which permits anyone to create and manage their distribution lists.

Aside from the Exchange Online admin roles inherited from admin roles, several other Exchange Online role groups also exist for granularly assigning rights to different users within your organization:

- **Compliance Management** – members of this role group can manage Data Loss Prevention, Information Rights Management, and Retention. In addition, members of this role group can view audit logs as well as all configuration and recipient attributes in the organization. Managing features such as DLP and IRM is often part of a more general security and information compliance role in an organization, not necessarily a duty performed solely by Exchange administrators, and this role group allows those users to be assigned the admin rights they need without having broader Exchange admin rights. Note that a separate set of roles govern access to the Purview compliance portal, discussed later.
- **Discovery Management** - members of this role group can perform discovery searches of mailbox data in Exchange Online, export search results to PST files, and configure legal hold on mailboxes. Microsoft has deprecated the eDiscovery experience in EAC (and it isn't even present in the modern EAC) in favor of using the Purview toolset, so you may not run into this role much in the future.
- **Help Desk** – members of this role group can view and manage the same individual recipient attributes that the users can view and manage for themselves. For example, users can log in to Outlook Web App and change personal details such as their phone number and street address or

update their password. Help Desk role group members can therefore also change the user's phone number, street address, or reset their password. This role group is suitable for low-level support staff, or for service accounts that automatically synchronize attributes such as phone numbers and street addresses from other systems such as an HR database.

- **Hygiene Management** – members of this role group have view-only access to Exchange configuration and recipients, and they have permission to manage some aspects of the transport system, mostly settings for anti-spam and anti-malware filtering, which are actually now in the Defender portal. You'll normally assign this role to people who manage your anti-malware appliances or services, and perhaps to select members of your organization's security team.
- **Organization Management** – members of this role group have access to manage all features of Exchange Online, except for the rights that the Discovery Management role group allows. As discussed earlier, Global admins and Exchange service administrators are automatically made members of this role group. You can also assign membership of this directly in Exchange Online, but you may find it better to assign the Exchange service administrator role instead as this will force the addition of an alternative email address for lost password recovery.
- **Recipient Management** – members of this role group have access to manage recipient objects in Exchange Online. Recipient objects include mailbox users, shared mailboxes, group mailboxes, resource mailboxes, distribution lists, contacts, and other mail-enabled objects. Recipient Management members can also perform mailbox migrations, message traces, and reset passwords. This role group is ideal for day-to-day tenant administration and is often assigned to first and second-level support teams. However, it is also suitable for general use by higher-level administrators who want to separate their admin accounts into low and high privilege use. For example, a Global admin or Exchange service administrator may have a separate admin account that is only granted Recipient Management rights. They can use the low privilege account for general administrative tasks, running report scripts, and so on, and only log on to the higher privilege account for the less frequent tasks that require those higher admin rights.
- **Records Management** – members of this role group are granted rights to perform compliance-related tasks such as configuring or viewing audit logs, configuring journaling, performing message traces, managing retention policies, and configuring transport rules.
- **Security Administrator, Compliance Management, Global Reader, Security Operator, and Security Reader** – although you will see these role groups in the Exchange admin center, they only appear here because the role groups are used across multiple Microsoft products (including Information Protection and the Purview portal). The Security Reader role group grants read-only access to Azure AD, Azure Identity Protection, Azure Privileged Identity Management, and all audit logs and sign-in reports. The Security Administrator role group grants the same access as Security Reader, plus the ability to configure security services. You won't manage membership in these role groups yourself; instead, you'll assign access to the services, and when you do, the role group membership will be automatically updated.
- **UM Management** – members of this role group can manage the Unified Messaging settings for the organization, as well as configure UM properties on mailboxes. This role group was only used for the now-deprecated Unified Messaging features.
- **View-Only Organization Management** – this role group gives read-only access to all recipient properties and configuration settings for the organization. This role group is ideal for anyone who needs broad visibility of the organization without the right to make changes. One scenario where this level of access is useful is for reporting scripts and tools that need to be able to read a wide range of information about the organization.

As you can see, a wide range of pre-configured administrative roles exists in Exchange Online to suit different requirements. Chapter 7 describes these roles, and their uses, in more detail.

## Assigning Security- and Compliance-Related Roles

The security and compliance features in Microsoft 365 support a set of permissions defined in RBAC role groups that allow the separation of responsibilities across different sets of users. The role groups are like those used by Exchange Online, except that these role groups only apply to the functionality available through the security and compliance portals and have no connection to the role groups used by Exchange Online, even though some of them share the same name.

To see the definitive version of the available roles and to assign or remove roles to users, go to the **Permissions** section of the Microsoft Purview Compliance portal or the **Permissions** section of the Microsoft 365 Defender admin center (see Chapter 4 if you need a refresher on the difference between these two). Each of these admin centers shows the same set of roles; there are also some roles associated with Azure AD that will be displayed in these admin centers but must be managed through the Azure AD admin center.

Here's a partial list of the roles you can assign to users:

- **Reviewer** - members of this role group can view the list of e-discovery cases that currently exist. Case managers can assign specific documents in an eDiscovery case for a reviewer to analyze or use a limited set of the analysis features in Microsoft Purview. Members of this group can see only the documents that are assigned to them.
- **Records Management** – members of this role group can manage retention and compliance features related to how long user data is stored and what policies are applied and enforced to archive or expire it.
- **Security Operator** – members of this role group can manage security alerts, view security reports, and view security settings.
- **Insider Risk Management** – members of this role group can manage and control the insider risk policies applied to users in a tenant.
- **Compliance Data Administrator** – members of this role group are typically IT administrators who are responsible for configuring security and compliance settings and policies, such as mobile device management, DLP, and preservation. Compliance administrators can also manage settings for reports in the Compliance portal.
- **Security Administrator** - membership of this role group is not intended to be managed by you directly through the security or compliance portals (although you can do so). Instead, Microsoft intends you to grant users access through the Security Administrators role in Azure AD. This role group may include Microsoft Support or external partners.
- **Security Operator** - this allows holders to read security reports and view all security settings, but they can't change any of them.
- **Compliance Administrator** – identical to the Compliance Data Administrator role, except that holders cannot manage data loss prevention settings.
- **Security Reader** - like the Security Administrator role, this role group is not intended to be managed by you. It gives read-only access to security and compliance features.
- **Supervisory Review** - members of this role group can create supervisory review policies for the organization, which determine the type of content that will be subject to supervisory review, and who will be performing the reviews.
- **Organization Management** - members of this role group can perform the same tasks as Compliance Administrators. They can also configure permissions for the Compliance portal and configure audit logging for the organization. Global admins are automatically added to this role group.
- **Mailflow Administrator** - this role group grants its members access to the mail flow dashboard. Users who are members of the Global admin or Exchange admin role groups will already have access

to this dashboard, but if you grant membership in Mailflow Administrator to users they will only have access to the dashboard, not to other Exchange or Microsoft 365 management features.

- **eDiscovery Manager** - members of this role group can perform searches across workloads, and apply holds to Exchange Online mailboxes, SharePoint Online sites, and OneDrive for Business libraries. An eDiscovery Manager has access to specific eDiscovery cases. When you grant a user the eDiscovery Manager permission, you can also grant them eDiscovery Administrator privileges, which allows the user to view and edit all eDiscovery cases for the organization.
- **Data Investigator** – a new role group whose holders can perform mailbox, SharePoint, and OneDrive searches using the new investigation features in the Compliance portal.
- **Service Assurance User** – members of this role group can access the Service assurance section in the Compliance portal, where they'll find reports, documents, and audit reports related to how Microsoft handles customer data.
- **Global Reader** – this role group allows the holder to read all reports, alerts, and settings for security and compliance features, similar in concept to how Exchange View-Only Administrator works.
- **Quarantine Administrator** - This role allows holders to manage mail quarantine settings and items in the Compliance portal.

You can create custom role groups. The **Create** button on the Permissions page allows you to create a new role group by specifying the base roles you want to use for the new role group, then assigning it to a set of users. It's important to know that your custom role group will get all the privileges assigned to the base roles you pick—there's no way to remove a specific privilege from the roles you create.

Like any other role assignment, the tenant administrator should review what accounts hold which roles regularly to ensure that users do not retain privileged access for longer than required. The eDiscovery Manager role group is particularly powerful because of the two sub-roles that it contains. eDiscovery Managers have access to all their assigned cases. eDiscovery Administrators have access to all cases and can assign themselves to a case. Although dependent on the volume of eDiscovery workload within the tenant, the usual situation is to have only a few eDiscovery Administrators and more eDiscovery Managers. However, in a small tenant, the same people might comprise the two groups.

## Administrative Units

[Administrative units](#) (AUs) are containers for objects, like organizational units on-premises. AUs allow you to group your users into logical units, then scope management tasks to the AU instead of the entire organization. For example, you could create an AU for all the users in the Sales department, or all users working in Slovakia; then you could grant permission for an ordinary user (let's call her Dagny) to perform user management tasks only on the users in a specific AU.

Like on-premises OUs, the real power of AUs is the ability to give Dagny permission to perform various tasks against *only those users* who are in the scope of an AU. This is very important for large organizations. The limited-scope admin roles (such as Exchange administrator) that exist control *what* their holders can do, but they don't restrict *which users* those role holders can manage.

The simplest way to create AUs is to use the Azure portal, as [described here](#). One nice optimization to this process is that you can create role assignments scoped to the group at the time you create the group itself. Of course, you can also create those role assignments separately by editing the properties of the AU in the Azure portal.

After you've created the group, you need to [add members](#) (which can be users or device objects). Device support is still in preview, though.

Finally, you need to add role assignments that are scoped to the AU. The result is that Dagny, or whoever else has been granted permissions on the AU, can exercise the privileges of the assigned role, but only on

users who are in the targeted administrative unit. The Microsoft 365 Admin center will filter out any users who aren't in the selected AUs, so Dagny won't even see them. If Dagny uses PowerShell instead, she will see users who are out of scope for her AU permissions, but she can't modify them; role scoping only applies to write operations.

If you prefer to manage your AUs with PowerShell, you can; the documentation linked above has examples for each of the cmdlets you'll need to use.

## Managing Privileged Accounts

Best practice in the Windows community is to perform administrative actions with dedicated administrator accounts rather than assigning the necessary permissions to user accounts. This is done to restrict permissions to a limited set of accounts rather than allowing for user accounts to gather a proliferation of permissions, sometimes for doubtful reasons. It also ensures that you do not have to perform regular reviews of highly permissioned accounts to remove permissions from accounts that do not need them. Inside the Microsoft data centers, very few administrators have elevated permissions, and great care is taken to ensure that permissions are only granted when needed and for a minimum period and that they are only used from privileged access workstations (PAWs), dedicated computers specifically configured for high security.

You can certainly bring the practice of using dedicated accounts for administrative tasks forward into Microsoft 365. However, given that the service is a very different environment that hosts many different workloads and the need for permissioned access is reduced because there are fewer administrative tasks to perform (no server management, for instance), perhaps this is a good time to consider whether a better approach could be taken.

### Privileged Identity Management

Microsoft offers two features to manage privileged access to user information. The first is [Azure AD Privileged Identity Management \(PIM\)](#), a framework for the management of privileged access to applications that depend on Azure AD, such as Microsoft 365. Using PIM, you can manage temporary elevation of user accounts to grant administrative privileges, such as Global Administrator. A user can be permanently eligible to elevate their permissions, or you can require the user to request approval each time they require elevated permissions. PIM will elevate the user's permissions for the time needed, then revoke the permissions afterward. Added controls such as enforcing multi-factor authentication (MFA) are also available in PIM, as well as an audit trail and activity report. PIM is available for customers who have Azure AD Premium P2 licenses, whether bundled with the Enterprise Mobility + Security E5 license or purchased separately.

### Privileged Access Management

The second feature is [Privileged Access Management \(PAM\)](#), which is generally available for Exchange Online. Other workloads will be added over time. Microsoft considers PAM to be a compliance feature now, so each user who requests or responds to a PAM request must have either an E5-equivalent (the Office 365 or Microsoft 365 E5, A5, G5, etc. SKUs) or an E3 license plus the Microsoft 365 E5 Compliance or Microsoft 365 E5 Insider Risk Management licenses.

PAM works on the basis that administrators create requests for authorization when they want to perform privileged tasks. Policies controlling individual cmdlets or RBAC roles or role groups state whether requests receive automatic approval or need manual review. These requests are routed to a set of approvers (a mail-enabled security group). Anyone in the group can approve a request through the Microsoft 365 Admin Center or with PowerShell. Once approved, the requester can execute the task for as long as the approval still is valid (the default is four hours).

In addition to considering who has permissions, you should also take steps to check the use of permissions through auditing. The audit log collects a vast array of events for administrative and user actions. You can use the audit log search in the Compliance portal to examine these events and export them for later analysis should the need arise.

## Managing Exchange Online Mailboxes

Early email systems only supported user mailboxes. Today's email systems are a lot more sophisticated and support many different mail-enabled objects designed for different purposes. Exchange Online supports the following recipient types:

- User mailboxes.
- Shared mailboxes.
- Room and resource mailboxes.
- Discovery mailboxes.
- Distribution lists.
- Public folders and public folder mailboxes.
- Mail contacts and mail users.
- Group mailboxes (used by Groups and Teams).
- Scheduling mailbox (for internal use only).

Exchange Online uses but does not allow tenant administrators access to system mailboxes. These mailboxes include arbitration mailboxes, such as that used to generate the Offline Address Book, the health mailboxes used by Managed Availability probes, and mailboxes created for system test purposes. On the other hand, the Exchange Online mailboxes you can manage have some unique characteristics not found on-premises, many of which we will meet here.

### The Link Between EXODS and Azure AD

Mail-enabled objects, including user mailboxes and groups, exist in both the Exchange Online directory service (EXODS) and Azure AD. This arrangement allows Exchange to control some extra properties for mail-enabled objects over and above the set available in Azure AD, but it also means that many attribute changes have to be written twice—once to each copy. Background processes synchronize EXODS and Azure AD. These processes use an identifier called *ExternalDirectoryObjectId* stamped on EXODS objects to link to Azure AD. The identifier does not feature in on-premises Exchange, but its importance is high enough for Exchange Online to display it as a default property when you run the *Get-ExoMailbox* cmdlet:

```
[PS] C:\> Get-ExoMailbox -Identity TRedmond

ExternalDirectoryObjectId : eff4cd58-1bb8-4899-94de-795f656b4a18
UserPrincipalName         : Tony.Redmond@office365itpros.com
Alias                     : Tony.Redmond
DisplayName                : Tony Redmond
```

Exchange cmdlets accept the *ExternalDirectoryObjectId* as a valid identity. In other words, this works:

```
[PS] C:\> Get-ExoMailbox -Identity eff4cd58-1bb8-4899-94de-795f656b4a18
```

*ExternalDirectoryObjectId* can also be used with cmdlets from other PowerShell modules. For example, this snippet retrieves the *ExternalDirectoryObjectId* for a mailbox and uses it to fetch information about the Azure AD user account to which the mailbox belongs.

```
[PS] C:\> $ObjectId = (Get-ExoMailbox -Identity Kim.Akers).ExternalDirectoryObjectId
Get-MgUser -UserId $ObjectId
```



The same technique works for Groups. In this example, we use the *ExternalDirectoryObjectId* to list the members of a group.

```
[PS] C:\> $ObjectId = (Get-UnifiedGroup -Identity ExchangeGoms).ExternalDirectoryObjectId
[array]$Members = Get-MgGroupMember -GroupId $ObjectId
ForEach ($Member in $Members) {Get-MgUser -UserId $Member.Id | Select -ExpandProperty DisplayName}
```

When an *ExternalDirectoryObjectId* is unavailable for an Exchange object, it means that the object is specific to Exchange and doesn't exist in Azure AD. Discovery mailboxes used to hold the results of old-style Exchange eDiscovery searches are an example of such objects.

## User Mailboxes

An on-premises Exchange administrator who begins working with Exchange Online mailboxes won't notice much difference between cloud mailboxes and their on-premises counterparts. You can't manage databases or move mailboxes around because Microsoft takes care of these activities, but the essentials of managing mailbox properties are similar. Many of the techniques used to work with mailboxes through the EAC or PowerShell are the same on both platforms.

Some features available on-premises are not in Exchange Online. The cmdlet extension agent is an example. This agent often runs in on-premises deployments to automate the population of mailbox properties during the creation of a new mailbox. For instance, you might set the time zone and language for a mailbox so that its owner has a seamless introduction to OWA the first time they access their mailbox.

Because Microsoft 365 is a massive multi-tenant infrastructure, it is logical that Microsoft imposes some throttles and controls over the resources that an individual user can consume, the size of mailboxes, and the volume and type of messages that the system handles. These limits exist to protect the integrity and performance of Exchange Online and are [documented online](#). Microsoft reviews the limits regularly and updates them in line with experience and user demand. You should acquaint yourself with these limits as they might influence the details of your deployment.

**Data Caching:** Exchange Online is a very different environment to an Exchange on-premises deployment, so it would be unreasonable to expect that everything will work the same. Caching is an example. Sometimes a change that you make to an object takes a few seconds – or even a few minutes – to be effective throughout Exchange Online. It might even take some time for a new object to appear because of the need for synchronization across different parts of Microsoft 365. This is quite normal and is a side-effect of the caching of data to improve performance and responsiveness within the service. The change or new object will appear eventually. Just have faith!

## Creating a New Mailbox

Because cloud mailboxes must have a Microsoft 365 account, the EAC does not include the ability to create a new mailbox. You can:

- Create a new account through the Microsoft 365 admin center and assign it an Exchange Online license. The mailbox is available a few moments after you create the account.
- Create mailboxes on-premises and synchronize them with Azure AD using a tool like AADConnect (hybrid mailboxes).
- Create a new account and mailbox through PowerShell.

Apart from instructing Exchange to create a mailbox for an account, the Exchange Online license assigned to an account controls some mailbox settings, such as its storage quota and access to other products such as SharePoint Online. [More information](#) is available online about the consequences of assigning or removing licenses to or from accounts. If you create accounts with PowerShell, you must make sure that the new

accounts are fully provisioned and licensed. Accounts created without a license will not be able to access any applications until they become licensed.

Once created, you can change the settings for Exchange Online mailboxes through the Microsoft 365 admin center, EAC, or PowerShell. Management of the settings for hybrid mailboxes always happens through the on-premises environment. One major difference between on-premises and cloud mailboxes is that Exchange Online does not apply naming policies to new objects. In practical terms, this means that you cannot control the format of display names. If your company organizes address lists using the last name, and first name convention, you must input the display name according to that policy or run a fix-up script afterward to ensure that all mailboxes follow the same naming convention.

**Display Names and Avatars:** There are many places in Microsoft 365 applications where user photos (avatars) are displayed, if available in the user account. If not, initials taken from the display name serve as a fallback. For instance, the avatar for the user Paul Robichaux will display PR. However, if you use the last name and first name convention for display names, the avatar will display RP. The apparently “wrong” choice of initials can be hard for users to understand.

## Editing a Cloud Mailbox

Many of the properties of Exchange Online mailboxes, including permissions, forwarding settings, and visibility in the GAL, can be set through the Microsoft 365 admin center. If you need access to the full set of mailbox properties, click the **Edit Exchange properties** link in the properties dialog to open the EAC mailbox properties page. At this point in the evolution of Exchange Online, you might rarely if ever go to the classic or modern EAC for anything, as most of the things you might commonly manage are available elsewhere.

**A note on mailbox aliases:** As its name implies, an alias is another name or way for Exchange to recognize a mailbox. Microsoft’s help suggests that *“The user’s alias is the portion of the email address on the left side of the @ symbol. It must be unique in your organization.”* However, although the advice is to be unique, Exchange Online currently allows you to create multiple mail-enabled objects with the same alias. This is a habit to be avoided because it can create confusion when processing objects, especially with PowerShell. Microsoft [announced that this will change](#), but they pushed back the original April 2022 implementation date—it is now due to be service-wide by the end of July 2022 but there haven’t been any updates to that date in a while.

## Creating a Mailbox with PowerShell

PowerShell is often used to script the creation of new user accounts and mailboxes because it allows companies to tie the creation of an account into other processes, such as the creation of a new HR record and access records, the printing of an employee badge, and so on. On-premises administrators are familiar with the concept of creating and updating mailboxes with PowerShell because the *New-Mailbox* cmdlet has been used for this purpose for well over ten years. The *New-Mailbox* cmdlet exists in Exchange Online but the environment in which it functions is very different, largely because of the multi-tenant nature of Microsoft 365 and the need to license users before they can access functionality. Exchange Online strictly enforces the need for mailboxes to be licensed and will remove unlicensed mailboxes if they don’t get a license within 30 days of creation. Thus, you cannot take scripts used to create mailboxes on-premises and expect to be able to use them with Exchange Online. Invariably, some adjustment is necessary, if only to assign a license to a new mailbox immediately after its creation.

One point of difference you need to be aware of is the difference between a mailbox and a remote mailbox. You know what a mailbox is: it’s homed either in the cloud or on-premises, in the same location as its

associated account object. A remote mailbox is a cloud mailbox associated a mail user object stored in an on-premises Active Directory. You'll create remote mailboxes when you want your users anchored in on-premises AD but their mailboxes in Exchange Online. Most likely, this will already have been done for you as part of your migration to Exchange Online.

It is not the intention to discuss how to use PowerShell to create Exchange Online mailboxes or how to write a bulk mailbox creation process (the Microsoft 365 admin center includes [an option for bulk account creation](#)). A scan of the Internet will provide many examples of code that you can examine and repurpose to suit your needs, including some from Microsoft (for instance, [a script to create multiple mailboxes](#) is available as is one to [assign licenses to multiple mailboxes](#)). Instead, this walkthrough will help you understand the differences that exist between Exchange on-premises and Exchange Online.

Remember that if you run a hybrid deployment, the on-premises environment is always the master, and Azure AD is the replica. Mailboxes and user accounts are created on-premises and then synchronized to Azure AD. However, even if you run a hybrid deployment and will never create cloud-based mailboxes, it's good to understand what happens when a new cloud mailbox is created from scratch. The basic steps in the process are as follows:

1. Make sure that your PowerShell session loads the necessary modules. You need both the Exchange Online and Azure AD cmdlets to create mailboxes and manipulate the underlying Azure AD objects. See Chapter 23 if you need more background on this.
2. Run the *New-RemoteMailbox* cmdlet to create a new user account and its Exchange Online mailbox.
3. Run the *Set-User* cmdlet to update the on-premises directory with the organizational and personal settings for the user object.
4. Run the *Update-MgUser* cmdlet to set the correct country code (location) for the new mailbox in the cloud. You assign a country to a user account to ensure that Microsoft 365 makes the services designated for that country available to the user. For example:

```
[PS] C:\> Update-MgUser -UserId Kim.Akers@Office365itpros.com -Country Germany
```

5. Assign a license to the user.

Creating a new mailbox with the *New-Mailbox* cmdlet on an on-premises Exchange server usually completes in a matter of seconds. The same is not true for Exchange Online. Although the cmdlet and syntax are the same, the creation of the new object across EXODS and Azure AD takes some time to synchronize across the directories.

**Generating passwords with PowerShell:** The need to use secure passwords that satisfy policy exists when you create new mailboxes. You can find lots of ways to approach the task of generating a secure random password in PowerShell on the web. Always be careful to test and check the code in any PowerShell scripts that you download from the web before using them in production.

## Updating Mailbox Attributes

Some of the attributes commonly populated for mailboxes are not set with the *New-Mailbox* or *Set-Mailbox* cmdlets. These are attributes controlled by Azure AD and are common to all Exchange recipient objects, whether or not they have a mailbox, and are updated with the *Set-User* cmdlet. As it happens, many of these attributes are useful when creating filters to create dynamic distribution lists or address lists, so it is important to give some attention to making sure that they hold the correct values. For example, the code shown below updates the user object for the mailbox that we just created with the organizational and personal information that you might expect to find in the corporate directory. The example shown here updates several properties, including the *Manager* property with a value that points to the name of the new user's direct manager. The update will fail if Azure AD cannot find the manager.

```
[PS] C:\> Set-User -Identity "Kim Akers" -City "Dublin" -CountryOrRegion "Ireland" -Department
"Marketing Operations" -Title "VP Marketing (New Products)" -Manager "Paul.Robichaux" -Office
"Dublin HQ" -Company "Popular Books"
```

Once the new mailbox exists, you can update its settings with *Set-Mailbox* and other cmdlets. For instance, here is how to enable an archive mailbox with the *Enable-Mailbox* cmdlet.

```
[PS] C:\> Enable-Mailbox -Identity "Kim Akers" -Archive
```

## Add User Photos to Mailboxes

Given the graphic nature of applications, it's a good idea to update a new mailbox with a suitable photo after it's created. The photo then shows up in the GAL, the people card, and apps like Teams which display user avatars. Several PowerShell cmdlets are available to administrators to update user photos.

- The Exchange Online *Set-UserPhoto* cmdlet updates the photo data in a mailbox. *Set-UserPhoto* can also update a photo for a group mailbox (be sure to specify the *GroupMailbox* switch). You cannot use *Set-UserPhoto* to update other mail-enabled objects, like distribution lists or mail contacts.
- The Teams *Set-TeamPicture* cmdlet updates the photo data for a team. This is analogous to running *Set-UserPhoto* to update the photo for a group mailbox. In most cases, it's best to use *Set-UserPhoto* to avoid the need to load another module.
- The Azure AD *New-MgUserPhoto* and *Update-MgUserPhoto* cmdlets writes photo data to an Azure AD user account. Use this cmdlet when you wish to update photo data for an Azure AD account that doesn't have an Exchange Online mailbox, like a guest account. As the cmdlet name suggests, the cmdlet processes thumbnail (small) photos. It does not generate the larger size photos which look better in Teams meetings. For this reason, you should always use *Set-UserPhoto* to upload photos for tenant accounts.

Exchange Online and Azure AD synchronize photo data to make sure that user accounts have the latest picture. After a short delay to allow the apps to refresh their caches, an updated photo will be active across the ecosystem.

Image files for user photos can be in JPEG or PNG format and should be 648 x 648 pixels. This is the largest resolution supported. Behind the scenes, Exchange Online generates smaller 64 x 64 and 96 x 96-pixel thumbnails for apps to use when small thumbnails are appropriate. Most digital photos are much larger (in pixels) so some resizing is needed. Square photos are best as they won't be cropped. Usually, the best results are obtained when the user faces directly into the camera. The photos must also be less than 500KB in size.

Although it can take 30 seconds or more to update a picture for a mailbox, running *Set-UserPhoto* is simple:

```
[PS] C:\> Set-UserPhoto -Identity James.Smith@office365itpros.com -PictureData
([System.IO.File]::ReadAllBytes("c:\Temp\James.Smith.jpg")) -Confirm:$False
```

If you want to check if a mailbox already has a picture (to avoid overwriting it), use the *Get-UserPhoto* cmdlet. This cmdlet returns *\$Null* if a mailbox does not hold photo data.

```
[PS] C:\> If (Get-UserPhoto -Identity Kim.Akers@Office365ITpros.com) {Write-Host "Kim has a photo"}
```

If you make a mistake and upload the wrong image, you can restart by removing the image with the *Remove-UserPhoto* cmdlet:

```
[PS] C:\> Remove-UserPhoto -Identity James.Smith@office365itpros.com -Confirm:$False
```

An example of how to scan user mailboxes to [find mailboxes without photos can be downloaded from GitHub](#).

## Allow Users to Update Their Photos

Exchange Online, Teams, and the Office portal give users the ability for users to update their photos, but this capability is controlled through the *SetPhotoEnabled* setting in the OWA mailbox policy assigned to the mailbox. The use of OWA mailbox policies means that organizations can decide to allow some users to manage their photos while barring others from doing so. In the latter case, the organization takes responsibility for updating user photos by:

- Building a connection to a system holding suitable photos, like an HR system.
- Building a special app to manage user photos (although Microsoft is deprecating the relevant Exchange Web Services APIs).
- Using a commercial tool, like [Photos for Office 365 from Code Two Software](#).

By default, the *SetPhotoEnabled* setting is *\$true*, meaning that users can upload a photo from apps. If this setting is off (as it may be in a legacy tenant) users will see a message such as “*picture options are disabled by policy*” if they try to change their photo. To allow users to upload and update their photos, either:

- Update the OWA mailbox policies so that *SetPhotoEnabled* is *\$True* in all policies, or:
- Create or update an OWA mailbox policy with *SetPhotoEnabled* set to *\$True* and assign this policy to the mailboxes of accounts you want to allow to upload photos.

For example, to update an OWA mailbox policy, run the *Set-OWAMailboxPolicy* cmdlet:

```
[PS] C:\> Set-OWAMailboxPolicy -Identity OWAFullAccess -SetPhotoEnabled $True
```

To assign an OWA mailbox policy to a mailbox, use the *Set-CASMailbox* cmdlet:

```
[PS] C:\> Set-CASMailbox -Identity Chris.Bishop -OWAMailboxPolicy OWAFullAccess
```

Changes to an OWA mailbox policy take up to 30 minutes before they are effective.

## Mailbox Plans

When you create a new Exchange Online mailbox, the new mailbox inherits many of its settings from a mailbox plan. Four mailbox plans are available within a tenant to accommodate the different Exchange Online plans included in Microsoft 365 and Office 365 products. To see the set of mailbox plans, run the *Get-MailboxPlan* cmdlet:

```
[PS] C:\> Get-MailboxPlan | Format-Table DisplayName, IsDefault, Name
```

DisplayName	IsDefault	Name
ExchangeOnlineEnterprise	True	ExchangeOnlineEnterprise-8fc1c029-5e32-485e-9810-179fb4701447
ExchangeOnlineDeskless	False	ExchangeOnlineDeskless-bc1e76cc-4c0b-491c-a518-3a0a43cbf78e
ExchangeOnline	False	ExchangeOnline-12c139bc-eafa-4a43-b4d2-e285f83e075d
ExchangeOnlineEssentials	False	ExchangeOnlineEssentials-1a1bf516-90d5-4c4b-a047-5b3544ad9826

The role of the mailbox plan is to be a template holding settings for mailbox properties. When you create a new mailbox, the new mailbox inherits settings from the mailbox plan chosen by Exchange Online. Most mailboxes are created along with new accounts via the Microsoft 365 admin center. When this happens, Exchange Online uses the license assigned to the account to select the mailbox plan to apply to the new mailbox. Table 5-1 lists the Office 365 and Microsoft products and the associated mailbox plans.

Products	Mailbox Plan
Exchange Online Kiosk Microsoft 365 F3 Office 365 F3	<i>ExchangeOnlineDeskless</i>

Exchange Online Plan 1 Microsoft 365 E1 Office 365 E1	<i>ExchangeOnline</i>
Exchange Online Plan 2 Microsoft 365 E3/E5 Office 365 E3/E5	<i>ExchangeOnlineEnterprise</i>
Microsoft 365 Business Basic	<i>ExchangeOnlineEssentials</i>

Table 5-1: Licenses and Mailbox plans

In the output for the *Get-MailboxPlan* cmdlet shown above, the Exchange Online Enterprise plan is marked as the default. If you create a user mailbox without a license, Exchange Online uses the default plan to populate its settings. Mailboxes that don't need licenses, like shared and resource mailboxes, use the Exchange Online mailbox plan. An administrator can specify the mailbox plan to use when creating a new mailbox with the *New-Mailbox* cmdlet.

To check how many mailboxes have each mailbox plan, we can check the plan registered for each mailbox. Note that the filter used to find mailboxes requires the distinguished name for the mailbox plan.

```
[PS] C:\> $Report = [System.Collections.Generic.List[Object]]::new()
$MbxPlans = Get-MailboxPlan
ForEach ($Plan in $MbxPlans) {
    $Dn = (Get-MailboxPlan -Identity $Plan.Name).DistinguishedName
    [Array]$Mbx = Get-ExoMailbox -Filter "MailboxPlan -eq '$Dn'" -Properties MailboxPlan -ResultSize
    Unlimited # Find mailboxes with the plan
    If ($Mbx) {
        ForEach ($M in $Mbx) {
            $ReportLine = [PSCustomObject][Ordered]@{
                Name           = $M.DisplayName
                UPN             = $M.UserPrincipalName
                Plan            = $Plan.DisplayName }
            $Report.Add($ReportLine) }
        }
    } #End ForEach
$Report | Group Plan | Format-Table Name, Count
```

Name	Count
ExchangeOnlineEnterprise	43
ExchangeOnline	17

The *Set-MailboxPlan* cmdlet configures settings in mailbox plans while the *Get-MailboxPlan* cmdlet reports the settings. Because the idea behind mailbox plans is to configure basic mailbox settings, not every property configurable with the *Set-Mailbox* cmdlet is available in a mailbox plan. The settings cover:

- Mailbox quotas and warning thresholds.
- Message send and receive size.
- Deleted items retention period.
- Mailbox retention policy.
- User role assignment policy.

In this example, we use *Set-MailboxPlan* to update the Exchange Online enterprise plan to update the largest supported message size for send and receive to 125 MB, change the deleted item retention period from 14 to 30 days, and assign a new default mailbox retention policy.

```
[PS] C:\> Set-MailboxPlan -Identity ExchangeOnlineEnterprise -MaxSendSize 125MB
-MaxReceiveSize 125MB -RetainDeletedItemsFor 30.00:00:00 -RetentionPolicy "General Mailbox
Retention Policy"
```

Somewhat frustratingly, although *Get-MailboxPlan* returns a large set of mailbox properties and values, *Set-MailboxPlan* is unable to update most settings. If you want to update a mailbox property outside the set supported by mailbox plans, you must run *Set-Mailbox* after creating the mailbox. For instance, you might want to write a value into one of the custom attributes.

Modifying the settings of a mailbox plan does not affect existing mailboxes. If you want to change settings for existing mailboxes, you'll need to run the *Set-Mailbox* or *Set-CASMailbox cmdlets*. However, if the license assigned to a user mailbox changes, Exchange Online applies the settings for the relevant plan to the mailbox (this doesn't happen immediately as it takes some time for Exchange to detect and react to the license change).

Each mailbox plan has a corresponding CAS mailbox plan. This mimics the relationship between *Set-Mailbox* and *Set-CasMailbox* where the first cmdlet updates essential mailbox settings while the second deals with connectivity. In this instance, the *Set-CASMailboxPlan* cmdlet allows administrators to control the following settings.

- Enabling Exchange ActiveSync.
- Enabling IMAP4 and POP3.
- OWA mailbox policy.

Here's an example of disabling the POP3 and IMAP4 protocols in all mailbox plans:

```
[PS] C:\> Get-CASMailboxPlan | Set-CASMailboxPlan -PopEnabled $False -IMAPEnabled $False
```

You can check the protocol settings by running the *Get-CASMailboxPlan* cmdlet to return the different protocol settings:

```
[PS] C:\> Get-CASMailboxPlan -Identity ExchangeOnlineEnterprise | Format-List DisplayName,
ImapEnabled, PopEnabled, MapiEnabled, ActiveSyncEnabled, OwaEnabled, OutlookMobileEnabled
```

```
Name                : ExchangeOnlineEnterprise
ImapEnabled         : False
PopEnabled          : False
MAPIEnabled         : True
ActiveSyncEnabled   : True
OWAEnabled          : True
OutlookMobileEnabled : True
```

## Recipient Limits

The default recipient limit for an Exchange Online mailbox is 500. This means that the mailbox owner can send messages addressed to up to 500 recipients. The limit exists to ensure that Exchange Online mailboxes do not consume large quantities of resources by sending messages to large numbers of recipients. You can update the recipient limit for a mailbox to anything from between 1 to 1,000 through the EAC or with PowerShell. For example, this command sets the limit to 900 for the chosen mailbox:

```
[PS] C:\> Set-Mailbox -Identity James.Ryan -RecipientLimits 900
```

To set a new recipient limit for every user mailbox in the tenant, use a command like this:

```
[PS] C:\> Get-ExoMailbox -RecipientTypeDetails UserMailbox -ResultSize Unlimited | Set-Mailbox -
RecipientLimits 900
```

Although the usual request is to increase the recipient limit to deal with situations such as the distribution of internal newsletters, a use case exists to reduce the limit to restrict the ability of specific users to communicate with large numbers of recipients.

If you want a new default recipient limit to apply for new mailboxes, you can update the setting in the appropriate mailbox plans. For example, this command updates the value for enterprise mailboxes:

```
[PS] C:\> Set-MailboxPlan -Identity ExchangeOnlineEnterprise-8fc1c029-5e32-485e-9810-179fb4701447 -RecipientLimits 750
```

## Multi-Geo Mailboxes

When Exchange is configured for multi-geo operation, a tenant can distribute mailboxes across the home region (for example, the United States) and one or more satellite geo locations (for example, the United Kingdom, France, and Norway), depending on where users are and the need to satisfy local data sovereignty. In multi-geo tenants, Azure AD user accounts have a *PreferredDataLocation* property to record the data center region where a user's data should be stored. Exchange Online synchronizes the value of *PreferredDataLocation* into the *MailboxRegion* property for mailbox objects stored in EXODS. This property is blank for mailboxes in single-geo tenants. Exchange Online uses *MailboxRegion* to know where user mailboxes and archive mailboxes are stored (both primary and archive mailboxes must be in the same region).

## Configuring a Mailbox-Enabled Account

Creating a new mailbox with the *New-Mailbox* cmdlet also creates a new user object in Azure AD. Creating a new mailbox doesn't assign a license and until the account is assigned a license for Exchange Online, the user won't be able to sign into the mailbox. To allow sign-on, we must update the account with a Microsoft 365 location code for the country where the account is located and assign a license. A license must be assigned to a new account within 30 days to avoid it being blocked and placed into Microsoft 365's automatic deletion routine.

The Microsoft 365 admin center shows the country to which a user is assigned. You can use the same country names with PowerShell, but the country code can also be passed – FR for France, DE for Germany, IE for Ireland, and so on. A tenant can accommodate users in multiple countries, and you can update a user's location as many times as you like, and the multi-geo capability available to large tenants allow them to distribute mailboxes across multiple data center regions. You can't update a tenant's location after it is created because the location determines factors such as billing and the data centers used to hold user data.

Two country properties exist for user accounts. The *Country* property is used for display and sort purposes. You can assign anything you like to the *Country* property for an account because the *Update-MgUser* cmdlet does not validate the input to ensure that the country exists. However, because the *Country* property is used for sort and display purposes, including the filters used for objects like dynamic groups, these values must be accurate and consistent.

The *UsageLocation* property, which uses the 2-character [ISO 3166 code for a country](#), is much more important because it controls the services that can be provided to the account. For this reason, *Update-MgUser* checks input to ensure that a valid country code is assigned to an account. Instances exist where certain functionality is unavailable in a specific country, as in the case of a voice calling plan. It is therefore essential to be accurate when you set the *UsageLocation* value for an account. Here's an example of setting the two country properties:

```
[PS] C:\> Update-MgUser -ObjectId Kim.Akers@Office365ITPros.com -Country "Ireland" -UsageLocation IE
```

As explained previously, it can take some time before updates to the new account synchronize across all workloads. You need to take this factor into account if you plan to script the creation of new mailboxes. After we update the user object with a valid location code, we can assign it a license. We do this by running the *Set-MgUserLicense* cmdlet. See Chapter 23 for information about how to manipulate licenses through PowerShell, including how to perform this task using the Azure AD cmdlets.



After you assign a license, we should have a licensed user with a functioning mailbox. You can confirm the assignment of the license to the user by examining their account through the Microsoft 365 admin center. It is easy to assign licenses to accounts through the Microsoft 365 admin center, especially when dealing with the finer points of turning on or off different services licensed by a plan. However, it is often more efficient to use PowerShell when you want to check the licenses that are already assigned to accounts or to reallocate licenses to accounts, as might be the case if you choose to upgrade a plan for selected users.

## Email Primary and Proxy Addresses

Mail-enabled objects have one or more addresses that allow the Exchange transport service to route messages to them. Exchange Online mailboxes usually have several SMTP proxy addresses (also known as email aliases), any of which can be used to send email to the mailbox. A mailbox always has an address for the onmicrosoft.com service domain used by the tenant. Administrators can assign other addresses to mailboxes for any of the domains owned by the tenant. The set of addresses is stored in the *EmailAddresses* attribute for mail-enabled objects. For instance, here's a typical set for a mailbox:

```
[PS] C:\> Get-ExoMailbox Jane.Sixsmith | Select -ExpandProperty EmailAddresses
smtp:Jane.Sixsmith+Amazon@office365itpros.com
SMTP:Janey@office365itpros.com
smtp:Jane.Sixsmith@office365itpros.com
SPO:SPO_edc2687c-5939-41d1-9ab1-24a7dc43ac6e@SPO_b662313f-14fc-43a2-9a7a-d2e27f4f3478
SIP:Jane.Sixsmith@office365itpros.com
smtp:Jane.Sixsmith@office365itpros.onmicrosoft.com
```

Each address is of a certain type. In this case, there are SMTP addresses, an address used by SPO (SharePoint Online) to synchronize information about documents created by the mailbox's owner with the Microsoft 365 substrate. You should not remove this address as internal Microsoft processes manage its creation and removal. The SIP address is used for SIP messaging communications, such as Teams calls.

One of the four SMTP addresses has a capitalized prefix. This is the primary address, which means that it is the one Exchange inserts as the reply address for outbound messages so that recipients use the address when they reply to messages. It is best practice to use the same address for both the User Principal Name and primary SMTP address for a Microsoft 365 account. For example, to set the primary SMTP address for a mailbox with PowerShell, run the *Set-Mailbox* cmdlet to set the *WindowsEmailAddress* attribute.

```
[PS] C:\> Set-Mailbox -Identity Jimmy.Jones -WindowsEmailAddress Jimmy@office365itpros.com
```

If the address selected as the primary is not already present in the mailbox properties, Exchange will add the address and make it the primary.

Except for plus addresses (see below), only administrators can assign multiple email proxy addresses to mailboxes, distribution groups, shared mailboxes, group mailboxes, mail contacts, and mail users using tools including the Microsoft 365 admin center, EAC, and PowerShell.

### Sending Email Using Proxy Addresses

Exchange Online supports the ability for users to send email using any of the secondary SMTP proxy addresses (otherwise known as email aliases) assigned to their mailbox. This is especially useful when a tenant supports multiple domains, as in the case of a corporate merger, and users need to send messages using proxy addresses for different domains. To allow users to send email using proxy addresses, update the Exchange organization configuration using PowerShell (below) or through the Mail flow settings in the EAC:

```
[PS] C:\> Set-OrganizationConfig -SendFromAliasEnabled $True
```

After updating the configuration, it can take several hours for all the mailbox servers used by the tenant to receive the change. Once this process completes, users can use the From field in the message compose form to select and use a proxy address. OWA users must select the proxy addresses they wish to use through the

Compose and Reply section of OWA settings to make addresses available for use. Outlook desktop and Outlook mobile users can insert proxy addresses in the From field for new messages. After sending, the proxy address appears for recipients in place of the sender's primary SMTP address.

## Plus Addressing

An SMTP address is composed of a local part and a domain. The domain routes messages to the server identified in the domain's MX record in DNS. The local part is the user address used to identify the eventual recipient. Exchange Online supports plus addressing, meaning that users can add a suffix (an arbitrary tag chosen by them) to the local part of their SMTP address. A plus sign divides the suffix from the local part. For example:

*Kim.Akers@Office365itpros.com* is a "normal" SMTP address.

*Kim.Akers+SomeValue@Office365itpros.com* is an SMTP address with a plus suffix.

When the transport service processes an inbound email with a plus address, it removes the plus sign and suffix and uses the remainder to deliver the message.

The idea behind plus addresses is that you can use them to find out if companies are sharing (or selling) your email address for marketing (or spamming) purposes. For instance, if a website asks for an email address before granting access to some content, you can create a plus address and use the name of the website as the suffix. If you later find that spam or other unwanted email arrives using that plus address, you know that the site shared your address. You can then use inbox rules to filter or block messages from that address.

Consumer mail systems like Outlook.com and Gmail support plus addressing. You can use plus addressing in two ways:

- **Administrator controlled:** The tenant assigns plus addresses as SMTP proxy addresses to mailboxes. These addresses are persistent like any other SMTP proxy address, which means that they can be used with features like sending using a proxy address described above. Administrators can assign plus addresses to mailboxes and groups through the EAC or PowerShell, but not through the Microsoft 365 admin center.
- **User-initiated:** Users can create their own plus addresses as needed by adding the plus sign followed by whatever text they want to use as a suffix to their regular email address when they give email addresses to other organizations. For instance, if doing business with Contoso, you could tell Contoso to send you an email at *First.Last+Contoso@tenant.com*.

Currently, the *AllowPlusAddressInRecipients* setting in the Exchange organization configuration controls how Exchange Online processes plus signs found in email addresses. The value can be:

- **\$True:** The Exchange transport service uses the plus sign to indicate that it can remove the tag after a plus sign and deliver the message using the remaining address. For example, Exchange will take the address *Tony.Redmond+eCommerce@office365itpros.com* and deliver the message to *Tony.Redmond@office365itpros.com*. This is the default for any new tenant.
- **\$False:** This is the default for older tenants where the possibility exists that valid email addresses exist containing the plus sign. This setting means that Exchange treats the plus sign as a literal character that's part of the email address and will only deliver messages if it can find a match in the proxy addresses assigned to a recipient. For example, if an administrator assigns *Tony.Redmond+eCommerce@office365itpros.com* as a proxy SMTP address to a mailbox, Exchange can deliver the message to that mailbox. It will not strip the "+eCommerce" portion off and attempt to deliver to [Tony.Redmond@office365itpros.com](mailto:Tony.Redmond@office365itpros.com).

The *AllowPlusAddressInRecipients* organization setting must be set to True to allow users to create their own plus addresses as described above. You can update the setting through the Mail flow settings in the EAC or by running the *Set-OrganizationConfig* cmdlet:

```
[PS] C:\> Set-OrganizationConfig -AllowPlusAddressInRecipients $True
```

When plus addressing is active, any intermediate gateway which processes inbound email must focus on the domain for routing. If the gateway attempts to do directory lookups to check recipient addresses and cannot handle plus addresses, the lookups might fail and cause the rejection of messages.

In April 2022, Microsoft removed the organization setting and make plus addressing available by default in all tenants. Any proxy addresses containing plus characters remaining for mail-enabled objects will not work as before because Exchange Online will no longer attempt to match the complete address (including the plus sign and tag) on inbound messages against its directory. This could lead to the non-delivery of email. If your organization wishes to opt out of plus addressing, run the command:

```
[PS] C:\> Set-OrganizationConfig -DisablePlusAddressInRecipients $True
```

See [this article](#) for details of PowerShell code to locate and remove plus addresses from mail-enabled recipients.

## Custom Attributes

Exchange Online supports two sets of custom attributes or properties for mail-enabled objects. Tenants can update these attributes as they wish to support their own needs. For example, you could use a custom attribute to store a staff number or an office location code. Note that these custom attributes are properties of the mailbox object associated with a user.

The two sets are:

- *CustomAttribute1* through *CustomAttribute15*: These attributes can store a single value.
- *ExtensionCustomAttribute1* through *ExtensionCustomAttribute5*: These attributes can store multiple values.

Apart from their numbers, the difference between the fifteen custom attributes and the five extension attributes is that the custom attributes can be updated through EAC. Extension attributes are only accessible through PowerShell.

To illustrate the difference between the two types of custom attributes, if you run this command:

```
[PS] C:\> Set-Mailbox -Identity "Kim Akers" -CustomAttribute5 "Custom", "Attribute", "5"
```

Exchange Online stores the string "Custom Attribute 5" in the attribute. But if you run this command:

```
[PS] C:\> Set-Mailbox -Identity "Kim Akers" -ExtensionCustomAttribute5 "Custom", "Attribute", "5"
```

Exchange Online returns an array of "5, Attribute, Custom." You can therefore do this to fetch the attribute and put it into an array variable:

```
[PS] C:\> $Values = Get-ExoMailbox -Identity Kim.Akers | Select -ExpandProperty  
ExtensionCustomAttribute5
```

And because it is an array rather than a string, we can use the information contained in the attribute differently. For example, to add an item to the array:

```
[PS] C:\> Set-Mailbox -Identity "Kim Akers" -ExtensionCustomAttribute5 @{Add="Great"}
```

Or to remove an item:

```
[PS] C:\> Set-Mailbox -Identity "Kim Akers" -ExtensionCustomAttribute5 @{Remove="Attribute"}
```

The custom attributes (but not the extension attributes) support server-side filtering for fast access. Thus, we can do this:

```
[PS] C:\> Get-EXOMailbox -Filter {CustomAttribute1 -ne $Null} -Properties CustomAttribute1 |
Format-Table DisplayName, CustomAttribute1
```

One note of caution. If you are synchronizing objects from an on-premises Active Directory and use these attributes, make sure to include the attributes in the synchronization set.

## Administrator Access to User Mailbox Settings

Occasionally, users need help from an administrator to update mailbox settings. Two methods are available:

1. Click the user avatar in the top right-hand corner of EAC and select **View another mailbox** from the menu. EAC displays a picker dialog to select the mailbox that you want to manage. The selected user's settings open in a browser window using a modified form of OWA Options. The user making the changes and the mailbox to which the changes are applied are clearly shown. Apart from needing to be a member of at least the Recipient Management role group, no special access rights need to be granted to the administrator to allow them to change settings for a user in this way. It is also possible to call up OWA Options directly by inputting a link to take you to the user mailbox. For example:

<https://outlook.office365.com/ecp/Frank.Clonan@Office365ITPros.com>

2. Through PowerShell. Exchange Online supports a set of cmdlets to allow administrators to manipulate different personal mailbox settings, or options usually selected by users. You can use these cmdlets to configure a mailbox on behalf of a user or to check that certain settings are in place. An account that performs maintenance on user accounts through PowerShell should be assigned the Recipient Management role. If the need exists to view the data in the mailbox, the account needs *Full Access* permission.

Using a browser to change user settings is easiest when you only have one or two mailboxes to update. PowerShell is the preferred method when changes need to be applied to multiple mailboxes, or when you want to update mailbox settings in a script that creates a new mailbox.

Administrators whose accounts are granted the RBAC User Options role can manage settings for any user mailbox. This example grants Tony Redmond full rights over the mailbox owned by Kim Akers.

```
[PS] C:\> Add-MailboxPermission -Identity 'Kim Akers' -User 'Tony Redmond' -AccessRights FullAccess
-AutoMapping $False
```

Note the use of the *AutoMapping* parameter. If omitted, Autodiscover will open the mailbox because it recognizes that the user has full access to the mailbox. As in the situation where you are configuring a mailbox for another user, this is probably not what you want to happen. Setting *AutoMapping* to *\$False* instructs Autodiscover to ignore the mailbox when it builds the list of resources available to Outlook. The user can edit their mailbox settings to open any extra mailboxes that they have access rights to at any time. For more information, see the section on *AutoMapping* later. To remove *FullAccess* rights, run the *Remove-MailboxPermission* cmdlet:

```
[PS] C:\> Remove-MailboxPermission -Identity 'Kim Akers' -User 'Tony Redmond'
-AccessRights FullAccess
```

# Securing the Data of Ex-employees

Over time, some employees will leave the organization. Many departures are expected and planned for, some are amicable, and some will be immediate and potentially traumatic. For legal reasons, you might need to preserve the employee's mailbox in the latter case. In an on-premises environment, you can preserve mailboxes by disabling the user account. If needed, you can reenable the account to restore a mailbox to full running order. The same is true for in the cloud, but you probably do not want to pay the monthly license for an unused account only kept in case the organization needs some of the information in its mailbox in the future. Removing the need for licenses is why inactive mailboxes are so valuable.

When you delete a user account through the Microsoft 365 admin center (Figure 5-7), the following points in the removal workflow need consideration:

- **License:** Cancel the license held by the deleted account or keep it for later reassignment to another account.
- **Mailbox:** Exchange Online removes all the proxy email addresses (primary address and any secondary addresses). In addition, Exchange removes all delegate permissions from the mailbox, but you can grant permission to allow another user to access the mailbox content and retrieve information. As shown in Figure 5-7, the reassignment cannot happen until after the removal of any holds on the mailbox. Alternatively, if the account has an Office 365 E3 or E5 license, you can place a hold on the mailbox before deleting the account. Exchange Online will then make the mailbox inactive (see Chapter 6) and retain it in that state until the hold lapses.
- **OneDrive for Business:** To review and retrieve important information held in the user's OneDrive for Business account, you can grant access to another user. Apart from documents, the account might hold information shared with other users through applications, like Teams meeting recordings, whiteboards, and Loop components. By default, access lasts for 30 days after which OneDrive permanently removes the information, so the person granted access must review and retrieve information in that period. A tenant can increase the retention period for deleted OneDrive accounts by updating the setting in the SharePoint admin center from 30 (the default) up to 3650 days (ten years).

The Microsoft 365 admin center can remove only cloud accounts. On-premises accounts are only manageable with an on-premises administration tool. In addition, the workflow does not address access to information in other workloads. For instance, if the user you want to delete is the sole owner of some Groups or Teams, the deletion process does not make sure that the groups/teams remain with at least one owner. The same is true for traditional SharePoint sites where the deleted user might be the sole administrator. In other words, treat the user removal workflow as something that deals with the user's data held in cloud mailboxes and OneDrive for Business accounts while ignoring the user's activity in other workloads. We'll come back to this topic shortly.

Remember that it might be a good idea to create an autoreply to let senders know that the person is no longer with the organization. This is a separate action that the Microsoft 365 admin center does not perform when removing an account.


Microsoft's account removal process is reasonable and suits the needs of many organizations. However, sometimes you might want to use a different process, especially if some element of discord is present when someone leaves. For this reason, we should discuss the different steps you can take to preserve all the data that a person might have access to when the time comes for them to leave the organization.



## Delete Andy Ruth (Director)

You can restore deleted users and their data, except for calendar items and aliases, for up to 30 days after you delete them. Data on their connected devices will be removed, as well as the following:

 Office 365 E3, Power BI (free), Microsoft Stream Trial will be unassigned and available for other users

Email aliases will be removed   
[View email aliases](#)

Mailbox delegate permissions will be removed   
[View mailbox delegate permissions](#)

Give another user access to Andy Ruth (Director)'s OneDrive files for 30 days after the user is deleted

  Kim Akers 

Give another user access to Andy Ruth (Director)'s email   
Andy Ruth (Director) has been placed on legal hold. Data can't be assigned to other users until the legal hold is removed

Delete user

Figure 5-7: Steps to delete a user account when removed through the Microsoft 365 admin center

## Blocking a User Account

The first and most essential action is to secure access to the user's account to prevent any unauthorized removal of data. We can then put their mailbox and other data into a suitable status for long-term preservation. Normally, when a user authenticates to connect to a Microsoft 365 application, they create a session with that application. The session receives an access token and a refresh token from Azure AD. Usually, an access token is valid for an hour. When that period elapses, an automatic reauthentication process kicks in to get a new access token to allow the session to continue. This interaction happens if the refresh token is still valid, and the account credentials are the same. The refresh token has a defined lifetime. This architecture leads to the occasional result that a user with a blocked or deleted account still has a valid refresh token and can still access some workloads. Because each session has a separate refresh token, you might have sometimes noticed that a user might be able to still sign in on one device after a password change or account blockage while other devices disallowed sign-in.

With the introduction of continuous access evaluation (CAE), as described in the Azure AD chapter, individual services can receive notifications that tell them when an account is blocked, has expired, or has critical security events posted against it. Because the client may still present an unexpired, valid-looking token even after events such as password changes, CAE enables the workloads to reject the token based on notification of a password change, account deletion, or similar event. Exchange Online and SharePoint Online may also receive risk events from Azure AD that indicate that the client's tokens shouldn't be trusted.

When the Microsoft 365 admin center blocks an account, it also sets the account's *RefreshTokensValidFromDateTime* property to the current date and time (this is the same as the forced sign-

out option for an account available in the Microsoft 365 admin center). Forcing apps to sign out means that any existing refresh tokens used by apps are invalidated and cannot be renewed. Because the forced sign-out invalidates the refresh tokens, the next time an app attempts to use its refresh token to renew its access, it discovers that the token has expired and so forces the user to reauthenticate. Because the account is blocked, the user cannot reauthenticate. In addition, the CAE mechanism will immediately revoke the refresh and access tokens.

You can manually block sign-ins, initiate forced sign-outs, and reset account passwords from the Microsoft 365 admin center. Any of these actions will flag the account for CAE, meaning that in nearly real-time the workloads will reject any previously issued tokens for that account. To block an account, select the account from the list of active users, edit its settings, and select **Block this user** from the set of icons under the user's name. Blocking an account from signing in through the Microsoft 365 admin center is the right approach when you manage the departure of a single employee. Using PowerShell commands is better when you need to secure a set of accounts quickly or you want to automate the process to secure an account for a departing employee.

To manually block an account using PowerShell, you can do the following:

- Block the account by running `Update-MgUser`. (This will trigger a CAE event, so you could stop here if you like.)
- Force a sign-out from apps by invalidating refresh tokens with a Graph API request. This will also trigger a CAE event.
- Create and set a new password for the account to allow administrator access once the time comes to recover information from apps. Naturally, the new password only works after unblocking the account by running `Update-MgUser -AccountEnabled:$True`. Password changes also trigger CAE events.

```
[PS] C:\> $UserId = (Get-MgUser -UserId Lotte.Vettler@Office365itpros.com).Id
Update-MgUser-UserId $UserId -AccountEnabled:$False
$Uri = "https://graph.microsoft.com/v1.0/users/$($UserId)/microsoft.graph.revokeSignInSessions"
Invoke-MgGraphRequest -Uri $Uri -Method Post
$NewPassword = @{}
$NewPassword["Password"] = "!NewYorkCity2022?"
$NewPassword["ForceChangePasswordNextSignIn"] = $True
Update-MgUser -UserId $UserId -PasswordProfile $NewPassword -AccountEnabled:$True
```

Some consequences flow from disabling a user account, including that Teams removes the disabled account from its membership rosters. While you might expect this to happen for org-wide teams, it often surprises administrators when they discover the removal of disabled accounts from other Teams. A Teams background process performs the removal, and this could take place several days after disabling the account. If an administrator reenables the account, Teams adds the account back to membership rosters. Again, this process can take some time. Other issues exist, as [documented here](#). An alternative to blocking an account is to change its password and force a sign-out from all sessions. The account remains active but is inaccessible unless a person has the new password.

## Permanently Removing User Accounts and Mailboxes

When you delete a user account, it remains in the Azure AD recycle bin for 30 days. You can remove the account permanently (hard deletion) beforehand by:

- Selecting the account in the Deleted users section of the Azure AD admin center and deleting it there.
- Running the `Remove-MgDirectoryObject` cmdlet

Permanent removal means that *no recovery is possible*. Azure AD synchronizes the removal to workload-specific directories, such as EXODS, to ensure the removal of any workload-specific objects such as the user's mailbox. Microsoft cannot recover user objects after permanent removal and cannot recover any workload-

related data either. Remember, Microsoft does not take backups of Azure AD or Exchange Online. It is therefore clear that deciding to make a user account irrecoverable should only happen when you are certain that it's safe to remove the object and its data.

## Putting a User Mailbox on Litigation Hold

After blocking the user account, you should place its mailbox on litigation hold to ensure that no one can remove information from the mailbox. As an example, this command places a mailbox on litigation hold and updates the *RetentionComment* property with details of who applied the hold and the date of its application.

```
[PS] C:\> Set-Mailbox -Identity 'Bad Employee' -LitigationHoldEnabled $True  
-RetentionComment ("Employee Terminated on " + (Get-Date) + " by "  
[System.Security.Principal.WindowsIdentity]::GetCurrent().Name)
```

You can create a core eDiscovery case (Chapter 18) whose only purpose is to manage a hold on mailboxes for ex-employees. Create a hold within the eDiscovery case and add the mailboxes that you want to keep to the hold. You can also add the OneDrive for Business sites belonging to ex-employees to the hold if you want to preserve their contents.

Risk always exists that a soon-to-be-terminated employee might hear of their impending fate in advance and not react well to the news. They might then decide to remove information from their mailbox and other data repositories (shared mailboxes, group mailboxes, Teams conversations, SharePoint Online document libraries, their OneDrive for Business personal site, and so on) before management makes the formal announcement. This is a tricky situation because of the lack of backups for cloud applications, but it is somewhat mitigated by the ability to assign retention policies to SharePoint Online sites to stop people from removing information.

Placing the mailboxes of affected employees on hold as soon as possible after management decides that the employees are leaving will preserve the mailbox contents no matter what steps their owners take to remove information, but it does create the possibility that employees will learn about their impending departure when HR asks IT to place mailboxes on hold. One way to fix this issue is to create a special RBAC role that allows HR representatives to place mailboxes on hold. This will not stop a maliciously-minded individual from erasing data from SharePoint Online or their OneDrive for Business account, but you can ask Microsoft support to help recover information removed from SharePoint using the backups that they take.

## Long-term Mailbox Preservation

Two choices exist as to how to preserve the mailbox for the long term. You can convert it into a shared mailbox or make it an inactive mailbox.

- **Convert to a shared mailbox:** An EAC option exists for this purpose. Converting to a shared mailbox removes the need for a license (unless the new shared mailbox has an archive mailbox, you place the shared mailbox on hold, or assign it a larger quota than the default 50 GB) and keeps the mailbox contents online so that whoever needs to access the mailbox can open it. Note that this option does *not* guarantee data preservation or immutability, as anyone with permission to access the shared mailbox can modify it unless you have applied a hold. You should check that full delegate auditing is enabled for the mailbox (see Chapter 21) so that Exchange records whatever actions the delegates take with items in the mailbox. Because the purpose of the shared mailbox is to preserve information, you should prevent people from sending messages to the mailbox by changing its SMTP address and hiding it from the GAL. You also need to remove the mailbox from the membership of any distribution lists and Groups to which it belongs and revoke its access to other SharePoint Online sites. Converting the mailbox of a departed employee to a shared mailbox is a simple and effective way to preserve mailbox contents that is preferable if you think that someone



will need to access the information in the mailbox in the short term. If you have the necessary licenses, you should put the mailbox on hold if you need to retain its contents for compliance purposes.

- **Convert to an inactive mailbox:** If no need exists for short-term access to the mailbox contents, you might prefer to “warehouse” the mailbox by making it inactive. As explained earlier, two prerequisites exist: first, the mailbox must be on hold before the account is removed. Second, the account must have a license that supports retention policies to allow the mailbox to be put on hold. Once the mailbox is on hold, it is safe to remove the user account because Exchange Online will keep the mailbox until the hold lapses, or the tenant removes all holds which apply to the mailbox. No one will be able to log into the mailbox and no need exists for a license. When a mailbox is made inactive, Exchange removes it from address lists and distribution lists. A mailbox can stay in the inactive state for a sustained period, but you can recover or restore an inactive mailbox if needed. Information held in the mailbox remains available for eDiscovery searches.

Exchange Online Plan 2 (included in Office 365 E3 and E5) or the Exchange Online Archiving add-on license are needed for retention policies. Because of this, mailboxes belonging to accounts with lower-cost licenses like frontline workers can't be made inactive and therefore their mailboxes must be converted into shared mailboxes if the need exists to retain these mailboxes for compliance purposes.

As noted earlier, you can run the *Remove-MgUser* cmdlet to remove a user account. Removing a user account releases the licenses assigned to the account. You can either reassign the licenses to other accounts or reduce the number of licenses that you pay for each month.

The steps needed to disable and preserve an employee's account in a hybrid deployment are different. Remember that in this scenario, the on-premises Active Directory is the master and all changes to accounts and mailboxes must happen on-premises and then synchronize to Azure AD, which usually creates a delay before a guaranteed block exists for an account.

## Dealing with Other Workloads

Of course, because information exists in many other places in Microsoft 365 workloads than the user's mailbox, administrators need to do some added work to look for and recover anything considered valuable. These include:

- The equipment was issued to the employee. PSTs and documents might be on the hard drive of the user's PC and hold information invisible to administrators unless they can gain physical access to the drive. In terms of the potential for information leakage, PSTs are especially vulnerable as users can remove them from the organization on easily portable USB thumb drives along with copies of documents downloaded from SharePoint and OneDrive for Business sites. You can scan the audit log to figure out whether the ex-employee downloaded an excessive number of files during their last weeks of employment, but it is impossible to know whether they copied messages from Outlook to a PST and took that PST with them. Copies of documents and messages protected with sensitivity labels cannot be accessed when someone leaves the organization because the ex-employee needs to authenticate using a valid account before they can open the content. This is a good reason to use sensitivity labels to protect confidential information.
- Given the widespread use of mobile devices, a terminated employee probably used a personal mobile device to access their mailbox. Because of caching, it might be possible to use cached credentials to connect to the mailbox with a mobile device even after the employee leaves. Given that a user might have multiple mobile devices, it is best to issue a remote wipe for all ActiveSync devices registered with the mailbox. Any future attempt to access the mailbox via ActiveSync will wipe the device that issues the connection request. Depending on how the mobile device vendor has implemented the ActiveSync protocol within the email client, the wipe might affect some or all

personal data. To wipe all the ActiveSync devices connected to a user account, run the following PowerShell command:

```
[PS] C:\> Get-MobileDevice -Mailbox 'Bad Employee' | Clear-MobileDevice
```

- Microsoft Endpoint Manager has more selective wipe capabilities for its managed devices than ActiveSync has. See [this page](#) for information about how to perform a full or selective wipe of a mobile device with Microsoft Endpoint Manager.
- The employee's OneDrive for Business site. Because this site stores personal information belonging to the removed user, its contents are out of sight to administrators unless you take steps to recover data after the removal of a user account. By default, unless the site is on hold, after the removal of a user's account, their OneDrive for Business site becomes a candidate for removal after 30 days by a background process (the My Sites Clean Up timer job). A message goes to the user's manager to warn that this will happen, and a second reminder goes 3 days before SharePoint Online removes the data. However, if the information in Azure AD does not allow SharePoint Online to decide who is the user's manager, SharePoint obviously cannot send the notification. If 30 days is too short a retention period, you can increase it to anything up to 3,650 days. For instance, this command sets a retention period of one thousand days.

```
[PS] C:\> Set-SPOTenant -OrphanedPersonalSitesRetentionPeriod 1000
```

- You can ensure continued access to the OneDrive sites owned by ex-employees by configuring the SharePoint Online settings to automatically assign a secondary owner to OneDrive for Business sites. Go to the **More features** page of the "new" SharePoint admin portal, open the **User Profiles** item, click **Setup My Site** under **My Site Settings**, verify that the **Enable access delegation** box is checked, and then select a user account to be the secondary owner. If Azure AD does not hold details of a manager for the removed user, SharePoint sends the reminders to preserve data to the secondary owner, who can then arrange for a review of the data and the capture of anything that needs preservation to another location (for example, another user's OneDrive for Business site or a SharePoint Online team site).
- SharePoint Online team sites and traditional sites that the user manages. The other users who have access to these sites can continue to work with the information held in the libraries and lists and the SharePoint Online administrator can grant ownership to these sites to a different user.
- Groups and Teams. If the user is the sole owner of a group or team, you should assign that role to another user. If you want to capture information about the conversations that an employee has participated in, you can run a content search against the group mailboxes for the teams they belong to and their mailbox for items. Once the search is complete, you can then export the results to a PST or ZIP file.
- The user might have created some videos in Stream that the organization might want to keep for reuse. When you remove an account, the videos owned by the account remain accessible in Stream. However, it's best to assign ownership for those videos to another person so that they can maintain permissions and settings. A Stream administrator can reassign video ownership.
- The user might have created some whiteboards to share ideas and concepts, including in Teams meetings. When you remove an account from Azure AD, any whiteboards the account owns are also removed. The [Invoke-TransferAllWhiteboards](#) cmdlet can be used to transfer ownership of whiteboards from one user to another.
- You can recover Microsoft Forms created by an ex-employee and transfer them to another user. See [this page](#) for information.

It is sensible to review the list of data sources annually as the potential places where users can store valuable corporate data might grow as the feature set available within Microsoft 365 expands.

**Removing Calendar Events:** When someone leaves the organization, their mailbox may be the organizer of events that exist in other peoples' calendars. In many cases, you might want to remove those events from calendars to allow someone else to reschedule replacement events (or not, as the case might be). The *Remove-CalendarEvents* cmdlet cancels future events organized by a specific mailbox and sends out cancellation notices. For example, this command cancels all meetings organized by Nancy Anderson from today's date, which is probably what you would do for an employee leaving the organization. Remember to run the cmdlet before you remove the mailbox.

```
[PS] C:\> Remove-CalendarEvents -Identity "Jim.Doe@Office365itpros.com" -CancelOrganizedMeetings
```

On the other hand, if the user is going away for an extended period and will eventually return, you can run the cmdlet to cancel events for a date range. See [this page](#) for more information about the cmdlet.

## Handling Inbound Email for Ex-Employees

When you convert a user mailbox to a shared mailbox, the mailbox keeps all the assigned email addresses and can continue to receive emails sent to the mailbox. Someone will need to process messages that arrive in the mailbox to let the senders know that the intended recipient no longer works at the company.

Alternatively, you can change the assigned email addresses so that Exchange will "bounce" (send a non-delivery notification) any new messages sent to the mailbox.

You can let people who try to contact the now-departed employee receive the normal non-delivery notification or you can create a better experience by telling them why the person they tried to contact is no longer available. This PowerShell code sets up internal and external auto-replies for the mailbox and adds a MailTip that is visible to internal users. We also hide the mailbox from all address lists and enable mailbox auditing to track any access that occurs to the mailbox.

```
[PS] C:\> Set-MailboxAutoReplyConfiguration -Identity "Terminated Employee" -InternalMessage "The person you are emailing no longer works for us. Please refer communications to Mr. Manager" -ExternalMessage "Mr. Terminated Employee is no longer an employee of this company." -AutoReplyState Enabled
```

```
Set-Mailbox -Identity "Terminated Employee" -MailTip "Terminated Employee has left the company. Please do not send any more mail to their mailbox" -HiddenFromAddressListsEnabled $True
```

Auto-replies and MailTips inform users about people that are no longer with the company if the mailbox still exists in an active state. However, we might want to remove the mailbox completely after harvesting any useful data in it. Exchange Online does not have a way to inform correspondents that a mailbox is no longer in use, but we can do this by creating a shared mailbox to hold the email addresses previously assigned to removed mailboxes or the old addresses of user mailboxes that have been converted to shared mailboxes. In effect, you use the shared mailbox (which does not need a license) to redirect inbound messages so that the senders will receive some information to inform them that their correspondent is no longer available.

Although Exchange Online limits the number of SMTP proxy addresses that you change assign to a mail-enabled object, a shared mailbox can easily hold 400 email addresses. If you need to keep a higher number of addresses for departed employees, you can spread the addresses over several mailboxes. One technique is to create a shared mailbox for each department so that external senders can receive an autoreply giving them details of a new contact within the department.

Of course, the shared mailbox will act as a black hole if you simply add the email addresses of departed employees to it. To complete the process, you should create an autoreply for the shared mailbox so that Exchange Online will respond to the senders after it delivers inbound messages for the addresses assigned to the shared mailbox. Ideally, the autoreply will tell the sender that they should not use the address in the future.

Although there might be some value gained by receiving email in the shared mailbox, often companies do not want to accumulate messages from people who have left. It can be an onerous task, not to mention a potential breach of personal privacy, if someone accesses the shared mailbox to process and respond to the messages found there, so the best solution is often to institute an automatic bounce mechanism to suppress inbound messages. This can be achieved with the combination of a distribution list and a transport rule. Here's how:

- Create a normal distribution list and add the shared mailbox (or mailboxes if necessary) to its membership. Make sure that the group owner (by default, the administrator who creates the group) is not added to the membership as they will not be able to receive new mail once the transport rule created in the next step implements a block for group members.
- Now create a transport rule to intercept messages sent to the members of the distribution list and return a rejection notice to the senders. We provide suitable text to explain why the rule rejects messages.

```
[PS] C:\> New-TransportRule "Block Email to Disabled Mailboxes" -SentToMemberOf "Disabled Mailboxes" -RejectMessageReasonText "Unfortunately the person with whom you attempted to communicate is no longer with our company." -Enabled $True
```

Once enabled, the rule will reject any message sent to one of the SMTP addresses assigned to the shared mailboxes in the distribution list. It is a simple but effective way to provide a better user experience.

## When Someone Dies

All of us will die someday. In large organizations, statistics show that the likelihood exists that one or more employees will die per 10,000 annually, depending on the average age of employees. Given the era we live in, when we do pass on, we will leave behind many digital assets, among which might be a corporate mailbox. It's a good idea to have a procedure to preserve the mailboxes and other information belonging to dead employees. The steps that you might take are like those that you use to preserve content in mailboxes for terminated employees and include:

- Disabling or blocking the account.
- Deciding if the mailbox should become an inactive mailbox or remain online. Alternatively, if you run a hybrid deployment, you could move the mailbox back to an on-premises database that is reserved for this purpose. The purpose of keeping the mailbox is to allow for the retrieval of any valuable information from the mailbox within a retention period decided by HR and/or the legal department and in compliance with applicable regulations such as GDPR. Some arrangements should be put in place to extract personal information from the mailbox and give it to the employee's family. For instance, some people store passwords and other valuable information in their mailboxes that might be needed by their family following their death. Personal and corporate data might also need to be recovered from the user's OneDrive for Business site. After retrieving whatever information is considered valuable from the mailbox, you might move it into an inactive state and keep the mailbox for a further period.
- Removing the mailbox from any distribution lists and the Groups/Teams to which it belongs. You can discover what groups a mailbox belongs to by looking at the mailbox properties via EAC or by running some PowerShell code (an example is in Chapter 11). If you make a mailbox inactive, its membership in distribution lists and groups is automatically canceled.
- If the mailbox remains online, remove the mailbox's access to any shared mailboxes. You can consider adding the mailbox to a special "Departed Employees" distribution list, which is hidden from the GAL. Some companies like to use a transport rule that blocks any incoming messages sent to the members of the "Departed Employees" group. The rule might also generate a customized NDR back to senders to inform them of the sad demise of the intended recipient. As described

above, you can use the *RejectMessageReasonText* parameter for the *New-TransportRule* cmdlet to create a thoughtful response to messages sent to a deceased employee. An alternative is to set an auto-reply message on the user's mailbox by either logging onto the mailbox or using the *Set-MailboxAutoReplyConfiguration* cmdlet (explained in Chapter 6).

Eventually, the retention period for policies applying to the mailbox or retention labels assigned to mailbox items will elapse, and you will either remove the user's account or cancel the hold that keeps their mailbox inactive.

Remember that a user is likely to have responsibility for documents and other information in other places across the service and that some effort is necessary to track down this information and transfer it to the safekeeping of another user.

**Humane effectiveness:** Although it is good to have a well-documented process to handle what happens when users die, it is also good if an organization can show some humanity. Some large businesses delegate the authority to manage the process of securing the digital assets of deceased employees to joint HR/IT teams who can flex and alter the process as necessary to meet the needs of any specific circumstances that might arise. The IT members of the team ensure that the technical processes are followed while HR ensures that everything is done humanely and thoughtfully. It is a good approach to follow.

## Compromised Accounts

A compromised account is one where someone outside the tenant (an attacker or hacker) manages to gain access to the account resources. Usually, this is because the attacker obtains the credentials necessary to sign in to the account, perhaps because their account credentials are compromised through a breach of another site.

Signs of unusual activity such as missing data, new rules or a forwarding address appearing in the mailbox that the user can't remember setting up, or strange emails from the mailbox might be indications of a compromised account. The steps necessary to secure the account and prevent further unauthorized access are:

- Block the account.
- Reset the account password and enable it for multi-factor authentication if not already done.
- Check all the mailbox settings to ensure that inbox rules, sweep rules, and forwarding addresses are valid (if a forwarding address exists outside the organization, ask if a business purpose exists for forwarding email to that address).
- If the account has administrative access, check what data might have been compromised through this access, validate that the access is needed, and remove it if not.
- The audit log can help you find what data the account has accessed recently. You should check any documents the account uploaded to SharePoint or OneDrive for Business to ensure that they don't contain malware.
- When the account is completely checked out, unblock it and give the new password to the user.

Microsoft's advice on the topic is in [this support article](#).

# Chapter 6: Managing Exchange Online

**Tony Redmond**

Two base workloads are the critical underpinning for much of the work within a tenant: Exchange Online for email, and SharePoint Online (including OneDrive for Business) for document management. Access to these workloads is reason enough for many companies to move to the cloud. This chapter reviews Exchange Online and discusses how it differs from its on-premises counterpart when it comes to what's stored inside online mailboxes and how to manage these objects.

## Exchange Online

For many companies, email is the first workload that they move into the cloud. Given the popularity of Exchange since its introduction in 1996, it was therefore unsurprising that many of the early organizations that moved to the cloud were migrations from on-premises Exchange servers.

Exchange Online is the cloud-based version of Exchange. Or is it? Both products share common roots and common functionality, but the version of Exchange that runs inside the cloud is very different from its on-premises counterpart. As we will see, the difference is somewhat inevitable given that Exchange Online must function inside a massive multi-tenant infrastructure while Exchange on-premises is designed to be able to serve the needs of many different customers, each of whom might adopt a different method to deploy and manage Exchange.

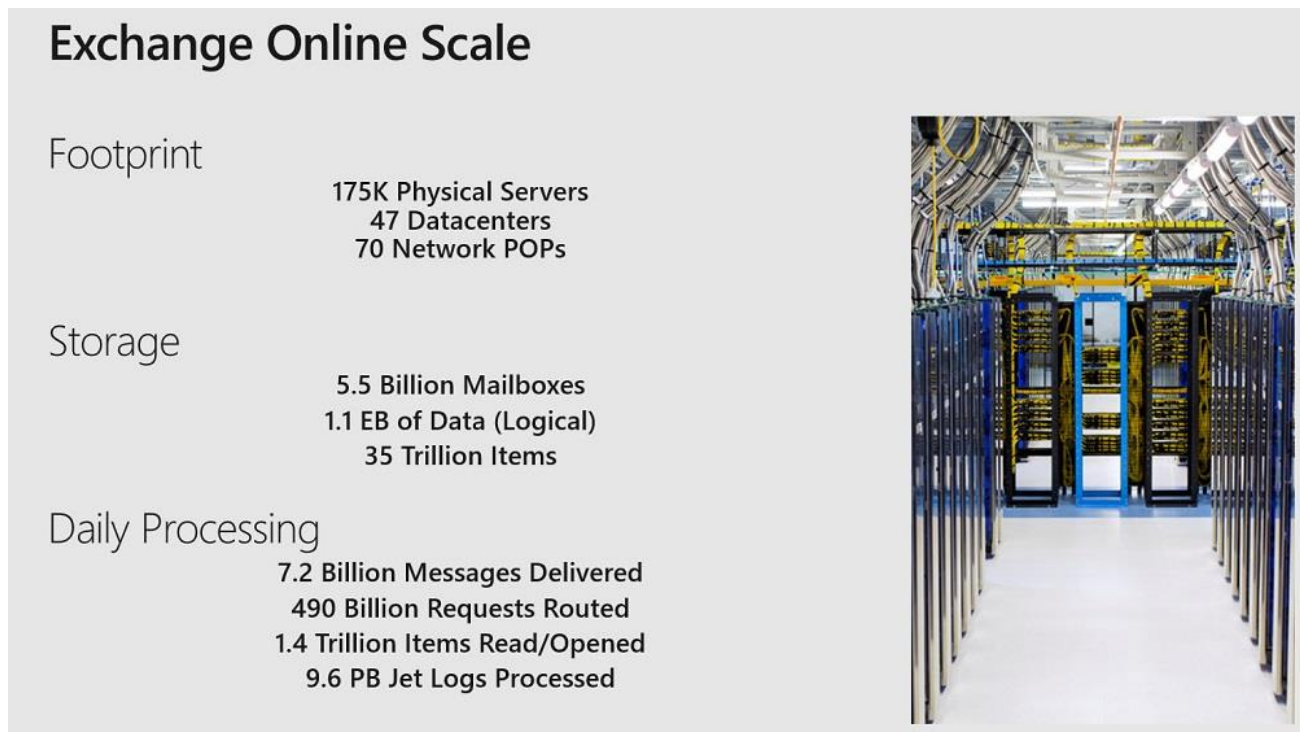


Figure 6-1: The scale of Exchange Online (source: Microsoft – Ignite 2018 conference)

The number of Exchange Online servers grows in line with the number of Office 365 accounts. At the Ignite conference in October 2018, Microsoft said that Exchange Online used 175,000 physical servers (Figure 6-1). Two years later, at the TEC 2020 conference, Microsoft reported that the number of Exchange Online servers had increased to 275,000, the number of data centers supporting Exchange Online had grown to 56, and 155 network point of presence (POPs) existed to carry client traffic to Microsoft's data center network. The email infrastructure delivered 10 billion messages daily while blocking two billion spam messages. Microsoft hasn't published updated numbers for Exchange Online since 2020, but given the growth in demand for Office 365 accounts, you can imagine how the numbers reported above have grown.

Microsoft's claim for Exchange Online to span five and a half billion mailboxes seems enormous in the context of 345 million paid Office 365 subscriptions, but the figure includes Outlook.com users ([400 million switched over to use the Exchange Online infrastructure in 2017](#)), shared mailboxes, group mailboxes, resource mailboxes, special cloud-only mailboxes used for compliance, and a very large number of system mailboxes used by the Microsoft 365 substrate. In other words, user mailboxes are a small minority of those active inside Exchange Online.

## Consumer Email

Exchange Online and Outlook.com share the same physical infrastructure. Outlook.com users receive service from mailboxes running on Exchange Online servers and connect to their mailboxes using a modified version of OWA. The functionality exposed in the consumer version of OWA is controlled to ensure that consumer accounts do not have the same level of functionality that's available to Exchange Online users; it's also segmented in terms of the functionality available to premium and free accounts. However, functionality developed for one platform does show up in the other. For example, the Sweep feature first appeared in Outlook.com before Microsoft decided to make it available to Exchange Online users. On the other hand, the Encrypt email feature first appeared in Exchange Online before it showed up in Outlook.com. Sharing the same physical infrastructure makes it very easy for Microsoft to switch features between the consumer and business platforms.

## Exchange Online Data Center Operations

Unlike the other Microsoft 365 services, Exchange Online still uses physical servers. The servers in the available pool receive updates on a rolling basis. Up to 10% of the servers are offline at any point. When the time comes to upgrade a server, Microsoft puts the server into maintenance mode, messaging queues are drained, and any active databases hosted by the servers move to other servers. The servers then leave the active pool. The refresh applies the latest version of the operating system, any required patches, and the latest build of Exchange to create a fully up-to-date server, which is then reintroduced into the operating pool. The passive copies of databases hosted by the server use the normal log shipping and replay mechanism to update themselves when the server comes back online. Eventually, the upgraded server becomes a candidate for database activation and the update cycle completes. Microsoft uses the wipe-and-load technique because it is a much simpler process than trying to keep servers updated with Windows and Exchange patches, service packs, and other fixes. At any point in time, Exchange Online servers run a variety of software versions because of their position in the refresh cycle. This can account for why some tenants see slightly different functionality than others. Given the number of servers in use, it is inevitable that some servers will not function as smoothly as they should. When this happens, Microsoft often finds that it is more cost-effective and simpler to shoot the server and remove it from the service than to try to perform a potentially complex hardware diagnosis and fix.

An on-premises administrator who paid a visit to a data center and reviewed how Microsoft runs Exchange Online would see some very familiar things, but many operation details are very different. The fundamentals of running a messaging system stay the same – users authenticate against Azure AD to connect to mailboxes

to receive and send emails. Messages exist in databases on Exchange mailbox servers and are transported to their destination via transport servers. Connectors link servers together and to the rest of the Internet via SMTP. Database replication hums along in the background to ensure that copies of user mailboxes are in four databases spread across at least two data centers. Apart from the sheer size of the infrastructure to support both Exchange Online and Outlook.com, the structure of mailboxes in databases running in Database Availability Groups spread across mailbox servers is very recognizable.

**The Largest Migration?** In February 2021, the [UK National Health Service](#) (NHS) completed the migration of 2.1 million mailboxes from on-premises servers to Exchange Online. About 2.3 petabytes of data was moved over six months at rates of up to 83,000 mailboxes daily. The NHS is the largest single Office 365 tenant. As far as we know, this is the largest mailbox migration to Exchange Online performed to date.

## Native Data Protection

Many people are struck by the fact that Microsoft does not backup the Exchange Online mailbox databases. Instead, Native Data Protection, a resiliency strategy formed by many different features incorporated into the Database Availability Groups, Information Store service, and transport service, protects data and removes the need for backups.

### Database Availability Groups

The Database Availability Group (DAG) is the cornerstone of Exchange Online storage. Mailboxes are stored in mailbox databases deployed in a DAG. Each DAG has sixteen mailbox servers spread across the data centers within a region. For instance, within the EMEA region, DAG member servers are in the Amsterdam, Dublin, Helsinki, and Vienna data centers. Each mailbox database has four copies distributed across servers in the data centers. The active copy is accessed by clients while the three other copies, one of which is lagged (kept seven days behind the active copy), synchronize with the active copy using log shipping and replay. Exchange captures transactions in transaction logs. Log shipping copies transaction logs generated by the active database copy to the servers holding passive copies. The target servers check the transaction logs as they arrive, and if the logs are valid, replay their contents to update the passive copy.

Because the DAG member servers are connected using Microsoft's high-speed data center network, the copy and replay operations are almost instantaneous. If a problem occurs with the active server, Exchange decides which of the passive copies to activate and switches it to service client connections. The previously active server then becomes a host for a passive copy.

Exchange Online servers use the ReFS file system (available in Windows Server 2012 or later) to gain better protection for data and to verify the integrity of data written to disk. The Exchange servers use a mixture of SSD (for the metacache) and JBOD RAID arrays and can automatically swap a replacement disk into an array should one fail.

### Database Engine Checks

Apart from log shipping and replay (which include consistency checks), the DAG includes other features to ensure that physical or logical corruption does not creep into databases. These include:

- **Single-bit correction:** The database engine detects and fixes single-bit CRC errors that result from hardware problems. The errors are corrected and flagged so that Microsoft's operations team can investigate.
- **Database consistency checker:** A background process reads and checks database pages for checksum failures. Any failed page is automatically fixed. The scheme used makes sure that every page is checked in a database every seven days.



- **Lost flush detection:** Lost flushes happen when the operating system (or disk) reports that a write happened, but the write operation did not complete (or is written to the wrong place). This represents logical corruption. To prevent this from happening, the database engine checks pages as it writes them to passive database copies. If a problem is found, it is fixed through single page patching.
- **Single page patching:** This is a process to replace bad database pages with good pages from other database copies in the DAG. If the problem is detected in a passive copy, the database engine copies the good page from the active database using the transaction log shipping and replay mechanism. If the problem is in the active database, the good page can come from any passive copy.

These features also exist in the on-premises version of Exchange.

## Other Exchange Resiliency Features

In addition to the DAG and database engine features, Exchange builds resilience into the transport service so that messages are always sent, even when failures occur. Shadow Redundancy means that the transport service takes a (shadow) copy of each message it receives. If Exchange ever thinks that a message might not have been processed, it can resubmit the shadow copy. The Safety Net queue captures copies of messages as they pass through the transport pipeline. If a failure occurs in the transport service or with a database, Exchange can replay messages from the Safety Net queue to the active database copy. Any redundant copies discovered during the replay are suppressed and not exposed to users.

Behind the scenes, Exchange Online servers are rebalanced on an ongoing basis to ensure that each takes approximately the same amount of user load. In addition, Microsoft withdraws mailbox servers from service to update their software as the need arises. These operations force mailboxes to be moved between databases. The Mailbox Replication Service (MRS) moves mailboxes between databases. During the moves, MRS checks mailbox contents for problems and if it detects problem items, it attempts to correct the items. If this isn't possible, MRS skips the corrupt items.

## Tenant-Controlled Resiliency Features

Tenants can control some mailbox management features that assist in resiliency. These include:

- **Single Item Recovery (SIR):** This feature ensures that Exchange Online retains messages that users purge (hard delete) in the Recoverable Items folder of mailboxes until the deleted item retention period expires. The maximum deleted item retention period is 30 days.
- **In-place and Litigation holds:** Holds make sure that Exchange keeps select items or all items in a mailbox for the set retention period. If a mailbox is deleted when it is subject to a hold, it becomes inactive, meaning that it is retained until the hold elapses. Holds can also be set by retention policies.
- **Large mailbox quotas:** Many enterprise licenses include 100 GB mailbox quotas. In effect, many users do not need to delete messages. If they do, the messages go into the Deleted Items folder and stay there until the user empties the folder. Even then, the deleted items go into the Recoverable Items folder (which can have another 100 GB quota) from where they can be recovered if needed for up to 30 days (the deleted items retention period).
- **Expanding archives:** Mailboxes can be archive-enabled. An archive mailbox is designed for long-term storage and its basic quota is 100 GB. However, tenants can opt for expandable archives, which means that the archive is composed of 50 GB "chunks" linked into a logical entity.
- **Retention labels and policies:** To ensure that important information is kept for defined periods, users can apply retention labels to folders and items. Exchange Online does not remove items under retention control until the retention period elapses. See the Compliance chapter for more information about Microsoft 365 retention policies and processing.

The combination of the DAG, database copies, database engine features, transport copies, and tenant-controlled resiliency features are enough for Microsoft to conclude that they do not need backups for Exchange Online mailbox data. Your situation might be different, but in many cases, tenants do not need to invest in a third-party backup service either and full consideration should be given to how to maximize the use of out-of-the-box functionality to decide if Native Data Protection meets the perceived need for backups.

## Managing Mailboxes

Exchange Online mailboxes hold a mixture of default, system, and user-created folders. The default folders are the set of well-known folders like the Inbox that exist in every mailbox. Some people only ever use a small set of default folders – Inbox, Sent Items, and Deleted Items – while others are dedicated filers and store items away in carefully-selected folders. Colloquially, the two types of users are “pilers” and “filers.” If its limits are respected, Exchange doesn’t care how data are organized in a mailbox. Limits for internal mailbox structures include:

- Maximum number of items per mailbox folder: 1 million
- Maximum number of items in the Recoverable Items folder: 3 million
- Maximum number of subfolders per mailbox folder: 10,000 (including the root folder)
- Maximum folder hierarchy depth: 300

Most won’t encounter a mailbox’s internal limits, but many have exceeded their storage quota.

You’ll find that we use PowerShell to manage many mailbox settings. To run these commands, you must connect a session to the [Exchange Online Management](#) endpoint.

### Unique Mailbox Identifiers

In June 2022, Exchange Online changed the way that it generates the Name and Distinguished Name properties for new mailboxes. Instead of using the *MailNickName* property from the Azure AD account as the basis for the *Name* and *DistinguishedName* properties, Exchange uses the Azure AD object identifier for the account when it creates new mailboxes. EXODS stores the object identifier (also called EDOID) for Azure AD accounts and groups in the *ExternalDirectoryObjectId* property.

The aim is to guarantee uniqueness for the *Name* and *DistinguishedName* properties to avoid issues when account data synchronizes from Azure AD to Exchange Online when accounts are homed on-premises. The problem arises because Active Directory creates objects in multiple organizational units whereas Azure AD creates objects in a single organizational unit named after the tenant. Thus, an Active Directory object with a *Name* property of John.Smith will cause a synchronization conflict if a similarly named object exists in a different organizational unit. Because Exchange Online stores the EDOID in the *Name* property for new mailboxes, including those created as remote mailboxes from the on-premises EAC, Azure AD no longer synchronizes the *Name* property with Active Directory.

The example below illustrates the new format for the *Name* and *DistinguishedName* properties:

```
[PS] C:\> Get-ExoMailbox -Identity b67c8bd7-a8d3-4358-b42f-cd51821f7ba3 -Properties Name
ExternalDirectoryObjectId : b67c8bd7-a8d3-4358-b42f-cd51821f7ba3
UserPrincipalName : Sue.P.Pickett@office365itpros.com
Alias : Sue.P.Pickett
DisplayName : Sue Pickett
Name : b67c8bd7-a8d3-4358-b42f-cd51821f7ba3
DistinguishedName : CN=b67c8bd7-a8d3-4358-b42f-cd51821f7ba3,
OU=Office365itpros.onmicrosoft.com,OU=Microsoft Exchange Hosted
Organizations,DC=EURPR04A002,DC=prod,DC=outlook,DC=com
```

Exchange Online will not update the properties of mailboxes created before the change became effective. However, if you want, you can update mailbox properties. For example:

```
[PS] C:\> $ExternalDirectoryObjectId = Get-ExoMailbox -Identity Kim.Akers@Office365itpros.com |
Select -ExpandProperty ExternalDirectoryObjectId
Set-Mailbox -Identity $ExternalDirectoryObjectId -Name $ExternalDirectoryObjectId
```

After you update the Name property, Exchange Online updates the Distinguished Name property to match.

## Calendar Permissions

While granting delegate access via Outlook or OWA are the normal ways for users to allow other people access to a folder like their calendar, administrators can run the *Add-MailboxFolderPermission* cmdlet to do the same. For example, this command gives delegate access to Michael Harty for the calendar of Kim Akers. The Editor access right is needed to create and edit items in the folder, and the sharing permission flags tell us that the delegate can see private items in the target calendar. In this case, *SendNotificationToUser* is specified to send a sharing notification to the new delegate to tell them that they can now access someone else's calendar. In some cases, you will not want to generate a sharing notification as the potential exists that the recipient might unwittingly refuse the invitation.

```
[PS] C:\> Add-MailboxFolderPermission -Identity Kim.Akers@office365itpros.com:\Calendar -User
Michael.Harty@office365itpros.com -AccessRights Editor -SharingPermissionFlags Delegate,
CanViewPrivateItems -SendNotificationToUser $True
```

FolderName	User	AccessRights	SharingPermissionFlags
Calendar	Michael Harty	{Editor}	Delegate, CanViewPrivateItems

To remove delegate access from a mailbox, run the *Set-MailboxFolderPermission* cmdlet and set the *SharingPermissionFlags* parameter to None.

Exchange holds details of delegate access in a hidden item in the user mailbox. Sometimes this item can become corrupted, and the user will no longer be able to add or remove delegates. In this case, you should run the *Remove-MailboxFolderPermission* cmdlet with the *ResetDelegateUserCollection* parameter to force Exchange to recreate the hidden item.

```
[PS] C:\> Remove-MailboxFolderPermission -Identity Kim.Akers@office365itpros.com:\Calendar
-ResetDelegateUserCollection
```

This action has the side-effect of removing the flags which enable delegate access. To complete the fix, you must recreate the delegate settings with *Set-MailboxFolderPermission*. For example, this command re-establishes Michael Harty as a delegate to manage the calendar in the mailbox of Kim Akers.

```
[PS] C:\> Set-MailboxFolderPermission -Identity Kim.Akers@office365itpros.com:\Calendar -User
Michael.Harty@office365itpros.com -AccessRights Editor -SharingPermissionFlags Delegate,
CanViewPrivateItems
```

## Regional Settings

When you create a new Exchange Online mailbox, none of its regional properties are set. These are the properties that define the language used to interact with the mailbox, the time zone that the mailbox owner is in, and their preferred date format. The first time someone logs into a mailbox with OWA, the client prompts them for this information, sets the properties, and creates a set of default folders in the mailbox based on the language setting. For example, mailboxes that use the English language have an *Inbox* folder, while mailboxes configured for French have *Boîte de réception*. If you connect to a brand-new mailbox with an Outlook desktop or a mobile client, the mailbox takes the language setting of the client and creates the default folders based on that value, but the other regional settings might not be set. To see the default regional configuration for a mailbox, we run this command:

```
[PS] C:\> Get-MailboxRegionalConfiguration -Identity 'Rob Young'
```

If necessary, you can run the *Set-MailboxRegionalConfiguration* cmdlet to tweak the regional settings. In this example, the mailbox language, time zone, and date format match the settings for a Dutch user working in the Netherlands. Notice the use of the *LocalizeDefaultFolderName* parameter, set to *\$True* to force Exchange Online to create the default folder names in the target mailbox using specified language:

```
[PS] C:\> Set-MailboxRegionalConfiguration -Identity 'Rob Young' -Language n1-NL
-TimeZone 'W. Europe Standard Time' -DateFormat 'd-M-yyyy' -TimeFormat 'HH:mm'
-LocalizeDefaultFolderName:$True
```

Exchange Online can be picky about the date and time formats used when updating mailbox regional configurations. The formats must be valid for the selected language. Sometimes it can be difficult to know what acceptable values are for date and time formats. A practical approach is to use OWA to change the regional settings for a mailbox and then examine the values for the mailbox's regional configuration. You can then reuse those values with other mailboxes. To know the values for the timezone setting, you can run this PowerShell code to report the values stored in the system registry:

```
[PS] C:\> $TimeZone = Get-ChildItem "HKLM:\Software\Microsoft\Windows NT\CurrentVersion\Time zones"
| ForEach {Get-ItemProperty $_.PSPath}; $TimeZone | Sort-Object Display | Format-Table PSChildname,
Display -Auto
```

Users can select their language and region settings through OWA. The selected language controls the language used in the OWA interface. It has less effect when using Outlook because the user interface is set by the version of Outlook installed on the workstations. A user can switch their language setting for OWA at any time, which can lead to interesting results if someone selects a language that they don't understand.

## Auto replies and Out of Office Notifications

Out of Office or OOF ([Out of Facility](#)) are names used for the autoreply feature which allows mailboxes to send automatic replies after the arrival of new messages. The *Get-MailboxAutoReplyConfiguration* and *Set-MailboxAutoReplyConfiguration* cmdlets retrieve and set the autoreply settings for a mailbox, including shared mailboxes. Users receive autoreply messages when they send emails to mailboxes that have an autoreply message configured and enabled. Exchange Online also displays a recipient's autoreply in clients like Outlook desktop when adding the recipient to a message. To discover which mailboxes have autoreply set and the time the autoreply lapses, you can use the following command:

```
[PS] C:\> Get-ExoMailbox -RecipientTypeDetails UserMailbox | Get-MailboxAutoReplyConfiguration |
Where {$_.AutoReplyState -eq "Scheduled" -or $_.AutoReplyState -eq "Enabled"} |
Format-List MailboxOwnerId, StartTime, InternalMessage
```

The first check looks for auto-replies scheduled for a specific period, the second finds mailboxes where the autoreply is enabled without dates. The *InternalMessage* property reveals the HTML-formatted text of the auto-reply message for internal correspondents (*ExternalMessage* holds the text seen by external correspondents).

### Setting Autoreply for Mailboxes

You can use the *Set-MailboxAutoReplyConfiguration* cmdlet to create an autoreply for a user who has gone on vacation and forgotten to let anyone know. Another scenario is when public holidays occur, and you want to set the autoreply for customer-facing shared mailboxes to let anyone who sends an email to the company know that a response will be delayed. In both instances, you enable autoreply, set a time limit, and create separate messages for internal and external audiences. You can also instruct Exchange Online that auto-replies go only to external people who are contacts of the mailbox owner (select *All* instead of *Known* if you want the auto-reply to go to anyone external who sends a message to the mailbox).

If you create an autoreply for a certain period, make sure that you set the *AutoReplyState* parameter to be "Scheduled" rather than "Enabled" as failure to do this will enable the autoreply for the mailbox at once rather than in the future. Exchange Online servers run using UTC, so if you specify a time (rather than just a date) in the start and end times, make sure that you convert the time into UTC. For example, the command below starts auto-replies at 19:30 UTC and ceases at 17:00 UTC on the respective dates.

```
[PS] C:\> Set-MailboxAutoReplyConfiguration -Identity "Kim Akers" -StartTime "04-Nov-2019 19:30"
-AutoReplyState "Scheduled" -EndTime "08-Nov-2019 17:00" -InternalMessage "Kim Akers is attending
the Microsoft Ignite event in Orlando and will respond to your message after she returns on November
10" -ExternalMessage "Kim Akers is on vacation" -ExternalAudience 'Known'
```

Note that when you view the autoreply configuration for a mailbox, PowerShell converts the times from UTC into the time zone applied to the local workstation. To turn off autoreply for a user, set the *AutoReplyState* property to "Disabled":

```
[PS] C:\> Set-MailboxAutoReplyConfiguration -Identity "Kim Akers" -AutoReplyState Disabled
```

Here's another example. In this case, we want to add auto-replies to all shared mailboxes to cover the period of a public holiday so that anyone who sends a message to the mailboxes will receive a reply to tell them that the user's out of the office. Note the use of HTML tags to add some basic formatting to the external message.

```
[PS] C:\> $HolidayStart = "04-Aug-2020 17:00"
$HolidayEnd = "6-Aug-2020 09:00"

$InternalMessage = "Expect delays in answering messages to this mailbox due to the holiday between
<b> + $HolidayStart + "</b> and <b> + $HolidayEnd + "</b>"
$ExternalMessage = "Thank you for your email. Your communication is important to us, but please be
aware that some delay will occur in answering messages to this mailbox due to the public holiday
between <b> + $HolidayStart + "</b> and <b> + $HolidayEnd + "</b>"

$Mbx = (Get-ExoMailbox -RecipientTypeDetails SharedMailbox | Select DisplayName, Alias,
DistinguishedName)
ForEach ($M in $Mbx) {
# Set auto reply
Write-Host "Setting auto-reply for shared mailbox:" $M.DisplayName
Set-MailboxAutoReplyConfiguration -Identity $M.DistinguishedName -StartTime $HolidayStart -
AutoReplyState "Scheduled" -EndTime $HolidayEnd -InternalMessage $InternalMessage -ExternalMessage
$ExternalMessage -ExternalAudience 'All' -CreateOOFEvent:$True }
```

## Autoreply Settings for Calendar Processing

OWA includes some settings to process incoming calendar requests during periods when autoreply is active for a mailbox. These settings are not yet visible to Outlook, but because they are acted upon by the server, their effect is felt when set by OWA. The options are:

- **Block my calendar for this period:** Other users will see this user's calendar as blocked out if they try to schedule a meeting during the period when autoreply is active. The *CreateOOFEvent* switch set by *Set-MailboxAutoReplyConfiguration* determines if Exchange creates a calendar event corresponding to the OOF period.
- **Automatically decline new invitations for events that occur during this period:** If new event invitations arrive, they can be automatically declined. The *AutoDeclineFutureRequestsWhenOOF* property controls this setting.
- **Decline and cancel my meetings during this period:** This option scans the user's calendar for meetings that are in place for the period when the autoreply will apply. Meetings that the user was invited to attend will be declined while meetings that they set up will be canceled. The *DeclineAllEventsForScheduledOOF* property controls this setting.

The settings can be controlled with PowerShell using the *Set-MailboxAutoReplyConfiguration* cmdlet.

## Calendar Configuration

The *Get-MailboxCalendarConfiguration* and *Set-MailboxCalendarConfiguration* cmdlets manage calendar settings. For example, this command configures the calendar to use Greenwich Mean Time (GMT) as the time zone with a starting time for the workday of 8:30 A.M.:

```
[PS] C:\> Set-MailboxCalendarConfiguration -Identity "Kim Akers" -WorkingHoursTimeZone "GMT Standard Time" -WorkingHoursStartTime 08:30:00
```

The cmdlet also controls the appearance of a user's calendar when viewed through OWA (Outlook uses different settings based on the time and time zone configured for the PC). For example, this command makes Monday the first working day of the week, starts a new year on the first day of the year, changes the default time increment from 30 minutes to 15 minutes, sets the weather unit to be Celsius, and defines the user's location to be County Dublin, Ireland. Figuring out the right longitude and latitude for a user's location might seem hard, but online tools help (like [this example](#)) or you can experiment by inputting different locations into the Weather section of the OWA Calendar options and noting what values are set.

```
[PS] C:\> Set-MailboxCalendarConfiguration -Identity "Kim Akers" -WeekStartDay Monday -FirstWeekOfYear FirstDay -TimeIncrement FifteenMinutes -WeatherUnit Celsius -WeatherLocations "{LocationId:9480;Name Dublin, County Dublin;Latitude:53.348;Longitude:-6.248}"
```

You can also use *Set-MailboxCalendarConfiguration* to control the scheduling of Teams online meetings by default by Outlook clients. A mailbox setting is available to override the organization setting created using *Set-OrganizationConfig*, but only for OWA and Outlook Mobile clients. Outlook desktop uses a different method to control if online meetings are the default for an individual account. This example configures online meetings as the default for the specified mailbox:

```
Set-MailboxCalendarConfiguration -OnlineMeetingsByDefaultEnabled $True -Identity Kim.Akers
```

### Automated Processing for Room Mailboxes

New room mailboxes receive a default value of *AutoAccept* for the *AutoProcessing* calendar property (before the default was *AutoUpdate*). The idea behind the change is to speed acceptance of meeting requests for room mailboxes. If this change doesn't fit with your corporate policies, make sure that you update new room mailboxes after creation. For example:

```
[PS] C:\> Set-CalendarProcessing -Identity NewRoom -AutomateProcessing AutoUpdate
```

For more information about how to control the processing of booking requests for room mailboxes, see [the online documentation](#).

### Viewing Details of User Availability

By default, Exchange Online makes limited free and busy information for mailboxes available to other tenant users to allow them to see when other people might be able to attend a meeting. The default setting is *AvailabilityOnly*, meaning that a user can see when someone is busy, but can't see any details about the reserved time slot. Many organizations choose to upgrade the setting to allow people to see details of slots reserved in other users' calendars. There's no organization-wide setting to control how the calendar works. Instead, you must update the permission granted to the special Default user for each calendar. The *Set-MailboxFolderPermission* cmdlet can update the permission. For instance, this code looks for mailboxes to update and runs *Set-MailboxFolderPermission* to update the access rights for the calendar folder for the Default user to *LimitedDetails*. The *LimitedDetails* setting allows other users to see the title, location, and status (out of office, tentative, etc.) for time slots.

```
[PS] C:\> # Find mailboxes that we have not yet reset the default sharing view
[array]$Mbx = Get-ExoMailbox -RecipientTypeDetails UserMailbox, RoomMailbox -ResultSize Unlimited -Filter {CustomAttribute13 -ne "Open" -and CustomAttribute13 -ne "Blocked"}
```

```

$CalendarName = "Calendar" # English language calendar folder
ForEach ($M in $Mbx) {
Write-Host "Processing" $M.DisplayName
# You can hard-code the calendar name (above) or try and find a local language value.
# This is one way to look for local values...
# $CalendarName = (Get-ExoMailboxFolderStatistics -Identity $M.UserPrincipalName -FolderScope
Calendar |?{$_.FolderType -eq "Calendar"}).Name
# Either way, you need to end up with a valid calendar folder reference
# - Like Tony.Redmond@office365itpros.com:\Calendar
$CalendarFolder = $M.UserPrincipalName + ":\\" + $CalendarName
Set-MailboxFolderPermission -Identity $CalendarFolder -User Default -AccessRights LimitedDetails
Set-Mailbox -Identity $M.ExternalDirectoryObjectId -CustomAttribute13 "Open"
} # End Foreach

```

Users in an organization might not use English-language versions of OWA or Outlook, in which case their Calendar folder might have a different name. The code above handles the situation by using the *Get-ExoMailboxFolderStatistics* cmdlet to look for the calendar folder and fetch its name. This cmdlet is expensive in terms of processing overhead, which is why the call is commented out. However, if you need to deal with multiple languages, you'll need to uncomment the command and take the hit.

After updating the permission, the code uses *Set-Mailbox* to update the *CustomAttribute13* attribute so that the next time it runs, it won't process this mailbox. In addition, the code ignores any mailbox with Blocked in *CustomAttribute13* to handle the situation where you don't want to share calendar details for some confidential mailboxes.

## Automatic Meeting Shortening

Meeting shortening means reducing the time assigned to an event by a set amount depending on its desired length. The idea is to allow users to have a buffer between meetings to have an opportunity to recharge before the next event. Outlook and OWA allow individual users to set how long they would like to reduce short (under an hour) and long (over an hour) meetings and whether the buffer should be at the start or end of the period. Tenants can apply default settings by updating the Exchange Online organization configuration with PowerShell. It's critical to understand that once a user selects their settings, the organization defaults do not apply to them.

Three organization-wide settings are available to control event shortening:

- **ShortenEventScopeDefault:** Sets whether event shortening is in effect (0 or none) or applies to ending meetings early (1 or *EndEarly*) or starting later (2 or *StartLate*). This parameter must be set to 1 or 2 before you can amend the periods.
- **DefaultMinutesToReduceShortEventsBy:** The number of minutes to shorten events by if they are scheduled for one hour or less. The default is five.
- **DefaultMinutesToReduceLongEventsBy:** The number of minutes to shorten events by if they are scheduled for over one hour. The default is 10.

To turn on event shortening for the organization and select to end events early, we run:

```
[PS] C:\> Set-OrganizationConfig -ShortenEventScopeDefault EndEarly
```

Using *Get-OrganizationConfig* to examine the settings afterward shows the current configuration:

```
[PS] C:\> Get-OrganizationConfig | fl defaultmin*, short*

DefaultMinutesToReduceShortEventsBy : 5
DefaultMinutesToReduceLongEventsBy   : 10
ShortenEventScopeDefault             : EndEarly
```

Like any organization-wide setting, some time is necessary to allow clients and servers to pick up new values. For now, there's no way for administrators to use PowerShell to update these settings for individual mailboxes as Microsoft hasn't upgraded the *Set-MailboxCalendarConfiguration* cmdlet.

## Controlling Calendar Requests for Mailboxes

The *Set-CalendarProcessing* cmdlet controls how the Resource Booking Assistant processes meeting requests arriving in user and room mailboxes. For example, this command sets the properties controlling the processing of forwarded meeting notifications and external meeting requests:

```
[PS] C:\> Set-CalendarProcessing -Identity "Kim Akers" -RemoveForwardedMeetingNotifications $True -ProcessExternalMeetingMessages $True
```

It's common to have conference rooms that the organization reserves for specific people. To accomplish the goal, we update the calendar processing configuration to block accepting requests from anyone except members of a distribution list:

```
[PS] C:\> Set-CalendarProcessing -Identity "Room 101" -BookInPolicy "Senior Leadership Team"
```

Users who attempt to book a room receive email responses to tell them if the calendar assistant has accepted or denied their request or when the request awaits approval by a delegate who manages reservations for the room. You can update the meeting responses by setting the *AddAdditionalResponse* switch to *\$True* and providing the text to insert into the responses in the *AdditionalResponse* property. This example shows how to add HTML-formatted text, including an emoji:

```
[PS] C:\> Set-CalendarProcessing -Identity "Room 101" -AddAdditionalResponse $True -
AdditionalResponse '<h2>Welcome to the Corporate Room Scheduling System</h2><p>We have a few basic
rules for you to follow.</p><ol><li>Please keep the room tidy and remove rubbish at the end of your
meeting.</li><li>Please do not change the settings of the AV equipment.</li><li>Please clean the
whiteboard before you leave.</li><li>Advise Corporate Meetings if you have any problems by sending
email to mailto:<a href="mailto:corporatemeetings@office365itpros.com">Corporate
Meetings.</a></li></ol><p><strong>Room 101</strong> can hold up to <strong>12</strong> people.
Please do not exceed this capacity.</p><p>If you need <strong>catering</strong>, please contact
Trina at 147-1497.</p><p>Thanks for meeting with us!&nbsp;</p><p>&nbsp;</p>'
```

## How Outlook Processes Inbound Meeting Updates

The *VisibleMeetingUpdateProperties* setting in the tenant Exchange Online organization configuration controls how the Outlook (for Windows) client processes meeting updates. In the past, each time a meeting organizer updated any property of a meeting, like its title, location, date, or body (description), email notifications go to all meeting attendees. The attendees then had to process the update. Outlook can auto-process meeting updates, meaning that Outlook automatically applies updates to the event in attendee calendars without the need for any human interaction. Notifications are still emailed but are moved to the Deleted Items folder after Outlook processes the updates.

In general, updating meeting updates without requiring user intervention is welcome. However, some of the updates might contain information that users want to see. For instance, the default configuration is to only show update notifications to users if the meeting location changes or a change is made to any detail of the meeting within 15 minutes of its start. Outlook processes (and hides) any other change, such as an update to the meeting subject or the body of the meeting notice, which is where details such as agendas are often published. Outlook always shows a meeting notification if any of the following conditions are true:

- A change is made to the meeting date, time, or recurrence pattern.
- The notification is for a delegated shared calendar.
- The recipient is @mentioned in the meeting body.
- The recipient has not yet responded to the meeting (in effect, the notification acts as a prompt for them to respond).

You can change how Outlook behaves by running the *Set-OrganizationConfig* cmdlet to update the settings in *VisibleMeetingUpdateProperties*. For example, this command forces Outlook to display updates for any change



to the meeting location or body 120 minutes or less before it starts, or a change at any time to the online (Teams or Skype for Business Online) details or meeting subject.

```
[PS] C:\> Set-OrganizationConfig -VisibleMeetingUpdateProperties Location:120, Body:120,
OnlineMeetingLinks, Subject
```

The setting applies to all mailboxes in the tenant. You can't change how Outlook works on a per-mailbox basis.

## Generating Automatic Calendar Events from Email

When it processes inbound email, a background Exchange agent scans messages to figure out if the messages relate to events generated by airlines, hotels, and other sources like booking agencies. If an event is detected, user settings control how Exchange processes the event. The settings to control automatic event detection are found in the **Calendar** section of OWA Options. Go to **Events from email** to select how Exchange should process each of the event types. You can choose to:

- **Don't show event summaries in email or calendar:** This choice stops processing these events.
- **Only show event summaries in an email:** Exchange delivers event information in an email but won't create a calendar event.
- **Show event summaries in email and calendar:** Exchange delivers event information in email and uses that information to create a calendar event based on the information in the email, such as the time and date for a flight together with the airline booking reference, the name of the departure, and destination airports. Details of the extracted event are added to the message. If several events (such as flights in a single reservation) are detected, separate calendar events are created for each event.

Event settings can be manipulated with PowerShell by running the *Set-EventsFromEmailConfiguration* cmdlet (this used to be done with the *Set-MailboxCalendarConfiguration* cmdlet). For example, to set rental car event processing to Email only while making sure that flight reservations are created in the calendar, run:

```
[PS] C:\> Set-EventsFromEmailConfiguration -Identity Kim.Akers@Office365itpros.com -
FlightReservationProcessingLevel Calendar -RentalCarReservationProcessingLevel Email
```

Unlike other Exchange Online cmdlets, you must identify the target mailbox using one of the SMTP proxy addresses assigned to the mailbox. You cannot pass an alias, display name, distinguished name, or user principal name. This means that regular pipeline processing is not possible. For instance, you can't do this to have flight reservations created as calendar events:

```
[PS] C:\> Get-ExoMailbox -RecipientTypeDetails UserMailbox | Set-EventsFromEmailConfiguration -
FlightReservationProcessingLevel Calendar
```

Instead, you can do this:

```
[PS] C:\> $Mbx = (Get-ExoMailbox -RecipientTypeDetails UserMailbox)
$Mbx.ForEach( { Set-EventsFromEmailConfiguration -FlightReservationProcessingLevel Calendar
-Identity $_.PrimarySmtpAddress } )
```

To disable a setting for an event, set it to *Disabled*. For instance, here's how to disable event creation for all events:

```
[PS] C:\> Set-EventsFromEmailConfiguration -Identity Kim.Akers@Office365itpros.com
-FlightReservationProcessingLevel Disabled -LodgingReservationProcessingLevel Disabled
-ParcelDeliveryProcessingLevel Disabled -RentalCarReservationProcessingLevel Disabled
```

As you can see, you must disable each event type separately.

To check the current event processing settings for a mailbox, run the *Get-EventsFromEmailConfiguration* cmdlet:

```
[PS] C:\> Get-EventsFromEmailConfiguration -Identity Kim.Akers@Office365itpros.com
```

```

DisableReason          : None
CreateEventsFromEmailAsPrivate : False
EntityTypeProcessorLevelSettings : {RentalCarReservation, EventReservation,
FlightReservation...}
Identity               :
IsValid                : True
ObjectState            : New

```

This reveals the option to create events as private entries in the calendar (off by default) but doesn't tell us what each of the event type settings is. We find out with:

```
[PS] C:\> Get-EventsFromEmailConfiguration -Identity Kim.Akers@office365itpros.com | Select -ExpandProperty EntityTypeProcessorLevelSettings
```

Name	Value
RentalCarReservation	Disabled
FlightReservation	Disabled
LodgingReservation	Disabled
ParcelDelivery	Disabled

You can reset to default values (all events set to email with no automatic calendar events created) as follows:

```
[PS] C:\> Set-EventsFromEmailConfiguration -Identity Kim.Akers@Office365itpros.com -ResetSettings
```

Before Exchange can process an email to extract event information, Microsoft must recognize the sending organization (like an airline). The full list of recognized senders [is available online](#). Microsoft updates this list on an ongoing basis, with organizations joining when they realize the value of having their email create events in user calendars.

## Room Mailboxes, Workspaces, and Room Lists

Room mailboxes are a special form of resource mailbox marked for use as a meeting location. You create new room mailboxes through the Resources section of the EAC or with PowerShell. As described below, the Outlook Places service helps users find rooms for meetings, but it's still a good idea to include some location information in the display name for a room. Here's how to create a new room mailbox with PowerShell:

```
[PS] C:\> New-Mailbox -Name "SF Executive Meeting Room" -Displayname "San Francisco Executive Meeting Room" -Alias SFExecMeeting -Room
```

A workspace is a form of room mailbox to represent a place like an individual desk, a small meeting room used for calls, or other space where people work. The first step to creating a workspace is to create a room mailbox and set its type to be a workspace:

```
[PS] C:\> Set-Mailbox -Identity "Floor 1 Desk 17" -Type Workspace
```

A room list is a special form of distribution list composed exclusively of room mailboxes. No other type of recipient (including resource or equipment mailboxes) can be a member of a room list. The idea behind room lists is that they are a convenient way to segregate the different conference or meeting rooms available within an organization so that you have a room list per building or location. The room lists can then be used to select the best location when scheduling a meeting from Outlook or OWA. This is probably not a feature that is of much interest to small companies, but it can be valuable in large campus scenarios (like Microsoft's own Redmond HQ) where many multi-floor buildings can host meetings.

Room lists are managed with PowerShell. Here is an example of how to create a new "HQ Rooms" list. Note the use of the *IgnoreNamingPolicy* parameter to override the distribution list naming policy in force for the organization.

```
[PS] C:\> New-DistributionGroup -Name "HQ Rooms" -Members "Room 101", "Room 102", "Room 103"
-RoomList -IgnoreNamingPolicy
```

You can use this command to discover the room lists that already exist within a tenant:

```
[PS] C:\> Get-DistributionGroup -RecipientTypeDetails RoomList
```

Editing the membership of a room is done using the *Update-DistributionGroupMember*, *Add-DistributionGroupMember*, and *Remove-DistributionGroupMember* cmdlets. You will see an error if you try to add a recipient that is not a room mailbox to a room list.

The *Remove-DistributionGroup* cmdlet removes a room list.

```
[PS] C:\> Remove-DistributionGroup -Identity "Old HQ Rooms"
```

## The Outlook Places Service and Location Metadata for Room Mailboxes

The Outlook Places service helps users find suitable locations when scheduling meetings using the Room Finder feature in OWA and Outlook (both clients use the same component). Outlook Places uses metadata for room mailboxes together with room lists. Each room list represents a building, which the Room Finder groups into cities based on the City property for room mailboxes (Figure 6-2). After the user selects a building, they can choose a room or workspace to schedule. The list of rooms includes room characteristics using icons for capacity, video, audio, display, and wheelchair access (workspaces don't support room characteristics). You can create custom characteristics for a room using the Tags parameter for the *Set-Place* cmdlet. For example, you could use Tags to show that the room needs a special key to gain entrance.

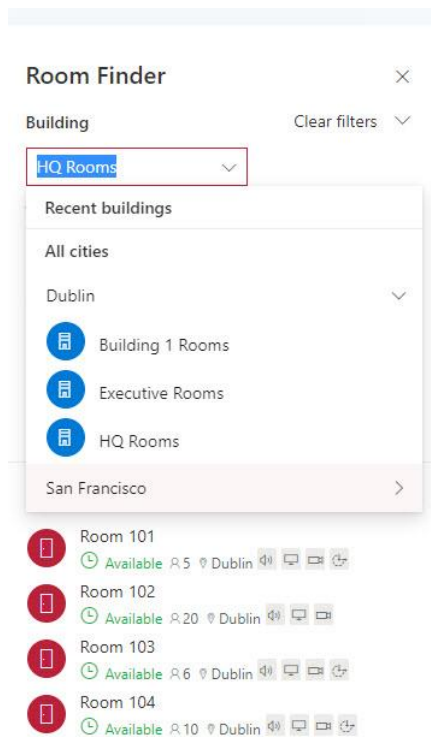


Figure 6-2: Outlook's room finder uses location metadata

The *Set-Place* cmdlet updates the location metadata for a room mailbox. For example, this snippet updates most of the properties available for a location:

```
[PS] C:\> Set-Place -Identity "SF Room 101" -CountryOrRegion "United States" -City "San Francisco" -
Floor 1 -Capacity 54 -Street "10 Sutter Street" -GeoCoordinates "37.790507; -122.400274" -Building
"Western HQ" -State CA -PostalCode 94104 -Phone "+1 206 177 4151" -Label "Training" -VideoDeviceName
"Crestron Flex UC-M150-T" -Tags "Training room"
```

It can take up to a day before updated metadata is available to clients. The *Get-Place* cmdlet retrieves information about a location. For example:

```
[PS] C:\> Get-Place -Identity DublinConfRoom@Office365itpros.com | Format-List
```

If the *PlacesEnabled* setting in the OWA mailbox policy assigned to a user mailbox is *\$True*, OWA displays location information to users in the room card when meetings are scheduled or viewed. If geocoordinates are available for a location, the Directions link calls the Bing Maps Locations API to generate directions to the location. Exchange Online uses a specific format to store geocoordinates for a location that is different from the format used by other applications (for instance, Google Maps uses a comma instead of a semi-colon to separate the latitude and longitude data). Outlook Mobile also consumes geocoordinates (if available) to show a map to a location, but only when scheduling a new meeting.

## General Mailbox Configuration

The *Get-MailboxMessageConfiguration* and *Set-MailboxMessageConfiguration* cmdlets retrieve and set the general properties of a mailbox. These settings control how OWA behaves. Although some are also respected by Outlook, you'll have to use a Group Policy Object or the [Office cloud policy service](#) to exert any real control over Outlook settings, including the roaming client settings supported by Outlook for Windows which are stored in mailboxes.

An example of how to use the *Set-MailboxMessageConfiguration* cmdlet is to create an autosignature for OWA to apply to new messages. This code defines some basic text for the autosignature and sets the default format for new messages created with OWA to "Plain text" (rather than the default HTML):

```
[PS] C:\> Set-MailboxMessageConfiguration -Identity "Kim Akers" -AutoAddSignature $True  
-SignatureText "From the desk of Kim Akers" -DefaultFormat PlainText
```

The equivalent command to create an HTML-format signature is shown below. In this instance, some simple HTML code creates the autosignature. The last parameter suppresses autosignatures for replies.

```
[PS] C:\> Set-MailboxMessageConfiguration -Identity "Kim Akers" -AutoAddSignature $True  
-SignatureHTML "<h3>From the Desk of Kim Akers</h3>" -DefaultFormat HTML -AutoAddSignatureOnReply  
$False
```

This is a simple example of creating a signature for OWA. You can download [a more comprehensive script from GitHub](#) to create a customized signature for every mailbox in the tenant. The signature includes user properties (name, title, etc.), a company logo, a clickable link for the user's email address, and links for Facebook and Twitter. Only OWA uses the signature defined by *Set-MailboxMessageConfiguration*; Outlook or Outlook Mobile use different autosignatures. If you want to from stop users changing the signature using OWA options after updating it in their mailboxes, you can do so using the technique [explained in this article](#).

**Autosignatures:** Apart from allowing users to create autosignatures using their client of choice or using the *Set-MailboxMessageConfiguration* cmdlet to create autosignatures for OWA with PowerShell, two other methods exist to control autosignatures. First, you can use an Exchange transport rule to insert an autosignature. This method has the advantage that it works for all clients and can use data extracted from Azure AD to populate the autosignature. Second, you can use a commercial ISV product to manage autosignatures. Examples of ISVs that provide this service include [Code Two Signatures](#), [Exclaimer](#), and [Outlook Signatures](#). ISV products cost, but they are very capable and provide the ability to manage autosignatures much more easily than is possible with either PowerShell or transport rules.

**Note:** Microsoft is in the process of introducing roaming signatures for Outlook for Windows with the goal of general availability in October 2022 (Microsoft 365 roadmap item 60371). Roaming signature means that Outlook stores the signature information in user mailboxes. One consequence of this initiative is that all Outlook clients will support multiple signatures in the same way that Outlook for Windows already does.

This feature is already available in OWA, and it can cause problems for scripts written to update user signatures with *Set-MailboxMessageConfiguration*. No PowerShell cmdlet or Graph API is currently available to manage roaming signatures.

Other common settings that you might consider updating for mailboxes include:

- **AlwaysShowBCC:** Display the *BCC:* control when composing new messages.
- **AlwaysShowFrom:** Display the *From:* control when composing new messages.
- **CheckForMissingAttachments:** If *True*, OWA checks messages before sending to check for the presence of attachments based on message text.
- **EmailComposeMode:** Set to *Inline* (default) to compose new messages in a pane within the same windows or *SeparateForm* to always launch a separate window for new messages.
- **EmptyDeletedItemsOnLogoff:** Controls whether OWA empties the Deleted Items folder when the user logs out.
- **HideDeletedItems:** Controls whether deleted items appear in conversation views. By default, clients show deleted items, so this property is set to *False*. Both Outlook and OWA respect this property.
- **IsReplyAllTheDefaultResponse:** Controls whether reply-all is the default response for messages. If *True*, a response to a message includes all recipients (this is the default value). Set the property to *False* to force OWA to create responses only addressed to the sender of the original message.
- **ReadReceiptResponse:** Controls how OWA generates read receipts for new messages delivered to the mailbox. The values are *DoNotAutomaticallySend* (prompt), *AlwaysSend*, and *NeverSend*. This setting applies only to OWA. See [this article for an explanation](#) of how read receipts work and how to control the relevant settings for OWA and Outlook.
- **NewItemNotification:** Set to *All* (default) to instruct OWA to signal the arrival of new messages (divided into the categories of email, fax, and voicemail) in every way that it can. Other values include *Sound* (a tone announces the arrival) and *EmailToast* (a pop-up “toast” notification signals the arrival).
- **PreviewMarkAsReadBehavior:** Controls how OWA sets the read status of messages. The default is *OnSelectionChange*, meaning that OWA marks a message as read if the user selects another message. If set to *Delayed*, OWA marks the item as read if the user spends more than the time specified in the **PreviewMarkAsReadDelayTime** property (by default, 5 seconds). You can also set this value to *Never*, meaning that the unread status of a message never changes no matter what the user does in the preview pane.
- **MailSendUndoInterval** sets how long OWA waits before sending a message. The interval allows the sender to stop the message if they've made a mistake or need to add something.

Microsoft updates PowerShell to allow scripting control over OWA options as new features appear in Exchange Online.

## MailTips

A MailTip is some informational text displayed by Exchange when a recipient is added to a message in Outlook desktop and OWA. Some MailTips are system-generated, and you can't affect the text they display, such as those generated when messages are addressed to moderated recipients or when a recipient's mailbox is full (a situation more unusual in the cloud than on-premises given the size of Exchange Online mailboxes). Others are controlled through the organizational configuration.

Although Exchange Online includes [protection against email reply-all storms](#), this only works for relatively large tenants. MailTips help users to avoid doing something bad like creating a reply-all email storm by inadvertently sending a message to a large distribution list. Several settings in the Exchange Online organization configuration control how clients see MailTips. To view the settings, run this command:

```
[PS] C:\> Get-OrganizationConfig | Format-List mailtip*
```

```
MailTipsAllTipsEnabled : True
MailTipsExternalRecipientsTipsEnabled : True
MailTipsGroupMetricsEnabled : True
MailTipsLargeAudienceThreshold : 10
MailTipsMailboxSourcedTipsEnabled : True
```

The values have the following meanings:

- **MailTipsAllTipsEnabled:** Set to True to instruct clients to use MailTips.
- **MailTipsExternalRecipientsTipsEnabled:** Set to True to have clients highlight messages addressed to external recipients.
- **MailTipsGroupMetricsEnabled:** Set to True to have Exchange Online calculate the number of members in distribution lists. This data is used to establish if the large audience threshold is exceeded in messages. A background process checks distribution lists periodically, so don't expect the numbers used to be 100% accurate.
- **MailTipsLargeAudienceThreshold:** The default is 25. It's lower here as a reminder that there might be a better way to share information with large audiences, such as a post in a Teams channel. Of course, if you don't use Teams, you could increase the threshold right up to the maximum number of recipients an Exchange Online user can address in a message (1,000).
- **MailTipsMailboxSourcedTipsEnabled:** Set to True to have Exchange Online look at mailbox data such as auto-reply messages to generate MailTips.

## Custom MailTips

Custom MailTips are text messages of up to 175 characters assigned to any valid mail-enabled recipient type including mailboxes, shared mailboxes, group mailboxes, mail contacts, distribution lists, and dynamic distribution lists.

For example, these commands set custom MailTips for a mailbox, a Microsoft 365 group, and the mail user object created for a guest account:

```
[PS] C:\> Set-Mailbox -Identity Oisin.Johnston -MailTip "Working 9 to 12 at present. Ping me on Teams if urgent"
[PS] C:\> Set-UnifiedGroup -Identity BankingTeam -MailTip "Messages to this Group are delivered to external guest members"
[PS] C:\> Set-MailUser vasil_michev.org#EXT# -MailTip "Be Careful with Vasil"
```

You can also look for objects with custom MailTips. For instance, here's how to do it for mailboxes:

```
Get-ExoMailbox -ResultSize Unlimited -Properties MailTip | ? {$_.MailTip -ne $Null} | Format-List
DisplayName, MailTip

DisplayName : Oisin Johnston
MailTip      : <html>
              <body>
                Working 9 to 12 at present. Ping me on Teams if urgent
              </body>
              </html>
```

Exchange Online stores MailTips in HTML format.

Users or the owners of distribution lists or groups cannot set MailTips unless their account holds the Exchange Online administrator role. Like many mailbox settings, it takes a few hours before clients pick up MailTip changes.

## Translated MailTips

Exchange Online operates in a multinational, multilingual environment. When you create a custom MailTip, it becomes the default for all languages. To create language-specific translations for a MailTip, you update the multi-valued *MailTipTranslations* property for an object. For example, this command sets up Spanish, French,

and German translations for a mailbox. Clients configured in these languages display the language-specific value. If no value exists for the client language, the default MailTip is used.

```
[PS] C:\> Set-Mailbox -MailTipTranslations @{Add="ES: Buzón no en uso activo", "FR: Boîte aux lettres non utilisée", "DE: Mailbox nicht aktiv genutzt"} -Identity CServices
```

## MailTips for AutoReplies

Exchange generates automatic MailTips for mailboxes that have an autoreply set to inform users of the current mailbox status. You can't control the generation of autoreply tips. Users might continue to see an autoreply MailTip even after the autoreply expires for a mailbox or is explicitly disabled. This is because Exchange Online caches autoreply information to prevent it from having to look up mailbox autoreply data every time a user adds a recipient to a message. The cached data is refreshed every hour. Exchange Online supports MailTips in the same way as on-premises Exchange.

## Folder Level Permissions

You can add folder-level permissions to allow a user to access a folder in someone else's mailbox. For instance, you might want someone to check new messages that arrive when you are on vacation. In this example, we permit Marc Vigneau to access the Inbox folder for Kim Akers:

```
[PS] C:\> Add-MailboxFolderPermission -Identity "Kim Akers:\inbox" -User "Marc Vigneau" -AccessRights Reviewer
```

Automapping is the process Outlook uses to connect to mailboxes automatically. This does not apply to folder-level permissions. The user who receives the permission must add the other user's folder as a shared folder. In addition, it can take several hours for Exchange to propagate the new permission and make it effective.

## Controlling Email Sent by Delegates

Delegate settings refer to the ability to give a user the right to access another person's mailbox. These rights include the ability to send messages on behalf of the mailbox owner or as the mailbox owner, or the *SendOnBehalfOf* and *SendAs* permissions. By default, when messages are sent by a delegate using these permissions, the outbound messages are stored in the Sent Items folder of the delegate's mailbox. This is because the *MessageCopyForSentAsEnabled* and *MessageCopyForSendOnBehalfEnabled* settings are set to *\$False*.

Mailbox owners often want to have copies of messages sent to them by a delegate. To force Exchange to create copies of these messages for the mailbox owner, set the properties to *\$True*. For example:

```
[PS] C:\> Set-Mailbox -Identity TRedmond -MessageCopyForSendOnBehalfEnabled $True -MessageCopyForSentAsEnabled $True
```

Similar control can be exerted over messages sent by delegates for a shared mailbox.

## Junk Mail Settings

The *Get-MailboxJunkEmailConfiguration* and *Set-MailboxJunkEmailConfiguration* cmdlets handle how to process junk email for a mailbox. A hidden Inbox rule (called Junk E-mail rule) holds the settings, which control the delivery of messages to the Junk Email folder based on the spam confidence level (SCL) and the safelist defined for the mailbox. Exchange Online applies the rule to all inbound messages – but only if it is enabled. It's therefore important to ensure that the rule is enabled for all mailboxes. To check, we can run a command like this:

```
[PS] C:\> $Mbx = Get-ExoMailbox -RecipientTypeDetails UserMailbox -ResultSize Unlimited
ForEach ($M in $Mbx) {
```

```
If ((Get-MailboxJunkEMailConfiguration -Identity $M.Alias).Enabled -eq $False) {
    Write-Host $M.DisplayName "mailbox has junk email disabled" }}
```

New mailboxes get a default rule to handle junk mail processing, but only after the mailbox is opened in Outlook (in cached mode) or OWA. The most interesting of the rule settings are the trusted senders and blocked sender lists. Both are multivalued properties. The following example adds an entry to the blocked senders and domains list (because we do not want to receive messages from this domain), specifies that Exchange should always treat the user's contacts as safe senders, and replaces the set of trusted domains.

```
[PS] C:\> Set-MailboxJunkEMailConfiguration -Identity "Kim Akers"
-BlockedSendersAndDomains @{Add="Badgirls.com"} -ContactsTrusted $True
-TrustedSendersAndDomains "Microsoft.com", "Office365.com", "Outlook.com"
```

Note that the rule supports a maximum of 1,024 contacts, but only if the *ContactsTrusted* setting in the mailbox's junk mail rule is set to *\$True*.

## Viva Insights Mailbox Settings

Viva Insights (previously MyAnalytics - see Chapter 6 in the Companion Volume) generates messages delivered to user mailboxes which you might want to control. These are system messages because they do not travel through the Exchange Online transport service. Instead, Microsoft 365 inserts the messages directly into mailboxes.

Several Viva Insights features can be individually enabled or disabled using the *Set-MyAnalyticsFeatureConfig* cmdlet. The same cmdlet also controls whether mailboxes opt into the collection of Microsoft Graph signals for email and calendar activity used to generate insights. The individual features are:

- **Email-digest:** If enabled, a mailbox receives a monthly digest generated from their activity containing insights for collaboration, network, focus, and wellbeing.
- **Dashboard:** If enabled, a mailbox has access to the Viva Insights dashboard.
- **Add-in:** If enabled for a mailbox, Outlook loads the Insights add-in.

The default values for these settings come from the MyAnalytics (Viva Insights) tenant settings in the Org settings of the Microsoft 365 admin center. New mailboxes are enabled automatically for Viva Insights. Users can disable Insights for their mailbox through the [Viva Insights home page](#). Alternatively, administrators can run the *Set-MyAnalyticsFeatureConfig* to control the settings. For instance, to disable Viva Insights for a mailbox, use a command like:

```
[PS] C:\> Set-MyAnalyticsFeatureConfig -Identity Vasil.Michev@office365itpros.com -PrivacyMode "opt-out"
```

Opting out from Insights disables all features. Mailboxes can be enabled for Insights but disabled for individual features. For example, many users don't like the monthly digest arriving in their inbox. To stop the digest, run *Set-MyAnalyticsFeatureConfig* to set the enabled state for the Email-digest feature to False:

```
[PS] C:\> Set-MyAnalyticsFeatureConfig -Identity Vasil.Michev@office365itpros.com -PrivacyMode "opt-in" -Feature Email-Digest -IsEnabled $False
```

You cannot enable or disable multiple features in one command. If you want to disable both the email digest and Outlook add-in, you must run *Set-MyAnalyticsFeatureConfig* twice.

Organizations can disable the daily briefing by updating the **Briefing email** settings in the Org settings section of the Microsoft 365 admin center. On an individual level, users can opt-out by unsubscribing, or an administrator can run the *Set-UserBriefingConfig* cmdlet. For example:

```
[PS] C:\> Set-UserBriefingConfig -Identity Oisin.Johnston@office365itpros.com -Enabled $False
```



Unlike most cmdlets in the Exchange Online Management module, the *Set-MyAnalyticsFeatureConfig* and *Set-UserBriefingConfig* cmdlets accept the user principal name for an account as the input identity instead of the display name, alias, object identifier, or distinguished name as with other Exchange Online cmdlets.

## Controlling Email Forwarding

When an organization assigns an Exchange Online mailbox to a user, they probably want that person to use the mailbox for email and keep messages in the mailbox for compliance purposes. It is possible that some users prefer to use another email system and will forward messages to that system using either a forwarding address created using OWA settings or a rule. No matter how forwarding occurs, messages leave the control of retention and other data governance policies, which is a bad thing. It's also the case that attackers often plant a mail forwarding rule to capture email from accounts that they wish to learn more about before launching a business email compromise attack. For instance, they might try to insert a mail forwarding rule or set up a forwarding address in the CFO's mailbox to discover information about the trading patterns of a company.

For historic reasons, Exchange Online supports two methods to forward email from a mailbox. The methods use different mailbox attributes, but both instruct the transport service to redirect messages to another SMTP address with the option to also deliver a copy to the original mailbox.

- A user can set the *ForwardingSmtpAddress* through OWA options or an administrator can set email forwarding for a mailbox through the Microsoft 365 admin center. Setting the *ForwardingSmtpAddress* attribute is the preferred approach. Because the mailbox owner can set up email forwarding through OWA, if an administrator creates a forward for a mailbox, its existence is known to the mailbox owner.
- An administrator can set the *ForwardingAddress* attribute through EAC through the **Manage email forwarding** option or by running the *Set-Mailbox* cmdlet. This redirect is invisible to the mailbox owner.

A significant difference between the two attributes is that *ForwardingSmtpAddress* supports any valid SMTP address, including those belonging to external domains. Microsoft recommends that you should use *ForwardingSmtpAddress* whenever possible. *ForwardingAddress* only supports addresses that are known to the tenant, including mail-enabled contacts pointing to external addresses. The properties of a mailbox can have both forwarding attributes set with different SMTP addresses. When this happens, Exchange Online will forward copies of all inbound messages to both addresses.

Here is an example of using the *Set-Mailbox* cmdlet to set a forwarding address for a mailbox. In this instance, the *DeliverToMailboxAndForward* property is set to *\$True* to instruct Exchange Online to forward a copy of any inbound messages to the supplied address and keep a copy in the mailbox.

```
[PS] C:\> Set-Mailbox -Identity "Andy Ruth" -ForwardingSmtpAddress Andy.Ruth@yandex.com -DeliverToMailboxAndForward $True
```

When you set up email forwarding for a user through the Microsoft 365 admin center, the forwarding address is written into the mailbox's *ForwardingSmtpAddress* attribute, and any value found in the *ForwardingAddress* attribute is cleared. Apart from checking that the input address is formatted properly, no attempt is made to confirm that messages can be redirected to the email address input as the forwarding destination. The administrator is notified that *"the mailbox owner will be able to view and change these forwarding settings"* (because they will be able to see the redirect address through OWA Options). If the need exists to hide forwarding, the administrator should use EAC or PowerShell to set up the redirect through the *ForwardingAddress* attribute.

Both the *Set-Mailbox* cmdlet and the Microsoft 365 admin center flag warnings if redirect addresses are detected in both attributes. *Set-Mailbox* will allow you to write redirect addresses into the two attributes, but

the Microsoft 365 admin center insists on removing the redirect contained in the *ForwardingAddress* attribute before it will update the email forwarding settings.

## Blocking Email Forwarding

Where organizations once allowed all users to forward messages, it's now more common to find restrictions in place. The easiest way to apply central control over email forwarding is to use the organization's [outbound spam filter policy](#) (see the Mail Flow chapter) to disable forwarding for mailboxes. If you allow some users to forward messages, you should check what email users forward and why and implement further blocks where necessary. Three approaches are available.

1. **Check for and stop users forwarding messages.** The *Get-ExoMailbox* command below lists all mailboxes with a populated forwarding address. We check both *ForwardingSmtpAddress*, which can redirect to an external address and *ForwardingAddress* (which only accepts internal addresses).

```
[PS] Get-ExoMailbox -RecipientTypeDetails UserMailbox -Filter {ForwardingAddress -ne $null -or ForwardingSMTPAddress -ne $Null} -ResultSize Unlimited -Properties ForwardingSmtpAddress, DeliverToMailboxAndForward | Format-Table DisplayName, ForwardingSmtpAddress, DeliverToMailboxAndForward -AutoSize
```

DisplayName	ForwardingSmtpAddress	DeliverToMailboxAndForward
Vasil Michev (Technical Guru)	smtp:vasil@contoso.com	True
Ståle Hansen	smtp:stale.hansen@fabrikam.com	True
Eoin Redmond	smtp:Eoin@contoso.com	True

After checking who's forwarding email outside the organization, we can remove the automatic forwarding with a variant of the command:

```
[PS] C:\> Get-ExoMailbox -RecipientTypeDetails UserMailbox -Filter {ForwardingSmtpAddress -ne $Null} -ResultSize Unlimited | Set-Mailbox -ForwardingSmtpAddress $Null
```

2. **Remove the ability of users to create rules to forward messages to external domains.** When forwarding is disabled for a domain, Exchange Online detects that this is the case and delivers new inbound messages to the original mailbox. You can block forwarding for all external domains as follows:

```
[PS] C:\> Get-RemoteDomain | Set-RemoteDomain -AutoForwardEnabled $False
```

Because good business reasons might exist to justify forwarding to certain external domains, a more granular approach can be taken to disable the ability for users to forward messages to selective external domains. This approach is often used to stop users from forwarding messages to consumer email services. For instance, let's assume that you don't want users to be able to forward messages to Yahoo.com accounts. You need to define Yahoo.com as a new external domain and then block forwarding to that domain. Here's how:

```
[PS] C:\> New-RemoteDomain -DomainName "*.yahoo.com" -Name "Yahoo.com"
[PS] C:\> Set-RemoteDomain -Identity "Yahoo.com" -AutoForwardEnabled $False
```

3. **Remove the ability of users to set forwarding through OWA options.** This approach needs you to update the default user role assignment policy because the OWA options are controlled by the policy. If you'd like the challenge of figuring out how to make the necessary changes, much of the same technique is used later to allow users to edit but not create distribution lists. You can find a [complete description of the required steps](#) online.

In the past, transport rules were commonly used to block forwarded messages to specific domains. The implementation of the block in the outbound spam policy removes the need to use these rules and they should be retired.

Power Automate offers several templates to allow users to forward new messages moved into a folder. These flows can forward messages to external addresses. None of the standard Exchange Online management tools stop Power Automate forwarding messages. However, [a transport rule](#) can block the forwarding of messages using Power Platform (flow).

Naturally, as in the case of any action that removes a facility that some people might be using, it's a good idea to capture the business reasons why a policy exists for email forwarding and to communicate why blocks are being enforced and how the changes might affect users before anything is done.

## Reporting Forwarding

If the outbound spam policy for the organization allows some users to forward messages, the Auto forwarded message report available in the Reports section of the EAC gives an insight into forwarding activity. In addition, the Microsoft 365 Defender admin center includes a default alert policy called *Creation of forwarding/redirect rule* to flag alerts when users create rules to forward messages outside the tenant. To round things out, it's possible to use PowerShell to check which mailboxes are forwarding messages. The steps required are:

- Creates a collection of user and shared mailboxes.
- Checks if the mailbox has a forwarding address set and reports it if found. We don't check the *ForwardingAddress* property because although it can forward messages to an address known in the directory, we are more concerned about mail going outside the tenant.
- Checks if any rules exist in the mailbox. If rules exist, check if any forward messages (including forward as an attachment).
- Check the forwarding addresses to see if the recipient is known to the tenant directory (including guest accounts) and report any forwarding address found.
- Optionally, if the forwarding address is unknown (does not exist in the directory), remove the rule from the mailbox.

You can [download an example script from GitHub](#).

Even if you did not remove the offending rules, you could run the script periodically to discover whether people are attempting to forward emails outside the organization, and if so, what are the target domains.

## Mailboxes Forwarding Email Must be Licensed

If you're in a situation where someone leaves the business and you want to keep their mailbox active and forward new messages to someone else for processing, remember that the mailbox (account) must remain licensed to allow forwarding to happen. In these scenarios, it might be better to convert the mailbox to a shared mailbox. A license is only needed then if the mailbox is larger than 50 GB or has an archive.

## Maximum Message Size

Exchange Online allows mailboxes to send and receive messages up to a maximum of 150 MB. The actual size of messages supported by a mailbox is set by the *MaxSendSize* and *MaxReceiveSize* properties, which can be changed using the *Set-Mailbox* cmdlet or by editing mailbox properties through the EAC (the values are available through the mailbox features page). For example, this command sets the send and receive size to 100 MB for the Kim Akers mailbox:

```
[PS] C:\> Set-Mailbox -Identity 'Kim Akers' -MaxSendSize 100MB -MaxReceiveSize 100MB
```

Although it is great to be able to send large messages, it is quite another matter to make sure that the recipient will be able to receive them as you probably have zero influence over the connectors and configuration of the email systems involved in the transfer of the message after it leaves Exchange Online. This is especially important in a hybrid situation as it is probable that the on-premises servers support

message sizes significantly smaller than the values supported by Exchange Online. It is also important to understand that the message size used for this purpose is the size of the message after it is coded into BASE64/MIME format to allow it to be accepted by other mail systems. This process can add up to a third to the size of a message, so a 60 MB message as seen by the user might become an 80 MB message when presented to the transport system for transmission. In turn, this might end up exceeding the permitted threshold and cause some bewilderment to the user.

## Enabling Third-Party Cloud Attachments

Working inside Exchange Online, it's natural to use "cloudy attachments" stored in SharePoint Online or OneDrive for Business document libraries. Your company might use other cloud-based document storage repositories like Dropbox or Google Drive, and you might want to allow users to add attachments to messages from these repositories. The *AdditionalStorageProvidersAvailable* setting in OWA mailbox policies controls access to both first-party and third-party storage providers. By default, this setting is *\$True*. If you block access to third-party storage providers, you must update OWA mailbox policies to ensure that they exert the same control over users. For example:

```
[PS] C:\> Get-OWAMailboxPolicy | ? {$_.ThirdPartyFileProvidersEnabled -eq $False} | Set-OWAMailboxPolicy -AdditionalStorageProvidersAvailable $False
```

OWA mailbox policy settings only apply to OWA clients and do not affect Outlook desktop or mobile. After making the change to the policy, wait an hour or so to allow the cached policy to be refreshed.

## Mailbox Quotas

The size of the assigned storage quota is a major difference between on-premises mailboxes and their cloud counterparts. Although it is still common for on-premises mailboxes to have relatively small quotas of between 2 and 10 GB, Exchange Online assigns a basic 100 GB quota to mailboxes with enterprise E3 and E5 plans, 50 GB to E1 and education plans, and 2 GB to frontline worker plans. Competitive pressure is one reason why Microsoft offers very large mailbox quotas for Exchange Online. For instance, Google Workspace plans include between 30 GB and "as much as you need" storage. Table 6-1 lists the current quotas assigned to [different types of mailboxes](#).

	<b>Frontline worker (F3)</b>	<b>Enterprise E1 (and Gov/Edu Equiv.)</b>	<b>E3 and E5</b>
Primary mailbox size	2 GB	50 GB	100 GB
Archive mailbox size	N/A	50 GB	Unlimited
Shared mailbox size	N/A	50 GB	50 GB
Resource mailbox size	10 GB	50 GB	100 GB
Group mailboxes	50 GB	50 GB	100 GB
Public folder mailboxes	N/A	50 GB	100 GB

Table 6-1: Exchange Online mailbox sizes

If a shared mailbox has an Exchange Online Plan 2 license, its quota increases to 100 GB, and it can have an archive. In the past, some unlicensed shared mailboxes had a 100 GB quota. These mailboxes can keep the erroneous quota if their status does not change (for example, administrators do not convert a shared mailbox to a user mailbox and back to become a shared mailbox). New shared mailboxes receive a 50 GB quota.

You do not have to assign the full quota to mailboxes and can restrict users to lower amounts. As you can see in the PowerShell example below, three properties control how Exchange applies a mailbox quota:

- The *IssueWarningQuota* property tells Exchange the point at which nagging messages should be sent to the mailbox owner to tell them that they are approaching the quota limit.

- The *ProhibitSendQuota* property marks the limit at which Exchange will no longer accept new outbound messages from the mailbox.
- The *ProhibitSendReceiveQuota* property tells Exchange when to cut off both outbound and inbound service to the mailbox.

For example, this code sets a warning limit of 23 GB, stops the user from sending messages at 25 GB, and stops the mailbox from receiving messages when the mailbox size reaches 30 GB.

```
[PS] C:\> Set-Mailbox -Identity TRedmond -ProhibitSendQuota 25GB -ProhibitSendReceiveQuota 30GB -IssueWarningQuota 23GB
```

You can scan for user mailboxes that have a certain quota and increase their quota with a command like:

```
[PS] C:\> $Mbx = (Get-ExoMailbox -RecipientTypeDetails UserMailbox -PropertySet Quota -ResultSize Unlimited)
[double]$QuotaCheck = 75161927680 # bytes - 70 GB in this case
Foreach ($M in $Mbx) {
    [double]$Quota = $M.ProhibitSendQuota -replace "(.*\(|,| [a-z]*\)", "" # value in bytes
    If ($Quota -lt $QuotaCheck)
    { Write-Host "Updating" $M.UserPrincipalName "quota of" $M.ProhibitSendQuota "to 75GB"
      Set-Mailbox -Identity $M.UserPrincipalName -ProhibitSendQuota 74GB -ProhibitSendReceiveQuota 75GB -IssueWarningQuota 72GB}
}
```

Clients display available quotas in different ways. The mailbox settings section in Outlook's "backstage" (available from the File menu) tells users how much quota they have consumed and how much remains. OWA gives more comprehensive information in the Storage section of its settings. OWA also displays a warning message to users at the bottom of the folder list once they consume 90% of their mailbox quota. The message contains a link to the storage section of settings to allow the user to remove some messages from the mailbox to free quota.

### Extra Space Available to Exchange Online Mailboxes

In addition to their basic quota, Exchange Online mailboxes also have a recoverable items quota of 30 GB, which is automatically increased to 100 GB if the mailbox is put on hold. The added quota accommodates the need to keep held items in mailboxes for extended periods. If necessary, it is possible to increase the recoverable items quota past 100 GB, but only by filing a support request with Microsoft. Altogether, the total available storage available per enterprise user mailbox is between 230-300 GB made up of the primary mailbox, primary archive, and recoverable items.

Exchange Online mailboxes also store data created by Exchange, other Office 365 applications, and the Microsoft 365 substrate in hidden folders. This data is inaccessible to users. The amount of system data can exceed user data, meaning that the overall storage occupied by a mailbox can be much larger than anyone expects. We'll dive into this aspect later.

**Moving very large on-premises mailboxes to Exchange Online:** The 100 GB mailbox quota assigned to enterprise Exchange Online user mailboxes is more than enough to store email for most people. Some on-premises users might have mailboxes larger than 100 GB, and if this is the case, you'll need to shrink these mailboxes before the Exchange Mailbox Replication Service (MRS) can move them to Exchange Online. Simply because of the amount of data to transfer, large mailboxes take longer to move and are somewhat more problematic than smaller mailboxes. With this in mind, it's a good idea to review the sizes of all large on-premises mailboxes before starting the migration process to make sure that they are at least 10 GB under the 100 GB quota (to allow for some growth) and if not, to reduce their size. It's also important to remove any items larger than 150 MB as items larger than this aren't supported by Exchange Online. If a mailbox is on litigation hold, a large proportion of the mailbox might be occupied by held items, in which case the Recoverable Items folder will be quite big. Users are unlikely to want to prune items out of a massive mailbox and will probably take too long to make a serious dent in the size, so the best approach is

to assign a mailbox retention policy to the oversized on-premises mailboxes to have the Exchange Managed Folder Assistant remove older items in the background (for instance, by deleting any message older than 2 years or moving old messages into the archive mailbox). Once the assistant has processed the mailboxes, you should be all set to move them to the cloud.

## Offline Storage

In a practical sense, once a quota is larger than 20 GB, adding more storage is nearly meaningless for the average user because it takes so long to fill the remaining quota, even if they move large amounts of old mail from their on-premises mailbox. From a user perspective, modern clients hide many of the complexities of large mailboxes. Outlook for Windows, for instance, supports up to 100,000 items per folder. Some degradation is likely when approaching any client limit, so users should be coached to keep the number of items in a folder well under this amount. Clients that work offline can also control the amount of information that they download from a mailbox. Here are some points to remember about offline storage:

- Outlook desktop's OST slider allows users to choose a synchronization period. Outlook desktop supports a range from 3 days to "everything." Outlook uses the synchronization period to decide what to download from the user-visible folders in the mailbox on the server to their offline folder replica file (OST). Virtual desktop clients often use smaller synchronization periods to reduce the amount of synchronized data. For laptops and other non-shared PCs, restricting the synchronization period to a year or less is an effective way to ensure that the OST stays a reasonable size (in the 4 GB to 10 GB range, depending on the number of items synchronized to the OST). The performance of larger OST files can suffer even on PCs equipped with fast SSDs.
- The OST slider does not affect the synchronization of the Contacts or Calendar folder. The assumption here is that users will want access to their complete calendar and all their contacts when working offline.
- Registry settings can be set by group policy or manually to control [the "slider window"](#) and whether Outlook synchronizes the contents of [public folder favorites and shared mailboxes](#). It is a good idea to review the set of shared mailboxes configured for use with Outlook periodically to remove mailboxes that are seldom accessed.
- Adding extra calendars or having too many items in calendar folders can [cause synchronization problems](#) with any version of Outlook for Windows.
- The [maximum size of an OST file is 50 GB](#). If your mailbox is larger, move the slider to reduce the amount of synchronized data. Remember that if you synchronize shared folders and other mailboxes, they also must fit within the 50 GB limit.
- Outlook must run in cached Exchange mode to access Groups.

Tenants considering a virtual desktop infrastructure (VDI) deployment need to pay extra attention to the amount of data cached by Outlook and perhaps consider OWA as a replacement client in these environments. If Outlook must be used, it's a good idea to consider VDI-specific solutions like the [FSLogix Office Container](#) (acquired by Microsoft in November 2018) to achieve good performance for users.

## Folder Associated Information

Administrators sometimes comment that the number of items reported in a folder by a client differs from the server-side data that they see when running a cmdlet like *Get-ExoMailboxFolderStatistics*. The difference is the hidden items (folder associated information or items, known also as FAIs) that Exchange stores in folders for different purposes such as to hold configuration settings (like the town chosen for a weather display in Outlook's calendar), RSS feeds, and retention policies. The FAIs are usually small and do not take up much space in the context of an overall mailbox, but they are very important to Exchange and its clients. The Inbox

folder is the location for most FAs. A test revealed that Outlook reported 8,734 items in the Inbox, but more were found with the following command:

```
[PS] C:\> Get-ExoMailboxFolderStatistics -Identity TRedmond -FolderScope Inbox | Select
ItemsInFolder

ItemsInFolder
-----
          9585
```

FAs account for most of the difference between the numbers reported for the Inbox by Outlook and the cmdlet. You can check this by using the MFCMAPI utility to examine the folder's associated items table.

## Other System Items in Mailboxes

A mailbox holds both items accessible to users through clients and system items hidden in folders that clients never reveal. System folders exist to hold data for application purposes. On-premises mailboxes have far fewer system folders than cloud mailboxes do because the Microsoft 365 substrate and applications use Exchange Online mailboxes to store a broader range of information. Most of the system data is in the Non-IPM part of the mailbox, which email clients do not access. Their focus is on content stored in the IPM (interpersonal messaging) section under a root folder called Top of Information Store. Folders in the IPM part of the mailbox include well-known folders like Inbox, Sent Items, and Calendar. Some applications, like To-Do, store data in IPM folders because the data is also exposed and used in email clients.

A good example of system data storage is how the substrate captures Teams compliance records in the *TeamsMessagesData* folder in the non-IPM section of group mailboxes (for channel conversations) and user mailboxes (for personal chats). System folders are usually hidden and do not appear in clients. The substrate creates these items to support Microsoft Search, machine learning, artificial intelligence, and other centralized services.

To illustrate just how much invisible data exists in a mailbox, let's use the *Get-ExoMailboxFolderStatistics* cmdlet to reveal what's stored in the *Non-IPM* part.

```
[PS] C:\> $Folders = Get-ExoMailboxFolderStatistics -Identity TRedmond -FolderScope NonIPMRoot
$Folders.Count
304
[PS] C:\> $Folders[0] | Select Identity, ItemsInFolderAndSubFolders, FolderAndSubFolderSize

Identity  ItemsInFolderAndSubFolders FolderAndSubFolderSize
-----
TRedmond\                451900 20.24 GB (21,727,913,262 bytes)
```

304 folders holding 451,900 items and occupying 20.24 GB. By comparison, here's what is in the visible folders in the same mailbox.

```
[PS] C:\> $Visible = Get-ExoMailboxFolderStatistics -Identity TRedmond -FolderScope All
$Visible.Count
92
[PS] C:\> $Visible[0] | Select Identity, ItemsInFolderAndSubFolders, FolderAndSubFolderSize

Identity  ItemsInFolderAndSubFolders FolderAndSubFolderSize
-----
TRedmond\Top of Information Store 38360 4.907 GB (5,268,429,234 bytes)
```

So, 92 folders against 304 and only 38,360 items in 4.907 GB against the whopping 451,900 stored in the invisible folders. If we look at the data in more detail and review the ten largest folders based on number of items, we gain some insight into the situation (you might find different folders in different mailboxes):

```
[PS] C:\> $Folders | Sort ItemsInFolder -Descending | Select -First 10 | Format-Table -Property
@{e="Name"; width=30}, @{e="FolderSize"; width=30}, ItemsInFolder
```

Name	FolderSize	ItemsInFolder
AllItems	3.901 GB (4,189,011,735 bytes)	24244
NoArchiveTagSearchFolder853...	3.869 GB (4,154,389,586 bytes)	24122
Audits	145.1 MB (152,185,788 bytes)	22781
TeamsMessagesData	1.667 GB (1,789,639,942 bytes)	16402
SPOOLS	1.155 GB (1,239,922,031 bytes)	6517
SpoolsSearchFolder	1.155 GB (1,239,922,031 bytes)	6517
Inbox	1.398 GB (1,501,199,889 bytes)	5270
Unified Inbox	1.367 GB (1,467,696,501 bytes)	5133
Sent Items	786.9 MB (825,137,524 bytes)	4731
EdgeSyncEntities	39.61 MB (41,531,884 bytes)	4675

Microsoft does not document how Exchange or other workloads use the items in hidden folders, so a little guesswork is necessary. Among the folders you might find when looking through the non-IPM part of a mailbox are:

- System folders are hidden from users. For instance, The **Audits** folder holds mailbox audit records (which are also transmitted to the audit log). The **Calendar Logging** folder holds change details for calendar items. Folders are present to store compliance data, such as **TeamsMessagesData** (Teams compliance records) and **Yammer**. OWA [stores items used for its People Favorites feature](#) in hidden “persona” folders while the [Planner integration with the Microsoft 365 message center](#) uses a hidden folder to hold calendar reminders for assigned tasks.
- Data stored by apps: Forms stores forms as PDFs and responses to forms as CSV files in the **c9a559d2-7aab-4f13-a6ed-e7e9c52aec87** folder under the **ApplicationDataRoot** folder. Likewise, Sway uses the **905fcf26-4eb7-48a0-9ff0-8dcc7194b5ba** subfolder to store any files created by the user in HTML format. In both cases, the application attaches its files to email messages to avoid the need to create extra MAPI message classes. Applications often store information in Exchange Online to include their data in search indexes and expose it to content searches.
- Search Folders used by different features. MAPI search folders don’t store copies of mailbox items, so they don’t occupy any storage. Instead, they store links to items chosen by applying search criteria (such as “All PDF attachments”). The real items exist in other folders. Even so, *Get-ExoMailboxFolderStatistics* reports the full item size. For example, the **GraphFilesAndWorkingSetSearchFolder** folder records details of attachments and files accessed by a user and is used by the OWA Files feature; the folder starting with **OwaFV15.1AllFocused** holds items in the *Focused* view for the Focused Inbox (another folder starting with **OwaFV15.1AllOther** holds items in the *Other* view).

To discover the set of mailbox folders that consume storage quota, you can use a command like this:

```
[PS] C:\> Get-ExoMailboxFolderStatistics -Identity James.Ryan -Folderscope NonIPMRoot | where
{($_.TargetQuota -like 'User') -and ($_.FolderSize -like "*GB*" -OR $_.FolderSize -like "*MB*")} |
Sort Name | Format-Table Name, FolderSize, ItemsInFolder
```

Administrators don’t need to worry about Exchange mailbox storage in the same way as they do on-premises. This underscores a key point about the cloud: you don’t worry about how a service is delivered if the service works reliably and meets your needs.

**Folder Limit for Cmdlets:** The *Get-ExoMailboxFolderStatistics* and *Get-MailboxFolderStatistics* cmdlets both can return data for a maximum of 1,000 folders. Although it’s uncommon to find mailboxes with more than a thousand folders, the storage of data (“digital twins”) in mailboxes by the Microsoft 365 substrate on behalf of other applications has increased over time. Users aren’t aware of the issue because these are system folders stored in the non-IPM part of the mailbox, but scripts can run into problems if they encounter a mailbox with more than a thousand folders.



## User Role Assignment Policies

Exchange Online controls access to features using role-based access control (RBAC), a method of ensuring that people who need to do something have the right to do it. By default, Exchange Online has a single-user role assignment policy to control several aspects of user functionality. Elsewhere in the book, we update a user role assignment policy to control whether users can create distribution lists and add personal retention tags to their mailbox retention policy. For now, all we need to do is introduce the concept that user role assignment policies exist. Each policy divides into roles, each of which controls some aspect of functionality, such as:

- **MyContactInformation:** Enables users to update their contact phone numbers and address.
- **MyProfileInformation:** Enables users to update their name information, such as their display name.
- **MyDistributionGroups:** Controls what actions users can take with distribution lists.
- **MyDistributionGroupMembership:** This allows users to control their membership in the distribution lists they join.

Not every tenant uses all these roles. Some, like text messaging, are remnants of technology that was once more important than it is now. Descriptions of the [full set of user roles](#) are available online.

## Autodiscover

Microsoft originally introduced Autodiscover to stop users from having to create a “profile” with the details of the server and database hosting their mailbox before they could connect MAPI clients like Outlook to Exchange. Since its introduction, Autodiscover has proven to be extremely valuable and has expanded to give information to clients about other resources (such as the location of public folders, the OAB, and Exchange Web Services). Equipped with this information, clients know about these resources and how to access them when needed. Clients that use Autodiscover (desktop Outlook, Outlook for Mac, and many Exchange ActiveSync clients) can connect to Exchange Online to retrieve all the information necessary to configure settings in the user profile. The technical details about where their mailbox is located are invisible to the user. See this post for a description of [how to use Autodiscover in scripts](#).

**Autodiscover and Teams:** Exchange 2016 CU3 and later on-premises servers run Autodiscover V2. Teams clients need this version to connect to Exchange on-premises servers to retrieve calendar events and other information from mailboxes.

Simplifying the first connection to a mailbox is an important contribution to making it easier to onboard new users. All a user needs to know is their Microsoft Online Services identifier (usually the same as their SMTP email address) and password. With this information, Autodiscover can connect to Exchange Online, retrieve information about the services provided by Exchange Online, and return the information to the client as an XML-formatted manifest. You can see the contents of the manifest by using Outlook’s “Test Email Auto-configuration” function. Do not include the Guessmart and Secure Guessmart Authentication methods as they only work with IMAP4 and POP3 servers. The Autodiscover process is a little more complicated in hybrid deployments because the first connection goes to the on-premises organization and then to Exchange Online.

## Recovering Deleted Mailboxes

When you remove an account, a clock starts ticking and the Azure AD object for the user account stays in a soft-deleted state for a 30-day retention period. During this time, you can recover and restore the data belonging to the user, including their mailbox. Once the retention period elapses, Azure AD removes all traces of the user account, and the account becomes irretrievable. You can use the **Deleted users** view in the Microsoft 365 admin center (Figure 6-3) to see the set of deleted users that are still within the 30-day

retention period. The list includes any deleted guest accounts. You can select a user from this list and restore their account, including their mailbox, at any time until their retention period expires.

To restore a user account, select their entry to view the details of the account. If you're sure that you want to restore it, click **Restore user** and give details of what the password should be for the restored account (you can assign one, force the user to create one when they first sign in, or have a password auto-generated). Azure AD will then restore the user account to its pre-deletion state, including the mailbox (if they had one) as well as memberships of any distribution lists, Groups, and Teams to which the account belonged. It takes between 15 minutes and 30 minutes before the system restores an account fully. After the restore operation finishes, the user should be able to log on and use their account as before. See the PowerShell chapter for information about how to restore deleted accounts with PowerShell.

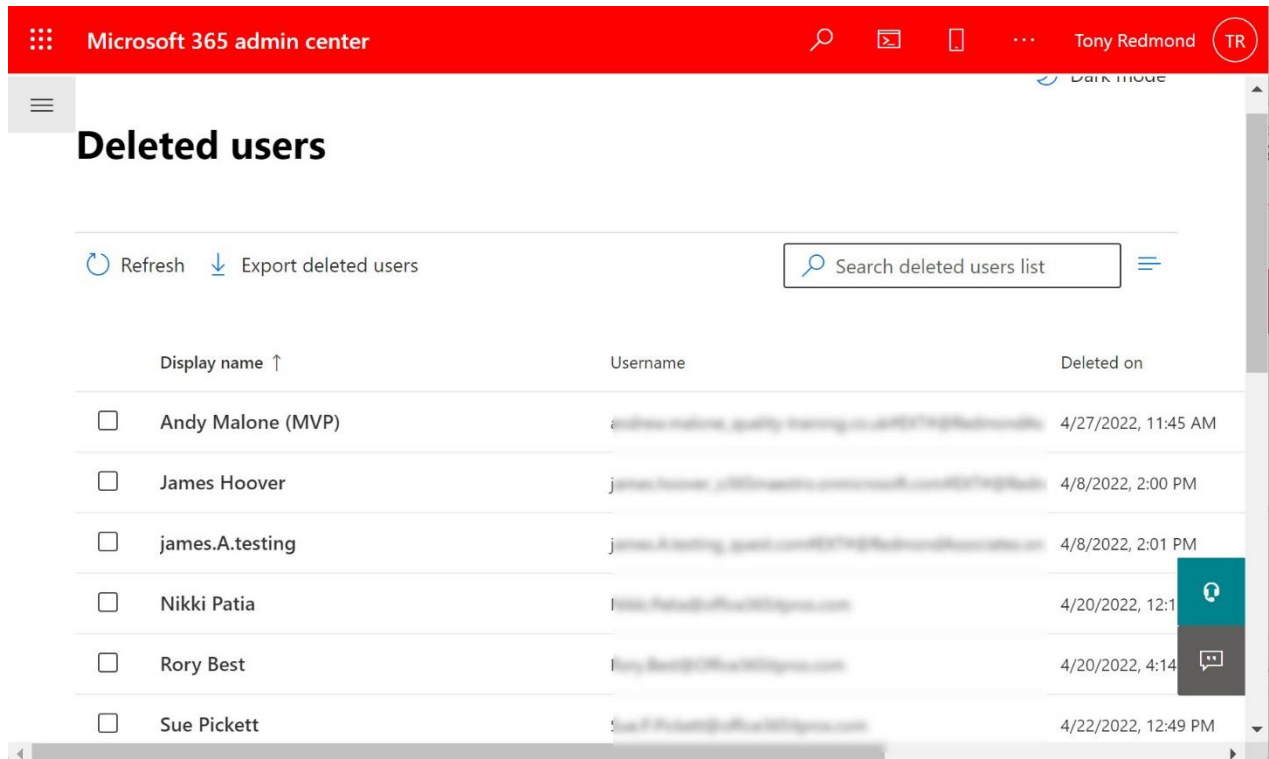


Figure 6-3: Viewing a deleted user in the Microsoft 365 admin center

Restored accounts might not receive licenses automatically and you should check the assigned licenses (and apps enabled or disabled for plans like Office 365 E3) after restoring an account to make sure that the account recovers full functionality. If you don't assign a license to a restored account, Exchange Online removes its mailbox in 30 days. The exception is if the mailbox is subject to a hold and you delete the user account which owns the mailbox. In this case, it becomes an inactive mailbox.

**Mailbox Recovery Troubleshooter:** Situations arise when administrators have discovered that they have removed the wrong mailbox and do not know what they should do to execute a recovery. To reduce issues in this area, Microsoft created a [mailbox recovery troubleshooter](#), a walk-through guide as to what an administrator needs to do to recover a mailbox. It is not a wizard and processing a successful recovery needs some skill in PowerShell and knowledge of the various PowerShell modules that are involved, but at least it's a step forward towards automated recovery.

## Removing an Unlicensed Mailbox

The usual methods to remove a mailbox from a Microsoft 365 account are to delete the account or take the Exchange Online license away from the account. If you remove the Exchange license from an account, Exchange notices the absence of the license and deprovisions the mailbox, putting the mailbox into a

disconnected state. The mailbox remains in this state for 30 days. During this time, you can reconnect the mailbox by restoring the account or adding the license back to the account. After the 30 days elapses, the Managed Folder Assistant permanently deletes the mailbox from its database.

When Exchange Online acknowledges the license removal, it sets the mailbox's *SkuAssigned* property to *\$False*. At this point, you can accelerate the removal process by running the *Disable-Mailbox* cmdlet (the cmdlet won't work if Exchange thinks the mailbox is licensed). Data in unlicensed mailboxes is not discoverable.

It's usually best to leave disconnected mailboxes for the Managed Folder Assistant to deal with and only use *Disable-Mailbox* when an urgent need exists to purge a mailbox immediately. Given the obvious compliance issues that might arise when a mailbox is purged especially when retention holds are in place, it's wise to document the reasons why this action is necessary and seek management approval before proceeding.

## Inactive Mailboxes

Sometimes you want to remove a user account because the account holder does not work for the company any longer. However, you need to keep the information in their mailbox for an extended period for legal or regulatory purposes or because it is a source needed for an eDiscovery case. In the on-premises world, you could disable the user's Active Directory account and leave their mailbox online. When the time comes to remove the account, you could export the mailbox to a PST if necessary to keep a copy of the mailbox.

Keeping unused mailboxes is not as straightforward in the cloud. The big difference is that you pay a monthly license fee for each account, so although you can disable (block) an account, you continue paying as long as the account and its mailbox remain online.

To address the problem of how to retain information in mailboxes potentially needed for eDiscovery after removing their owners' accounts, Microsoft created the concept of inactive mailboxes. These are online mailboxes belonging to accounts that no longer exist in Azure AD. The deciding factor to remove a mailbox when removing its account or making it inactive is whether any holds existed on its owner's account. The [holds that force Exchange Online to retain an inactive mailbox](#) include both org-wide holds (that apply to all mailboxes), litigation holds, and the holds placed by retention policies. Common ways to place a hold on a mailbox include:

- Put the mailbox on litigation hold. The hold covers the entire contents of the mailbox.
- Include the mailbox within the scope of a Microsoft 365 retention policy that keeps items covered by the policy for a defined retention period. The hold covers all content matching the hold criteria (which might be the complete mailbox and its archive).
- The presence of retention labels with a retain then delete or retain action on one or more items in a mailbox. Users can assign labels to mailbox items manually or the system can do so using an auto-label policy. In this instance, the hold is on the labeled items and once their retention period expires, the mailbox becomes available for deletion.
- Include the mailbox in a Core or Premium eDiscovery case that includes an in-place hold that includes the mailbox within its scope. eDiscovery cases allow you to use a query to select and hold a subset of the content in a mailbox. Do not input any search conditions if you want to hold everything in the mailbox (the equivalent of a litigation hold).
- Older in-place holds belonging to Exchange eDiscovery searches remain valid. Given that Microsoft deprecated these searches in 2020, the number of such holds remaining is now minimal.

Adaptive scopes support inactive mailboxes. Some organizations use this capability to ensure that they retain inactive mailboxes for a defined period. To do this, they create an adaptive scope to find inactive mailboxes and use this with a retention policy. The retention period set in this policy makes sure that the expiration of

other retention policies or labels do not cause the removal of inactive mailboxes. In effect, the policy acts as a retention backstop. For more information about retention policies, see the chapters covering eDiscovery and Compliance.

When an administrator deletes an Azure AD account, a 30-day countdown starts. During this time, administrators can recover the deleted account through the Deleted Users section in the Microsoft 365 admin center. The mailbox is inactive at this point because of the deletion of its corresponding user account. After 30 days, Azure AD permanently removes the user account and evaluates if the mailbox continues to be inactive because one or more holds exist on the mailbox. Remember that the personal data belonging to an account usually spans much more than a mailbox. Extra steps are necessary to secure all the information belonging to an account. We'll get to this point soon.

Because an inactive mailbox is no longer associated with an Azure AD account, it does not need a license. You can't sign into an inactive mailbox. If you want to access its contents, you must restore or recover the mailbox. You can also export the data from an inactive mailbox by running a content search and exporting the search results to a PST.

Exchange Online retains an inactive mailbox until the termination of the last hold on the mailbox. When this happens, Exchange Online puts the mailbox into a soft-deleted state and no longer considers it to be inactive. Exchange Online keeps mailboxes transitioning from the inactive state for 183 days before proceeding to permanent removal. The retention period (the soft-deleted date) commences when the mailbox moves out of the inactive state and not when the holds lapsed.

**Inactive Mailboxes and SMTP Addresses:** It's possible to have an inactive mailbox with the same SMTP address as an active mailbox. This works because the inactive mailbox is invisible for routing purposes, so the transport service only ever delivers emails sent to the address to the active mailbox. However, it's a horrible idea to allow this situation to occur because it's bound to cause confusion. One way around the problem is to remove all SMTP addresses from a mailbox and to assign it a new SMTP address before deleting its account. Ideally, the new SMTP address should be something that will never be used in production, like *Inactive.Andy.Ruth@Office365itpros.com*. This must be done before the mailbox becomes inactive because once it is inactive, you can't update its properties.

## Hybrid Inactive Mailboxes

Some problems exist in hybrid configurations where:

- The user account is on-premises.
- Their primary mailbox is on-premises or in the cloud.
- The mailbox is archive-enabled, and the archive is in the cloud.

In this scenario, you cannot make the mailbox inactive by deleting the user account. Microsoft has acknowledged the problem, but no solution is currently available. Two workarounds exist:

**Convert the mailbox into a shared mailbox.** This approach allows the retention of the archive. The shared mailbox must have at least an Exchange Online Plan 1 license. To make the shared mailbox more like an inactive mailbox, hide it from Exchange address lists and replace the email addresses for the mailbox with something that people are unlikely to guess.

**Transfer all the content from the archive back into the primary mailbox.** This can be done using Exchange Web Services ([here's a script](#) to show how), but only if the resulting size of the primary mailbox remains under 100 GB. When the transfer is complete, disable the archive and wait for Exchange to process the command (allow an hour or so), and then delete the user account.

Taking everything into consideration, the best solution is to convert the mailbox into a shared mailbox.

## Finding Inactive Mailboxes

The Data lifecycle section of the Microsoft Purview Compliance portal includes a tab to show details of inactive mailboxes. Although you can see what mailboxes are inactive and some details of the inactive mailboxes, you can't recover, restore, or do much else with the information.

The *Get-ExoMailbox* cmdlet can list any inactive mailboxes that exist in a tenant. This command retrieves the list of inactive mailboxes. The *WhenSoftDeleted* property tells us how long Exchange has kept the mailbox. Remember, this date commences upon the original deletion. The fact that some inactive mailboxes are present well after 30 days since their deletion tells us that holds remain on these mailboxes.

```
[PS] C:\> Get-Mailbox -InactiveMailboxOnly | Sort WhenSoftDeleted -Descending | Format-Table
DisplayName, WhenSoftDeleted
```

DisplayName	WhenSoftDeleted
-----	-----
Jack Smith	17/06/2021 15:37:53
Sanjay Patel	26/11/2020 14:10:56
Nancy Anderson	03/10/2021 13:14:05
Boris Johnstone	29/05/2022 09:23:00

If you check for soft-deleted mailboxes, you'll find that inactive mailboxes are in the returned set. This is because inactive mailboxes are technically in a soft-deleted state. The difference between the two sets is that Exchange Online permanently removes soft-deleted mailboxes which aren't inactive 30 days after they enter the soft-deleted state while inactive mailboxes remain untouched until the removal of the last hold. This command generates a list of soft-deleted mailboxes:

```
[PS] C:\> Get-ExoMailbox -SoftDeletedMailbox -Properties WhenSoftDeleted | Format-Table DisplayName,
WhenSoftDeleted
```

Although inactive mailboxes aren't in use, they need some management. Most of the time, these tasks amount to responding to an occasional need to recover or restore an inactive mailbox. But it's also a good idea to keep an eye on the set of inactive mailboxes and know why they are in that state. Inactive mailboxes are not visible within EAC (an [inactive mailboxes page](#) in the Data lifecycle management section of the Microsoft Purview Compliance portal lists inactive mailboxes), and it is easy to miss the fact that holds are in place for some deleted (and now inactive) mailboxes. For this reason, it is sensible to update the account properties of inactive mailboxes so that they become more obvious to administrators. For instance, you could update the display name for inactive mailboxes to mark their status. Hopefully, the visual reminder is enough to stop administrators from making embarrassing mistakes.

## Making Inactive Mailboxes with Holds

Inactive mailboxes remain in a soft-deleted state because they come within the scope of one or more holds. The type of hold isn't important: what is important is that the hold must exist before removing the Azure AD account. It is the combination of deleted account and an on-hold soft-deleted mailbox that makes an inactive mailbox. Data in inactive mailboxes remain indexed and discoverable for eDiscovery, but retention policies process inactive mailboxes to remove the information not required by the holds.

Although an inactive mailbox does not need a license, the account owning a mailbox must have a suitable license to allow the assignment of a hold on the mailbox before the user account is removed. Before administrators can assign retention policies to a mailbox, the owning account must have an Exchange Plan 2 or Exchange Online Archiving license. The Office 365 E3 and E5 plans include Exchange Plan 2.

You can put an inactive mailbox on litigation hold. This is useful when a hold associated with an eDiscovery search is about to lapse and you do not want this action to force Exchange Online to remove the inactive mailbox. To put an inactive mailbox on litigation hold, use the command:

```
[PS] C:\> Set-Mailbox -Identity "Jill Smith" -InactiveMailbox -LitigationHoldEnabled $True
```

## Removing Org-Wide Holds from Inactive Mailboxes

When a mailbox is inactive, a mixture of org-wide and specific holds might apply to it. The presence of any hold is enough to retain an inactive mailbox. As the tenant creates new org-wide holds, those holds apply to both active and inactive mailboxes (if you use an adaptive scope, it can be set to apply to mailboxes in a specific state, such as inactive). The net effect is that the number of org-wide holds that apply to inactive mailboxes can grow over time, which might then mean that some inactive mailboxes exist for longer than they should because the tenant applied extra org-wide holds after the mailboxes became inactive. This goes against the principle that organizations should be able to control the retention of information.

To solve the problem, the *Set-Mailbox* cmdlet supports the *ExcludeFromOrgHolds* and *ExcludeFromAllOrgHolds* parameters.

- **ExcludeFromOrgHolds:** Takes one or more GUIDs pointing to org-wide holds as input and excludes these holds from the evaluation of whether to keep an active mailbox.
- **ExcludeFromAllOrgHolds:** Excludes all org-wide holds from the evaluation of whether to keep an inactive mailbox.

When you exclude org-wide holds, Exchange will only keep an inactive mailbox if a specific hold exists on the mailbox or an org-wide hold is present. If no other holds exist, Exchange removes the mailbox.

For example, let's assume that you have many inactive mailboxes and want to clean up the set. To remove org-wide holds from the evaluation used by the Managed Folder Assistant to decide if an inactive mailbox is still subject to a hold, we retrieve the identifiers for the holds by running the *Get-OrganizationConfig* cmdlet:

```
[PS] C:\> Get-OrganizationConfig | Select -ExpandProperty InPlaceHolds
mbx9696959111f74ecda8a40aef97edd2c2:1
grp703105e3b8804a1093bb5cb777638ea8:1
mbx19200b9af08442529be070dae2fd54d3:1
grp6a1654abdba4712a43c354e28a4d56c:1
mbx703105e3b8804a1093bb5cb777638ea8:1
mbxc1e2d6f1785d4bf8a7746a26e58e5f66:1
```

Holds applying to user mailboxes are prefixed by "mbx." We can pass the identifiers in the *ExcludeFromOrgHolds* parameter as a comma-separated list. The example below passes the identifiers for two org-wide holds. The values used for the hold identifiers are in the same format as those reported by *Get-OrganizationConfig*. Notice that Exchange Online uses distinguished names to identify inactive mailboxes to ensure that it can find the mailboxes.

```
[PS] C:\> $InactiveMbx = (Get-Mailbox -InactiveMailboxOnly -ResultSize Unlimited)
ForEach ($Mbx in $InactiveMbx) {
  Write-Host "Removing Specific Org-Wide holds from" $Mbx.DisplayName
  Set-Mailbox -Identity $Mbx.DistinguishedName -ExcludeFromOrgHolds
  "mbx9696959111f74ecda8a40aef97edd2c2:1", "mbx19200b9af08442529be070dae2fd54d3:1" -Confirm:$False
  -Force}
Write-Host "Checking inactive mailboxes again"
$InactiveNow = (Get-Mailbox -InactiveMailboxOnly -ResultSize Unlimited)
Write-Host "$InactiveMbx.Count - $InactiveNow.Count" "mailboxes processed to exclude selected org-
wide holds "
```

In this example, we exclude all org-wide holds, so there's no need to pass any hold identifiers. Because you're now removing all org-wide holds from the evaluation of inactive mailboxes, it's even more important to make sure that you are happy for Exchange to remove all inactive mailboxes not covered by a specific hold.

```
[PS] C:\> $InactiveMbx = (Get-Mailbox -InactiveMailboxOnly -ResultSize Unlimited)
ForEach ($Mbx in $InactiveMbx) {
  Write-Host "Removing Org-Wide holds from" $Mbx.DisplayName
  Set-Mailbox -Identity $Mbx.DistinguishedName -ExcludeFromAllOrgHolds -Confirm:$False -Force}
```

```
Write-Host "Checking inactive mailboxes again"
$InactiveNow = (Get-Mailbox -InactiveMailboxOnly -ResultSize Unlimited)
Write-Host ($InactiveMbx.Count - $InactiveNow.Count) "mailboxes excluded from all org-wide holds"
```

If you examine the properties of an inactive mailbox after running *Set-Mailbox* to exclude some or all org-wide holds, you'll see that the GUIDs for the excluded holds are present in the mailbox's *InPlaceHolds* property and that Exchange prefixes each hold with a minus sign. This indicates that the Managed Folder Assistant should exclude the hold when evaluating the mailbox. For instance:

```
[PS] C:\> Get-ExoMailbox -Identity David.Pelton | Select -ExpandProperty InPlaceHolds

-mbxc1e2d6f1785d4bf8a7746a26e58e5f66
-mbx703105e3b8804a1093bb5cb777638ea8
-mbx19200b9af08442529be070dae2fd54d3
-mbx9696959111f74ecda8a40aef97edd2c2
-mbx6a1654abda4712a43c354e28a4d56c
UniH26c5d797-0fd3-496d-92ac-4f405700c917
```

The hold that is keeping this inactive mailbox is the last one on the list (it doesn't have a minus sign). This identifier points to a specific in-place hold.

Don't exclude org-wide holds from inactive mailboxes without understanding exactly what holds are keeping mailboxes in the inactive state as you cannot retrieve an inactive mailbox after Exchange removes it.

## Removing the Azure AD User Object for an Inactive Mailbox

Normally you must wait 30 days for a deleted mailbox to become fully inactive because its matching user account remains in the deleted objects section of Azure AD to allow administrators to recover the account and reconnect the mailbox during that period. You can discover which inactive mailboxes are in this state by running this command:

```
[PS] C:\> Get-Mailbox -InactiveMailboxOnly | ?
{![string]::IsNullOrEmpty($_.ExternalDirectoryObjectId)} | Format-Table DisplayName,
ExternalDirectoryObjectId

DisplayName ExternalDirectoryObjectId
-----
Imran Khan b8eef43d-6854-4d77-9e03-745cf2e11e11
```

If necessary, you can remove a user account by first deleting the account and then removing the deleted object. For example:

```
[PS] C:\> $UserId = (Get-MgUser -ObjectId David.Jacobs@Office365itpros.com).Id
Remove-MgUser -UserId $UserId
$Uri = "https://graph.microsoft.com/beta/directory/deleteditems/" + $UserId
Invoke-MgGraphRequest -Method Delete -Uri $Uri
```

Alternatively, you can remove a soft-deleted account through the Users section of the Azure AD portal. Select the Deleted Users view, then select the account you want to remove, and then click the **Delete permanently** button.

Usually, it's best to let nature take its course and let user accounts move through the 30-day account retention process until Azure AD purges the accounts and their mailboxes become inactive, but you never know when you might want to accelerate the process.

## Restore or Recover Inactive Mailboxes

In addition to being able to export the contents of inactive mailboxes through eDiscovery searches, you can also [retrieve information by restoring or recovering an inactive mailbox](#). When you restore an inactive mailbox, Exchange Online merges the contents of the mailbox into another mailbox. This might be done when

a user needs to work with the information contained in a mailbox that belonged to an ex-employee. Restoring data from an inactive mailbox leaves the inactive mailbox intact and still available for eDiscovery.

## Restore an Inactive Mailbox

The first step is to run the *Get-Mailbox* cmdlet to return a list of inactive mailboxes and identify which mailbox to restore. Several of the inactive mailboxes might share the same or a similar display name or other attributes, so we need a unique value for the mailbox to pass to Exchange Online to restore the mailbox. The Distinguished Name is best for this purpose, so we'll use that.

```
[PS] C:\> $Inactive = (Get-Mailbox -InactiveMailboxOnly -Identity "Jill Smith").DistinguishedName
```

We can now set up the restore with the *New-MailboxRestoreRequest* cmdlet to ask Exchange Online to fetch the data from the inactive mailbox and move it into a target mailbox. Remember, a restore operation leaves the inactive mailbox intact, so this is in effect a copy operation. The *AllowLegacyDNMismatch* switch allows the *New-MailboxRestoreRequest* cmdlet to process the restore request even though the distinguished names (DN) of the inactive and target mailboxes do not match. Normally, to safeguard against the misdirection of items to a mailbox that they don't belong to, *New-MailboxRestoreRequest* would refuse to copy items into a target mailbox if a DN mismatch existed. We proceed even though we know a mismatch exists, so we override the natural caution of the cmdlet by telling it that it's OK to go ahead and copy the items:

```
[PS] C:\> New-MailboxRestoreRequest -SourceMailbox $Inactive -TargetMailbox
Abrus@Office365ITPros.com -TargetRootFolder "Jill Smith Old Mailbox" -AllowLegacyDNMismatch
```

Name	TargetMailbox	Status
MailboxRestore	Abrus	Queued

The restore operation continues in the background. You can check it by running the *Get-MailboxRestoreRequest* cmdlet. When the status is "Completed," all the data from the inactive mailbox should be in the target mailbox. The target root folder specified in the restore request ("Jill Smith Old Mailbox") is in the target mailbox with all the folders and items that belonged to the inactive mailbox underneath that root.

If the inactive mailbox owns an archive, you can restore items out of the archive and direct them to either the archive of a target mailbox or to the target mailbox itself. *New-MailboxRestoreRequest* supports the *SourcesArchive* switch to control whether to copy items from the primary (the default) or archive mailbox and the *TargetsArchive* switch to control whether to restore items to the primary mailbox of the target or into its archive. Restoring items to an archive mailbox has the value of creating a clear separation between the restored items and the owner's items in the target mailbox.

## Recover an Inactive Mailbox

Recovering an inactive mailbox means the transformation of a mailbox into one that is usable by a new user. To make this happen, we run the *New-Mailbox* cmdlet to create a new mailbox and instruct Exchange Online to populate that mailbox with the inactive mailbox.

```
[PS] C:\> $Inactive = (Get-Mailbox -InactiveMailboxOnly -Identity "Jill Smith").DistinguishedName
```

```
[PS] C:\> New-Mailbox -InactiveMailbox $Inactive -Name "Joe Healy" -FirstName Joe
-LastName Healy -DisplayName "Joe Healy" -MicrosoftOnlineServicesID "Joe.Healy@Office365ITPros.com"
-Password (ConvertTo-SecureString -String "Testing123!" -AsPlainText -Force)
-ResetPasswordOnNextLogon $True
```

In this case, we take the inactive mailbox of Jill Smith and use it to create a new mailbox under the control of Joe Healy, a new account. After *New-Mailbox* completes, the old inactive mailbox is gone, and its content is in the Joe Healy mailbox. To complete the process and make the mailbox fully operational, you must assign a license to the new mailbox.



You cannot use the recover method for an inactive mailbox while its user account still exists in Azure AD (for 30 days following the removal of the account). During this period, you can use the standard Recover Deleted Users option to restore the account and reactivate the mailbox, but you can't recover the data to a new mailbox. To test whether the user object still exists for an inactive mailbox, run *Get-Mailbox* as shown below. In this case, Exchange returns a GUID, so you know that the object still exists.

```
[PS] C:\> Get-Mailbox -InactiveMailboxOnly -Identity 'Jill Smith' | Select ExternalDirectoryObjectID
ExternalDirectoryObjectId
-----
636578d1-89fd-42e6-8b1d-237c96635a95
```

If you recover an inactive mailbox, any holds that existed on the mailbox when it was removed are not active. Instead, Exchange enables single item recovery for the mailbox and puts a retention hold in place for 30 days to be sure that nothing is removed through the application of retention policies in that time.

## Automatic Mailbox Maintenance

Like on-premises Exchange, the Managed Folder Assistant (MFA) runs on a workcycle basis to apply retention policies to mailboxes and clean out deleted items that have exceeded their retention period. The SLA for the MFA calls for it to process mailboxes at least once weekly, but MFA often processes mailboxes more regularly. You can't affect when mailbox management happens because this happens automatically, but you can affect how MFA processes mailboxes by changing the mailbox properties that govern retention policy. To view details of a mailbox's assigned retention policy, open the *Manage mailbox policies* section of mailbox properties in EAC. The *Manage mailbox archive* section of mailbox properties shows if the mailbox has an archive and if so, how much quota the archive uses. Alternatively, you can run the *Get-ExoMailbox* cmdlet to retrieve the same information:

```
[PS] C:\> Get-ExoMailbox -Identity TRedmond -PropertySets Archive -Properties DisplayName,
RetentionPolicy | Format-Table DisplayName, RetentionPolicy, ArchiveName

DisplayName      RetentionPolicy      ArchiveName
-----
Tony Redmond     Management retention policy  {Grubby old stuff}
```

Exchange Online assigns the default Messaging Retention Management (MRM) retention policy to all mailboxes upon creation, including migrated on-premises mailboxes. In contrast, an on-premises administrator must make an explicit choice to assign a retention policy to a mailbox. The rationale for having a default retention policy in place for all mailboxes is that it allows MFA to exert some level of control over mailbox contents.

Figure 6-4 shows some of the retention tags included in the Default MRM Policy, with the "Default 2 year move to archive" tag selected. This is a default tag, meaning that the Managed Folder Assistant applies its action to all items not governed by a more explicit retention tag. The action is to move items to the archive once they are 730 days (2 years old). In effect, MFA checks the retention period on items each time it processes a mailbox and will move any older than 2 years to the archive mailbox. Logically, this action can only happen if a mailbox has been archive-enabled. If not, MFA ignores the directive contained in the default archive tag. Unlike other retention policies that you might be aware of, the default MRM policy used by Exchange Online does not include a default delete tag, so if items are not archived, they continue to accumulate in the primary mailbox unless the user deletes them.

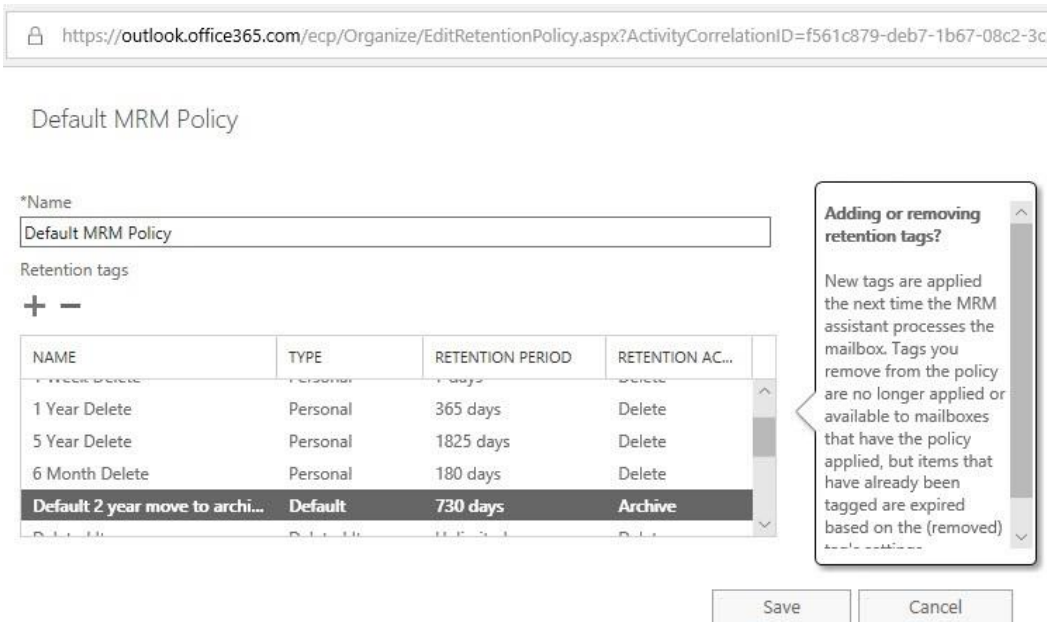


Figure 6-4: Retention tags in the Default MRM Policy for Exchange Online

## Mailbox Data Retention

Exchange Online uses what Microsoft calls Native Data Protection to protect data. Here are some aspects of mailbox management to consider:

- Each mailbox database has at least four copies spread across two data centers within the host Office 365 region. One of these is a lagged copy.
- As explained in the Compliance chapter, the default retention policy applied to Exchange Online mailboxes does not force the removal of items from the Deleted Items folder. You can change this behavior by changing the retention policy assigned to user mailboxes. If not, items stay in the Deleted Items folder until the user empties the folder or the items move to an archive mailbox.
- Exchange Online enables Single Item Retention (SIR) for every mailbox so that items moved into Recoverable Items stay in the database for the full retention period set on the mailbox. The default for the deleted items retention period used to be 14 days but (since 2017) it is 30 days, which is also the longest period you can set for this property. Exchange retains deleted Calendar items for up to 120 days.
- If you want to update older mailboxes, you can do so with this command:

```
[PS] C:\> Get-ExoMailbox -RecipientTypeDetails UserMailbox -Properties RetainDeletedItemsFor | ?
{$_ .RetainDeletedItemsFor -Lt 30} | Set-Mailbox -RetainDeletedItemsFor 30
```

If the deleted items retention period for a mailbox does not update when you run the command, check its *UseDatabaseRetentionDefaults* property. This must be *\$False* before you can update *RetainDeletedItemsFor*. Some older mailboxes have this property set to *\$True* (it's a legacy of the on-premises heritage of Exchange Online).

- Compared to the frequent access for items in primary mailboxes, archive mailboxes usually experience infrequent access. The default retention period moves items from the primary mailbox to the archive after they are two years old.

## Recovering Deleted Items

It is possible users will discover that they need to recover some items after they empty the Deleted Items folder, or a retention policy removes items from the folder. Exchange Online's Single Item Recovery feature ensures that deleted items remain in Recoverable Items for the full deleted items retention period configured for the mailbox.

Under Recoverable Items, two important sub-folders hold items:

- The **Deletions** folder stores deleted items removed from the Deleted Items folder (by emptying the complete folder or removing individual items) or items removed from other folders with the SHIFT+Delete command. While items remain in the Deletions folder, users can recover them by opening the Deleted Items folder and then using the **Recover items deleted from this folder** feature supported by Outlook or OWA (Figure 6-5).
- The **Purges** folder (which is invisible to clients like Outlook) stores items removed from Deletions. For example, if someone uses the Recover Items feature to find items and then purges them, Exchange removes the items from Deletions into Purges. The items remain there until the deleted items retention period passes or, if the mailbox is on hold, any relevant holds elapse.

As explained below, it's also possible for administrators to recover items on behalf of users by running PowerShell cmdlets or using an option in the EAC. Once the retention period elapses, the Managed Folder Assistant removes the items from the database the next time that it processes the mailbox. At this point, the items become irrecoverable.

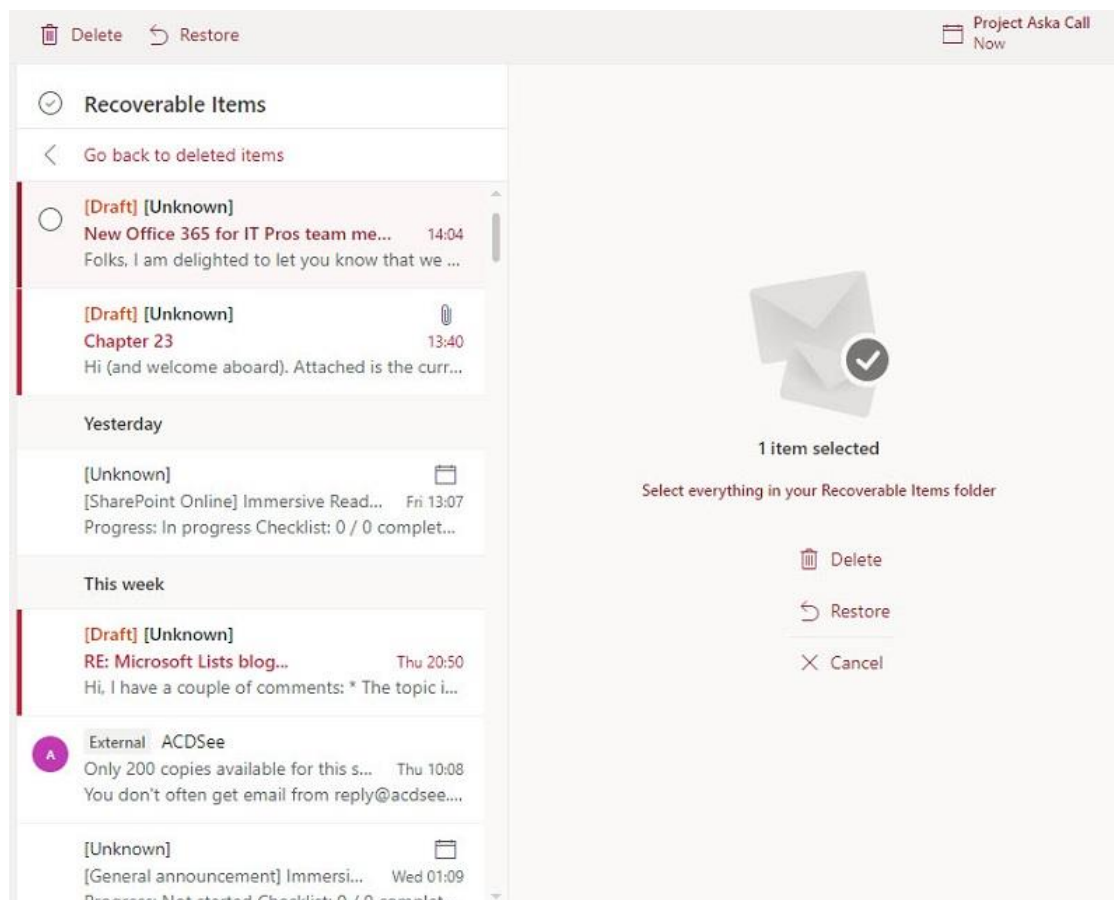


Figure 6-5: OWA lists recoverable items

Microsoft does not take backups of Exchange Online mailbox databases. Once Exchange removes an item from a database, it is permanently gone and no number of appeals to Microsoft Support will result in its

recovery. For this reason, some administrators place some mailboxes on a permanent in-place hold or a litigation hold. The logic is as follows:

- An in-place hold keeps deleted items needed by the hold in the hidden Purges subfolder of the Recoverable Items folder until the hold lapses. The in-place hold feature is not part of every plan and a mailbox needs to have at least an E3 or Exchange Online Plan 2 license before an administrator can put the mailbox on hold.
- Problems are more severe if items are irrecoverable for mailboxes belonging to executives or other critical personnel.
- Users will not remember to protect every piece of sensitive information – the system must do this for them.
- Exchange Online makes sufficient quota available to the Recoverable Item folder (and its sub-folders) to store deleted mailbox items to satisfy the requirements of Single Item Recovery or in-place holds.
- Deleted items under hold are invisible to the user and clients. An administrator can retrieve these items by executing an eDiscovery search. Once found, exporting the search results to a PST makes the items available for return to the user (or, in the case of an investigation, to investigators).

It is difficult to estimate how much of the Recoverable Items quota a hold will consume. Even a busy mailbox will remove less than 10 GB of unwanted messages annually and a hyperactive mailbox will consume perhaps 20 GB. The default quota is therefore capable of holding five years of deleted items – and probably more for less active mailboxes. However, when you have the combination of a busy mailbox and litigation hold (or multiple in-place holds), you need to keep an eye on the growth of Recoverable Items to ensure that a mailbox does not exceed its quota ([premium monitoring](#) includes this capability). Once the quota consumed approaches 85% of its limit, it's time to talk to Microsoft Support to see if they can assign some additional mailbox quota. If not, you might need to take another approach, such as:

- Use the Mailbox Folder Assistant (MFA) to remove duplicate items from Recoverable Items. Often a surprising amount of information accumulates that can be safely removed without compromising the effectiveness of a hold. To have the MFA scan and remove duplicate items, run the *Start-ManagedFolderAssistant* cmdlet and specify the *HoldCleanup* parameter. For example:

```
[PS] C:\> Start-ManagedFolderAssistant -Identity Kim.Akers -HoldCleanup
```

- Use retention tags and policies to [move data from Recoverable Items to the archive mailbox](#) (and exploit auto-expanding archives).
- After securing approval to do so, [delete some items from the Recoverable Items folder](#).

## Administrator Recovery of Deleted Items

Sometimes, users need help to recover items. In the past, an administrator had to sign in to the user mailbox to perform recovery on behalf of the user. The problem with this approach is that it compromises the privacy of the mailbox. For this reason, Exchange Online includes two cmdlets to help administrators recover user data without needing to access the mailbox with a client:

- *Get-RecoverableItems* searches the Deleted Items folder and the Purges and Deletions sub-folders within Recoverable Items of the target mailbox to find items. The administrator doesn't need to sign into the target mailbox, which can be a user or a shared mailbox.
- *Restore-RecoverableItems* finds and copies items from Deleted Items and the Purges and Deletions sub-folders of Recoverable Items to their original folders.

The basic idea is that you use *Get-RecoverableItems* to construct a search to find the desired items and then use the search results as input to *Restore-RecoverableItems* after you've found the right items. Before trying to run these cmdlets, make sure that the account you use to sign into PowerShell holds the Exchange "Mailbox

Import Export" RBAC role. To find out who has the role already, you can run the following command. In this example, the members of the Organization Management role group have the role as does the Administrator account.

```
[PS] C:\> Get-ManagementRoleAssignment -Role "Mailbox Import Export" | Format-Table RoleAssigneeName
RoleAssigneeName
-----
Organization Management
Administrator
```

An example of using *Get-RecoverableItems* to search a mailbox for items is shown below. The search looks for any item of type *Ipm.Note* (messages) moved into Recoverable Items during a certain period. This happens when the user or the Managed Folder Assistant moved the item into its current folder. For instance, if a mailbox has a retention policy that moves items from Deleted Items into Recoverable Items after 120 days, the Managed Folder Assistant might have processed the items found by the search above at least four months ago. A user can bypass Deleted Items and send an item direct to Recoverable Items by using the *SHIFT+Delete* key combination. In this case, the *LastModifiedTime* property (used for date filters) is the date when the user executed *SHIFT+Delete*.

```
[PS] C:\> Get-RecoverableItems -Identity Kim.Akers -SourceFolder RecoverableItems -FilterStartTime "2/16/2018 10:00:00" -FilterEndTime "2/20/2018 17:00:00" -FilterItemType Ipm.Note
```

## Location of Deleted Items

Administrators can recover deleted Items from three locations, each referred to as a source folder:

- **DeletedItems:** The Deleted Items folder from the user's mailbox.
- **RecoverableItems:** The Deletions subfolder in the Recoverable Items folder of the user's mailbox.
- **PurgedItems:** The Purges subfolder in the Recoverable Items folder. Deleted items kept due to a retention policy stay here until the retention period set in the policy expires.

These names are language-independent. You cannot search folders in the archive mailbox. To search for deleted items across all locations, don't pass a *SourceFolder* parameter to *Get-RecoverableItems*. For example, this search finds all deleted items for Kim Akers' mailbox:

```
[PS] C:\> $Items = Get-RecoverableItems -Identity Kim.Akers
```

Capturing the items returned by a search in an array makes it easier to process them afterward:

## Finding Deleted Items

Although you could go through every deleted item from all locations to find something to restore, it's better when the user for whom you're restoring items gives some hints about the items they want to recover. Users might not be sure when an item was deleted, but they should be able to tell you something about the message subject. Here's an example of using *Get-RecoverableItems* to find a message by subject.

```
[PS] C:\> $Items = (Get-RecoverableItems -Identity Kim.Akers -SourceFolder RecoverableItems -SubjectContains "disgruntled")
```

Be aware that a search based on *SubjectContains* finds any item which contains the provided string in its subject, so it's good to be as precise as possible. For instance, this search would find items with subjects like "Disgruntled at work" or "Handling disgruntled employees."

The example of the data returned for an item in Recoverable Items is shown below.

```
LastParentPath      : Junk Email
LastParentFolderID  : EBA28A7861EE1F4485DA85FE1279C88C000009CA9700
OriginalFolderExists : True
Identity            : Kim.Akers
```

```
MailboxIdentity      : b662313f-14fc-43a2-9a7a-d2e27f4f3478\ea58dd70-4581-4190-aeef-52075e470846
ItemClass           : IPM.Note
Subject            : How to Catch Attacks by Disgruntled Employees
EntryID            : 00000000E4D17F986EC65C4EB677E1EB8F1015F20700EBA28A7861EE1F4485DA85FE1279C
                   : 88C0000000001140000EBA28A7861EE1F4485DA85FE1279C88C0005228415330000
SourceFolder       : Recoverable Items\Deletions
```

Obviously, the more precise the search, the more likely you are to find the right item. For instance, it's possible that the user will be able to give an approximate period when an item was deleted, so you could use that to refine the search. For example:

```
[PS] C:\> $Items = (Get-RecoverableItems -Identity Kim.Akers -SourceFolder PurgedItems -
FilterStartTime "13-Jun-2020 00:17" -FilterEndTime "13-Jun-2020 00:35")
```

A date-based search only works against a single location.

## Searching for Specific Types of Deleted Items

You can specify different types of items to look for, but you cannot combine different item types in a search. Instead, if you want to find items of multiple types, don't pass a *FilterItemType* parameter and the search will return items of all supported items. If needed, you can then apply a filter using PowerShell to isolate the required items. The valid item types include:

- IPM.Note: A standard email message item.
- IPM.Appointment: A calendar meeting or appointment.
- IPM.Task: A task.
- IPM.Contact: A contact.
- IPM.File: A file stored in the mailbox. These include files created by Office 365 as the result of some processing.

## Restoring Deleted Items

Once you are happy that your search finds the right items, you can proceed to recovery. The *Restore-RecoverableItems* cmdlet takes the same search that you use to find items and restores each item to its original location. You can input the same search as you used to find the item, but if multiple items are returned and you only want to restore a single item, pass the *EntryID* (a unique identifier) for the item. For example, let's assume that the search returned twenty items. The items are stored in the *\$Items* variable, so we can pass a reference to the exact item we want to restore with:

```
[PS] C:\> Restore-RecoverableItems -Identity Kim.Akers -EntryID $Items[0].EntryID -SourceFolder
RecoverableItems
```

The response to a successful restore will confirm the folder the item has been restored to:

```
RestoredToFolderPath : Junk Email
RestoreToFolderId    : EBA28A7861EE1F4485DA85FE1279C88C000009CA9700
WasRestoredToOriginalFolder : True
WasRestoredSuccessfully : True
Identity             : Kim.Akers
MailboxIdentity      : b662313f-14fc-43a2-9a7a-d2e27f4f3478\ea58dd70-4581-4190-aeef-52075
                   : e470846
ItemClass           : IPM.Note
Subject            : How to Catch Attacks by Disgruntled Employees
EntryID            : 00000000E4D17F986EC65C4EB677E1EB8F1015F20700EBA28A7861EE1F4485DA85
                   : FE1279C88C000000001140000EBA28A7861EE1F4485DA85FE1279C88C00052284
                   : 15330000
SourceFolder       : Recoverable Items\Deletions
```

Of course, if you wanted to, you could create a processing loop to process a batch of mailboxes and restore the messages, possibly after an administrator makes a mistake and deletes messages in error using a content search purge. Here's an example:

```
[PS] C:\> $Mbx = (Get-ExoMailbox -RecipientTypeDetails UserMailbox -Filter {CustomAttribute1 -eq "IT"} | Select Alias, DisplayName)
Write-Host "Recovering items for" $Mbx.Count "mailboxes..."
ForEach ($M in $Mbx) {
    Write-Host "Checking mailbox" $M.DisplayName
    Restore-RecoverableItems -Identity $M.Alias -SourceFolder RecoverableItems -SubjectContains "Important and Critical Message" }
```

## Recover Deleted Items in EAC

The EAC includes a GUI for the *Get-RecoverableItems* and *Restore-RecoverableItems* cmdlets to recover items for user and shared mailboxes (Figure 6-6). The GUI leverages the cmdlets and can do much of the functionality described above. The differences are:

- EAC offers fixed date ranges (7, 14, and 30 days).
- EAC can process a single mailbox at a time.

Deleted on	Entry ID	Subject line	Item type	Folder type	Original folder
6/18/2020 12:39:57 AM	0000000064D17...	Office 365 IT Pros Blog - Digest	IPM.Note.ConnectorMessage	Recoverable Items/Purges	Inbox
6/18/2020 12:39:57 AM	0000000064D17...	Weekly digest Office 365 changes	IPM.Note	Recoverable Items/Purges	Inbox
6/18/2020 12:39:57 AM	0000000064D17...	You have been added to a team in Microsoft Teams	IPM.Note	Recoverable Items/Purges	Inbox
6/16/2020 8:51:17 PM	0000000064D17...	ANITA CLARK sent you "PO-09507 (0).pdf"	IPM.Note	Recoverable Items/Purges	Inbox
6/16/2020 8:51:17 PM	0000000064D17...	ANITA CLARK sent you "PO-09507 (0).pdf"	IPM.Note	Recoverable Items/Purges	Inbox
6/16/2020 8:51:17 PM	0000000064D17...	ANITA CLARK sent you "PO-09507 (0).pdf"	IPM.Note	Recoverable Items/Purges	Inbox
6/16/2020 3:14:11 AM	0000000064D17...	Kim, New Guide: Top 5 CRM Software With the Best Uti...	IPM.Note	Recoverable Items/Purges	Inbox
6/16/2020 3:14:11 AM	0000000064D17...	Spam Notification: 1 New Messages	IPM.Note	Recoverable Items/Purges	Inbox
6/16/2020 3:14:11 AM	0000000064D17...	Office 365 IT Pros Blog - Digest	IPM.Note.ConnectorMessage	Recoverable Items/Purges	Inbox
6/16/2020 3:14:11 AM	0000000064D17...	Kim, New Guide: Top Rated FrontRunners for Customer...	IPM.Note	Recoverable Items/Purges	Inbox
6/16/2020 3:14:11 AM	0000000064D17...	Office 365 Roadmap Watcher - Digest	IPM.Note.ConnectorMessage	Recoverable Items/Purges	Inbox
6/15/2020 2:06:19 AM	0000000064D17...	Office 365 Meeting Test	IPM.Appointment	Recoverable Items/Purges	Calendar
6/15/2020 2:06:19 AM	0000000064D17...	Office 365 IT Pros Blog - Digest	IPM.Note.ConnectorMessage	Recoverable Items/Purges	Inbox
6/15/2020 2:06:19 AM	0000000064D17...	You appeared in 3 searches this week	IPM.Note	Recoverable Items/Purges	Inbox

Figure 6-6: Recover deleted items in the Exchange admin center

When an administrator recovers items on behalf of a user by running the cmdlet in PowerShell or the EAC, Exchange Online generates a *Restore-RecoverableItems* audit record for each recovered item in the audit log. An example script showing how to report these audit records can be [downloaded from GitHub](#).

## Archive Mailboxes

The motivation to introduce archive mailboxes (also known as “in-place archives” or “online archives”) for Exchange arose for many reasons. Over the years, Microsoft did not like the fact that customers deployed third-party products like Symantec Enterprise Vault to offload information from user mailboxes to a separate repository. Often there was good reason to move data, especially when Exchange mailbox quotas were small, and databases ran on expensive SAN storage. Moving data to secondary storage managed by other products allowed Exchange mailboxes to continue without the need for increased quotas, albeit with “stubs” left behind in mailboxes as pointers to the actual items. If you plan to move previously archived items from external

repositories back into Exchange Online, it is important to ensure that the stubs are “rehydrated” (become fully complete Exchange items) as part of the process.

From a Microsoft perspective, moving information out of Exchange databases implied a certain loss of control as the data could be as easily migrated to a different server as moved back to Exchange. Another issue that has become increasingly important is that PST data are invisible to compliance and data governance features. An example of how important it is for companies to protect their commercial interests and business reputation is seen in [the Sony hack of December 2014](#) when hackers stole and shared valuable contractual information held in PSTs. The messages revealed to the public had details of negotiations, strategies, opinions about business partners, and so on. It would have been so much better had this data been securely kept in online databases rather than being vulnerable to hackers.

Small mailbox quotas encouraged the proliferation of PSTs and created a plague of insecure and potentially corruptible files holding valuable user information. Because email messages are commonly shared between multiple users, PST data usually includes a high percentage of duplicate items. The existence of so much duplicated information slows down the transfer of data from PSTs to online mailboxes. In addition, users can attempt to protect PSTs with passwords that range from very easy to very difficult to crack. If you gather PSTs from users for migration, you must remove passwords from the PSTs to make their data accessible for migration processing. Some third-party PST migration engines can deduplicate content extracted from PST before ingestion and crack open any password-protected files to extract the information contained within. These are valuable features that you should consider and evaluate during any PST migration (eradication) project.

Because PST migration is often tiresome and expensive, a better solution is to encourage users to keep their data online. The solution to allow online storage to replace PSTs came about in two parts. To make it feasible to provide very large mailbox quotas to users, Microsoft engineered the mailbox database engine to support JBOD. This effort included the introduction of the Native Data Protection features used within Exchange Online today. Archive mailboxes, first introduced in Exchange 2010, are the second part. At that time, the plan envisaged that larger mailboxes meant that no one would ever want to use a PST because their mailbox had so much available space. The accompanying archive mailboxes are the repository for long-term information and avoided the need for Exchange customers to buy a third-party archiving product. Mailboxes have become larger and archive mailboxes are in common use, but users have not yet discarded PSTs. It takes a long time and much effort for Outlook users to break working habits based on PSTs, including using these files as shared repositories when much better options exist like Groups or SharePoint team sites.

**Keeping everything in the primary mailbox:** [Enterprise plans \(E3 - E5\)](#) grant 100 GB primary mailbox quotas to users, which then creates a question of whether archive mailboxes are necessary. After all, 100 GB should be enough to hold as much data as anyone would want to keep. Although large mailboxes enable people to avoid using archive mailboxes, a workable case exists to separate information into the data which needs to be on hand and items users must keep for reference purposes. In this scenario, the items you need on-hand remain in the primary mailbox, and those needed for reference (or compliance) go into the archive. Good retention policies help users achieve the split by automatically moving items into the archive after a set period (two years is the default).

Archive mailboxes are now available to any Exchange on-premises or cloud mailbox except those using online frontline (kiosk) plans. Some plans [limit the combined storage for primary and archive mailboxes](#) but plans E3 and above allow “unlimited” storage. This used to mean unlimited in that auto-expanding archive mailboxes could grow to well over 2 TB. Such massive mailboxes caused operational difficulties, such as problems moving mailboxes between databases, so Microsoft decided to limit the size of archive mailboxes to 1.5 TB.

Microsoft offers Exchange Online Archiving to allow on-premises customers to use a cloud-based archiving service. In this scenario, mailboxes hosted on Exchange on-premises servers can connect to unlimited Exchange Online archive mailboxes without having to perform a full hybrid deployment. The reverse (archive



on-premises and mailbox in the cloud) is unsupported. Exchange Online Archiving is also available as an add-on plan for Exchange Online kiosk mailboxes.

An archive mailbox is an online-only extension of the primary mailbox. The link between the two is through the *ArchiveGuid*, a property of the user mailbox that points to the location of the archive. Other archive-related properties set when a mailbox becomes archive-enabled include:

- **ArchiveDatabase:** The Exchange Online database holding the archive mailbox.
- **ArchiveName:** A user-friendly name for the archive mailbox that shows up in a client's resource list. You can change the name to whatever value you like. For example, "Joe's Online Archive."
- **ArchiveQuota:** The current storage quota assigned to the archive. The current default for Exchange Online is 100 GB. Administrators can apply to Microsoft support to have the quota increased.
- **ArchiveWarningQuota:** The threshold for the user to receive warning messages to tell them that space is running out in the archive. The current default is 90 GB.
- **ArchiveStatus:** Set to "Active" when the archive mailbox is available to a user. It is set to "None" when an archive mailbox is not present.
- **ArchiveState:** Set to "Local" when the mailbox and archive are on the same platform (the case when Exchange Online hosts both).

To assign an archive to a mailbox, select the mailbox in the EAC, go to the Others section in mailbox properties, and use the Manage mailbox archive option to enable or disable the archive. Alternatively, edit the mailbox properties, go to the mailbox features section, scroll down to archiving, and click **Enable**. The equivalent PowerShell command is:

```
[PS] C:\> Enable-Mailbox -Identity 'Joe Smith' -Archive
```

Only Outlook desktop and OWA support client access to archive mailboxes. Outlook Mobile does not include this capability because the necessary API support is unavailable. The same is true for mobile clients based on the ActiveSync protocol, such as the native mail app client on iOS and Android devices.

Shared and room mailboxes can both be enabled with an archive. Although there are many reasons why you might need an archive for a shared mailbox, the case is much less obvious for a room mailbox. You can't enable an archive for a group mailbox. If you archive-enable a shared or room mailbox, remember that you need to assign at least an Exchange Online Plan 1 license to the mailbox.

You can check for mailboxes that have archives by running the *Get-ExoMailbox* cmdlet with the *Archive* switch. The first command below returns a list of archive-enabled mailboxes. The second returns a simple count of archive-enabled mailboxes.

```
[PS] C:\> Get-ExoMailbox -Archive -RecipientTypeDetails UserMailbox -PropertySet Archive -Properties DisplayName -ResultSize Unlimited | Format-Table DisplayName, ArchiveName
```

DisplayName	ArchiveName
Deirdre Redmond	{Deirdre's Online Archive}
Tony Redmond	{Grubby old stuff}
Jeff Guillet	{Jeff's Archive}
Kim Akers	{In-Place Archive - Kim Akers}

```
[PS] C:\> (Get-ExoMailbox -Archive -ResultSize Unlimited -RecipientTypeDetails UserMailbox).Count
15
```

Disabling an archive removes the ability of the mailbox owner to access the content held in the archive. The data in the archive remains there and does not move back to the primary mailbox. Disabling an archive means that you break the connection between the primary and archive mailboxes. To disable an archive with EAC, open mailbox properties and select the **Manage mailbox archive** link, and then set the state to Off. Exchange Online allows disablement of an archive only if no in-place or litigation holds exist for a mailbox. This

restriction exists to ensure that no one can remove data potentially needed for eDiscovery from Exchange Online. The PowerShell command to disable an archive is:

```
[PS] C:\> Disable-Mailbox -Identity 'Joe Smith' -Archive
```

Although disabling an archive prevents user access to the archive, it does not remove the content from its database. Instead, a 30-day retention period starts. During this time, you can recover the archive and reconnect it to the primary mailbox by re-enabling the archive. Exchange Online removes the archive mailbox after the 30-day deleted mailbox retention period expires. If you examine mailbox properties after disabling an archive, you'll see that the *ArchiveGuid*, a unique value used by Exchange to find the archive mailbox within a database, is a set of zeros. This is how clients know that they shouldn't try to open an archive for this mailbox. However, Exchange preserves the original value of the *ArchiveGuid* in the mailbox's *DisabledArchiveGuid* property, which means that it is easy to re-establish the link to the archive by running the *Enable-Mailbox* cmdlet. After the archive is re-enabled, the two GUID values should be identical.

```
[PS] C:\> Get-ExoMailbox -Identity 'Joe Smith' -PropertySet Archive | Format-List *ArchiveGuid
```

```
ArchiveGuid           : 00000000-0000-0000-0000-000000000000
DisabledArchiveGuid : d7c65fee-c983-4eac-8fa3-6381a8673212
```

```
[PS] C:\> Enable-Mailbox -Identity 'Joe Smith' -Archive | Format-List *ArchiveGuid
```

```
ArchiveGuid           : d7c65fee-c983-4eac-8fa3-6381a8673212
DisabledArchiveGuid : d7c65fee-c983-4eac-8fa3-6381a8673212
```

Several of the PowerShell cmdlets that are commonly used to work with mailboxes include an *Archive* switch to point them to the archive rather than the primary mailbox. The *Get-ExoMailboxStatistics* and *Get-ExoMailboxFolderStatistics* cmdlets are good examples:

```
[PS] C:\> Get-ExoMailboxStatistics -Identity 'Kim Akers' -Archive
[PS] C:\> Get-ExoMailboxFolderStatistics -Identity 'Kim Akers' -Archive
```

```
DisplayName           : Online Archive - Kim Akers
MailboxGuid           : afc1e472-0826-498e-b990-85de223e809d
DeletedItemCount     : 508953
ItemCount             : 75218
TotalDeletedItemSize : 40.34 GB (43,317,558,017 bytes)
TotalItemSize        : 9.129 GB (9,802,226,010 bytes)
```

Note that the *DeletedItemCount* and *TotalDeletedItemSize* properties reported by *Get-ExoMailboxStatistics* refer to information stored in the Recoverable Folders structure. The storage occupied by these items does not count against mailbox or archive quotas. They do apply against the overall 1.5 TB limit for an auto-expanding archive mailbox.

The Data lifecycle management section of the Microsoft Purview Compliance portal also includes a list of archive mailboxes. The same functionality is available there to set up and manage archive mailboxes as through the EAC.

**Naming archives:** The default name for an archive mailbox is *"In-Place Archive – User"* (for example, *"In-Place Archive – Tony Redmond"*). However, you can assign whatever name you like by running the *Set-Mailbox* cmdlet. The name can be at least 80 characters (the largest value I have tested). OWA is the only archive-capable client that will display the name that you assign as Outlook ignores the value and uses a normalized name of *"Online archive – primary SMTP address"*. There is no reason why the two client development teams decided to display different default names for archive mailboxes or why Outlook did not follow OWA in allowing the display of user-specific names. It is just another Exchange quirk.

## Moving Information into Archive Mailboxes

After a mailbox is archive-enabled, it appears automatically in client interfaces. OWA recognizes the archive the next time the user connects to Exchange Online, and the Outlook desktop client learns that the archive exists when it refreshes the list of available resources through the Autodiscover process. When the archive appears, the user can move information to it by creating folders and dragging and dropping items from their primary mailbox.

Exchange Online uses mailbox plans to assign default values to new mailboxes. One of the values in a mailbox plan defines the retention policy for the mailbox. This isn't a Microsoft 365 retention policy. Instead, it's an Exchange mailbox retention policy that can contain a default archive tag to control how items move from the primary mailbox into the archive (if a mailbox is archive-enabled). For instance, if the default archive tag defines a retention period of two years, the Managed Folder Assistant moves items that don't have another retention tag from the primary to the archive mailbox once they reach two years old. The items go into a folder with the same name.

Using a default archive tag means that archive mailboxes grow on an ongoing basis as items move from the primary mailboxes. To remove items from the archive, the retention policy can include a default delete tag. For instance, if the retention period for the default delete tag is five years, the MFA will:

- Obey the retention period specified in the default archive tag and move items from the primary mailbox to the archive after two years.
- Obey the retention period specified in the default delete tag to remove items from the archive after a further three years (the items reach the five-year limit).

In addition, when a mailbox is archive-enabled, MFA moves items captured because of retention processing in the Purges, Versions, and DiscoveryHolds sub-folders under the Recoverable Items folder one day after their creation (depending on when the MFA processes a mailbox, the actual move might take place two or three days after creation). This action frees up space in the Recoverable Items quota assigned to mailboxes (100 GB when a mailbox is archive-enabled) to allow the primary mailbox to store more time-critical data such as the items in the Deletions folder.

Of course, most users are blissfully unaware of the boring details of retention policies. They might not even notice the movement of items from the primary mailbox to the archive mailbox until they look for something and can't find it because the item is not where they thought it should be – in their mailbox. Of course, the item is still available, and it is, in a roundabout way, still in their mailbox. It is just invisible in some respects because the user doesn't know where it is or understand how it got there.

Microsoft defines archive mailboxes as personal repositories. [Microsoft explicitly prohibits](#) the use of transport rules, journal rules, auto-forwarding, or other methods to move information into a mailbox from multiple sources for archiving purposes. However, it is permissible to import data from multiple PSTs that might have originated from multiple users into an archive mailbox. The difference between the two is that using an archive mailbox as an archive destination creates the potential that the archive mailbox will expand on an ongoing basis to cope with the items moving into the archive. On the other hand, if you perform a once-off import to an archive mailbox, its content will probably not change all that much in the future as the archive is essentially historical rather than live. In other words, it's a question of dealing with "hot" data that is constantly growing or "cold" data imported once and hardly ever accessed thereafter.

## Effective Use of Archive Mailboxes

Because primary user mailboxes can hold up to 100 GB, many users can quite happily get along by just using their primary mailbox and never need to use an archive. Consider the following:

- It will take most users several years to accumulate 10 GB of mailbox data. Some users accumulate 20 GB+ of new mailbox content annually, but they are exceptional. On the other hand, people have had quite a while to accumulate information in their mailboxes at this point and many of the mailboxes that move to Exchange Online are already up at the 20 GB mark (or higher).
- In the past, many caveats were expressed about the desirability of holding more than 2 GB in the primary mailbox. Apart from the cost of storage, most problems related to offline access and the need to synchronize such a large amount of data to an OST file whose internal structure was never designed to cope with large volumes. These issues have largely been dealt with today due to faster networks and smarter Outlook synchronization. The larger amounts of data synchronized to clients do slow down OST access speeds, a factor that can be handled by equipping laptops with SSDs that effectively disguise the inefficiency of the OST.
- Archive mailboxes can only be accessed online. Outlook does not synchronize any archive folder into the OST. While the bulk of data moved into the archive might never be accessed again, the potential that someone might need an archived item when a network connection is not available does exist.
- Searches can find items stored in archives but only if the user specifies that Outlook should search "All Mailboxes." Searching the archive in OWA needs the user to be positioned within an archive folder. Neither of these operations is especially intuitive to someone who might be unaware that they have an archive.
- ActiveSync clients cannot access archive mailboxes because the protocol does not support this type of resource. The same is true for Outlook for iOS and Android clients.
- Users seldom want to decide what items should be always available (and in their primary mailbox) and those that can be safely moved into the archive. Expecting users to decide on a data storage strategy is an act of folly.

With these points in mind, it should come as no surprise to find that some tenants avoid the use of archive mailboxes and tell users to exploit the storage now available in primary mailboxes. Indeed, some companies have reversed course and decided to only use primary mailboxes. In some cases, they have had to move information back from archives to primary mailboxes. Apart from dragging and dropping items from folder to folder using Outlook or OWA or exporting archive items to a PST and importing them back into the primary mailbox, there is no in-built method to do this. In most cases, the solution is to create a script to move items using a combination of PowerShell and Exchange Web Services.

On the other hand, some tenants consider archive mailboxes to be the perfect answer to the problem of unmanaged and proliferating PSTs and make strenuous efforts to import data from user PSTs into archive mailboxes so that everything is online and available for compliance. The availability of Microsoft's Import Service, which allows tenants to load PSTs into Azure for later import into user mailboxes (primary or archive) has spurred many companies to consider how they should deal with PSTs in the future. Moving information from PSTs into archive mailboxes exposes the data for compliance purposes, but this is an exercise that involves much more than ingesting data from individual PST files. Before importing anything, you must decide how to collect the PSTs, clean them up (remove corruptions), crack passwords set on the PSTs, and possibly deduplicate the content so that the import processes clean information. Remember the adage that rubbish in equals rubbish out.

Another effective use of archive mailboxes is to hold data migrated from older POP3 or IMAP4 systems. It is also fair to say that companies who have moved data back from third-party archiving solutions to Exchange Online find archive mailboxes to be a natural evolution. In summary, the decision to use or ignore archive mailboxes comes down to the circumstances and business conditions that exist within a tenant.

**PST Import Tools:** Many tools are available to help you move data from user PSTs into primary or archive mailboxes. Like any software utility, you should carefully test the software in your environment to discover which product best meets your needs. Once you have collected and prepared (made sure no corruption

exists) the PSTs, you can use the Import service to either upload the PSTs over the network or send them on hard drives to a Microsoft data center for processing. In either case, the PST data will be ingested into Azure and can be imported from there into user mailboxes.

## Auto-Expanding Archives

The mailbox quotas assigned in enterprise plans are more than enough to handle the storage requirements for most users. The low demand for user interaction makes it practical to allow archive mailboxes to have much higher storage limits than primary mailboxes. In turn, a greater storage capacity allows archive mailboxes to act as long-term repositories to replace PSTs ingested through the Import service. Microsoft's solution is the "auto-expanding archive," meaning that the archive mailbox can expand as it fills.

User or shared mailboxes with E3 and E5 licenses or with the Exchange Online Plan 2 license or the Exchange Online Archiving add-on can use auto-expanding archives. To figure out whether a mailbox is eligible, you can check its persisted capabilities, which expose the licenses assigned to a mailbox. In this case, the presence of *BPOS\_S\_Enterprise* means that the mailbox can use auto-expanding archives. The other values are *BPOS\_S\_ArchiveAddOn* (Exchange Online Archiving) and *BPOS\_S\_Archive* (Exchange Online Plan 2).

```
[PS] C:\> Get-ExoMailbox -Identity TRedmond -Properties PersistedCapabilities | Select -ExpandProperty PersistedCapabilities
```

```
BPOS_S_ThreatIntelligenceAddOn  
BPOS_S_EquivioAnalytics  
BPOS_S_CustomerLockbox  
BPOS_S_Analytics  
BPOS_S_Enterprise
```

In some respects, the technique used to expand archives borrows from the auto-split capability built into public folder mailboxes. When a public folder mailbox approaches 50 GB, the Mailbox Replication Service creates a new mailbox and (MRS) transfers folders to the new mailbox to balance the load.

**Auto-Expand Archive Limits:** A 1.5 TB limit for auto-expanding archives applies from November 1, 2021. While there are thousands of terabyte-plus archives in use, the limit is unlikely to affect many tenants. If the change affects your organization, you should contact Microsoft support. A script to report how close archive-enabled mailboxes are to the limit is [downloadable from GitHub](#).

[Microsoft's guidelines](#) say that only individual or shared mailboxes with a growth rate that does not exceed 1 GB per day support expandable archives. The reason for this restriction is to ban the use of archive mailboxes as targets for the migration of non-personal data from legacy services. It is perfectly acceptable to use an archive mailbox as a migration target for personal data such as PSTs.

### Enabling Auto-Expanding Archives

To enable auto-expanding archives, run the *Set-OrganizationConfig* cmdlet to update the tenant configuration:

```
[PS] C:\> Set-OrganizationConfig -AutoExpandingArchive
```

Once set, the new configuration ensures that all existing archives become auto-expanding, and any new archives are auto-expanding. The process of enabling existing archives can take some time to complete due to the need for the Managed Folder Assistant to process each mailbox. Only [a subset of Exchange Online clients](#) can access an auto-expanding archive. Other clients can access information in the primary archive but cannot access any data moved into an auxiliary archive.

If you do not want to enable auto-expanding archives for all mailboxes, you can control the capability on an individual basis using the *Enable-Mailbox* cmdlet. An archive must already exist for the mailbox before you can make it auto-expanding. For instance, to enable an auto-expanding archive for Kim Akers:

```
[PS] C:\> Enable-Mailbox -Identity "Kim Akers" -AutoExpandingArchive
```

In a hybrid deployment, you can enable auto-expansion for a cloud archive where the primary mailbox is on-premises. However, if you do this, you will not be able to move the archive back to Exchange Server.

Some thought and planning are necessary for the deployment of auto-expanding archives. Enablement is a one-way process. No method exists to collect the pieces of an auto-expanding archive together to make them into a single "standard" archive mailbox. In addition, you cannot move a mailbox with an auto-expanding archive to an on-premises server until Microsoft releases a version of Exchange Server which supports the transfer of these mailboxes. In the interim, if you want to move items out of an auto-expanding archive to transfer data to another mail system or an on-premises server, you will have to export the needed data to PSTs. Running an eDiscovery content search to find and export the data is an effective way to achieve this goal. You can make a mailbox with an auto-expanding archive inactive by putting it on hold and then removing the user account. Although this allows for the long-term preservation of mailboxes for compliance purposes, you should still review your procedures covering how to remove user accounts from the tenant to ensure that everything works as you expect.

When you enable auto-expanding archives for a mailbox, Exchange:

- Increases the normal 100 GB quota for the primary mailbox to 110 GB and modifies the quota warning threshold from 90 GB to 100 GB.
- If the mailbox comes within the scope of a hold, Exchange increases the Recoverable Items quota from 100 GB to 110 GB.

It can take up to 30 days before the process to enable auto-expansion for an archive completes. The changes to the primary mailbox quotas reflect the fact that expanding archives are more likely used by busy mailboxes, so it makes sense to give the primary mailbox a little extra headroom to ensure that the mailbox can continue operating while the process to enable the auto-expanding archive is active.

One limitation is that if you need to search auto-expanding archives with OWA or Outlook, you can only search within a specific folder. eDiscovery content searches can find information stored in any part of an auto-expanding archive, which also supports litigation and in-place holds as normal.

In many cases, it is preferable to enable auto-expanding archives for selected accounts instead of a complete tenant. Those accounts are usually those that have a genuine business need to keep massive quantities of information for certain periods.

## How an Archive Mailbox Expands

When you enable a mailbox for archiving, it starts with a single 110 GB archive mailbox. After the occupied space within the archive mailbox approaches the transition threshold (90 GB), a mailbox assistant automatically provisions a new archive mailbox. Exchange calculates the occupied size from the total size of folders in the archive mailbox with their *Movable* flag set to *\$True* or the *FolderType* set to *DeletedItems* or *RecoverableItems*.

To find the set of mailboxes with expandable archives, run the command:

```
[PS] C:\> Get-ExoMailbox -ResultSize Unlimited -RecipientTypeDetails UserMailbox, SharedMailbox | ? {$_.MailboxLocations -like "*AuxArchive*"} | Format-Table DisplayName, MailboxLocations
```

To gain an insight into the current state of the primary archive for an individual mailbox, we can use PowerShell to scan the folders in the primary archive to report the occupied space. Later in this section, we discuss how to retrieve the GUID for the primary archive to use as input for the *Get-ExoMailboxFolderStatistics* cmdlet.

```
[PS] C:\> $CheckUser = Read-Host "Enter User to check"  
$UserGuid = (Get-ExoMailbox -Identity $CheckUser | Select -ExpandProperty ExternalDirectoryObjectId)
```

```

$MLO = (Get-MailboxLocation -User $UserGuid | Select MailboxLocationType, MailboxGuid | ?
{$_MailboxLocationType -eq "MainArchive"})
$Folders = (Get-ExoMailboxFolderStatistics -Identity $MLO.MailboxGuid.Guid | Select FolderPath,
Movable, FolderType, Name, ItemsInFolder, FolderSize)
$NumFolders = 0; $TotalSize = 0
$Folders | Add-Member -MemberType ScriptProperty -Name FolderSizeInBytes -Value {$this.FolderSize -
replace "(.*\(|,| [a-z]*\)", ""}
ForEach ($F in $Folders) {
    If ($F.FolderType -eq "DeletedItems" -or $F.FolderType -eq "RecoverableItems" -or $F.Movable -eq
$True) {
        $TotalSize = ($TotalSize + $F.FolderSizeInBytes)
        $NumFolders++ }
}
$TotalSizeGB = [math]::round($TotalSize/1GB, 3)
$ThresholdPercent = ($TotalSizeGB/90).ToString("p")
Write-Host $NumFolders "movable folders found. Occupied space" $TotalSize "bytes or" $TotalSizeGB
"GB." "At" $ThresholdPercent "of 90 GB transition threshold"

Enter User to check: TRedmond
92 movable folders found. Occupied space 5834011419 bytes or 5.433 GB At 6.04% of the 90 GB
transition threshold

```

After provisioning, the new archive mailbox is a “chunk” or “shard,” or more correctly “auxiliary archive,” and joins the overall archive. Exchange links the GUID for the new auxiliary archive with the GUIDs of the existing auxiliaries and primary archive to form a chain or set of mailboxes that the Information Store considers a single logical entity. The expansion of an archive to form an archive chain occurs without user intervention and without affecting supported clients, which continue to query the Information Store for data and receive data back without knowing which part of the archive holds the data.

The Managed Folder Assistant coordinates the movement of information out of the primary archive to an auxiliary archive to reduce the size of the primary archive under the transition threshold. This exercise aims to move enough data to the auxiliary archive to get the primary archive under 50% of its current size. So, if the primary archive grows to 95 GB, the Managed Folder Assistant examines the folders in the primary archive and selects enough to move approximately 47.5 GB to the auxiliary archive. The Deleted Items and Recoverable Items folders are not movable, but the Managed Folder Assistant can create sub-folders under these folders in the target archive and move content there.

The Mailbox Replication Service moves the data from the primary to the auxiliary archive and takes care of ongoing synchronization to ensure that any changes made to data while the move progresses are in the moved data. The copying of content occurs in the background. To ensure that no data loss occurs during the rebalancing of the archive, the primary archive keeps the copied data for 30 days. When this period elapses, the Managed Folder Assistant flushes the copied data from the primary archive to release the space.

## Archive Links

Archives link to their primary mailboxes by storing the GUID pointing to the archive as a mailbox property. The GUID is enough for Exchange to find the archive in a database. The auto-expanding archive replaces the single GUID that connects the mailbox to the archive with a linked list of GUIDs. Each of the GUIDs points to a separate auxiliary archive of up to 50 GB, which Exchange Online combines with the other auxiliary archives and the primary archive to form a logical archive mailbox.

You can see the details of the GUIDs by running the *Get-ExoMailbox* cmdlet to examine a mailbox’s properties. If you look at the *MailboxLocations* property, you will see something like this:

```

[PS] C:\> Get-ExoMailbox -Identity TRedmond -Properties MailboxLocations | Select -ExpandProperty
MailboxLocations

1;0370f354-2752-4437-878d-cf0e5310a8d4;Primary;eurprd04.prod.outlook.com;d96ca5a2-340d-4c83-be33-
d4a7a8c9b1d6
1;afc1e472-0826-498e-b990-85de223e809d;MainArchive;eurprd04.prod.outlook.com;e46d4e31-3734-47dc-
801d-5d59f9988766

```

The information about mailbox locations reported by *Get-ExoMailbox* divides into two sections: one for the primary mailbox and the second for the archive. Only the primary mailbox and the primary archive are shown here. If other auxiliary archives are present, the archive section lists them as mailboxes 2, 3, 4, and so on. The information for the two mailboxes is:

Primary mailbox:

- The *ExchangeGUID* (which ties the mailbox back to a user account).
- "Primary" to show that this data refers to the user's primary mailbox.
- If, as in this case, the mailbox database is in Exchange Online, the name of the Exchange Online forest is noted (eurprd04).
- The GUID of the database holding the mailbox.

Archive mailbox:

- The *ArchiveGUID* (which is only present when a mailbox is archive-enabled).
- "MainArchive" to show that this data is an archive set. When auxiliary archives are part of the set, they are tagged with "AuxArchive."
- The name of the Exchange Online forest hosting the archive mailbox. This value is empty for on-premises mailboxes.
- The GUID of the database holding the archive.

In this case, the mailbox and the archive are in the same database (you can confirm this by using *Get-Mailbox* to examine the *Database* and *ArchiveDatabase* properties). And, as you would expect, both the primary and archive mailboxes are in the same Exchange Online forest. Another way of accessing information about these archives is with the *Get-MailboxLocation* cmdlet. For example:

```
[PS] C:\> Get-MailboxLocation -User TRedmond | Sort MailboxLocationType -Descending | Format-Table MailboxGUID, MailboxLocationType
```

MailboxGuid	MailboxLocationType
0370f354-2752-4437-878d-cf0e5310a8d4	Primary
afc1e472-0826-498e-b990-85de223e809d	MainArchive
bb131464-1461-147e-b774-41646ddadd11	AuxArchive

In this case, the *MailboxGuid* property for the user's primary mailbox and all the parts of the archive are more obvious. The *MailboxGuid* is needed if you want to check how much data Exchange has moved to an auxiliary archive. For example:

```
[PS] C:\> Get-ExoMailboxStatistics -Identity bb131464-1461-147e-b774-41646ddadd11 | Format-Table ItemCount, TotalItemSize
```

ItemCount	TotalItemSize
30225	4.398 GB (4,722,299,639 bytes)

We now know that Exchange has moved a certain amount of data to the auxiliary archive.

**Migrating Large On-Premises Archives:** On-premises archive mailboxes can grow past 100 GB. At this point, they can no longer be migrated to Exchange Online using the online migration tools. The reason is that even for auto-expanding archives, the initial quota assigned by Exchange Online is all that's available until the auto-expansion provisioning process completes. Even then, an auto-expanding archive cannot add storage dynamically in the middle of an online migration. For this reason, two options exist if you have large on-premises archive mailboxes to move to Exchange Online: reduce the archives to under 100 GB and use the online migration tools or find a different migration tool. One approach is to export archive data to PSTs and import the PSTs into Exchange Online using the Office 365 Import Service. A third-party



migration service might be able to automate much of the processing needed to get the data imported into Exchange Online.

## Outlook's Archive Folder

Exchange Online mailboxes include an Archive folder as one of the default folder set created inside all mailboxes. Apart from a user being able to apply a retention tag to the Archive folder to move items in the folder to the archive mailbox after a period, the Archive folder has no relationship to the "online archive." Instead, Microsoft envisages the Archive folder as a convenient place to move items from the Inbox after a user has finished processing the items but wishes to keep them for a period.

The advantage of putting items in the Archive folder rather than removing them or having Exchange Online move them to the archive mailbox is that items in the Archive folder are available to mobile clients. Items in the archive mailbox are inaccessible to mobile clients because the protocols used by these clients only support primary mailboxes. Items in the Archive folder are also available offline while items in the archive mailbox are only accessible when a network connection is available. Sometimes that network connection is not fast enough to allow easy access to archived items, including searching those items. Because local searches can process cached copies of the Archive folder, searches complete faster, and items are more accessible.

Outlook, OWA, and Outlook Mobile support one-click (or swipe) options to allow users to easily move items to the Archive folders. The idea is that a user can quickly move through their Inbox to triage items by reading, removing, or archiving items as required. Because it is a mailbox folder, you can synchronize the Archive folder to other mobile clients, but one-click access will not be available unless the developer of the client incorporates the necessary GUI.

## Shared Mailboxes

Shared mailboxes have been part of the Exchange product since Exchange 2000. They meet the need to have a mailbox to handle messages that a team of people might process, such as the team staffing a support desk. The implementation of shared mailboxes in Exchange Online is like that found on-premises, with the disabled user object used to instantiate the shared mailbox created in Azure AD. Shared mailboxes have many uses, including:

- To allow groups of users shared access to functional email. For example, all the support agents who staff a help desk can use a shared mailbox to review and respond to messages sent to the help desk by users asking for help. Sometimes these mailboxes are called functional mailboxes and are often used to ensure that incidents can be managed effectively by multiple people across shift boundaries. A major attraction of this use is that messages sent in response come from the shared mailbox rather than the individual user.
- To allow users access to the mailbox of a colleague who has left the company. This is done by changing the personal mailbox into a shared mailbox and assigning access to the mailbox to those who need access. This technique is used extensively in on-premises organizations.

Mailbox delegation goes together with shared mailboxes because there is not much point in creating a shared mailbox if you cannot then access the mailbox. Because a shared mailbox is linked to a user object in Azure AD, access to its contents must be gained by delegating or granting permissions over the mailbox to other users. The *Get-ExoMailbox* cmdlet enables us to discover the set of shared mailboxes known in a tenant:

```
[PS] C:\> Get-ExoMailbox -RecipientTypeDetails SharedMailbox | Format-Table Name, DisplayName, Alias
```

Name	DisplayName	Alias
----	-----	-----

<b>Customer Services</b>	<b>Customer Services</b>	<b>CServices</b>
<b>Book Feedback</b>	<b>Book Feedback</b>	<b>BookComments</b>
<b>Help Desk</b>	<b>Corporate Help Desk</b>	<b>HelpDesk</b>

A shared mailbox can be hosted by Exchange Online or, in a hybrid environment, on-premises. In this instance, the shared mailbox is known as a remote shared mailbox when accessed from the other platform. It is best to keep shared mailboxes on the same platform as used by the accounts that have delegated access to the mailboxes. In other words, if you want to move some shared mailboxes to Exchange Online, move the mailboxes that access those shared mailboxes to Exchange Online too. Exchange Online places no limit on the number of shared mailboxes that you can create within a tenant.

## Licensing, Quotas, and Limitations of Shared Mailboxes

By default, a shared mailbox does not need a license. However, once a feature described in Table 6-2 is enabled for a shared mailbox, you should assign a license. Go to the Active users view in the Microsoft 365 admin center, select the shared mailbox, edit its properties, and assign the license under **License and Apps**.

<b>Feature</b>	<b>License</b>
Enable an archive for a shared mailbox.	Exchange Online Plan 1 or Exchange Online Archiving.
Place a hold on a shared mailbox (this includes the mailbox being a custodian in a Microsoft 365 Advanced eDiscovery case).	Exchange Online Plan 2.
Exceed the default 50 GB mailbox quota.	Exchange Online Plan 2.

Table 6-2: Licensing requirements for shared mailboxes

Microsoft documentation has always said that the default quota for shared mailboxes is 50 GB. Because of some errors in the provisioning process, many shared mailboxes received a 100 GB quota without the need for a license. These mailboxes keep their 100 GB quota unless the mailbox state changes. For example, if you convert a shared mailbox to be a user mailbox and then convert it back again, the shared mailbox object then needs a license to keep its 100 GB quota. New shared mailboxes receive the correct 50 GB quota.

Here's a PowerShell script to report on the set of shared mailboxes in a tenant. The report shows the current number of items in each mailbox, the size of the mailbox, the assigned quota, whether it is licensed, and if it has an archive. In this case, only one shared mailbox is licensed, one was previously licensed but had the license removed, and the others have never been licensed. All the mailboxes existed before Microsoft put the new provisioning process in place, so they all have 100 GB quotas.

```
[PS] C:\> $SMbx = Get-ExoMailbox -RecipientTypeDetails SharedMailbox -ResultSize Unlimited -
PropertySet All
$Report = [System.Collections.Generic.List[Object]]::new()
ForEach ($S in $SMbx) {
    $Archive = "Disabled"
    If ($S.ArchiveName -ne $Null) { $Archive = "Enabled" }
    $Quota = $S.ProhibitSendReceiveQuota.Split("(")
    $Stat = (Get-ExoMailboxStatistics -Identity $S.UserPrincipalName | Select ItemCount,
TotalItemSize)
    $ShMbxSize = $Stat.TotalItemSize.Value
    $Size = $ShMbxSize.ToString().Split("(")[0]
    $ReportLine = [PSCustomObject]@{
        Mailbox      = $S.DisplayName
        TotalItems   = $Stat.ItemCount
        MailboxSize  = $Size
        Quota        = $Quota[0]
        Licensed     = $S.SkuAssigned
        Archive      = $Archive }
    $Report.Add($ReportLine)}
$Report | Format-Table Mailbox, TotalItems, MailboxSize, Quota, Licensed, Archive -AutoSize
```

Mailbox	TotalItems	MailboxSize	Quota	Licensed	Archive
Customer Services	213	7.484 MB	100 GB	False	Enabled
Book Feedback	301	7.278 MB	100 GB		Enabled
Redirect for Removed Mailboxes	198	814 KB	100 GB		Disabled
Redmond Shared Events	3777	190.5 MB	100 GB		Disabled
Company Information	252	897.4 KB	100 GB		Disabled

If a shared mailbox without a license exceeds 50 GB, Exchange stops delivering email to the mailbox. You can assign a license to the mailbox to restore delivery. Because Exchange caches mailbox information, the newly licensed status, and the 100 GB quota, take about 15 minutes to become effective, after which email delivery recommences.

If you convert a user mailbox holding more than 50 GB to a shared mailbox (often done to keep mailboxes for ex-employees), Exchange continues to deliver mail to the mailbox while the mailbox is licensed. In most cases, people convert user mailboxes to shared mailboxes to free up licenses, so if you go ahead and do this, Exchange detects that the mailbox holds more than 50 GB and ceases email delivery. You can restore delivery by removing content from the mailbox to get it under the 50 GB quota or by assigning a new license.

Moving shared mailboxes from on-premises Exchange organizations will fail if they hold more than 50 GB and the MailUser object for the target shared mailbox in Exchange Online is unlicensed. To allow migrations to succeed (up to the 100 GB quota), assign a license to the MailUser object. This limitation exists for all migrations performed by the Mailbox Replication Service (MRS) or third-party migration utilities.

An unlicensed shared mailbox can be placed on hold or archive-enabled as Exchange Online does not perform a licensing check before enabling the feature. These are technical licensing breaches that will only come to light in an audit. Even if the mailboxes continue to work without a license, the danger always exists that Microsoft might turn on code to enable restrictions to stop unlicensed shared mailboxes from working, so it's best to make sure that the correct licenses are in place.

Although companies often use shared mailboxes for functional purposes, apart from inbox rules, Exchange doesn't have any out-of-the-box way to automate the processing of messages arriving in a shared mailbox, so a mailbox delegate must sign into the mailbox to handle new emails. Performance for a shared mailbox can be problematic if more than twenty users try to access it concurrently. Finally, users who have an Exchange Online frontline/kiosk license cannot add delegates to their mailbox, but they can access a shared mailbox if someone makes them a delegate for that mailbox.

## Creating a New Shared Mailbox

New shared mailboxes are created through the **Groups** section of the Microsoft 365 admin center. Not much information is needed to create a shared mailbox (Figure 6-7). After you click Add, Exchange Online needs a short delay to provision the mailbox and make it ready to add the users who will access and use the mailbox. Remember that users do not log into a shared mailbox as happens with a regular mailbox and that access to its contents occurs by opening the mailbox as a secondary resource within Outlook or OWA.

✕

## Add a shared mailbox

Email can be sent to and from the name and email address of the shared mailbox, rather than an individual. After you create the shared mailbox, you can add members who can read and reply to email.

Name \*

Email \*

 @  

[Save changes](#)

Figure 6-7: Creating a new shared mailbox with the Microsoft 365 admin center

## Mailbox Delegates and Permissions

After creating a shared mailbox, you must assign permission to the mailbox to allow access to those who need to work with it (its delegates). You can assign permissions when creating a new shared mailbox or by editing the mailbox object in the Microsoft 365 admin center or through PowerShell. Mailboxes support three delegate permissions:

- **Full Access** permission allows a user to open the shared mailbox and access its contents. Full Access does not mean that a user has *SendAs* permission. Microsoft 365 refers to this permission as “Read and manage mail to this mailbox.”
- **Send As** permission allows a user to send messages from the shared mailbox that appear as if the messages came from the shared mailbox rather than the user. In effect, a user with this permission can impersonate the mailbox. This is the best situation when you want replies to conversations started from the mailbox to flow back to the shared mailbox.
- **Send On Behalf Of** permission means that a user can send an email on behalf of the mailbox. Unlike the Send As permission, the name of the user who sends the message is obvious. You can only set this permission for a shared mailbox through PowerShell.

Full Access permission allows a user to work with all folders and items inside the shared mailbox. Users need both Full Access and Send As permissions to work with a shared mailbox as if it was their personal mailbox. In addition to shared mailboxes, which is why these are the two permissions featured in the admin centers, Exchange Online supports the assignment of the Send As and Send on Behalf of permissions to user and group mailboxes.

After creating a shared mailbox with the Microsoft 365 admin center or EAC, you can use the **manage mailbox permissions** option (Microsoft 365 admin center) or **Manage mailbox delegation** option (EAC) to define permissions for the shared mailbox. As noted above, the following permissions are available for shared mailboxes:

- Read and manage (Full Access).
- Send As.

Figure 6-8 shows that two users have full access to a shared mailbox (oddly, three have the Send As permission). Because Exchange Online caches permissions for better performance, it can take up to an hour for the new permission to be effective. In addition to a refresh of the server caches, clients must learn that

users have permissions for a shared mailbox. For example, Outlook clients learn about access to a shared mailbox through its Autodiscover process, which runs every 15 minutes for this purpose.

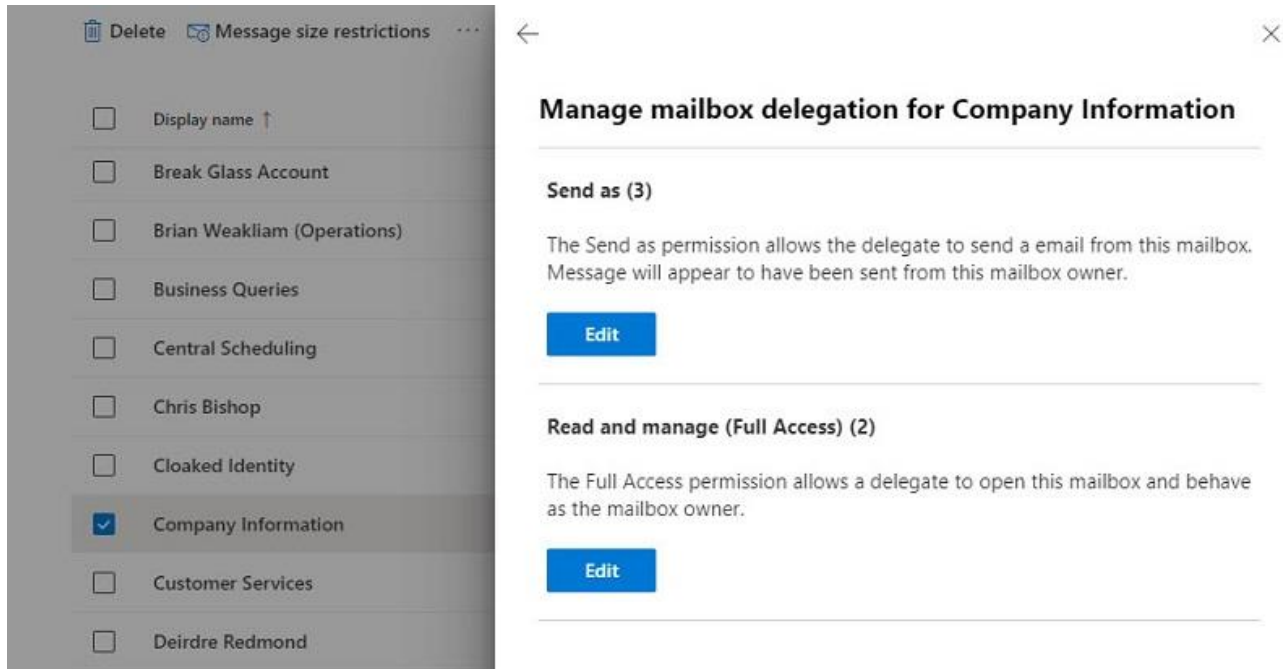


Figure 6-8: Assigning delegate permissions for a shared mailbox (EAC)

If you need to assign access to a shared mailbox to multiple users, it is often more convenient to assign permission to a distribution list, which must be a security-enabled group. You cannot assign permissions for a shared mailbox to a normal distribution list because these groups are not security principals, and you cannot assign permissions to a group or a dynamic distribution list either.

You can edit the delegate permissions for a shared mailbox to add or remove members or work with individual permissions. Select the mailbox in the Microsoft 365 admin center or EAC and use **Manage mailbox permissions** or Manage mailbox delegation to assign or remove individual permissions. Once again, it will take a little while before the amended permissions become fully active across the tenant.

**Shared mailbox or Groups:** On the surface, Groups seem like a better and more modern choice as the basis for team sharing than a shared mailbox. That statement might be true if the requirement is to collaborate based on shared documents, meetings in the group calendar, and threaded conversations with perhaps a link to a plan managed by Planner. However, shared mailboxes still have some unique strengths. For instance, shared mailboxes allow access to a full range of folders rather than the limited set used in a group mailbox. In addition, shared mailboxes support shared contacts and tasks while these are unavailable to groups. Shared mailboxes also support categories, rules, and Outlook add-ins, all of which can be very important to customer support or sales teams. The point is that the two types of mailboxes are suitable for different purposes. Think about how people need to share information and what type of information they need to work with before you select which type to use.

## Handling Messages Sent from Shared Mailboxes

If you use a shared mailbox for customer communications or a similar purpose, you probably want to keep any replies sent by people using the *SendAs* or *Send on Behalf of* permission in the mailbox. The default behavior is that Exchange keeps messages in the mailbox of the person who sends a message, even if they are replying as or on behalf of the shared mailbox. This is fine for personal mailboxes, but not so good for shared mailboxes. Consider the example of a mailbox used to receive customer comments or complaints where a team of support agents accesses the mailbox. If an agent answers a message, their mailbox stores the reply

and none of the other team members know that the customer has received a response (or what that response said).

You can control these settings with PowerShell by running the *Set-Mailbox* cmdlet. For example, this command tells Exchange to retain copies of messages sent using the Send As and Send on Behalf Of permissions in a shared mailbox.

```
[PS] C:\> Set-Mailbox -Identity SharedMailbox -MessageCopyForSentAsEnabled $True
-MessageCopyForSendOnBehalfEnabled $True
```

To ensure consistency and enable these settings across all shared mailboxes in the tenant, the command is:

```
[PS] C:\> Get-ExoMailbox -RecipientTypeDetails SharedMailbox -PropertySet All | ?
{$_ .MessageCopyForSendAsEnabled -eq $False -or $_ .MessageCopyForSendOnBehalfEnabled -eq $False} |
Set-Mailbox -MessageCopyForSendOnBehalfEnabled $True -MessageCopyForSentAsEnabled $True
```

Somewhat along the same vein is the situation that occurs when a member removes items from a shared mailbox. Logically, you might think that these items would go into the Deleted Items folder of the shared mailbox. However, Outlook moves such items into the Deleted Items folder of the delegate's mailbox. To change this behavior, update the system registry by creating a new DWORD value at:

```
HKEY_CURRENT_USER\Software\Microsoft\Office\[xxx]\Outlook\Options\General\DelegateWastebasketStyle
```

Replace [xxx] with 15.0 for Outlook 2013, and 16.0 for both Outlook 2016 and Outlook 2019. Set the value to 4 (four) and restart Outlook. Deleted items will now stay in the Deleted Items folder of the shared mailbox, which is really where they should be.

## Creating and Managing Shared Mailboxes with PowerShell

The PowerShell command to create a new shared mailbox is straightforward as all you must specify is the name. However, it's good practice to specify an alias, and primary SMTP address (which must belong to the tenant domain) to make sure that they have the expected values. Other important mailbox properties such as the MRM policy, RBAC role assignment policy, and mailbox quotas are set automatically. This command is an example of how to create a new shared mailbox.

```
[PS] C:\> New-Mailbox -Shared -Name "End User Services" -Alias Shared.EndUserServices
-PrimarySmtpAddress "EndUserServices@Office365ITPros.com"
```

Remember to create a unique alias for the new shared mailbox. In the example above, the alias is created by prefixing "Shared" and a period in front of a value derived from the mailbox's display name. This should be enough to ensure uniqueness.

After creating the new shared mailbox, we need to add some members. Somewhat confusingly, we need to run two different cmdlets to grant the full set of permissions over the mailbox. The reason why this situation exists is that we need to grant a user permission to open the mailbox and work with folders and then the right to send as (impersonate) the mailbox. In other words, the first is an access right, the second is the right to do something for the mailbox. In this example, we use the *Add-MailboxPermission* and *Add-RecipientPermission* cmdlets to grant Kim Akers *FullAccess* rights to the mailbox and then the *SendAs* permission.

```
[PS] C:\> Add-MailboxPermission -Identity 'Book Feedback' -User 'Kim Akers' -AccessRights FullAccess
```

```
[PS] C:\> Add-RecipientPermission -Identity 'Book Feedback' -Trustee 'Kim Akers' -AccessRights
SendAs -Confirm:$False
```

You can remove any extraneous permissions with the *Remove-MailboxPermission* or *Remove-RecipientPermission* cmdlets. For example:

```
[PS] C:\> Remove-MailboxPermission -Identity 'Book Feedback' -User 'Kim Akers' -AccessRights
FullAccess -Confirm:$False
```

```
[PS] C:\> Remove-RecipientPermission -Identity 'Book Feedback' -Trustee 'Kim Akers' -AccessRights SendAs -Confirm:$False
```

It is sensible to conduct periodic reviews of the permissions that are assigned to shared mailboxes and to remove permissions that are no longer needed. You can use the *Get-ExoMailboxPermission* and *Get-RecipientPermission* cmdlets to check who has access to a mailbox and what they can do. Here is an example of using the *Get-RecipientPermission* cmdlet to view details of accounts that have the *SendAs* permission for a mailbox. The same command works for both shared and personal mailboxes.

```
[PS] C:\> Get-RecipientPermission -Identity 'Book Feedback'
```

Identity	Trustee	AccessControlType	AccessRights	Inherited
Book Feedback	NT AUTHORITY\SELF	Allow	{SendAs}	False
Book Feedback	Tony Redmond	Allow	{SendAs}	False
Book Feedback	Kim Akers	Allow	{SendAs}	False

The *Get-ExoMailboxPermission* cmdlet returns access permissions for a mailbox. In the past, the full set included many system accounts and management role groups that have access to mailboxes. Today, the set is limited to users and system entries like NT AUTHORITY\SELF. In this example, we use the *Get-ExoMailboxPermission* cmdlet to check the permissions for a mailbox and trim the returned set to only report delegated user accounts:

```
[PS] C:\> Get-ExoMailboxPermission -Identity 'Book Feedback' | ? {$_.User -like "*@*"} | Format-Table User, AccessRights
```

User	AccessRights
Tony.Redmond@office365itpros.com	{FullAccess}
Kim.Akers@office365itpros.com	{FullAccess}

A more complete script to create a report of all non-standard permissions (FullAccess, Send on Behalf Of, and SendAs) currently present on user and shared mailboxes is [available on GitHub](#).

We can also use the permissions on shared mailboxes to discover the list of shared mailboxes that a user can access:

```
[PS] C:\> Get-ExoMailbox -RecipientTypeDetails SharedMailbox | Get-MailboxPermission -User "Kim.Akers@Office365itpros.com"
```

## Send on Behalf Of permission

Exchange Online also supports assigning the *Send on Behalf Of* permission for a mailbox. The difference in terms of functionality between the *SendAs* and *Send on Behalf Of* permissions is the degree of impersonation implied. When a user sends a message on behalf of another person, the message is marked as such and you know that someone else has taken responsibility for composing the message. The *Send On Behalf Of* permission is intended to cover scenarios such as when an assistant processes emails on behalf of an executive. You know that the executive did not personally send the message (because that fact is clear in the message header), but the message has still come from their mailbox. In the world of letter-writing, using the *Send on Behalf Of* permission to send a message is the equivalent of signing a letter for someone and adding "pp" (per pro) beside your signature.

You can use the *Set-Mailbox* cmdlet to grant the *Send on Behalf Of* permission to a shared mailbox, just like you can grant the permission to send messages on behalf of distribution lists or dynamic distribution lists with the *Set-DistributionGroup* and *Set-DynamicDistributionGroup*, or indeed, for a Microsoft 365 Group using the *Set-UnifiedGroup* cmdlet. The same syntax is used in all cases. For instance, to grant the permission for the Customer Services shared mailbox to Jill Smith:

```
[PS] C:\> Set-Mailbox -Identity "Customer Services" -GrantSendOnBehalfTo "Jill Smith"
```

The command above overwrites any existing permission. To add someone to the list of those allowed to send on behalf of a mailbox, use this format to add the user.

```
[PS] C:\> Set-Mailbox -Identity "Customer Services" -GrantSendOnBehalfTo @{Add="Terry.Hegarty"}
```

You can pass a comma-separated list of user accounts to grant permission at the same time, as in this example:

```
[PS] C:\> Set-Mailbox -Identity "Customer Services" -GrantSendOnBehalfTo
@{Add="Ken.Bowers","James.Ryan"}
```

To view the accounts that hold *Send on Behalf Of* permission to mailboxes, you can run this command:

```
[PS] C:\> Get-ExoMailbox -RecipientTypeDetails SharedMailbox -Properties GrantSendOnBehalfTo | ?
{$_GrantSendOnBehalfTo -notlike ""} | Format-Table Name, GrantSendOnBehalfTo
```

Even if you grant an account the *FullAccess* and *Send on Behalf Of* permissions to a mailbox, this is different from giving the account the *Send As* permission. *Send As* is a different permission to assign users the ability to send messages as if the sender is the mailbox owner. If a user has both the *Send As* and *Send on Behalf Of* permissions for a mailbox, the client applies the *Send As* permission when they send a message.

Remember to change the behavior for Sent Items storage as described earlier if you want to keep copies of messages sent using the *Send As* or *Sent on Behalf Of* permissions in the mailbox.

## Mailbox Automapping

Like their on-premises counterparts, Exchange Online shared mailboxes can be auto-mapped as an Outlook resource who have full access to the mailboxes. Automapping means that Autodiscover returns details of the shared mailbox in the XML manifest generated when an Outlook client queries the server to discover the set of resources available to it. Because the shared mailbox is part of its resource set, Outlook automatically opens the shared mailbox. Automapping occurs every time the Autodiscover component refreshes information for Outlook, so it will occur shortly after Outlook starts and every 60-90 minutes thereafter. After Outlook refreshes its set of resources, the shared mailbox appears in the set of resources just like any other mailbox.

An example of the XML from an Autodiscover manifest to define a shared mailbox as a resource is below.

```
<AlternativeMailbox>
  <Type>Delegate</Type>
  <DisplayName>Editing Team</DisplayName>
  <SmtpAddress>EditingTeam@Office365ITPros.com</SmtpAddress>
  <OwnerSmtpAddress>EditingTeam@Office365ITPros.com</OwnerSmtpAddress>
</AlternativeMailbox>
```

Including a shared mailbox in the Autodiscover manifest means that the Outlook desktop client opens the mailbox no matter what workstation the user uses to run Outlook. Although automapping usually is convenient, you might not always want to automap a shared mailbox for all users. If so, you must remove the *FullAccess* permission for the mailbox from that user's account and then grant it again, this time specifying that auto-mapping should not occur. No options are available to control automapping in the administrative portals, so it must be done through PowerShell. Another way of handling the issue is to use the *Remove-MailboxPermission* cmdlet. For example, this command removes all automapping for the shared mailbox passed in the identity parameter. After running the command, users who have Full Access to the shared mailbox will continue to have access but automapping will not happen.

```
[PS] C:\> Remove-MailboxPermission -Identity 'Customer Services' -ClearAutoMapping
```



You can clear the complete set of mailbox permissions from a mailbox by running *Remove-MailboxPermission* with the *ResetDefault* switch. This instructs Exchange to reset the mailbox back to its default state by removing all delegated permissions. The *Send As* and *Send on Behalf Of* permissions are unaffected.

```
[PS] C:\> Remove-MailboxPermission -Identity 'Customer Services' -ResetDefault
```

Note that you cannot simply flip the automapping switch from *\$True* to *\$False* once a set of permissions exist. If you make a mistake, you must remove the permissions and then add them back to the mailbox, making sure that the correct automapping choice is in place this time. In this example, we use the *Remove-MailboxPermission* cmdlet to remove the permission from the mailbox and then the *Add-MailboxPermission* cmdlet to add the permission back again:

```
[PS] C:\> Remove-MailboxPermission -Identity "Shared Mailbox" -User "UserWithAccess"
-AccessRights FullAccess
[PS] C:\> Add-MailboxPermission -Identity "Shared Mailbox" -User "UserWithAccess"
-AccessRights FullAccess -AutoMapping:$False
```

The next time Outlook refreshes its list of resources from Autodiscover, it will remove the shared mailbox. Alternatively, you can close and re-open Outlook to accelerate the process.

## Delegated Folder-Level Permissions

Outlook and OWA support folder-level permissions to allow delegates access to specific folders in a mailbox. This is an older form of delegate access that has existed in Outlook for almost 20 years to support the classic manager-assistant scenario where the manager delegates access over their inbox and calendar to the secretary to allow that user to process inbound emails. You can find more information about how to use folder-level delegation in this [support article](#).

Folder-level permissions are set with the *Set-MailboxFolderPermission* cmdlet and retrieved with *Get-MailboxFolderPermission* (or *Get-ExoMailboxFolderPermission*). As an example of their use, [this script](#) generates a report of all folder-level delegated permissions in a tenant. As with mailbox permissions, it is wise to conduct periodic reviews of folder-level permissions to ensure that people don't hold permissions that they no longer need.

## Who Sent That Mail?

One of the questions that often occurs is who sent a certain message from a shared mailbox. The question does not arise when someone uses the *Send on Behalf Of* permission because their name appears in the message header, but *SendAs* takes the name of the shared mailbox and does not tell you who created and sent the message. If mailbox auditing is enabled for the shared mailbox (it is by default), Exchange captures audit records for actions such as *SendAs* and you can search those records to discover the sender. In this example, we run the *Search-MailboxAuditLog* cmdlet to find the audit records for the Customer Services shared mailbox for a selected five-hour period and then filter the records to focus on those for *SendAs* events. We then display the name of the user who signed into the shared mailbox and the date and time of the event. By comparing the date and time with the timestamp in the message header, you know who sent the message.

```
[PS] C:> Search-MailboxAuditLog -Identity "Customer Services" -LogonTypes Delegate -StartDate "1-May-2022 12:00" -EndDate "31-May-2022 17:00" -ShowDetails -Operations SendAs| Select
LogonUserDisplayName, LastAccessed
```

The Office 365 audit log ingests the mailbox audit records from Exchange alongside audit data generated by other applications. You can search these records using the audit log search in the Microsoft Purview Compliance portal or by running the *Search-UnifiedAuditLog* cmdlet. If you suspect you know who sent the message, you can refine the search by passing their email address as the value for the *UserIDs* parameter.

Otherwise, the search returns all audit records for all mailboxes where *SendAs* events occurred. See the Auditing chapter for an example of how to search the audit log for *SendAs* events.

## Accessing Shared Mailboxes

Once created, users can access the new shared mailbox through Outlook or OWA. If a user has *FullAccess* permission for a mailbox, they do not have to do anything specific as an auto-mapping process kicks in to inform Outlook to include the shared mailbox in the list of resources that it opens. For more information on opening and using shared mailboxes in Outlook, see [this support article](#).

You can override auto-mapping by including a shared mailbox in an Outlook profile. This is an explicit instruction to open the mailbox for every Outlook session and is the older method used to access a shared mailbox. In this case, you click **More Settings** when editing the profile, go to the Advanced tab and click **Add** to enter the names of the shared mailboxes that you want Outlook to open. You can use the SMTP address, alias, or display name to tell Exchange which shared mailbox to use.

If you use Outlook in cached Exchange mode, Outlook depends on the OAB to check unknown addresses (except when using SMTP addresses). The new shared mailbox will not be present in the OAB until after Exchange Online has updated the OAB files and distributed them to clients. This process can take 24 hours or more, so if you want to send messages as the shared mailbox in the interim, use the online Global Address List (which has the shared mailbox because the mailbox joins the address list once after its creation) to lookup the name that you enter in the "From:" field. If you do not do this, Exchange Online cannot deliver the messages because the address you enter does not include the necessary information to allow Exchange to check it against its address lists, and you will receive a non-delivery notification containing the following error:

**This message could not be sent. Try sending the message again later, or contact your network administrator. Error is [0x80070005-00000000-00000000].**

### Opening Shared Mailboxes in OWA

Although OWA does not use Autodiscover or have a profile, the client can also open and use shared mailboxes. To access a shared mailbox, expand the full set of folders in your mailbox and right-click **Folders** in the list of OWA resources, and select the "**Add shared folder**" option from the menu. Then input all or part of the name of the shared mailbox to search for it in email contacts and the directory (Figure 6-9). OWA validates that you have the correct permissions and if all is in order, will open the mailbox and display it in the folder list. This operation will not affect the set of resources shown in Outlook, just like auto-mapping does not influence the set of folders available to OWA. Both clients function independently of each other.

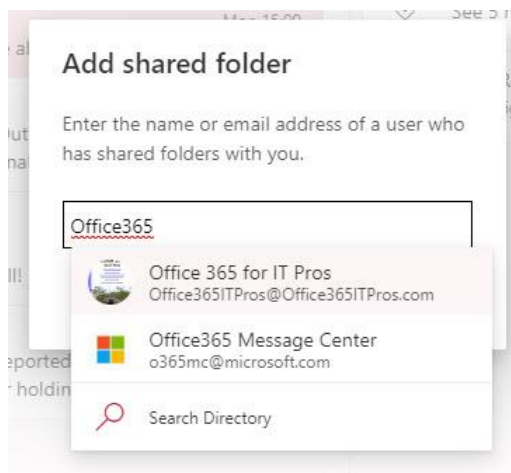


Figure 6-9: Choosing a shared mailbox to open with OWA

A good way of gaining access to a shared mailbox that exposes all the functionality available to the mailbox owner is to click the mailbox icon in the top right-hand corner of the menu bar and then **Open another**

**mailbox.** This method opens another session with the shared mailbox. Once a shared mailbox is open, you can access its contents in the same way as any other mailbox. You can send messages as if they came from the mailbox if you have the *SendAs* permission. You can select the mailbox to use by inputting the name of the shared mailbox into the "From:" field of the message so that the message comes from the shared mailbox rather than you.

The same technique works to open individual folders assigned to delegate users with folder-level permissions.

**Logging on to a shared mailbox:** The usual course of events is to grant access to a shared mailbox and have users open the shared mailbox as an added resource in either Outlook or OWA. However, at the time of writing, it is possible to change the password for the account for the shared mailbox and log on to the mailbox directly, without going anywhere near a user's account. Using this method is frowned upon for several reasons. First, it depends on password sharing, which is never a good thing, and second, according to Microsoft's Terms and Conditions, once you perform a direct logon to an account, you need to license that access because the mailbox is considered as no longer shared.

## Converting a Shared Mailbox to a Regular Mailbox

The EAC includes a method to convert a shared mailbox to a regular mailbox (and vice versa). Once the conversion is complete, you must assign a license and reset the password. The license is necessary to allow the newly regularized mailbox to have full functionality while the reset password allows the person who takes ownership of the mailbox to access it. Shared mailboxes have user accounts that no one logs into (unless they assign a license to the account), so they do not need a password. Instead, access to the shared mailbox is gained by giving permissions to the users who need to access the mailbox.

The EAC converts a shared mailbox by running the *Set-Mailbox -Type 'Regular'* command. The reverse is also possible, in which case the EAC runs the *Set-Mailbox -Type 'Shared'* command. In neither instance does the EAC do anything to assign a license or remove a license, nor does the EAC handle conversions of several mailboxes at one time. If you want to script the entire process, including the granting or removal of a license, you need to run some PowerShell commands. Here's how to convert a regular mailbox to a shared mailbox

```
[PS] C:\> Set-Mailbox -Identity "Shared@Office365ITPros.com" -Type Shared
```

If you want to convert a shared mailbox to a regular mailbox:

```
[PS] C:\> Set-Mailbox -Identity "Shared@domain.com" -Type Regular
```

After converting a mailbox to be a shared mailbox, you should remove any licenses it has. When you convert a shared mailbox to be a regular mailbox, you need to assign an Exchange Online license (either individually or part of a product like Office 365 E3). The commands to assign and remove licenses are covered in the PowerShell chapter.

## Redirecting Mail Sent to a Shared Mailbox

Companies often use shared mailboxes as the basis for customer communication. Things are a little more complicated when you have different teams of customer support agents, each using a shared mailbox, but you still want to have all communications recorded in a central mailbox for compliance or other purposes, such as integration with a CRM system that archives and categorizes customer interactions. You can achieve the goal by using a couple of mailbox properties.

Our scenario is as follows: our customer services team handles many different products. We want each product to have its own identity when email links are published on web pages and in other communications, so we create a set of shared mailboxes, one for each product, like this:

- Diapers.

- Cosmetics.
- Cars.

In addition, we have a central Customer Services shared mailbox.

Taking care of outbound emails is simple. The customer service agents simply make sure that they are sending messages as the Customer Services mailbox and that Outlook is configured to store the message in the shared mailbox rather than their personal mailboxes. The email address of the central mailbox will be stamped into the outbound message header so that any response from the customer will flow back into the central mailbox.

Inbound mail is redirected by setting two properties on each of the product mailboxes. The forwarding address is set to the central mailbox while the deliver and forward property is set to *\$True* to instruct Exchange Online to both forward a copy of the message to the central mailbox and keep a copy in the brand mailbox. The address selected for the forwarding address must belong to a mail-enabled object belonging to the tenant, including other mailboxes, public folders, mail contacts, shared mailboxes, and distribution lists.

A case can be argued not to forward a copy of the message to the brand mailbox (by setting the *DeliverToMailboxAndForward* flag to *\$False*) but this means that someone must process all the new mail arriving at the central mailbox and forward it on for attention. It is much better to have one copy captured centrally and another going direct to the people who must deal with the customer. In this scenario, the central copy acts as a customer contact record while the copy delivered to the "action" mailbox initiates the response.

**Sending an acknowledgment from a shared mailbox:** It is a common requirement to want to issue an acknowledgment for messages that arrive in a shared mailbox. For example, when people send an email to a customer services mailbox, it is good to have them receive a response to say that their email will be answered soon (or whatever is the right text). There is no obvious way to set up an auto-reply message for a shared mailbox, but it is easily done:

- Click your photo in the OWA menu bar.
- Select **Open another mailbox**.
- Enter the name of the shared mailbox and click OK.
- When OWA opens the shared mailbox, click **Options** (cogwheel) then all Outlook settings, and select **Mail**, then **Automatic replies**, and create the automatic reply.

Another way to connect is to include the email address of the shared mailbox in the URL for OWA. For example, <https://outlook.office.com/owa/bookcomments@office365ITPros.com/>. You can then go to Options as before. This procedure is the best approach if you want to use a text editor to format the content of an autoreply. Alternatively, if you are interested in basic text auto-replies, you can use the *Set-MailboxAutoReplyConfiguration* (explained earlier) to set internal and external auto-replies for shared mailboxes. The thing about auto-replies is that Exchange only sends a single response per correspondent. If you want responses for every email, you must use Outlook to create an Inbox rule that forces the server to create a reply to every message using a template.

## Using Mobile Devices with Shared and Delegated Mailboxes

Shared mailboxes are a very convenient method for teams to have access to information on which they need to work together. Given the highly mobile nature of today's workforce, it is natural to assume that you will be able to use shared mailboxes on smartphones or tablets, but this isn't the case. The Exchange ActiveSync protocol only supports access to personal mailboxes. Most mobile email clients bundled with Android and iOS phones use ActiveSync and are limited by the functions built into the ActiveSync protocol, which don't support shared mailboxes.

Current versions of Outlook Mobile (iOS and Android) use a different protocol called the Microsoft sync technology to connect clients to Exchange Online (and Outlook.com). The Microsoft sync technology includes

[support for shared mailboxes](#) and [delegate access to user mailboxes](#) along with other advanced features that will never be supported by ActiveSync. These features are available when the target (shared or user) mailbox and the mailbox of the delegate user both use Exchange Online.

To have delegated access to another person's mailbox with Outlook Mobile, a user must be assigned:

- Full access permission to the delegated mailbox. This means that the delegated user has unrestricted access to all folders in the mailbox.
- Either *Send As* or *Send On Behalf Of* permission to send messages for the delegated mailbox.

No other mobile client uses the Microsoft sync technology, which means that if you want to access shared mailboxes using a mobile client in a robust and supported manner, use Outlook Mobile.

## Mail Contacts and Mail Users

Exchange Online supports both mail contacts and mail users and includes the two object types in address lists (by default, All Contacts, the GAL, and the OAB), which makes the objects addressable by any Outlook client. Third-party clients can also access and use these objects with the appropriate code. The differences between the two objects are:

- A mail contact is a pointer to a user of an external email system.
- A mail user has an external email address but also has Microsoft 365 credentials and can sign in to access resources such as SharePoint Online or OneDrive for Business sites. Mail user objects are a relic of on-premises systems which only Exchange uses. Other applications such as Microsoft 365 Groups, Teams, SharePoint Online, Yammer, and Planner use Azure B2B Collaboration to enable guest access to their data.

Organizations often use mail contacts to provide users with a GAL (and OAB) entry that points to a known, valid email address. Often, mail contacts hold contact details for external business partners, such as a Public Relations agency. You can create a mail contact with the EAC or Microsoft 365 admin center by completing six fields:

- First Name: Optional.
- Last Name: Optional.
- Initials: Optional.
- Display Name: The name that appears in the GAL. It's a good idea to mark the object as being external to the company and to include a visual clue about which organization the contact belongs to.
- Alias: The alias must be unique, and it can't have any spaces.
- External email address: An SMTP address that is external to the tenant.

Remember that in a hybrid environment, management of mail contacts and mail users is on-premises with changes synchronized to Exchange Online.

## Using PowerShell to Create Mail Contacts

The *New-MailContact* and *Set-MailContact* cmdlets are available to create and update mail contacts. In this example we first create a new mail contact, setting their preferred mail format to be HTML, and then run the *Set-MailContact* cmdlet to enforce moderation and to inform senders that moderation applies to any message sent to this address. Note that you cannot create a mail contact with an SMTP address already used by another mail-enabled object, including guest accounts.

```
[PS] C:\> New-MailContact -ExternalEmailAddress "Danny.Flowers@contoso.com" -LastName "Flowers"
-DisplayName "Danny Flowers (Contoso)" -FirstName "Danny" -Name "Danny Flowers" -MessageBodyFormat
HTML -Alias Danny.Flowers
```

```
[PS] C:\> Set-MailContact -Identity Danny.Flowers -ModeratedBy "Kim Akers" -ModerationEnabled $True
-MailTip "Message will be moderated before dispatch"
```

After creating a mail contact in Exchange Online, a synchronization process creates a contact object in Azure AD. The *Get-MgContact* cmdlet retrieves details about Azure AD contacts.

```
[PS] C:\> Get-MgContact -OrgContactId (Get-MailContact -Identity
AlexW@o365maestro.onmicrosoft.com).ExternalDirectoryObjectId
```

You don't have to worry about the contact objects held in Azure AD. They exist for internal purposes and are not used for mail routing.

Some of the mail contact settings supported by Exchange on-premises, such as the ability to set the maximum message size, are unavailable in Exchange Online. Others, such as phone numbers and organizational information, are updatable with the *Set-Contact* cmdlet. In this example, we use PowerShell to read a simple CSV containing a set of records and create mail contacts for each object found. You can see how the *New-MailContact* cmdlet first creates the new object before the code uses *Set-Contact* to update some extended properties.

```
[PS] C:\> $InputContacts = import-csv c:\temp\inputcontacts.csv
Write-Host $InputContacts.Count "contacts found"
ForEach ($Contact in $InputContacts) {
    $Alias = $Contact.First + "." + $Contact.Last
    # Real simple code to make sure that we have an alias
    If ($Alias -eq $Null) { $Alias = $Contact.Name.Split(" ")[0] + "." + $Contact.Name.Split(" ")[1] }
    If ((Get-Recipient -Identity $Contact.EmailAddress -ErrorAction SilentlyContinue) -eq $Null) {
        # Recipient is not known, so we can add them
        Write-Host "Adding contact" $Contact.EmailAddress
        New-MailContact -Name $Contact.Name -ExternalEmailAddress $Contact.EmailAddress -Alias $Alias
        -FirstName $Contact.First -LastName $Contact.Last
        # Update country and phone numbers
        Set-Contact -Identity $Alias -MobilePhone $Contact.MobilePhone -Phone $Contact.WorkPhone -
        CountryOrRegion $Contact.Country -Company $Contact.Company }
}
```

You can even add a photo to mail contacts. Outlook displays these photos when users view contacts in the GAL. Here's how to use the *Import-RecipientDataProperty* cmdlet to add a photo. For best results, size the JPG file at 150 x 150 pixels (or smaller).

```
[PS] C:\> Import-RecipientDataProperty -Identity "John Contoso" -FileData ([Byte[]](Get-Content -
Path "c:\temp\DefaultGuestPicture.jpg" -Encoding Byte -ReadCount 0)) -Picture
```

## Creating Mail Users

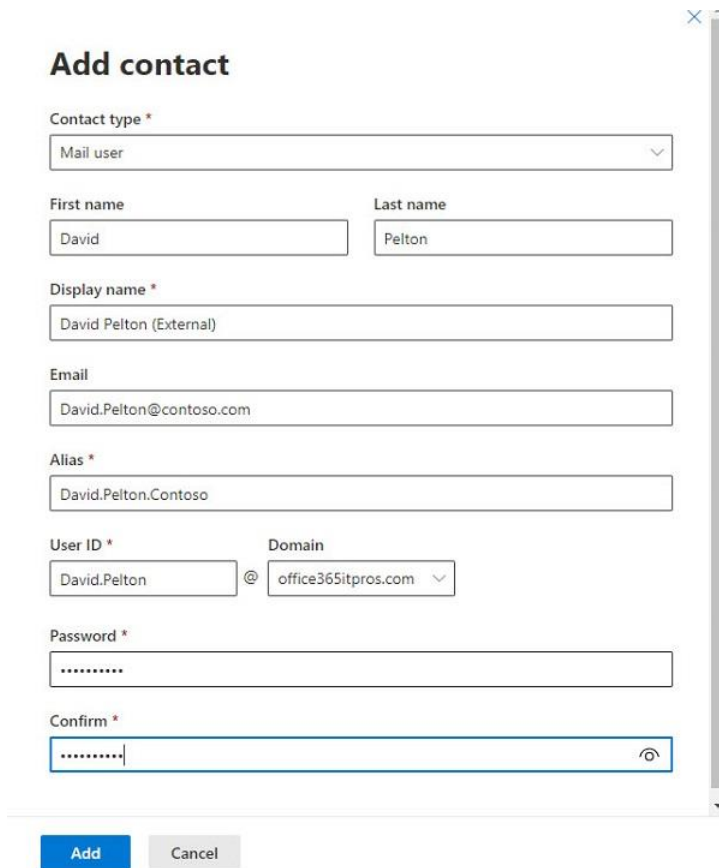
Although you can create the new mail user object by running the *New-MailUser* cmdlet, the easiest way to create a mail user is with EAC (you can't create mail users through the Microsoft 365 admin center). As shown in Figure 6-10, you must give a valid SMTP email address to point to the user's mailbox on an external email system.

Behind the scenes, Exchange Online creates a user object in Azure AD which appears in the Microsoft 365 admin center. However, the new user object must be assigned a license before the account can access any services.

Mail user objects can be manipulated with the *Set-MailUser* cmdlet. For example:

```
[PS] C:\> Set-MailUser -Identity "David Pelton" -CustomAttribute10 "External Recipient"
```

Azure AD creates mail user objects when apps like Teams and Planner add guest accounts (you'll see a record for the *New-SyncMailUser* cmdlet in the audit log). Exchange links the mail user objects to the guest accounts and will remove the mail user object on the deletion of the associated guest account.



**Add contact**

Contact type \*  
Mail user

First name: David      Last name: Pelton

Display name \*  
David Pelton (External)

Email  
David.Pelton@contoso.com

Alias \*  
David.Pelton.Contoso

User ID \*      Domain  
David.Pelton      @ office365itpros.com

Password \*  
.....

Confirm \*  
.....

**Add**      Cancel

Figure 6-10: Creating a new mail user

## Blocking Basic Authentication

Basic authentication means what it says - a basic mechanism to authenticate a connection to a service. Basic authentication is simple to use and simple to abuse, which is why attackers try to exploit its simplicity in exploits like [password spraying attacks](#). Exchange Online supports many different mailbox connection protocols from ActiveSync to POP3 to IMAP4 to MAPI. Supporting now-obsolete email connection protocols reflects the fact that Exchange has been around since 1996. On a positive note, supporting multiple protocols allows people to use their client of choice to connect to their mailbox. Unfortunately, the profusion of connection protocols creates a problem because each protocol must be secured to stop attackers from probing for protocol weaknesses that they can use to penetrate a tenant. To address the issue, Microsoft will disable basic authentication for connections to Exchange Online for a set of email protocols (see list below). The protocols are on the list because experience shows that they are vulnerable to attack. Microsoft's plan of record is to:

- Disable basic authentication for the target set in new tenants. This already happens.
- Disable basic authentication for protocols in the target set when telemetry shows that tenants do not use basic authentication with these protocols. Tenants get a 30-day notice in the message center in the Microsoft 365 admin center when Microsoft plans to block basic authentication for a protocol. This process is now happening.
- Begin to selectively block the connection protocols for randomly selected tenants for 12-36 hours starting in early 2022. The idea is to allow tenants to judge what effect a permanent block for these protocols will have.

- Block basic authentication for the listed protocols (except SMTP AUTH) [on October 1, 2022](#). This is a permanent block for Exchange Online. Microsoft is likely to remove the exemption for SMTP AUTH in the future.

To check if Microsoft has blocked some or all protocols, run the *Get-OrganizationConfig* cmdlet and examine the *BasicAuthBlockedApps* property. The value returned will be in a range from 0 (zero – no protocols are blocked) to 255 (all protocols are blocked):

```
[PS] C:\> Get-OrganizationConfig | Select BasicAuthBlockedApps
```

```
BasicAuthBlockedApps
```

```
-----  
255
```

The value of *BasicAuthBlockedApps* is a bitmask composed of values for each protocol. Zero (0) means that basic authentication is not blocked for any protocols. The values for each of the targeted protocols are:

- Exchange ActiveSync (EAS): 1
- Exchange Web Services (EWS): 2
- POP3: 4
- IMAP4: 8
- Remote PowerShell: 16
- MAPI over RPC (Outlook Anywhere): 32
- Offline Address Book (OAB): 64
- RPC: 128

The driving force behind the removal of basic authentication for email connections is to encourage organizations to deploy clients and applications which support modern authentication, such as any modern Outlook client. PowerShell connections can use MFA with the Exchange Management PowerShell module or move to [PowerShell within Azure Cloud Shell](#). In the interim, you can prepare for the deprecation of basic authentication by using [authentication policies](#) to [control the connectivity protocols available to mailboxes](#). Although Microsoft will block basic authentication for many connectivity protocols, authentication policies will remain active and useful afterward because they control basic authentication for other protocols that Microsoft won't block.

Azure AD conditional access policies are another way to block connection attempts based on protocol and should be used whenever organizations need to exert precise control over inbound connections to a tenant. In comparison to the granular control that can be exerted by conditional access policies, authentication policies are a basic and blunt weapon. However, authentication policies offer one great advantage: conditional access policies are processed post-authentication. If you deploy an authentication policy to block basic authentication, most attempted attacks to compromise accounts will be blocked at the first hurdle. This is the reason why tenants should create and use an authentication policy to block basic authentication for as many protocols as possible.

## Creating an Authentication Policy

The easiest way to create and manage the default authentication policy for a tenant is through the Modern Authentication section in Org Settings in the Microsoft 365 admin center. If an authentication policy doesn't already exist, it is created and populated with the settings you choose. Thereafter, if you edit the settings (for instance, to block basic authentication for a protocol), the default authentication policy is updated. Before you can block basic authentication, you must enable modern authentication for the tenant and be sure that users have clients that support modern authentication, like Outlook, as older clients will not be able to connect if you disable basic authentication for protocols through a policy.



You can have multiple authentication policies in a tenant, each of which allows basic authentication for a different set of protocols. Non-default authentication policies can only be managed through PowerShell. To create a new policy, run the [New-AuthenticationPolicy](#) cmdlet. For example:

```
[PS] C:\> New-AuthenticationPolicy -Name "No Basic Auth"
```

```
RunspaceId : fd030e40-053a-404c-90f9-3cf9f2c2dcef
AllowBasicAuthActiveSync : False
AllowBasicAuthAutodiscover : False
AllowBasicAuthImap : False
AllowBasicAuthLogExport : True
AllowBasicAuthMapi : False
AllowBasicAuthOfflineAddressBook : False
AllowBasicAuthOutlookService : False
AllowBasicAuthPop : False
AllowBasicAuthReportingWebServices : False
AllowBasicAuthRest : False
AllowBasicAuthRpc : False
AllowBasicAuthSmtplib : False
AllowBasicAuthWebServices : False
AllowBasicAuthPowerShell : False
AdminDisplayName :
ExchangeVersion : 0.20 (15.0.0.0)
```

You can see that the only protocol enabled for basic authentication is Log Export, which is probably not going to be used by an attacker.

If you want to change a policy setting to allow basic authentication for a protocol, run the [Set-AuthenticationPolicy](#) cmdlet. For example:

```
[PS] C:\> Set-AuthenticationPolicy -Identity "No Basic Auth" -AllowBasicAuthPop:$True
```

After you've created the authentication policies you need, you assign the policies to user accounts to control basic authentication connections for those accounts. If a user account is not assigned an explicit authentication policy, Exchange Online applies the organization's default policy.

If you decide to apply a non-default authentication policy, the best approach is to run the *Set-User* cmdlet to assign the policy to target accounts. At the same time, you should reset the refresh token for the account to the current date and time. This will force Exchange to request clients which use basic authentication for connections to reauthenticate using modern authentication. This code uses the *Get-Recipient* cmdlet to find all current mailboxes:

```
[PS] C:\> Get-ExoMailbox -RecipientTypeDetails UserMailbox -ResultSize Unlimited | Set-User -AuthenticationPolicy "No Basic Auth" -STSRefreshTokensValidFrom $([System.DateTime]::UtcNow)
```

Remember that if you want the non-default authentication policy to apply to new mailboxes, you must assign the policy to new accounts following their creation.

## Checking Policies Are Applied to Accounts

To check that the policy is assigned to accounts, run the *Get-User* command to fetch the set of user mailboxes and report the policy assigned to the account. You should see that each account is assigned the desired authentication policy and the account's refresh token is reset to the time when you ran the *Set-User* command.

```
[PS] C:\> Get-User -RecipientTypeDetails UserMailbox -ResultSize Unlimited | Format-Table DisplayName, AuthenticationPolicy, Sts*
```

DisplayName	AuthenticationPolicy	StsRefreshTokensValidFrom
Deirdre Smith	No Basic Auth	21 Oct 2018 14:30:42
Tony Redmond	No Basic Auth	21 Oct 2018 14:31:06

```
TempAdminAC No Basic Auth 21 Oct 2018 14:31:11
```

Authentication policies don't come with any sophisticated methods for tracking when they are applied or how effective they are. Apart from finding out whether people use obsolete clients to connect to mailboxes, the biggest issue you might face is that disabling basic authentication for PowerShell forces accounts to use multi-factor authentication when they connect to Exchange Online. This isn't a problem with the Exchange Online Management module because it includes native support for modern authentication.

Limiting basic authentication for connections using an authentication policy only affects Exchange Online and doesn't affect other Microsoft 365 workloads.

## Defining a Default Authentication Policy

Exchange Online assigns the default authentication policy for the tenant to new accounts. Unless you define a default authentication policy in the organization configuration, the value assigned is `$Null`, meaning that no policy is assigned. To change this, run the `Set-OrganizationConfig` cmdlet and set a default policy. In this example, we set the default authentication policy to be "No Basic Auth":

```
[PS] C:\> Set-OrganizationConfig -DefaultAuthenticationPolicy "No Basic Auth"
```

Exchange Online validates that the authentication policy exists and won't update the organization configuration if it can't find the specified policy. You can check the value of the current default authentication policy with the `Get-OrganizationConfig` cmdlet:

```
[PS] C:\> Get-OrganizationConfig | Format-List DefaultAuthenticationPolicy
```

```
DefaultAuthenticationPolicy : No Basic Auth
```

## SMTP Authenticated Submissions

Many SMTP-enabled devices like printers and scanners submit messages to Exchange Online using [SMTP AUTH client submission](#) (SMTP AUTH) by logging in with a simple username and password. These connections do not support modern authentication methods such as multi-factor authentication or certificate-based authentication, which means that any account set up to be used for SMTP AUTH might be a target for attacks that depend on finding a username and password combination to penetrate accounts.

If you don't need to use SMTP AUTH, you should disable these connections at either a tenant or mailbox level. To disable at the tenant level, run the command to update the Exchange transport configuration:

```
[PS] C:\> Set-TransportConfig -SmtClientAuthenticationDisabled $True
```

If you leave SMTP AUTH enabled in the transport configuration, you can disable the feature selectively at the mailbox level. To find out what mailboxes are enabled, run the command:

```
[PS] C:\> Get-CasMailbox | ? {$_.SmtClientAuthenticationDisabled -eq $Null -or  
$_.SmtClientAuthenticationDisabled -eq $False } | Format-Table DisplayName
```

To disable the feature, change the command to:

```
[PS] C:\> Get-CasMailbox | ? {$_.SmtClientAuthenticationDisabled -eq $Null -or  
$_.SmtClientAuthenticationDisabled -eq $False } | Set-CasMailbox -SmtClientAuthenticationDisabled  
$True
```

When a mailbox is blocked from SMTP AUTH, it can't be used to submit messages to Exchange Online by running the PowerShell `Send-MailMessage` cmdlet (as explained in the PowerShell chapter, PowerShell scripts can send messages using the Outlook Graph API instead). Many PowerShell scripts use `Send-MailMessage` to send emails for different purposes from welcoming new users to a tenant to reporting the results of a

background job. Until Microsoft delivers a more secure mechanism for email submission, you should minimize the number of accounts used by scripts and make sure that these accounts remain enabled for SMTP AUTH.

As noted earlier, Microsoft will disable basic authentication for email connectivity protocols for all tenants on October 1, 2022. Reflecting the widespread use of SMTP AUTH in applications and multi-function devices and prioritizing the elimination of basic authentication like POP3 and IMAP4 that attackers target most frequently, Microsoft has made an exception for SMTP AUTH. Here's the situation:

- Microsoft will not disable SMTP AUTH if the protocol is used by a tenant.
- If SMTP AUTH is not in active use, Microsoft will disable SMTP AUTH (using the transport configuration setting) along with other unused connection protocols in tenants starting in September 2021.
- If a tenant discovers that they need to reenable SMTP AUTH, they can do so by running *Set-CASMailbox* to enable the protocol for selected mailboxes.
- SMTP AUTH will remain available to tenants who choose to use it after the general block on basic authentication for email connection protocols effective October 1, 2022.

Eventually, Microsoft will disable SMTP AUTH across Exchange Online. To prepare for this eventually, organizations can replace SMTP AUTH with Graph API calls (see the example in the PowerShell chapter) or use [OAuth 2.0 for SMTP AUTH connections](#). Developers can use these capabilities to upgrade applications that use SMTP AUTH with basic authentication to use modern authentication instead.

# Chapter 7: Mail Flow

## Gareth Gudger

This chapter focuses on mail flow and managing the features available in Exchange Online, Exchange Online Protection, and Microsoft Defender for Office 365 to ensure that messages transit securely and reliably from senders to recipients. Managing mail flow covers much more than ensuring the successful delivery of messages. It is also about keeping you safe from threats such as malware, targeted phishing attacks, spoofing, spam, and (accidental) data loss.

Administrators perform the tasks discussed here through the Exchange admin center (EAC), the Microsoft 365 Defender portal, and PowerShell cmdlets. Many examples use PowerShell because it is often easier and quicker to update a setting with PowerShell and the commands are much less prone to change than any of the GUI-based tools. In any case, we will begin with a discussion about configuring mail flow.

## Configuring Mail Flow

Before Exchange Online accepts messages for your domains, you must add and confirm your email domains in the tenant. The process for adding domains to a tenant is [covered in Microsoft's documentation](#). After the domain(s) are added to the tenant, you must decide when to configure the domain's MX records to point to Exchange Online Protection (EOP) so that sending organizations know where to route email.

The information to set up the necessary DNS records for your domain(s) is automatically returned in the wizard while adding domain(s) to the tenant. For mail flow purposes, only the MX and SPF records are initially relevant. As soon as you configure extra protection features, like DKIM (explained later), additional DNS records might be required.

**Real-world:** If you do not add your custom domains to your tenant, Exchange Online only accepts messages for the tenant's default service domain, which is in the form of *tenantname.onmicrosoft.com*. This is not very useful if you want your tenant to process email for all your domains, but it allows you to set up a test tenant to verify certain features without registering a new domain.

## Mail Exchanger (MX)

When you configure your domain to use EOP, the MX record for that domain points to a hostname in the form of *domain-tld.mail.protection.outlook.com* rather than an IP address. The following example illustrates how to resolve an MX record for a domain configured in EOP:

```
[PS] C:\> Resolve-DnsName office365itpros.com -Type MX
```

Name	Type	TTL	Section	NameExchange	Preference
office365itpros.com	MX	3572	Answer	office365itpros-com.mail.protection.outlook.com	0

When a sending server queries the MX record for the domain *office365itpros.com*, the A record *office365itpros-com.mail.protection.outlook.com* is returned. Next, the sending server will try to translate that hostname into an IP address, for which it needs to perform an extra DNS query. Before responding to the query, Microsoft's servers perform an internal lookup to check the region the tenant belongs to. Once the region is known, the service responds with the IP addresses of the EOP systems within the same region of the tenant.

**Real-world:** When you add a new domain to your tenant, you are asked to configure various DNS records, including the MX record for your domain. In a greenfield deployment, it is probably okay to go ahead and configure the MX record to point to EOP. However, reconfiguring the MX record for an existing domain will break your mail flow.

Switching from your current solution to EOP should be done at the right time. When strictly depends on your migration approach. Typically, you reconfigure your MX record once your pilot for Exchange Online is complete, or you wait until most of your users have been migrated to Exchange Online. The time for changing your MX record is when you have ensured your configuration is ready to accept messages for on-premises recipients from EOP. For more information, see [this Microsoft article](#).

## Sharing SMTP Namespaces

Many organizations use one or more unique SMTP domain names (for example, office365itpros.com). Those organizations sometimes need to share a common SMTP domain name across multiple environments to present a single domain name to the outside world.

Sharing an SMTP namespace with another tenant is not possible because you cannot register a domain name in more than one tenant. Because of this, it is technically impossible to share a common email domain without an elaborate scheme. Because address rewriting is not possible with Exchange Online, we will not discuss sharing an address namespace between multiple tenants. However, we will discuss how you can share a namespace across a tenant and one or more on-premises Exchange organizations.

Sharing a domain name across an on-premises organization and Microsoft 365 is easier than sharing a namespace across multiple tenants. There are several ways you can do this:

- **Use a third-party or custom broker service** that "catches" all incoming emails and routes those messages to the backend system hosting the recipient mailboxes. For the broker service to determine what host to direct a message to, you must synchronize the recipient information from all connected systems to a specific location where the broker service can access it; sometimes, this is a directory of the broker service itself. An example of such a service capable of handling routing for various backend systems is Mimecast. The downside of this approach is that it introduces another component in the mail flow process, increasing your solution's complexity and cost.
- **Reconfigure domains in Exchange from Authoritative to Internal Relay.** By default, verified domains in Office 365 are added as authoritative in Exchange Online. This means that Exchange Online will accept and handle mail flow for the domain if a recipient exists in the tenant directory. If a recipient's email address does not exist, Exchange Online generates a Non-Delivery Report (NDR). When you reconfigure a domain as an Internal Relay domain, Exchange Online first tries to deliver the message locally. If no matching recipient is found, the message can be forwarded to another mail system through a connector. If no connector is found to match the address, Exchange performs an MX lookup to decide how to route the message.
- **Setup routing domains** and configure each environment to forward messages to specific recipients using the configured forwarding address (based on the specific routing domain). This is the approach used in a hybrid deployment.

It can quickly become tricky to manage this solution with multiple environments: the first environment (Exchange Online) must forward messages to the second environment. The second environment forwards messages to the third, and so on. The last environment to receive messages must then reject unknown recipients. If you do not reject unknown recipients and forward the emails back to any of the previous environments, you will introduce a mail loop. Mail loops occur when messages go to recipients that do not exist in any environment.

## Sender Policy Framework (SPF)

The Sender Policy Framework (SPF) record is part of the [SPF validation system](#), created to protect from the receipt of spoofed messages. The receiving server does this by evaluating the purported sender domain's SPF record. The receiving email system will use the information on the SPF record to determine whether the email was received from a messaging system authorized to send emails from that domain.

When a message is received from an external system, the receiving system examines the IP address of the prior server in the mail flow chain. It then looks for an SPF record for the sender's domain. Then, one of two things can happen:

1. If an SPF record is found, the IP address is verified against the values specified in the SPF record. Depending on how the record is configured, this will either generate a neutral result or a "soft" or "hard failure."
2. If no record is found, the SPF lookup is considered a "soft failure."

The suggested SPF record for Exchange Online looks like the example below. The record specifies that only the servers specified in the SPF record for `spf.protection.outlook.com` can send messages on behalf of this domain:

```
v=spf1 include:spf.protection.outlook.com -all
```

If another server not listed on the SPF record for `spf.protection.outlook.com` tries to send messages for this domain, a hard failure should be generated because the `"-all"` parameter is specified.

The syntax of the SPF record controls how the receiving system should treat a failure:

- **-all**. The minus sign shows that any SPF failure should be considered a hard fail. Whether the hard fail results in the message being rejected depends on how it is configured in the receiving server's anti-spam solution.
- **~all**. The tilde shows that any SPF failure should be treated as a soft fail; the message will be accepted unless the receiving system rejects soft fails.
- **?all**. The question mark indicates that the domain owner is neutral towards the use of SPF records. A failure will not generate a soft or hard fail. Instead, it will generate a neutral result. The message should be accepted regardless of the result.

If EOP is the only system that will send messages for your domain(s), the suggested SPF record will suffice. However, if other systems also send emails for your domain(s), you must manually tweak the suggested SPF record to include those systems. For instance, if a marketing platform sends an email from an address belonging to one of your custom domains, the SPF record should be modified to include that system. In the example below, we add the required entry for the marketing platform to our SPF record to `include:servers.office365itpros.org`.

```
v=spf1 include:spf.protection.outlook.com include:servers.office365itpros.org -all
```

In the scenario where an on-premises organization uses EOP to protect outbound email, you should also include the public IP address for the on-premises Exchange server in the SPF record. In this example, `1.2.3.4` is the public IP of the on-premises Exchange Server(s) that are sending emails to EOP:

```
v=spf1 include:spf.protection.outlook.com ip4:1.2.3.4 -all
```

**Note:** Like the examples shown above, some SPF records use the `include` statement to point to another SPF record. This tells the recipient to trust everything in your SPF record and trust everything in the SPF record identified by the `include` statement. The challenge is that this SPF record might also contain an `include` statement(s), which in turn point to additional SPF record(s), which might also contain another

layer of *include* statements and so on. This has the potential to cascade into many SPF records. Every time a hostname or pointer to another record is added to an SPF record (such as the *include* statement), a new DNS lookup is triggered to try and resolve the hostname to an IP address. The total number of lookups must not exceed ten, as this will cause the SPF check to fail. Therefore, when you add an *include* statement, it is prudent to review what the target SPF record permits on your domain's behalf. It is also worth noting that the SPF record behind the *include* statement can change at any time. A single SPF record is also limited to 255 characters.

Creating and updating SPF records is not a trivial task, mainly because most email administrators do not need to do it very often. However, various tools on the internet can help. For instance, [this site](#) helps generate the SPF record to configure in your external DNS, and [this site](#) can help you check you do not exceed the maximum of ten DNS lookups.

**Real-world:** Although not all organizations check for a sender's SPF records, it is important to configure the records for the domain to reduce the risk of encountering email delivery issues elsewhere. Organizations do not always know what other mail systems send messages on behalf of their domain. For instance, a company's marketing department might periodically send messages through a third-party solution. In such cases, you should change your domain's SPF record to include the third-party solution.

If you cannot obtain that information, the best practice is to use a different domain for the marketing department, such as `marketing.office365itpros.com`. Make sure that the namespace uses a different public IP address and SPF record than your corporate main SMTP namespace. That way, if the domain `marketing.office365itpros.com` (or the underlying public IP) gets blocked, it will not affect email delivery for the main corporate namespace.

It is recommended to add SPF records for all your internet domains, even those that do not send mail. Bad actors will use your domain regardless of whether you do. For domains that do not send mail, you can protect them with an SPF record that hard fails all senders. For example, `v=spf1 -all`.

While SPF records have long been a critical component of email security, they are no longer the only DNS records you need to implement to maintain a modern email system.

## Domain Keys Identified Mail (DKIM)

Domain Keys Identified Mail (DKIM) allows an organization to add a cryptographic signature to outgoing emails. This signature, which is based on the message's headers and body, is added in a DKIM-specific header called *DKIM-Signature*. The purpose of DKIM-signing a message is to allow a remote organization to verify whether the message originated from a platform authorized to send messages from a given domain. A DKIM signature does not encrypt the contents of the message body or any attachments. DKIM should be a part of the baseline configuration for your email message flow.

Figure 7-1 illustrates how DKIM works: The sending server signs the message with a DKIM signature (1). The DKIM signature contains both a selector name and domain name (2), which the receiving server uses to perform a DNS lookup against the sender's DNS zone (3). The returned DNS record contains the matching public key (4), which the receiving server uses to match the DKIM header in the email. The receiving server adds the DKIM result (pass or fail) to the message header (5). The message is then queued for delivery (6). DKIM not only confirms the origin of the email but also that no one has tampered with the message during transit.

DKIM supports multiple keys per domain. This is useful in situations where an organization operates multiple email platforms, wants to delegate control of DKIM signatures to individual departments within the organization, uses a third-party service to send messages on behalf of one of the organization's domains, or when administrators want to update one of the keys without affecting current mail flow. To support multiple

keys, a “selector” is prepended to the domain name and recorded as part of the DKIM-Signature header. For instance: “selector.\_domainkey.domain.com.” With the selectors in place, each system (or third-party solution) can use a different selector and still generate valid DKIM signatures.

By default, Exchange Online verifies incoming DKIM signatures and signs outbound messages – even when you have not configured DKIM records. Exchange Online uses the “selector1” and “selector2” selectors.

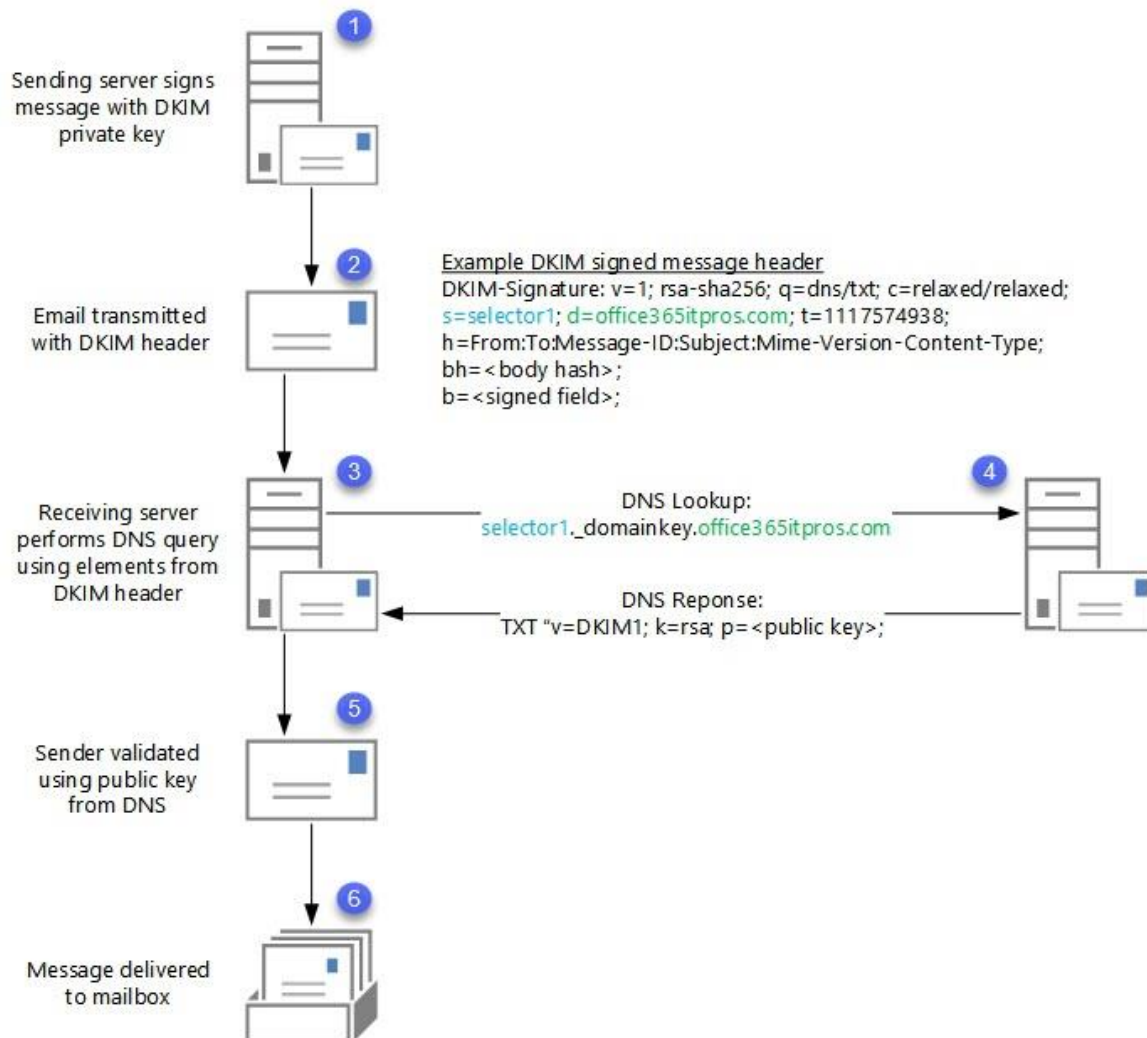


Figure 7-1: How EOP performs a DKIM lookup on incoming mail

## DKIM Verification

After an inbound message is processed, the result of the DKIM verification process is written to the *Authentication-Results* header of the message. Depending on the outcome of the test, the header will show either *dkim=pass*, *dkim=none*, or *dkim=fail*, followed by a more human-readable explanation of the result in parentheses. For instance, the following *Authentication-Results* header reveals that the sender did not sign the outgoing message:

```
authentication-results: office365itpros.com; dkim=none (message not signed) header.d=none;
```

However, if the message was signed and the DKIM signature is successfully verified, the header will look similar to this:

```
authentication-results: office365itpros.com; dkim=pass (signature was verified)  
header.d=office365itpros.com;
```



## Configuring DKIM Signing

Typically, several steps are necessary to enable the signing of outbound messages. However, as described in the section below, you do not necessarily have to perform these steps if you only use Exchange Online as a mail service. This is because Microsoft enables DKIM signing by default and uses a clever workaround to ensure that signatures are valid.

If you manually configure DKIM for your domain to associate signatures with vanity domains assigned to the tenant and not just the default service domain (*tenantname.onmicrosoft.com*), you must publish two DNS records for the vanity domain name for which you want to enable DKIM signing. Each record points to a specific target based on a combination of your domain and tenant name. For example, imagine you have a vanity domain called "office365itpros.com" and a tenant named "mycompany.onmicrosoft.com." The value of the CNAME records would then be the following:

```
selector1._domainkey.office365itpros.com CNAME
selector1-office365itpros-com._domainkey.mycompany.onmicrosoft.com
```

and

```
selector2._domainkey.office365itpros.com CNAME
selector2-office365itpros-com._domainkey.mycompany.onmicrosoft.com
```

The target of the CNAME record always starts with either *selector1* or *selector2*, followed by a hyphen and the domain MX key. The domain MX key is the first part of the MX record for your domain from the domains information page. The last part of the CNAME target is your tenant domain name. This is the default domain automatically created when you sign up for a Microsoft 365 tenant and is in the form of *<name>.onmicrosoft.com*. After publishing the DNS records, the easiest way to enable DKIM signing is through the Defender portal ([security.microsoft.com](https://security.microsoft.com)) under **Threat management > Policy > DKIM**. First, select the domain you want to enable DKIM and toggle the Sign messages for this domain with DKIM signatures slider to Enabled for the domain name you wish to enable DKIM.

Alternatively, you can enable DKIM signing with PowerShell:

```
[PS] C:\> New-DkimSigningConfig -DomainName office365itpros.com -Enabled $true
```

Domain	Enabled
-----	-----
office365itpros.com	True

If the CNAME records are unavailable or if the records were created incorrectly, the above command will emit the following warning:

```
WARNING: Config is created but cannot be enabled since CNAME records are not published. Please
enable this policy using Set-DkimSigningConfig once CNAME records are published.
```

If you have recently updated your DNS records and set the CNAME record, you might need to wait for DNS caches to update. After sufficient time has passed, you can attempt to enable the configuration again using the following command:

```
[PS] C:\> Set-DkimSigningConfig -Identity office365itpros.com -Enabled $true
```

If you continue receiving the warning message, try comparing the CNAME records you created with the values EOP expects. To retrieve the correct selector CNAME values, run the following command:

```
[PS] C:\> Get-DkimSigningConfig | fl *CNAME*
```

```
Selector1CNAME: selector1-office365itpros-com._domainkey.mycompany.onmicrosoft.com
Selector2CNAME: selector2-office365itpros-com._domainkey.mycompany.onmicrosoft.com
```

It can take up to one hour after you configure DKIM before the changes have successfully replicated to all EOP servers. Once the new DKIM settings are replicated, all messages are automatically signed as the vanity domain rather than the tenant domain. For example, the following output is from a test message sent to an external recipient when the source tenant has enabled DKIM signing. Note the *dkim=pass* reference:

```
Authentication-Results: spf=pass (sender IP is 1.2.3.4) smtp.mailfrom=office365itpros.com; dkim=pass (signature was verified) header.d=office365itpros.com; dmarc=bestguesspass action=none header.from=office365itpros.com; compauth=pass reason=109
```

```
DKIM-Signature: v=1; a=rsa-sha256; c=relaxed/relaxed; d=office365itpros.com; s=selector1; h=From:To:Date:Subject:Message-ID:Content-Type:MIME-Version; bh=WeG072W2hUk4jdiG4YKREVB5fA4PaMBqox2yZKnVYI=;
```

```
b=mSTzjicUsZiAJizFSwmW0EpaNYBjUFXDoFQNaAgFwpbDhCcun5P5W+MYPxRRx5//KANWT8hsv559VEU/E6PCENAQbRUnJ/CgbiQRcoh6+X5+vtisdXenC7gnc0TOXWDAK1sqh46mzW3Ls1vGVH4xdzFV6i7spZR0Mif2cd/qtU=
```

Along with the signature (shown prepended by *b=*), the DKIM header also includes additional information, represented as tags used during the DKIM verification process. For instance, it indicates which selector was used (*s=selector1*) and a list of headers used during the signing process (*h=*). Trying to decipher the headers can be extremely tedious and is usually unnecessary. To learn more about the finer details of the DKIM header and the various tags it can have, read [RFC6367](https://tools.ietf.org/html/rfc6367).

## Default DKIM Signing

DKIM is a vital tool to fight spoofed messages. As such, Microsoft automatically signs all outbound email traffic, even when the customer has not set up DKIM or configured any specific DKIM DNS records. You might wonder how Microsoft can do that, especially because they do not control the DNS zone for your domains? However, Microsoft controls one DNS zone directly linked to your Office365 tenant: the default tenant service domain *<tenantname>.onmicrosoft.com*. This allows them to create a DKIM signature based on the tenant domain and still correctly represent your organization. The following example shows a *regular* DKIM signature as you would expect when you have configured DKIM and the appropriate DNS records:

```
From: sender@office365itpros.com
DKIM-Signature: v=1; a=rsa-sha256; c=relaxed/relaxed; d=office365itpros.com; s=selector1; h=From:To:Date:Subject:Message-ID:Content-Type:MIME-Version; bh=<body hash>; b=<signed field>;
```

Now, take a look at the following example of a DKIM signature when you have not explicitly configured DKIM. Notice the selector (*s=*) and domain (*d=*) fields:

```
From: sender@office365itpros.com
DKIM-Signature: v=1; a=rsa-sha256; c=relaxed/relaxed; d=o365itpros.onmicrosoft.com; s=selector1-o365itpros-com; h=From:To:Date:Subject:Message-ID:Content-Type:MIME-Version; bh=<body hash>; b=<signed field>;
```

Based on the above information, you can reconstruct the actual DNS records being referenced to the following value: **selector1-o365itpros-com.\_domainkey.o365itpros.onmicrosoft.com**.

Using the *Resolve-DNSName* cmdlet, you can then verify that the record exists:

```
[PS] C:\> Resolve-DnsName selector1-microsoft-com._domainkey.microsoft.onmicrosoft.com -Type TXT
v=DKIM1; k=rsa;
p=MIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQCKHq3ztGI1R8a1D+7oZiaG5mTUttFdF01pKHRBZCPFG4sugV1EfF5F6Jpwb
JDzZmyI1qYfTgUkmY0vbHsoYvW7rddLKVTh+vE1SZ5P9coIHrw759hXbpPDSQ9JNP8aN+Bfng6YMEwnOGA+PL+ZpyvswcB0jz9M6
yMvowOxChv5QIDAQAB; n=1024,1435867504,1
```

The key returned by the record is used by the target system to verify the DKIM signature and is no different from the record otherwise used when manually configuring DKIM. After all, when setting up explicit DKIM signing for Exchange Online, you do not create a TXT record with the key; you create a CNAME record pointing to the same TXT record in Microsoft's DNS zone for your tenant illustrated in the example.

### Automatic DKIM Key Rotation

Microsoft uses two different keys for DKIM signing. As mentioned earlier, several reasons exist why two keys are used. Still, the most important one is that it allows Microsoft to rotate (update) the keys approximately every six months without affecting the DKIM verification process.

Here's how this works: messages are signed using one of the two selector keys at any given time. For instance, *selector1*. When Microsoft wants to update the key value behind *selector1*, they do not just update it. If they were to do that, messages in transit would be invalidated because their DKIM signature has not yet been verified. This is because the recipient's system queries the key value as it processes the message. As such, if the key value is updated between the moment a message is signed (with the old value) and checked (with the new value), the verification process will fail, and the DKIM Signature is considered invalid.

To overcome this problem, Microsoft uses a two-step process to update the key values. In the above example, Microsoft first switches to using the key value behind the second selector (*selector2*). Any new message sent by Exchange Online from then on is signed with this new key. When remote systems process those messages, they will query the key value for *selector2* instead of *selector1*. Messages still signed by the first selector key can still be verified because the key value is still publically available in DNS.

After about a week, Microsoft removes the DNS entry of the first selector key. This leaves ample room for messages to be delivered and verified, removing the need for that key to be in DNS. For the next six months, outbound emails are signed with *selector2*, and then the process repeats – a new *selector1* is created, email traffic starts to use the new selector value, and a week later, *selector2* is removed. Using a two-step process, Microsoft avoids invalidating messages still in transit and potential delays in DNS propagation.

### DKIM With Third-Party Email Filters

If you route outbound email via a third-party filtering or email hygiene service, you should ensure that this service does not invalidate Microsoft's DKIM signature on outbound messages. If the third-party service rewrites the email or adds new content to the message body, then the DKIM signature will be invalid, and your recipients will see DKIM failure errors (probably just in the headers, but they might reject your email because it fails). Therefore, it is often important to ensure that the last sending system is the system that applies the DKIM signing, remembering that this might not be EOP.

## Domain-based Message Authentication, Reporting, and Conformance (DMARC)

Domain-based Message Authentication, Reporting, and Conformance (DMARC) is a method for authenticating messages that build upon SPF and DKIM. It proves particularly useful in preventing phishing attacks. This is when an attacker spoofs a legitimate domain to trick the recipients into visiting a malicious website and possibly revealing personal information such as passwords. Spoofing a message means that the spammer places a value in the RFC5322 *From* field that is not their domain and uses a value from another domain to pretend that the email is from that organization instead. The RFC5322 *From* header is the one that is shown in clients such as Outlook. Because of this, spoofed messages are hard to differentiate from legitimate emails for end-users.

DMARC specifies what action the receiving server should take if SPF or DKIM validation fails for an email received from your domain. Without DMARC to inform the receiving system what the intent of the sending

system is (concerning quarantine or rejection), the receiver may still decide not to reject messages if either SPF or DKIM validation fails.

Similar to SPF and DKIM, DMARC requires a DNS TXT record that the receiving party can query. A typical DMARC record looks like the following:

```
v=DMARC1; p=reject; pct=50; rua=mailto:postmaster@office365itpros.com;
ruf=mailto:dmarcfailures@office365itpros.com; fo=1
```

- **V** specifies the DMARC protocol version.
- **P** recommends the action the receiving system should take if SPF and DKIM validation fail. Possible actions are:
  - **Reject**
  - **Quarantine**
  - **None** (monitor mode)
- **PCT** gives the percentage of messages subjected to the DMARC policy. The default is 100.
- **RUA** (or *reporting URI for aggregate reports*) contains the email address(es) to which aggregate reports should be sent.
- **RUF** (or *reporting URI for forensic reports*) contains the email address(es) to which forensic reports should be sent.
- **FO** (or *forensic options*) can have a value of 0 or 1. 0 is implied if the *fo* field is absent from the DNS record.
  - 1 = Send forensic reports when either DKIM or SPF fails.
  - 0 = Send forensic reports when both DKIM and SPF fail.
- **SP** (or *sub-domain policy*) is not shown in the above example. If this field is absent from the DNS record, all subdomains inherit the policy from the parent domain (defined with *p=*). If you want a different policy for all subdomains than the parent policy, include *sp=*. For example, if you know you have no subdomains, then add *sp=reject*. This will automatically reject all subdomains as you do not send emails from them. If you have specific subdomains, add a DMARC record for that domain with its policy to override the parent policy.

Aggregate reports are reports the receiver sends daily to the email address specified in the “*rua*” field, containing information such as how many emails have been received and if these messages passed SPF and DKIM tests. Forensic reports are better known as failure reports and are sent to the recipients specified in the “*ruf*” field each time DMARC fails. These reports are extremely useful for figuring out why messages fail DMARC processing. Note that most large email providers (for example, Exchange Online, Gmail, and Yahoo) do not send forensic reports regardless of if you request them in your DMARC record.

## How DMARC Works

The following steps outline how DMARC processes a message:

1. A user receives a message. For example, from *emails@office365itpros.com*.
2. The receiving server will verify if *office365itpros.com* has a DMARC policy (DNS record).
3. Does the message pass SPF/DKIM validation?
  - a. For SPF, does the RFC5322 *From* field and the envelope sender (RFC5321 *MailFrom* or RFC5321 EHLO domain) header match?
  - b. For DKIM, does the RFC5322 *From* match the DKIM signing domain (*d=header*).
4. If the SPF and/or DKIM alignment checks fail, the message fails DMARC processing, and the action specified in the DMARC policy is executed.

If we apply this logic to the DMARC example policy from above, this means that a spoofed message (which would fail SPF or DKIM) would be *rejected*, and a failure report would be sent to [dmarcfailures@office365itpros.com](mailto:dmarcfailures@office365itpros.com):

```
v=DMARC1; p=reject; pct=50; rua=mailto:postmaster@office365itpros.com;
ruf=mailto:dmarcfailures@office365itpros.com; fo=1
```

If you decide to create a DMARC record for your domain, it is best to configure DMARC in *monitor mode* first. This allows you to sift through the failure reports to understand better why any messages from your domain fail DMARC processing. DMARC processing does not always fail because a message is spoofed. Sometimes messages come from a legitimate server that is not included in the SPF record or does not do DKIM signing. The more messages your environment processes, the more failure reports you will receive. For large organizations, going through the failure reports can be very time-consuming. Luckily third-party services exist, which you can specify in the DMARC policy. These services will then process the failure reports for you and provide you with a summary of the failures and the reasons.

**Real-world:** If your domain is a target for many spoofed messages, which is often the case for more prominent and well-known organizations, enabling DMARC might decrease the amount of phishing and spoofed email messages your organization receives. However, if you are a smaller organization, the risk from spoofed emails is likely just as high, as lots and lots of organizations are getting business compromise attack emails, and the end-users are responding and revealing passwords or transferring money, etc. Therefore, we recommend that each domain starts with a policy of none; this allows time to monitor and remediate any legitimate mail failing DMARC. Once all legitimate messages have been remediated, these domains should be transitioned to a policy of quarantine and, finally, a policy of reject.

## DMARC Processing on Inbound Messages

Exchange Online performs DMARC processing on inbound messages if a DMARC policy exists for the sender's domain. Just like SPF and DKIM processing, the result of the DMARC test is written to the *Authentication-Results* header:

```
authentication-results: spf=pass (sender IP is 1.2.3.4) smtp.mailfrom=phishing.com; dkim=none
(message not signed) header.d=office365itpros.com; dmarc=fail action=quarantine
header.from=office365itpros.com; compauth=fail reason=000
```

In this example, the DMARC test failed because the alignment between the RFC5322 *From* address and the envelope sender did not match (phishing.com does not equal office365itpros.com). Therefore, the composite authentication result of SPF, DKIM, and DMARC together (the compauth value) shows a failure, and the reason code means the message failed explicit DMARC authentication.

If the SPF domain in the *mailfrom* header does not match the *From* header on a message entering Exchange Online, it is marked as junk. This is because the *mailfrom* header could match a valid SPF record but be unrelated to the domain that appears in the client application (the *From* header). This scenario results in spoofed emails being marked as valid in terms of SPF. Therefore, EOP requires that the *mailfrom* and *from* header values match to stop EOP considering the message a spoofed email.

Because many organizations employ complex routing instead of just EOP (third-party vendors or on-premises routing, for example), Microsoft cannot always guarantee that a DMARC failure constitutes an actual failure. For this reason, Microsoft will not reject the message even when the DMARC policy is configured with a reject action. Instead, EOP will add the "action=oreject" value to the *authentication-results* header (or "oquarantine" as another lesser seen example), which mail flow rules can use to override the verdict for such messages. Most emails in this category are rejected before reaching the end user's mailbox, but messages can also be rescued from the spam filter by a mail flow rule or allow lists.

## Configuring DMARC

DMARC looks easy to set up, and to an extent, it is – you just add a valid TXT record in public DNS, and the aggregate emails appear within 48 hours in the mailbox identified in the “rua” value of the TXT record. However, the processing of these aggregate emails is more complex. Several services on the internet, such as Dmarcian and Agari, will take these aggregate emails and produce the analytics for you.

Once you have the analytics for your domain, you can tweak your SPF and DKIM configuration to ensure that all your legitimate senders are within scope. Spoofing senders would be outside the scope of your SPF or DKIM records – you would then increase your DMARC policy to  $p=quarantine$  and then eventually  $p=reject$ .

Getting there is the hard part. To help customers with this, Microsoft partners with Valimail to [offer Valimail Monitor](#) as a free service for Microsoft 365 customers to get started with DMARC. This will help you monitor your mail flow and help you work towards enforcing compliance.

## DNS Authentication of Named Entities (DANE) for SMTP

When mail servers transfer mail, they first must agree to use TLS to protect the connection, and if so, which version of TLS. During this initial handshake, these packets are unencrypted. The servers then upgrade the connection TLS after negotiating a common TLS protocol. Mail then transfers over this secure TLS connection.

The challenge is the initial unencrypted handshake. During this unencrypted state, packets are subject to man-in-the-middle and downgrade attacks. To answer this problem, the industry created a new standard called [DNS Authentication of Named Entities](#) (DANE).

DANE allows an organization to publish details of the TLS protocols it supports through a special DNS record in their external DNS. These DNS records are known as TLSA records. The example below illustrates what a TLSA record might look like for the office365itpros.com domain.

```
_25._tcp.office365itpros.com. IN TLSA 3 1 1  
e1c362c8c03a15023fff83831a70d6fce33203d499a3f3d0b13243a1ac689088
```

TLSA records are not exclusive to mail flow. In the example above, the TLSA record exists for TCP 25, which is the port used by SMTP. Multiple TLSA records could exist in external DNS to serve several TCP ports. For example, an external web application that leverages HTTPS could publish TLSA records for TCP 443.

When a mail server that supports DANE wishes to send an email, it performs a DNS query against the recipient's external DNS to see if the TLSA DNS record exists. If the record exists, it is retrieved with DNSSEC, which secures the DNS records by signing the TLSA records in DNS using public-key cryptography. The existence of the certificates enables servers to know that an attacker has not tampered with the TLSA records using a man-in-the-middle attack.

Figure 7-2 illustrates how DANE for SMTP works: the sending server queries the external DNS of the recipient domain for their MX record (1), which is returned (2). The sending server then queries the external DNS of the recipient domain for a TLSA record for TCP 25 (3). If a TLSA record for TCP 25 exists, it is returned via DNSSEC (4). This returned record contains the certificate fingerprint. Next, the sending server initiates a TLS connection to the recipient server (5) using the previously retrieved MX record. The recipient server responds with its certificate fingerprint (6). Finally, the sending server matches the certificate fingerprint transmitted from the recipient server to the fingerprint received from the TLSA record (7). If the fingerprints match, the connection is established, and mail is sent (8). If the fingerprints do not match, the sending server drops the connection.

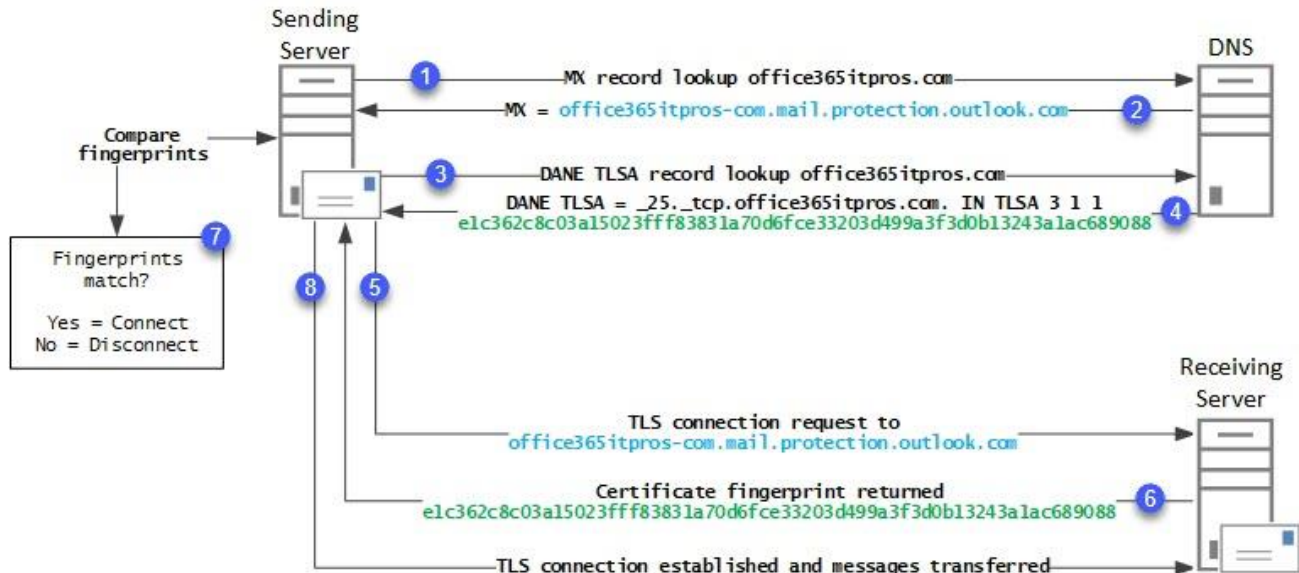


Figure 7-2: How DANE for SMTP prevents downgrade and man-in-the-middle attacks

**Note:** Outbound support for DANE is enabled in all tenants and will not require any tenant administrator action. Inbound support for DANE is slated for release in December 2022.

## MTA Strict Transport Security (MTA-STS)

MTA Strict Transport Security (MTA-STS) is a security technology developed by the Messaging, Malware, and Mobile Anti-Abuse Working Group (M3AAWG). MTA-STS combats the same problem as DANE: how to secure mail flow from man-in-the-middle and downgrade attacks? The main differences are that MTA-STS does not require DNSSEC and uses TXT records instead of TLSA records.

The benefit of MTA-STS is that it is easier to deploy. For example, if your external DNS provider does not support DNSSEC or TLSA records, you cannot use DANE. However, all DNS providers should support TXT records, the only DNS requirement for MTA-STS.

When checking MTA-STS, the sending server looks for a TXT record in the recipient domain with a value of `_mta-sts`. Using `office365itpros.com` as an example, the TXT record is `_mta-sts.office365itpros.com`. The example below illustrates what an MTA-STS TXT record might look like for the `office365itpros.com` domain.

```
_mta-sts.office365itpros.com. 3600 IN TXT v=STSV1; id=20220301000000Z;
```

`v=STSV1` identifies the version of MTA-STS used by the recipient domain. In the example above, this is version 1. Currently, version 1 is the only version of MTA-STS.

The `ID` value identifies the current version of the MTA-STS policy. If the ID value changes, the sender retrieves the latest policy from `https://mta-sts.<domain name>/.well-known/mta-sts.txt`. If the ID value is the same, the sender uses a previously cached copy of the recipient's policy file. This helps eliminate unnecessary HTTPS requests to the server hosting the policy file.

**Note:** Each time you publish a new MTA-STS policy, you should update the ID published in DNS. This alerts senders that your policy has changed and to retrieve a new policy file. Using an ID that specifies a date and time format is recommended, but this ID could be whatever you desire.

One can argue that the lack of DNSSEC means that the retrieval of the TXT record could be subject to man-in-the-middle attacks. However, even if a bad actor changed the ID's value before returning it to the sender, it would merely instruct the sender to download a new policy file from a known location that the recipient owns. The best a bad actor could do is return that no MTA-STS TXT record exists.

The known location used by MTA-STS is always `https://mta-sts.<domain>/.well-known/mta-sts.txt`. This path, file name, and file extension are mandatory for MTA-STS. In addition, the policy file is only retrieved over HTTPS. This means the webserver that hosts the policy file must have a valid third-party certificate. Using `office365itpros.com` as an example, the certificate must contain a subject name (or subject alternate name) of `mta-sts.office365itpros.com`. Alternatively, a wildcard certificate could also be used.

The example below illustrates what the contents of this policy file could contain.

```
Version: STSv1
mode: enforce
mx: office365itpros-com.mail.protection.outlook.com
max_age: 604800
```

**Version** identifies the version of MTA-STS in use. In our example, this is version 1.

**The mode** has three possible values: "enforce," "testing," and "none."

- **Enforce** instructs a sender not to transmit messages to any host that fails certificate validation. Or if the host does not support STARTTLS or TLS 1.2 (and greater).
- **Testing** instructs a sender to transmit messages (including to hosts that fail certificate validation) and provide reports to the recipient of any failures. This mode helps administrators identify and remediate any misconfigured legitimate mail exchangers before enforcing the policy.
- **None** instructs the sender to transmit all messages and treat the recipient domain as if it had no MTA-STS policy.

**MX** defines all MX records served by this policy. This could be a single MX record, like in our example above, or multiple MX records, each entered onto a separate line. Wildcards are also permitted. In the example above, we define Office 365 as a valid MX record for `office365itpros.com`. MTA-STS requires that the hosts defined in the MX records support TLS 1.2 or greater.

**Max age** defines how long (in seconds) a sender should cache this policy. In our example above, 604,800 is 7 days in seconds. This instructs the sender to discard the policy file after 7 days.

Figure 7-3 illustrates how MTA-STS works when in enforce mode: the sending server queries the external DNS of the recipient domain for their MX records (1), which are returned (2). The sending server then queries the external DNS of the recipient domain to see if a `_mta-sts` TXT record exists (3). If the `_mta-sts` TXT record exists (4), the sender checks the ID to see if it has changed since the last time it retrieved the `_mta-sts` record (5). If the ID has changed, the sender performs an HTTPS request to retrieve the MTA-STS policy file from `https://mta-sts.office365itpros.com/.well-known/mta-sts.txt` (6). If the ID is the same, the sender uses a previously cached copy of the MTA-STS policy. With a policy of "Enforce," the sender then checks all MX records defined in the policy, starting with the lowest priority MX record first (7). If the host passes certificate validation, the mail is transmitted (8). If the host fails certificate validation, MTA-STS proceeds to the next MX record in the policy. Mail is not transmitted if all hosts fail certificate validation (9).



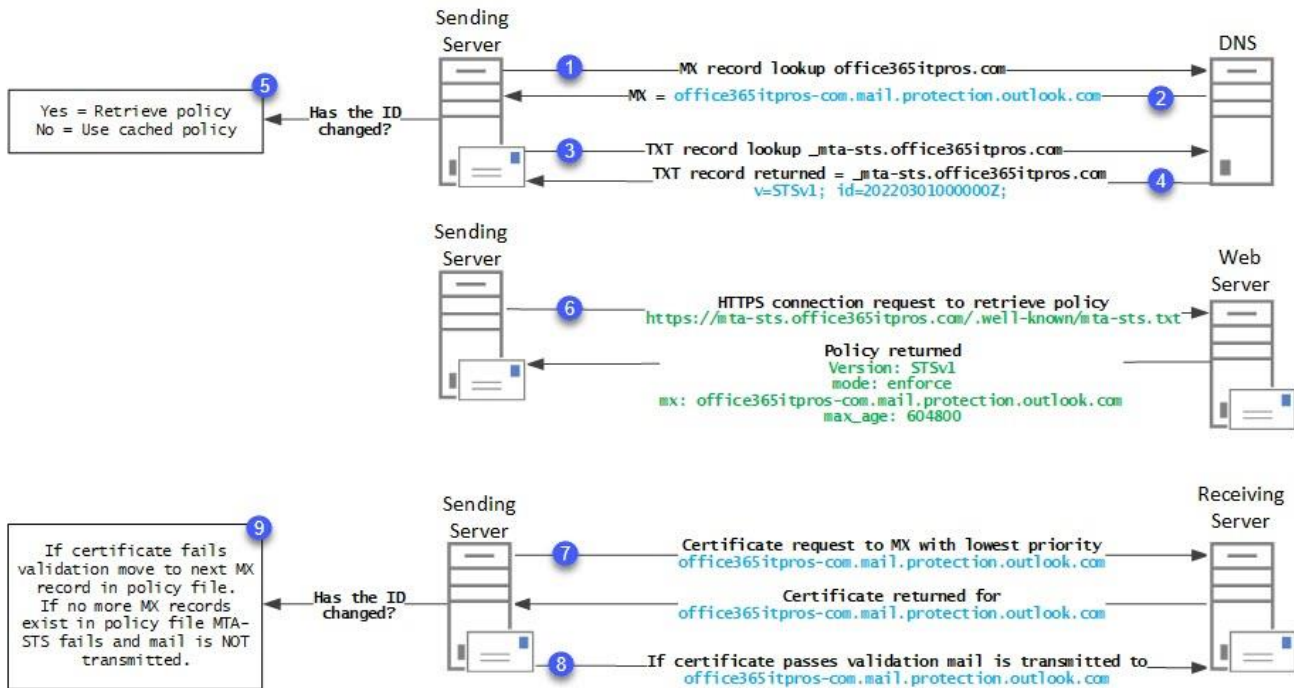


Figure 7-3: How MTA-STS prevents downgrade and man-in-the-middle attacks

**Note:** There is no detriment to having both MTA-STS and DANE configured on your domains. This allows you to cover more senders who may only support one of the technologies.

## Authenticated Receive Chain (ARC)

The challenge with methods used to authenticate outbound email is that they do not account for forwarding. Consider a scenario where `microsoft.com` sends a message to `office365itpros.com`, and a rule at `office365itpros.com` then forwards that message to `contoso.com`. When `contoso.com` queries the headers of the forwarded `microsoft.com` email, it will attribute the sender of the email as `office365itpros.com`. This will fail both SPF and DMARC because `office365itpros.com` is not a valid Microsoft sender.

Authenticated Receive Chain (ARC) combats this by adding a series of email headers that encrypt and store the original SPF and DMARC results to forwarded messages. The receiving system can then make decisions based not just on the SPF and DMARC information it calculates (for example, SPF might fail if the message was forwarded) but on the SPF and DMARC decisions of the original receiving system.

ARC works by utilizing three new headers. First, the *ARC-Authentication-Results* header contains the original SPF and DMARC results. The *ARC-Message-Signature* header stores information about the state of the headers as created by the forwarding system (so that the final receiver can trust what was recorded). Lastly, the *ARC-Seal* header snapshots the *ARC-Authentication-Results* and *ARC-Message-Signature* headers so that the receiver knows if they have been tampered with.

When the email arrives at the receiving server, the system reads the *ARC-Authentication-Results*. If the forwarding system is trusted, the server can base its decision to classify the email as junk or deliver the email on what the original receiver determined (i.e., did the message pass SPF and DMARC on original delivery even though it will fail, on either or both tests, upon forwarding). The *ARC-Message-Signature* and *ARC-Seal* are used to prove the validity of the *ARC-Authentication-Results* header. For example, the *ARC-Message-Signature* hashes some headers from the message (listed in the *h* tag) and stores that hash in the *bh* tag. Changes to the message headers or body that invalidate the ARC headers are easily detected.

Each forwarding of the message creates a new ARC header set starting with an incrementing number (i=2; i=3 etc.). Within the *ARC-Authentication-Results*, you will see that spf=pass, dmarc=pass, and dkim=pass for alignment purposes. An example of what the ARC headers look like in a message header is shown below:

<i>ARC-Authentication-Results</i>	<code>i=1; mx.microsoft.com 1; spf=pass smtp.mailfrom=office365itpros.com; dmarc=pass action=none header.from=office365itpros.com; dkim=pass header.d=office365itpros.com; arc=none</code>
<i>ARC-Message-Signature</i>	<code>i=1; a=rsa-sha256; c=relaxed/relaxed; d=microsoft.com; s=arcselector9901; h=From:Date:Subject:Message-ID:Content-Type:MIME-Version:X-MS-Exchange-SenderADCheck; bh=message-body-hash; b=hash-of-headers-in-h-tag</code>
<i>ARC-Seal</i>	<code>i=1; a=rsa-sha256; s=arcselector9901; d=microsoft.com; cv=none; b=hash</code>

The keys used to sign the hashes are in public DNS under the existing DKIM headers. In the above example, the selector (s=) is arcselector9901. The public key for this selector can be found in a TXT record called arcselector9901.\_domainkey.microsoft.com.

Microsoft 365 tenants trust each other. This means forwarding between tenants results in ARC headers, and the final receiving tenant can use the SPF and DMARC results of the original receiver.

### Adding trusted forwarders (ARC sealers) to your tenant

You may have a situation where you need to add your own trusted forwarders (ARC sealers) to your tenant. This could occur if Office 365 does not trust a forwarder that you need to trust, or you use a third-party message hygiene service in front of EOP that supports ARC sealing.

You can add a list of trusted forwarders (ARC sealers) using the Defender portal. From the **Defender portal**, navigate to **Policies and rules > Threat Policies > Email authentication settings**. From the *Email Authentication Settings* page, select the **ARC** tab. Click the **Add** button to add a trusted sealer. If you have previously added a trusted sealer and need to add more, click the **Edit** button to update the list of trusted sealers. You can also use the *Edit* button to remove trusted sealers from the tenant.

Check [this article](#) for more information on ARC and adding trusted sealers to your tenant.

## Sender Rewrite Scheme (SRS)

Authentication of your email messages with DMARC, DKIM, and SPF improves the reputation of your domain and brand by reducing spoofing and increasing the likelihood of detection of spoof attacks against the domain. However, some problems exist with these techniques. The primary issue is that SPF lists the IP addresses used for outbound emails. If your messages are delivered to another service that forwards them on again (using mail flow rules, distribution lists with external members, or inbox rules, for example), then the SPF information published for your domain makes the forwarding service appear to be a generator of spoofed email.

The Sender Rewrite Scheme (SRS) changes the P1 (or envelope from address) of auto-forwarded (or redirected) messages so that the receiving next-hop sees that the sender is the forwarding service and not the original sender. The goal is to try and avoid an SPF failure. SRS does this by changing the P1 header from the original sender's address and domain to an address that encapsulates the original sender's address but is really the forwarding service. Note, though, that this means that NDRs sent to these newly rewritten addresses will no longer return to the original sender but to the forwarding service, which in this case is Exchange Online. Exchange Online then unwraps the SRS rewritten header and sends the NDR back to the original sender. NDRs that cannot be unwrapped will go to a bounces@<default-accepted-domain> mailbox, which you need to create. The P2, or the *From* address the user sees in the email client, will not be changed. For example:

Before SRS is implemented, the P1 address is: *brian@office365itpros.com*

After SRS is implemented and the initial receiving tenant forwards messages from the original sender, the P1 address might be something like this:

```
sales+SRS=f2ss=IX=office365itpros.com=brian@itpros.com
```

This is best explained as:

```
<Forwarding Mailbox Username>+SRS=<Hash>=<Timestamp>=<Original Sender Domain>=<Original Sender Username>@<Forwarding Mailbox Domain>
```

Now that the forwarded email comes from the forwarding tenant and not the original system, it should not fail SPF. More details about the implementation of SRS are available on the [Exchange Blog](#).

## Managing Connectors

EOP uses connectors to control inbound and outbound mail flow. By default, Exchange Online uses hidden connectors to allow messages to be sent between the tenant and the internet. However, if you have a requirement, such as delivering mail to a smart host or forcing TLS, you need to leverage a custom connector.

Connectors do more than enable mail flow with external organizations; they also apply security settings such as transport layer security (TLS). The receiving server almost always specifies requirements for TLS. Exceptions to this rule can exist, for instance, when a connector is configured to require (or force) TLS on a specific connector. By default, EOP will always try to negotiate TLS with the receiving server. This is called opportunistic TLS. When TLS cannot be used, the sending server (EOP in this case) will revert to using an unencrypted connection. However, if the recipient's message system supports TLS, EOP uses TLS for the connection.

Messages destined for the on-premises Exchange environment in a hybrid deployment are, by default, always secured with TLS and require a valid, trusted digital certificate. The same is true for messages sent to domains protected by EOP. However, if you often communicate with a business partner or send highly confidential information, you might want to enforce specific security settings, such as requiring TLS encryption. There are several ways in which TLS encryption can be enforced. For instance, you can require the certificate to encrypt the traffic with a specific subject name. To configure specific TLS settings for a remote organization, you must create a custom connector.

### Creating a connector

You can use the *New-InboundConnector* and *New-OutboundConnector* cmdlets to create connectors via PowerShell. However, the EAC wizard is much easier to work with. You do not have to specify what type of connector you want to create (inbound/outbound). Instead, you specify the source and target systems, and the wizard will automatically create the right connector. To open the wizard, open the EAC and navigate to **Mail flow > Connectors > Add a connector**.

For example, suppose we need to route mail to a specific domain through a smart host (rather than delivering to its MX record). For this, we can create an outbound partner connector. To do this, launch the connector wizard from **Mail Flow > Connectors > Add a connector**. Select Office 365 as the source and Partner Organization as the destination. Click **Next**. Give the connector a **Name** and **Description**. On this page, you can also choose whether to enable the connector as part of the wizard. You may wish to uncheck **Turn it on** if you want to enable this connector later manually. Click **Next**. On the *Use of connector* page, add your target domains in the **Only when email messages are sent to these domains** field. Click **Next**. On the *Routing* page, check **Route email through these smart hosts** and enter a fully qualified domain name (FQDN) or IP.

Click **Next**. On the *Security Restrictions* page, specify whether to require TLS and the requirements for the certificate. Click **Next**. Validate the connector and click **Next**. Click **Create connector**.

According to Microsoft, many reported mail flow problems are caused by incorrectly configured connectors. To reduce the number of problems due to misconfiguration, when creating an outbound connector (from Office 365 to somewhere else), the wizard tries to validate the connector is functional before creating it. Then, an attempt is made to send a test message to a recipient you specify. The test is successful if the wizard can connect and deliver the message to the remote environment. If the validation fails, you can select to override the failure in the wizard and create the connector.

Using our example above, we can create this connector with PowerShell using the *New-OutboundConnector* cmdlet. We specify the scope of this connector with the *RecipientDomains* parameter and the delivery target with the *UseMXRecord* and *SmartHosts* parameters. The *TlsSettings* and *TlsDomains* parameters specify the TLS security when transferring mail over the connector.

```
[PS] C:\> New-OutboundConnector -Name "Contoso Outbound Connector" -ConnectorType Partner -
RecipientDomains contoso.com -TlsSettings DomainValidation -TlsDomain *.contoso.com -UseMXRecord
$false -SmartHosts mail.contoso.com
```

To review the properties of the connector (whether created in the EAC or PowerShell), you can use the *Get-InboundConnector* or *Get-OutboundConnector* cmdlet. In the example below, we retrieve the properties of the *Contoso Outbound Connector*.

```
[PS] C:\> Get-OutboundConnector "Contoso Outbound Connector" | Format-List

Enabled                : True
UseMXRecord            : False
Comment                :
ConnectorType          : Partner
ConnectorSource        : Default
RecipientDomains       : {contoso.com}
SmartHosts             : {mail.contoso.com}
TlsDomain              : *.contoso.com
TlsSettings            : DomainValidation
IsTransportRuleScoped : False
RouteAllMessagesViaOnPremises : False
CloudServicesMailEnabled : False
AllAcceptedDomains    : False
SenderRewritingEnabled : False
TestMode               : False
LinkForModifiedConnector : 00000000-0000-0000-0000-000000000000
ValidationRecipients   :
IsValidated            : False
LastValidationTimestamp :
Name                   : Contoso Outbound Connector
```

For more information on creating connectors with PowerShell, check the [New-InboundConnector](#) and [New-OutboundConnector](#) Microsoft articles.

## Mail Flow Rules

A modern messaging system must be able to handle the various requirements an organization might have. It is no longer enough for a server to send messages from point A to point B. As an organization evolves, its messaging system must offer the right feature set to support the changing requirements of that organization.

One could compare the changing behavior of interacting with a messaging system to how people interact with the post office. Where before it was good enough to deliver a written letter to its recipient, today the post office offers many additional services. For instance, when you move to a new house, you can request the post office to automatically forward mail addressed to your old location to your new home. This not only

gives you the time to inform all users about the address change, but more importantly, it prevents mail from being returned to sender or, worse: delivered at the wrong address.

The general idea behind mail flow rules is very like the above example. It provides Exchange Online with a way to dynamically handle incoming or outgoing messages based on one or more criteria. But, as you will see, you can use mail flow rules for much more than forwarding messages to a different email address. Mail flow rules also support Data Loss Prevention (Chapter 19) and Information Protection (Chapter 20).

The way mail flow rules work in Exchange Online is very similar to how they operate in on-premises Exchange, with the main exception that Exchange Online mail flow rules have more conditions and actions to choose from. As a result, an organization can create up to 300 different mail flow rules.

## Mail Flow Rule Conditions

Conditions (also known as predicates) define when a mail flow rule should be triggered. For instance, a condition can be used to check for a specific value in the sender's email address, or it can look for the existence of an attachment. The easiest way to build a new mail flow rule is to use the EAC, as you can select conditions from a drop-down list.

A rule can contain multiple conditions, in which case the message must match all individual conditions before any actions are triggered (logical "AND" configuration). You must create separate mail flow rules if you need to match several conditions in a logical "OR" configuration. If a single rule must apply for different values of the same condition, then add multiple values to a given condition. If multiple values are specified, the message must match one of the values assigned to the condition (logical "OR" configuration).

**Note:** It is also possible to create a "catch-all" rule. If you do not specify a condition (or you select *[Apply to all messages]*), the mail flow rule will be applied to every message that flows through the organization.

## Mail Flow Rule Exceptions

Exceptions are, just like conditions, used to scope the applicability of mail flow rules. Exceptions are used in conjunction with 'regular' conditions to exclude matches against (one of) the primary condition(s). Each condition has a matching exception.

Multiple exceptions can be configured for a given rule, in which case a single match is needed for the rule to be skipped (logical "OR" configuration). Similarly, you can specify multiple values for a given exception, and the rule will be skipped if any of the values are encountered.

**Note:** Any predicate can be used if messages are unencrypted when they are processed. S/MIME encrypted messages (not to be confused with messages protected with sensitivity labels or Office 365 Message Encryption) cannot be processed by mail flow rules based on conditions that inspect the contents of the message. In any other case, mail flow rules with conditions based on a message's envelope header will still work correctly.

For a comprehensive list of mail flow predicates and exceptions, check this [Microsoft article](#).

## Mail Flow Rule Actions

Actions ultimately carry out a task on the message. Each action uniquely affects the message, either by changing some of the message's properties or altering the routing behavior of the message. Amongst other actions, you can, for instance, re-route, reject, mark as spam, or silently redirect a message.

A mail flow rule can include multiple actions. Each rule has a priority number and the transport service processes applicable rules in ascending order of priority. This means that if you want to execute actions in multiple rules on a message, you must plan the rules so that the desired actions run in the correct order. In

determining the priority order for rules, you should consider the possibility that a rule might cause all further processing to stop after it completes because it includes the *Stop processing more rules* option. Rules also stop processing if a rule's action is to drop (delete) a message, as the next rule will not have a message to process. Something similar happens if the action is set to moderate the message (sent for approval): Exchange won't process lower priority rules until the message is approved. After approval, the transport service evaluates the remaining rules against the message.

**Note:** Over the years since the introduction of EOP, the options for spam and malware filtering have expanded. This means fewer reasons to use mail flow rules for spam/malware filtering. For example, the malware filter includes attachment protection settings, thus avoiding the need to look for attachments by file extension name in a mail flow rule.

For a complete list of all actions available for a mail flow rule, check [this article](#).

## Mail Flow Rule Properties

Apart from the building blocks detailed above, each mail flow rule also has properties that control various aspects of its processing. These include:

- **Priority** determines the order in which rules are processed. By default, Exchange Online orders mail flow rules by creation date, but you can override the processing order by adjusting the value of the priority property. Rules are processed from the lowest to the highest value.
- **Audit this rule with severity level** allows you to “tag” a rule as either *Low*, *Medium*, *High*, or *Not specified severity*, making it easier to filter or group the data in the mail flow rule reports.
- **Choose a mode for this rule** gives you the option to test a rule before turning it on (explained below).
- Optionally, you can toggle the **Activate this rule on the following date** and/or **Deactivate this rule on the following date** checkboxes to “schedule” the rule to be active only during a given time.
- The **Stop processing more rules** setting effectively stops mail flow rule processing after the given rule has executed its actions. Any mail flow rule with lower priority will not be affected.
- Turning on the **Defer the message if rule processing does not complete** option will cause the message to be resubmitted for reprocessing in the event of a failure to process the given mail flow rule.
- The selection under the **Match sender address in message** dropdown allows you to specify whether matches against any conditions or exceptions, including the sender address, are performed against the header value, the envelope value, or both.
- Lastly, the **Comment** field is available to store additional information about the rule, such as its creation date, the reason for a given modification, etc.

Apart from all the properties listed above, you can also toggle a mail flow rule On or Off by selecting the corresponding checkbox in the list of rules presented in the EAC or via the *Enable-TransportRule* or *Disable-TransportRule* cmdlets.

## Creating a New Mail Flow Rule

As mentioned earlier, the easiest way to create a new mail flow rule is to use the wizard in the EAC, as it allows you to select conditions, exceptions, and actions from a drop-down list. The EAC also allows you to select a pre-configured mail flow rule from a list of templates. These templates cover common scenarios, such as allowing a specific message to bypass spam filtering. The conditions, exceptions, and actions are already pre-selected depending on the template you selected. You only need to specify a unique value(s) for any of the specified properties. For instance, if you select the **Bypass spam filtering...** template, the **Set the spam**

**confidence level (SCL) to...** action has already been pre-selected, and you then select which SCL value you want to apply to the message.

You can build a mail flow rule from scratch using the EAC. By default, if you select **Create a new rule**, you will only see a subset of conditions and actions. However, clicking the **More options** link gives you the complete list of available conditions, exceptions, and actions.

## Testing New Mail Flow Rules

Before you put a new mail flow rule into production, it is wise to verify that the rule does what you expect it to. You can verify whether a mail flow rule works by enabling it in test mode. You will have the following options to choose from when creating a new mail flow rule:

- Enforce (default).
- Test with policy tips.
- Test without policy tips.

The **Enforce** option triggers the rule if all the conditions of the rule are met. However, **Test without policy tips** will stop the rule action from running because the rule is in test mode. But how do you know the rule has worked if it does not fire? The answer is to add the **Generate incident report and send it to** action.

The **Generate incident report and send it to** action sends an email to the selected recipient with information about the message that caused the rule to trigger. In test mode (with Test without policy tips enabled), Exchange Online generates and sends the incident report email. Still, the other actions specified in the rule, such as setting a disclaimer, are not performed. The incident report action allows you to select the properties of the message you want to include in the incident report. For example, if you are working with sensitive information or local legislation prevents you from accessing the contents of a message without notifying the user, you can opt not to include the original mail and make sure that the incident report only reveals non-crucial information about the message. An example of an incident report email appears below, where we see that the disclaimer rule “Corporate Disclaimer” was triggered. Because the rule was in test mode, the actual disclaimer was not applied and, therefore, not seen by the recipient.

```
Report Id: bbc7457a-64a9-44c5-a246-1b7d35db116b
This email was automatically generated by the Generate Incident Report action.
Message Id: <VI1PR06MB1183C94FCA77A1BFB979A927C9560@VI1PR06MB1183.eurprd06.prod.outlook.com>
Sender: john.smith@office365itpros.com
Subject: Incident Report Testing
Recipients: jane.doe@office365itpros.com
Severity: Low
Override: No
False Positive: No
Rule Hit: Corporate Disclaimer, Action: ApplyHtmlDisclaimer, GenerateIncidentReport
```

Incident Reports are an easy way to check the effectiveness of a mail flow rule and, at the same time, gather information about what kind of messages will trigger the rule. This will allow you to decide whether the mail flow rule meets your expectations. Once you decide that the mail flow rule is ready for production, you can edit the mail flow rule and choose to **Enforce** it. Do not forget to remove the action to send an incident report, or you will receive a report each time the rule is triggered!

## Monitoring Mail Flow Rule Usage

One measurement of the effectiveness of a mail flow rule is the number of times it triggers. Of course, there might be cases where you hope a mail flow rule never triggers, such as when a mail flow rule enacts an *ethical wall* to prevent two departments from communicating with one another. This is often seen in financial institutions where regulations define that market researchers cannot communicate with brokers to avoid a conflict of interest or (in-)voluntary market manipulation.

The Exchange Admin Center includes the “Exchange Transport Rule” report dedicated to mail flow rules. The report gives you a breakdown of all the individual mail flow rule matches in either graphical or table view. In addition, you can also group the results by the audit severity value, which we described earlier. Like many of the other Exchange-related reports, you can schedule this report to be emailed to specified recipients on a weekly or monthly basis. We will discuss this report and others later.

The *Get-MailDetailTransportRuleReport* cmdlet searches for specific transport-rule related events. For instance, it can be used to search for 'hits' within a given timeframe or to look for actions that have been performed because a mail flow rule was executed. The example shown below reports all mail flow rule hits in the past 5 days and lists the date and time, message subject, action taken, and rule that was executed:

```
[PS] C:\> Get-MailDetailTransportRuleReport -StartDate (Get-Date).AddDays(-5)
-EndDate (Get-Date) | Format-List Date, Subject, Action, *Rule*
```

If you need to search for hits of a specific mail flow rule, you can use the following command:

```
[PS] C:\> Get-MailDetailTransportRuleReport -TransportRule "Check for sensitive data"
```

The cmdlet has many more parameters that allow you to refine the results further. For instance, you can look for messages from a specific sender:

```
[PS] C:\> Get-MailDetailTransportRuleReport -StartDate (Get-Date).AddDays(-5)
-EndDate (Get-Date) -Sender Joe@office365itpros.com
```

The dashboard reports and the data generated by the PowerShell cmdlets come from Microsoft's reporting data warehouse. The only downside to this approach is that the information in the data warehouse isn't always up to date. It can take up to a day before these cmdlets return complete data. If you are looking for an immediate way to measure mail flow rules, adding an incident report is preferred.

## Mail Flow Rule Limitations

Microsoft limits the amount of transport and journal rules each tenant can create. These limitations protect the service by limiting the amount of routing logic applied to each message. Table 7-1 lists the rule limits.

<b>Feature</b>	<b>Limit</b>	<b>Additional Information</b>
Journal Rules	300	The maximum number of journal rules that can be created in a tenant
Mail Flow Rules	300	The maximum number of mail flow rules that can be created in a tenant

Table 7-1: Transport rule limits

## Common Use Cases for Mail Flow Rules

Because of the many available conditions, mail flow rules are useful in a myriad of scenarios. It would be impossible to list all the use cases here, but we will look at some of the more common scenarios and tasks. The examples from the scenarios described below should provide you with enough insights into the capabilities of mail flow rules so that you can produce viable solutions to any unique requirements you have.

### Organization-wide Disclaimers

Sometimes legal or regulatory requirements dictate the need for an organization to add disclaimer text to outbound messages. Mail flow rules can handle this requirement because one of the actions available to a mail flow rule is the ability to prepend or append a disclaimer to a message. The content of the message can be plain text, but it can also include HTML code to make it more dynamic.

Additionally, you can include certain variables based on the attributes of a user's Azure AD account to insert user-specific values in the disclaimer. For instance, by adding *%%DisplayName%%* to the disclaimer text, Exchange Online inserts the sender's display name into the text.



However, when you add a disclaimer to all outgoing messages, the user experience may not be what you expect. For example, in Figure 7-4, when you append a disclaimer to a message, the transport service adds the disclaimer text to the end of the message. Of course, this is precisely where you would expect it to be for a new message (1). However, if you reply to an email thread, you will notice that the disclaimer is added at the very bottom of the entire thread, not after the latest reply (2). This might be perfectly acceptable for a disclaimer, but it looks strange for signatures.

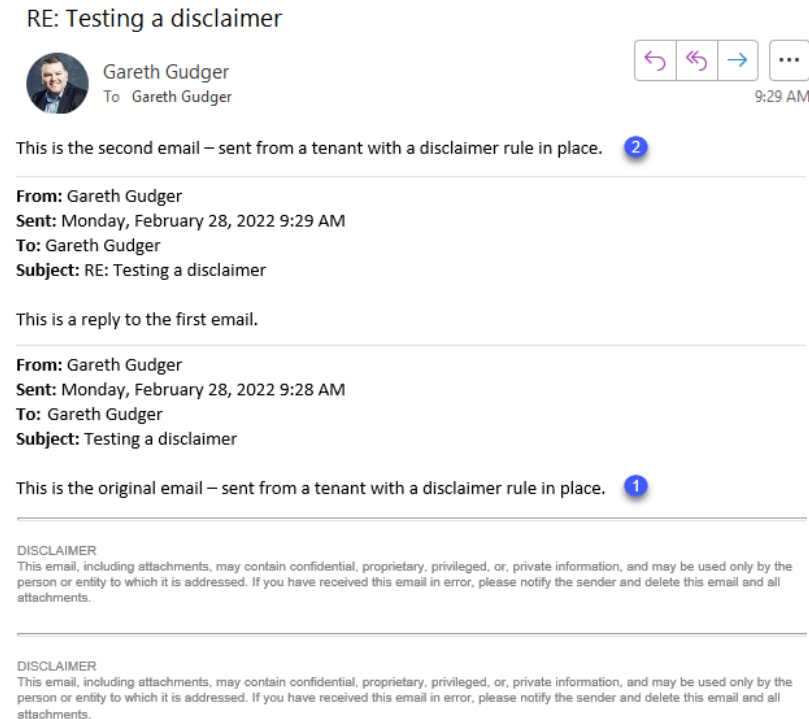


Figure 7-4: The disclaimer text added by a mail flow rule appears at the bottom of a message

Although you cannot make a mail flow rule add text at the end of the latest message, you can stop the rule by adding disclaimer text or signature if one already exists by adding an exception to the rule. This PowerShell example shows how:

```
[PS] C:\> New-TransportRule -Name "CompanyDisclaimer" -Enabled $True
-SentToScope "NotInOrganization" -ApplyHtmlDisclaimerLocation "Append"
-AppllyHtmlDisclaimerText "<br><hr><font face='Arial' color='Gray' size='1'><br>DISCLAIMER<br>This
email, including attachments, may contain confidential, proprietary, privileged, or, private
information, and may be used only by the person or entity to which it is addressed. If you have
received this email in error, please notify the sender and delete this email and all
attachments.<br></font>"<table border="0"> <tr><td>Column1</td> <td>Column2</td></tr><tr><td
colspan="2"><div style="font-size:10pt; font-family:'verdana'">This is the disclaimer text and it
cannot exceed 5,000 characters. </div></td></tr></table>"
-ExceptIfSubjectOrBodyContainsWords "This email, including attachments, may contain confidential,
proprietary, privileged, or, private informationThis is the disclaimer text, and it cannot exceed."
```

The first time the message is processed, the rule runs and adds the disclaimer because it's improbable that the phrase used by the exception is present in the body or subject. However, on any subsequent occasion, the message or its replies pass through the transport service, such as when a recipient sends a further reply, the rule exception prevents the insertion of the disclaimer because the body now contains the exact words added by the same rule the first time the message was processed.

**Note:** One important limitation of mail flow rules is that the total size of the disclaimer cannot exceed 5,000 characters. This includes the text and any HTML tags or Cascading Style Sheet (CSS) code you might add.

## Bypass Spam Filtering

As discussed later, EOP automatically collects allowed and blocked sender settings from the junk mail settings manipulated by users with Outlook and Outlook Web App. Sometimes, an organization might want to centrally control which senders receive automatic approval and those marked as spam. For instance, this will be necessary if you introduce a new corporate communications tool that sends messages from an external system.

To achieve this goal in Exchange Online, you can use a mail flow rule to add a message header set to a specific value. In this case, the header is the "SCL" header (Spam Confidence Level) which will be set to a value of "-1". The value of the SCL header determines if a message should be sent to the user's junk email folder at the end of the processing pipeline. Table 7-2 describes the values used for the SCL header in Exchange Online:

<b>SCL Value</b>	<b>Meaning</b>
-1	The message is deemed safe because the sender is safe-listed, the sender's server is on an IP safe list, or the recipient is on the safe-recipient list. Therefore, the content filtering engine does not process the message.
0, 1	After processing by the content filtering engine, the message was deemed to be clean.
5, 6	After processing by the content filtering engine, there is reasonable confidence to mark the message as spam.
9	After processing by the content filtering engine, there is no doubt (high confidence) that the message is spam.

Table 7-2: SCL header values in Exchange Online

In the following PowerShell example, a new mail flow rule is added to create a safe list that will automatically mark messages from the office365itpros.com domain as safe:

```
[PS] C:\> New-TransportRule -Name DomainSafeList -SenderDomainIs "office365itpros.com" -SetSCL "-1"
```

**Real-world:** Care should be taken with setting blanket safe list rules like those shown above. If the domains or senders in question have their email address spoofed so that a message does not come from the user or organization who owns the mailbox but instead from someone acting as that sender or domain, EOP will mark the email as safe. This is because the rule checks the domain or user and not doesn't do anything to validate the reliability of the sender. Anti-phishing rules are recommended as a replacement for blanket allow lists.

## Conditional Mail Routing

Conditional Mail Routing allows you to use mail flow rules to alter the default routing behavior for specific messages that match the conditions of a mail flow rule. However, if you want to control the behavior for all outbound messages, it is better to create a custom connector instead.

Consider the following scenario: You represent a large organization with multiple types of users—for instance, an educational institution with students, staff, and faculty. You use Exchange Online for both types of users but want email from staff and faculty to be routed to the internet through an on-premises appliance that offers various features like journaling, signatures, secure messaging, etc. Even though these features also exist in Exchange Online, you previously invested in the solution and do not want to lose out on those investments. If you created a regular outbound connector, all messages would be routed through the external appliance, including those from students. So instead, you create a new connector, only to be used by a mail flow rule which defines that outbound messages sent by faculty and staff mailboxes should be routed through the new connector.

This scenario is an excellent example of how conditional mail routing can help an organization meet regulatory requirements without making extensive changes to other parts of the configuration and leveraging

prior investments. Although conditional mail routing adds a layer of complexity to the entire solution by splitting the message routing logic, it can be extremely valuable if carefully planned and used for the right purposes.

**Real-world:** The most common use of conditional mail routing is to ensure that messages sent to specific users or applications are either encrypted with TLS or sent directly to a specific system. The latter option requires the outbound connector to be configured with a smart host. As such, the MX records for the recipient's domain are ignored, and messages are delivered directly to the specified smart host. This is also done to ensure messages are sent to servers in a specific region. Often the recipient's organization does not have the infrastructure to dynamically (and geographically) route incoming messages directly to servers in the same region as the mailbox. In such a case, the sender's organization can then use conditional mail routing to accomplish the task.

Conditional mail routing relies on mail flow rules to check the properties or content of a message. The mail flow rules will redirect the message to the selected connector if a specific condition is met. Thus, before creating a mail flow rule to redirect messages to a specific connector, you must first set up a connector and configure it for conditional mail routing. To do this, create a custom outbound connector and select the option **Only when I have a transport rule set up that redirects messages to this connector** when asked *'When do you want to use this connector?'*

In the following example, we will configure a mail flow rule to redirect all messages sent to "kim.akers@office365itpros.com" through the custom connector *Outbound to Office365ITPros*. We will only configure the basic settings to make this rule work. Start by opening the new mail flow rule wizard from the EAC, select **Mail Flow**, navigate to **Rules**, click on the plus sign to show the drop-down menu, and select **Create a new rule...**

By default, the list of actions displayed on the screen is limited and does not allow redirecting messages through a specific connector. To enable access to the more advanced options, you must first select **More options...** After clicking the link, the additional settings will be displayed.

1. Specify a descriptive name for the rule. This will ensure that you can easily find the rule in the list of rules.
2. Select a condition to which the message should apply. For this example, select **The recipient is...** If you haven't created a contact for the recipient, type it manually in the *check* names field and then hit Enter. For example, type *kim.akers@office365itpros.com*.
3. From the actions drop-down menu, select **Redirect the message to...** and **the following outbound connector**.
4. Select **Outbound to Office365ITPros** from the list of available connectors.

After saving the mail flow rule, it is enabled and will apply to any outbound message sent to kim.akers@office365itpros.com.

## Identifying External Senders

Adding a warning message to emails from external senders is a common use case for mail flow rules. For many organizations, this warning message is a core part of end-user training to identify and report potential phishing messages. The downside to using a mail flow rule is that the warning message is often unsightly and might be included in future replies to the sender. You can enable external email tagging in Exchange Online to transition away from a mail flow rule (if you have one). Exchange Online's external email tagging functionality is available in Outlook, OWA, Outlook Mobile, and Outlook for Mac.

To enable external email tagging, run the following command in Exchange Online PowerShell:

```
[PS] C:\> Set-ExternalInOutlook -Enabled:$true
```

When Exchange Online marks a message as external, it sets the *IsExternalSender* MAPI property for the message to True. Clients check for the property and display the warning if it is True.

You can specify external email domains that Exchange should not treat as "external." For example, you might have messages from a third-party training system that should appear to be from an internal sender. To add a domain to the list of allowed external senders by running the following command:

```
[PS] C:\> Set-ExternalInOutlook -AllowList @{Add="Office365ITPros.com"}
```

## Exporting and Importing Mail Flow Rules

Although mail flow rules can be enabled in test mode, it can sometimes come in handy to experiment with a few settings in a test tenant. Once testing is complete, you can export the rule(s) from a test tenant and import them into the production tenant. To export mail flow rules from a tenant, run the following command while connected to Exchange Online PowerShell:

```
[PS] Set-Content C:\ExportTransportRules.xml -Value ((Export-TransportRuleCollection).FileData -Encoding Byte)
```

**Be careful when importing:** The *Import-TransportRuleCollection* cmdlet overwrites the existing mail flow rules in the target tenant with the content of the XML file. This command is a rip and replace. In this example, this means the mail flow rules you export from the test tenant will overwrite the rules in the production tenant. To ensure you have a rollback option, export rules from the target tenant before import.

Next, connect to the target tenant and run the following command to import the ruleset:

```
[PS] C:\> Import-TransportRuleCollection -FileData (Get-Content -Path C:\ExportTransportRules.XML -Encoding Byte -ReadCount 0)
```

Confirm

Importing a rule collection will overwrite all existing rules in that collection. Do you want to continue?

[Y] Yes [A] Yes to All [N] No [L] No to All [?] Help (default is "Y"): Y

## Remote Domains

Any email domain not registered with your tenant is considered a remote domain. By default, you can deliver email to any remote domain. However, an organization might want to limit what messages get sent to remote organizations, or maybe they want to control the message format. This is where remote domain configuration settings are useful.

A remote domain allows you to control:

- **Out of Office replies.** Allow or deny replies to be sent to the domain. You can also control whether the internal or external out-of-office response is sent.
- **Automatic replies and forwarding.** You can control if a user-generated rule can automatically reply to or forward messages to the remote domain.
- **Message reporting.** This allows you to control whether message status notifications should be sent to the remote domain. For instance, you can prevent non-delivery reports from being sent to a remote domain or configure that meeting forward notifications are sent; by default, meeting forward notifications are only sent within the organization.
- **Format of the message.** Allows you to define if a message should be sent in rich text and what MIME character set should be used, if at all.

**Note:** The settings you specify for a remote domain override any user-specific settings, for instance, configured through the *Set-MailboxAutoReplyConfiguration* cmdlet. If you want to modify default settings for all other unlisted domains, you must modify the default remote domain.

You can add a new remote domain or configure an existing one by going to **Mail flow** and **Remote domains** in the EAC.

Automatic message forwarding is a common action that attackers take when they compromise a mailbox. Because this threat exists, Exchange Online does not allow automatic forwarding by default. You can amend the default policy to allow users to auto-forward email, but this is a terrible idea. Instead, if some users need to auto-forward email from their mailboxes, you can create a custom outbound spam filter policy. The custom policy specifies the email addresses of users who can auto-forward their email outside the organization.

## Device and app mail relay to Exchange Online

Even in the digital world, people still print many documents. Often, these documents are scanned again for digital archiving. For years, scanners could email scanned messages. Despite sending email messages, many of these devices do not support authentication and can only send plain, unauthenticated SMTP messages. Scanned messages mostly go to internal recipients anyway. However, the same is not always true for some line-of-business (LOB) applications. For example, the application might need to send a message to an external recipient when an order is updated.

This is not a problem for Exchange on-premises because you can scope your firewall to accept messages from specific device or application IP addresses. However, it's not that simple with EOP. Because EOP is publicly available on the internet, allowing unauthenticated SMTP would open the door to spammers to use EOP as an "open relay." Today, there are four possible solutions:

- If your device or application supports authentication, you configure it to use the credentials of an Exchange Online mailbox to send messages internally and externally. This method requires the device to support TLS 1.2 and send mail over port 587. Microsoft refers to this method as **SMTP Auth**.

**Note:** Using TLS 1.0 and TLS 1.1 with SMTP Auth is unsupported. Microsoft has stated they will block TLS 1.0 and TLS 1.1 to the smtp.office365.com endpoint in 2022. Therefore, any devices or applications that require these older protocols will need to leverage the new endpoint of smtp-legacy.office365.com. To activate this [legacy endpoint](#), tenant administrators must run *Set-TransportConfig -AllowLegacyTLSClients \$true*.

- If the device or application either does not support authentication, does not support TLS, or can only use port 25, you can use **Direct Send**. Note that this can only be used to send mail to internal recipients in your tenant. Any external recipient will be rejected. Mail will also be heavily scrutinized and be subject to anti-spam policies and could be flagged as spam.
- If your device needs to send mail to external recipients, another option is to use a **custom connector** to identify and authenticate all on-premises devices and applications that need to send email messages. This option will likely require network configuration changes to ensure that all outbound connections from these devices and applications to Exchange Online originate from a specific set of IP addresses.
- A custom or third-party SMTP relay solution is used to forward messages to Exchange Online on behalf of the SMTP device or application. For instance, the IIS SMTP service can be used to create a custom SMTP relay solution. The SMTP service will accept unauthenticated messages from devices and applications within your network and, in turn, forward the messages to Exchange Online using the credentials of an Exchange Online mailbox.

For information on configuring each of these relay methods, check this Microsoft [article](#).

**Managing hybrid mail flow:** While it is possible to eliminate your last on-premises Exchange server for management tasks, many organizations maintain Exchange servers for SMTP relay for on-premises devices and apps. In those scenarios, it is easier to keep routing those emails through the existing Exchange Servers configured for secure hybrid mail flow.

The other aspect is security. Many enterprises leverage firewall rules to govern what devices and apps can send outbound mail. For security-conscious organizations, they want to minimize the number of outbound SMTP connections to the absolute minimum. To use an analogy, they do not want to swiss-cheese their firewall. In this scenario, security dictates that only a handful of known and centrally managed Exchange servers (or 3<sup>rd</sup>-party relays) can send outbound mail.

## Exchange Online Protection (EOP)

Exchange Online Protection (EOP) is Microsoft's cloud-based email filtering service. EOP includes a set of message hygiene features to sanitize inbound and outbound mail flow and remove threats from messages. EOP protects all emails exchanged between Exchange Online (including between tenants) and external sources. Protection is in place to divert spam and malware and to guard against potential data loss. There are three configurations for EOP:

- **Exchange Online (cloud-only) deployment.** Although EOP is a separate feature, it is an integral part of and tightly integrated into Exchange Online. If you have an Exchange Online tenant, you automatically use EOP.
- **Standalone.** Organizations that do not use Exchange Online can route email traffic to EOP to use it as an email hygiene service. This can be for an on-premises Exchange environment or another email solution (hosted or on-premises).
- **Hybrid deployments.** In this scenario, EOP protects the traffic between the cloud and the on-premises servers. In a non-centralized mail flow approach, EOP can protect on-premises mailboxes as it does in a standalone deployment.

Exchange Online provides redundancy and load-balancing within a data center region. This also satisfies regulatory requirements that might need data to remain within a particular geographical area, such as the EU data processing guidelines. EOP running in country-level data centers (like Norway, France, or the United Kingdom) processes messages for tenants in those countries. EOP running in data centers in the United States (US) processes messages for US tenants. This means that it is vital to be aware of the correct namespace to use when routing messages so that you do not route to the wrong data center and have the message declined.

### How EOP Processes Email

The filtering system in EOP consists of multiple layers and processes to handle inbound and outbound messages. To understand how Exchange Online processes messages, let us review the diagram shown in Figure 7-5, which shows how EOP uses inbound mail filtering to protect Exchange Online mailboxes.

The transport service routes email from source to destination unless the message triggers some component of the service that deems the message unsafe, in which case the properties of the message are changed, or the message is routed to the quarantine or rejected outright.

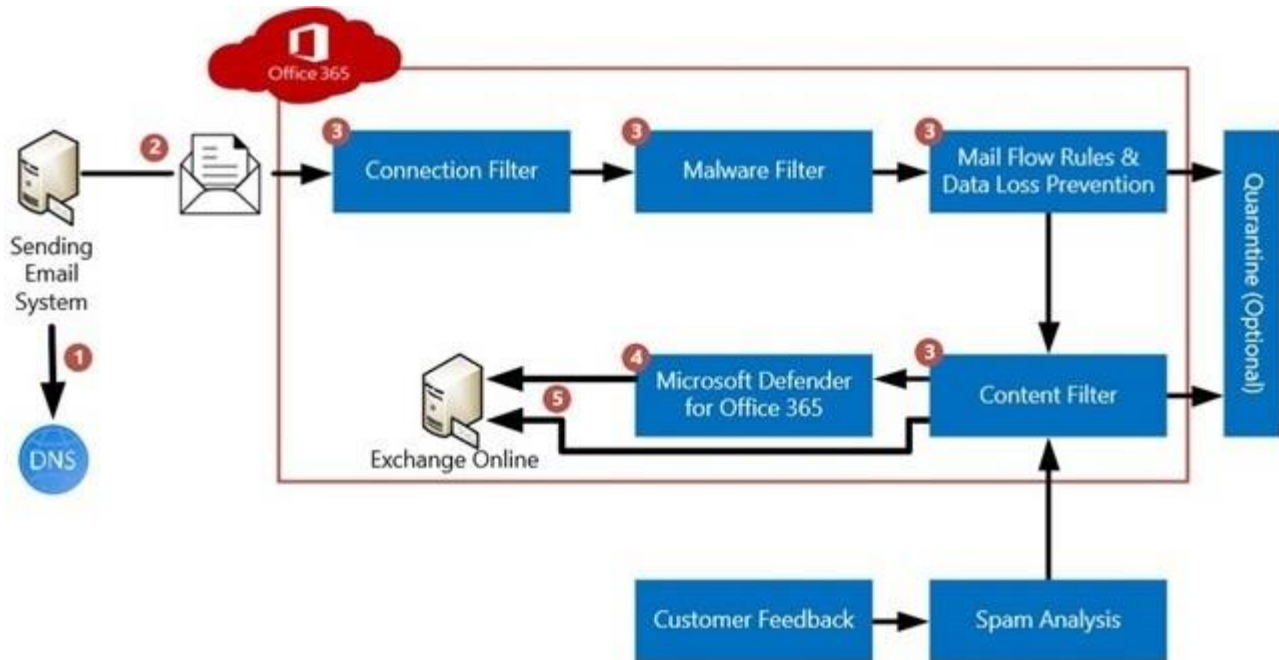


Figure 7-5: How EOP protects email

Figure 7-5 shows the route a message takes through EOP:

1. The sender's mail server looks up the domain MX record information (in our example, the MX record resolves to EOP).
2. The MX record resolves to the data centers in the primary region to which the tenant belongs. The sending mail server then tries to deliver the message to the MX endpoint retrieved earlier. To accept messages from the internet, you do not need to configure anything. The default connector in Exchange Online (which is invisible to the administrator) automatically accepts all messages.
3. Before the message is delivered to the recipient(s), it goes through several layers of filtering:
  - a. **Connection Filtering.** The connection filter is the first layer of defense. It checks several items, including the sender's reputation. In addition, an administrator can define an IP allow/block list to allow or block connections from specific IP addresses. Administrators can also opt into the safe list, a curated list of IPs Microsoft has deemed safe. By enabling this option, all safe list members bypass spam filtering.
  - b. **Anti-Malware filter.** This filter inspects the message for malware and viruses. If a message is deemed malicious, it is removed from the mail pipeline by default. However, it is not rejected at this stage because the sender does not get a notification or non-deliverable response that you might expect when a message fails to reach its target. Configuring anti-malware rules will be covered later.
  - c. **Policy Enforcement.** The policy filter checks messages against configured mail flow rules (including those created by data loss prevention (DLP) policies). If a message matches a rule, the action(s) configured for that rule or policy are applied to the message.
  - d. **Content Filtering** is where the content of messages goes through checks to determine whether the message is spam or phishing based. If a message is considered spam, phishing, bulk, or high-confidence phishing, it can be deleted, sent to the user's Junk Email folder, or quarantined within EOP, depending upon the settings and rules you have configured.
4. If the tenant has Microsoft Defender for Office 365 licenses and the right policies are in place, messages will receive additional scanning before moving on to the next step.
5. If the message was not dropped, rejected, or quarantined, it is delivered to the recipient(s).

**Hybrid mail flow** is handled slightly differently. Messages received from the on-premises organization are marked as "internal" and bypass content filtering (anti-spam). However, if the hybrid mail flow is not correctly configured, internal senders could be treated as anonymous and external to the organization. This could cause messages from internal senders to be incorrectly quarantined, moved to junk, or outright rejected.

Outbound messages are also scanned and evaluated:

1. Outbound messages first go through the **Anti-malware filter**, where they are scanned for known viruses.
2. Next, messages make their way through the **Policy Enforcement** engine. Here, messages are evaluated using configured policies such as mail flow rules, email encryption, and data loss prevention (DLP) rules. If a message matches one of the rules, the action from that rule is applied to the message.
3. The **Content Filtering** process exists to decide if sent messages are spam. A variety of techniques are used, and if a message is likely to be spam, two things can happen:
  - a. The message is routed through EOP's *High-Risk Delivery Pool*. This pool of EOP servers is used to send messages flagged as potentially being spam. This means that regular mail flow will not be affected if one of the IP addresses in the high-risk pool is blocked. Out-of-office messages are also sent via the high-risk delivery pool.
  - b. Depending on the configuration, and the characteristics of the email, it can also be quarantined.
4. If a message is considered safe, it is routed to its destination through the regular outbound pool. Unlike the high-risk pool, these servers are only used to send messages considered safe.

**Queuing:** Sometimes, the mail server to which EOP should deliver messages may be (temporarily) unavailable. Queuing will only happen for "transient errors" such as connection time-outs, refused connections, or other SMTP errors with 4XX error codes. Any "hard failure," such as invalid recipients, authentication mismatch, or failures shown by an SMTP 5XX error code, will generate a non-delivery report (NDR). In hybrid mail flow from on-premises, the Hybrid Wizard changes the mail flow configuration so that some 5XX error codes are downgraded to less severe 4XX error codes.

EOP will automatically try to re-deliver queued messages approximately every five minutes until it can successfully deliver the message. If the message cannot be delivered within 48 hours, the message expires, and an NDR is generated. Reports on messages that have been queued for over one hour are available in the Exchange admin center under **Reports > Mail Flow Report > Queued Messages Report**.

## Zero-hour Auto Purge (ZAP)

Despite all the efforts to fight spam and malware, malicious messages can escape the scrutiny of EOP's multiple scanning engines and be delivered to user inboxes. Many reasons could result in anti-malware code not classifying a message as spam or malware. Sometimes this is because the sender (still) has a good reputation, or perhaps because the spam or malware in the email was so new that the malware signature database in EOP is unable to recognize the vulnerability exploited by the attacker.

Zero-hour Auto Purge (ZAP) retroactively deals with malware, spam, and phishing that penetrates defenses by continuously monitoring EOP for signature updates.

- If a new malware signature is detected that matches a signature in a previously delivered message, ZAP retroactively removes the attachment or message from each user's mailbox.
- If a new phishing or spam signature is detected that matches a signature in a previously delivered message, ZAP retroactively takes the action of the anti-spam policy, which could include moving the



message to quarantine, moving the message to the user's junk email folder, deleting the message, or to take no action.

- The opposite is true for false positives: ZAP can move misidentified messages back into the user's mailbox.

ZAP is enabled by default and can be turned off through PowerShell via the *Set-HostedContentFilterPolicy* cmdlet (for phish and spam) or the *Set-MalwareFilterPolicy* cmdlets (for malware). ZAP will only work if the following conditions exist:

1. The user's junk email filtering is enabled. By default, this is true.
2. The spam filter policy is configured to apply any action except *Add X-Header*. The default is to move messages to the junk email folder.

Note that mail flow rules and end-user inbox rules take precedence over ZAP.

To verify if ZAP is enabled, you can use the following cmdlets:

```
[PS] C:\> Get-HostedContentFilterPolicy | Select PhishZapEnabled,SpamZapEnabled

PhishZapEnabled SpamZapEnabled
-----
                True           True

PS C:\> Get-MalwareFilterPolicy | Select ZapEnabled

ZapEnabled
-----
          True
```

Although the actions are transparent to the end-user, information remains for the administrator in the message's routing information. Look for the mention of **Zero-Hour Auto Purge (ZAP)** in the details of the message trace. If the information is there, it means ZAP intervened.

## Directory-based Edge Blocking (DBEB)

By default, EOP will reject any message addressed to recipients it does not recognize, that is, for which EOP cannot find a matching mail-enabled object in the directory. This feature is known as directory-based edge blocking (DBEB) and offers several benefits, especially when a domain is under a dictionary attack by a spammer. DBEB is enabled for each authoritative accepted domain within Exchange Online and blocks messages at the perimeter network.

When you register a new domain with Microsoft 365, the domain is automatically configured as *Authoritative*, designating Exchange Online as your primary (and only) mail system. All valid recipients should then be represented by an object within Exchange Online, giving the system confidence that every recipient not matched against an existing object should be considered invalid. In turn, DBEB will reject messages addressed to invalid recipients, even before the message reaches the filtering layers. For this reason, messages rejected by DBEB are not visible in the message trace logs.

In a hybrid deployment, the transport service must be able to deal with email-enabled recipients homed on-premises. The same is true in a standalone scenario where EOP protects an on-premises organization. Directory synchronization is then used to synchronize on-premises recipients to Exchange Online. However, directory synchronization does not synchronize all recipient types. For instance, dynamic distribution groups are not synchronized. As such, [without extra configuration](#), messages will be rejected. Two options exist to avoid the dropping of messages for missing recipients:

1. Create a mail contact for the missing recipients. This will ensure that a valid recipient exists in the tenant, and DBEB will accept the message, after which it is forwarded to the on-premises environment.

2. Reconfigure the domain as an *Internal Relay* domain instead of *Authoritative*.

In the latter scenario, DBEB will be disabled automatically. However, messages for recipients not found in Azure AD will be forwarded to the on-premises organization or, if you do not have a dedicated connector, to where the MX record points. Thus, you might also need to configure an outbound connector scoped to the domain name. For instance, if you have configured the domain *office365itpros.com* as an internal relay domain, you should also create an outbound connector, as illustrated in the following PowerShell example:

```
[PS] C:\> New-OutboundConnector -Name "To On-premises"  
-RecipientDomains "office365itpros.com" -ConnectorType "OnPremises"  
-SmartHosts "onpremises.office365itpros.com" -UseMXRecords $False
```

When running the *New-OutboundConnector* cmdlet, it is important to include the *OnPremises* connector type. In this example, the connector is scoped to include only the *office365itpros.com* domain. However, you might have multiple relay domains. If you want to forward all messages for all accepted domains to a specific smart host without specifying each domain individually, you can use the *-AllAcceptedDomains* switch instead. This will ensure that a single connector is used for all known accepted domains.

```
[PS] C:\> New-OutboundConnector -Name "To On-premises" -AllAcceptedDomains $True  
-ConnectorType "OnPremises" -SmartHosts "onpremises.office365itpros.com" -UseMXRecords $False
```

The benefit of using the *AllAcceptedDomains* parameter is that new accepted domains are automatically included in the scope of the outbound connector.

## Preset Security Policies

Preset security policies allow an organization to quickly deploy their message hygiene and Microsoft Defender for Office 365 policies with a few simple clicks. Microsoft provides two predefined security policies: standard and strict. Think of these as policies Microsoft is curating and managing on your behalf. All you need to do is select which users or groups of users you want to include in the preset policies. If your organization needs a custom policy, then the preset security policies will not work for your organization.

You can deploy these policies by navigating to **Policies and rules > Threat Policies > Preset security policies**. Clicking the **Manage protection settings** link (under either *Standard Protection* or *Strict Protection*) allows you to add users to either policy. Users can be added individually, through group membership, or through their primary domain. Once you complete the wizard, these policies are created automatically under Anti-Phishing, Anti-Spam, Anti-Malware, Safe Attachments, and Safe Links. However, unlike custom policies, these policies cannot be edited, deleted, or disabled through the regular policy screens. For example, in Figure 7-6, we can see that the strict protection policy is disabled. However, even though the strict protection policy is disabled, it is still present on the other policy screens. The strict protection policy is only deleted from the other policy screens when all users, groups, or domains have been removed from the policy.

To disable the policy, navigate to the **Preset security policies** page and toggle the switch to the left (disabled). If you would rather delete the policy, click the **Manage protection settings** link under the policy you wish to delete. From the wizard, set the **Apply Protection** checkboxes to **None** on the *Apply Exchange Online Protection* and *Apply Defender for Office 365 protection* pages. Click **Confirm**. The policy toggle will now be greyed out and removed from all policy screens.

You can learn more about preset security policies and how their management differs from Microsoft's [documentation](#).

Policies & rules > Threat policies > Preset security policies

## Preset security policies



<p><b>Standard protection</b></p>  <p>A baseline protection profile that protects against spam, phishing, and malware threats.</p> <ul style="list-style-type: none"> <li>✓ Balanced actions for malicious content</li> <li>✓ Balanced handling of bulk content</li> <li>✓ Attachment and link protection with Safe Links and Safe Attachments</li> </ul> <p><input checked="" type="checkbox"/> Standard protection is on</p> <p><a href="#">Manage protection settings</a></p>	<p><b>Strict protection</b></p>  <p>A more aggressive protection profile for selected users, such as high value targets or priority users.</p> <ul style="list-style-type: none"> <li>✓ More aggressive actions on malicious mail</li> <li>✓ Tighter controls over bulk senders</li> <li>✓ More aggressive machine learning</li> </ul> <p><input type="checkbox"/> Strict protection is off</p> <p><a href="#">Manage protection settings</a></p>
---	--

Figure 7-6: Preset Security Policies in EOP

## Anti-Spam Protection

As defined in Wikipedia, email spam is unsolicited, undesired, or illegal email messages. It should come as no surprise that the various anti-spam features in EOP work together to:

1. Reduce the amount of spam that is delivered to a user's inbox; and
2. Reduce the number of false positives; messages are marked as spam but are not.

When an organization uses Exchange Online, spam filtering for inbound messages is enabled by default and cannot be disabled. However, an administrator has a multitude of options to adjust the settings for the different anti-spam features, including the ability to create custom antispam policies for different groups of users.

### Anti-spam Inbound Policy

Inbound spam filtering can be accessed from the **Defender portal** and navigating to **Policies and rules > Threat Policies > Anti-spam policies**. The *Anti-spam policies page* lists the default three policies (which cannot be disabled) and any custom anti-spam policies you have defined.

To create a new policy, select **Create Policy > Inbound**. This will launch a wizard to guide you through the policy creation. The elements of that wizard are the same as if you edit a custom policy. We will cover that next.

To customize an existing policy, select that policy, and the properties pane will be displayed. The properties pane lets you customize what constitutes spam and what actions to take on that spam.

From the policy properties pane, you can perform the following actions:

- Enable or disable the policy with the **Turn on** or **Turn off** buttons. These buttons are absent for the default policy as it is always enabled.
- Change the processing order of the policy with the **Increase Priority** or **Decrease Priority** buttons. These buttons are absent from the default policy because it always has the lowest priority.
- Delete the policy with the **Delete policy** button. The default policy cannot be deleted.

- The **Edit description** link allows you to change the name and description of the policy. The description field is a great way to document any changes you have made to the policy. Note that the default policy cannot be renamed, so the name field will be greyed out if you are editing the default policy.
- **Edit users, groups, and domains** allow you to define the scope of the policy. This section is not present on the default policy because the default policy is the catch-all for anyone not defined in a custom policy. Custom policies have priority over the default policy and are always processed first. If you have multiple custom policies that scope the same set of users, then the policies are processed in order of priority, where the lowest number (zero) is processed first.
- **Edit spam threshold and properties** allow you to configure what constitutes spam or bulk email (we will cover this in the next section).
- **Edit actions** define what actions to take on messages marked as spam, high confidence spam, phishing, high confidence phishing, and bulk mail. These actions are as follows:
  - **Move message to Junk Email Folder** delivers the message to the user but deposits the message into the user's junk folder. This could be a helpful training tactic as it may prompt the user to scrutinize the source and content of the message.
  - **Prepend subject line with text** delivers the message to the user but modifies the subject line. This is useful if you want to provide a visual cue to the user by prepending the subject line with BULK or SPAM as an example. Like moving the message to the user's junk folder, this is also a helpful training tactic to get the user to scrutinize the message.
  - **Quarantine message** prevents delivery of the message to the user and instead delivers the message to a quarantine. A quarantine policy dictates how the user interacts with a message in the quarantine. You can define a different quarantine policy for different types of mail. For example, the quarantine policy on high confidence phish could be much more restrictive than the policy on bulk email. Quarantine policies are discussed later in this chapter.
  - **Redirect message to email address** prevents delivery of the message to the user and instead delivers the message to someone else. This option is helpful for egregious phishing that you want to redirect to a SecOps mailbox for analysis.
  - **Add X-header** adds a header to the message. This header can be targeted by mail flow or inbox rules for additional actions.
  - **Delete message** does precisely as the name implies. The message is deleted and cannot be recovered.
  - **No action** does precisely as the name prescribes. The message is forwarded to the user without modification.

**Note:** Not all actions are available to each spam or phishing type. For example, *No action* is only available to an email defined as bulk. On the other hand, the only actions for high confidence phish are Quarantine and Redirect messages.

- **Retain spam in quarantine for this many days** defines how long messages are kept in quarantine if the quarantine action and a quarantine policy have been selected.
- **Enable spam safety tips** define whether safety tips should be displayed to users in an Outlook client. This is useful if you deliver spam to the user's mailbox (e.g., prepend subject or move to junk) but want to provide additional warnings about the message.
- **Enable zero-hour auto-purge (ZAP)** allows Exchange Online to reach into a mailbox and take retroactive action on previously delivered mail. This is useful for zero-day exploits that are only identified after the delivery of mail to a user's inbox. You can also choose whether to

enable ZAP for phishing messages, spam messages, or both. It is best practice to do the latter.

- **Edit allowed and blocked senders and domains** allow an administrator to define safe and blocked senders. These can be in the form of individual senders or an entire domain (we will cover this in a later section).

## Spam Thresholds and Properties

As mentioned in the previous section, an administrator can fine-tune what constitutes spam and bulk mail by clicking the **Edit spam threshold and properties** link in an Anti-spam inbound policy.

The first control is a **Bulk email threshold** slider. This slider dictates at what level bulk email should require an action to be taken. The recommended value is 7, on a scale of 1 through 9. This means any bulk email with a rating of 7 or higher is subject to the bulk email action defined in the policy.

The rating is also known as a Bulk Complaint Level (BCL). EOP applies a BCL rating on every message and applies that rating in the **X-Microsoft-Antispam** header in a field called "BCL." The field can have the values described in Table 7-3.

<b>BCL Value</b>	<b>Meaning</b>
0	The message is not from a bulk email sender.
1,2,3	The message is from a known bulk sender but is unlikely to generate many complaints.
4,5,6,7	The message is from a known bulk sender and can generate many complaints.
8,9	The message is from a known bulk sender and is likely to generate a high number of complaints.

Table 7-3: BCL header values

The inbound spam policy lets you increase the spam score on messages if any of the following are enabled:

- Image links to a remote website (rather than using embedded images)
- URLs containing IP addresses (rather than using a registered domain name)
- URLs that redirect to another port (rather than using the standard HTTP or HTTPS ports)
- URLs that use a top-level domain of BIZ or INFO

The inbound spam policy lets you mark messages as spam if any of the following are enabled:

- Messages with no content
- HTML code with IFRAME, FORM, EMBED, or OBJECT tags
- Messages with Javascript or VBScript
- Offensive words in the subject or body (this is a curated list by Microsoft that cannot be modified)
- Senders that hardfail SPF
- Senders that hardfail sender ID
- Backscatter (we discuss Backscatter later in the chapter)
- Messages received in specific languages or from specific countries

Each of these settings can be configured as **On** or **Off**. You should consider each of these options with care. Changing a setting might increase the likelihood of a message being marked as spam which tends to have a more significant impact on the productivity of your end-users than the occasional spam message.

Some settings also have the option to be configured in **Test** mode. Configuring a setting in test mode can take additional actions such as adding an X-Header or BCC the message to a mail recipient. Testing is a great way to determine what these settings are operating in the way you expect before fully enabling them. While a setting is in test mode, the original email will be delivered to the user's inbox.

## Spam Allow and Block Lists

You may have a scenario where you need to allow or block a sender or domain globally for the entire organization or a group of users. To allow or block a sender or domain globally for the entire organization, you can modify the **Anti-spam inbound policy (Default)**. From the pop-out window, select **Edit allowed and blocked senders and domains**. You can add entries to the allowed senders, allowed domains, blocked senders or blocked domains list. You can also use the *Set-HostedContentFilterPolicy* cmdlet for this purpose.

If you want to scope these allow or blocklists to a specific group of users and not the entire organization, you will need to create a new antispam policy. From the **Threat Policies > Anti-spam policies** page, click **Create Policy > Inbound**. Give the policy a name and description and click **Next**. The *Users, groups, and domains* page allows you to scope this new policy to either one or more users, one or more groups, or one or more of your accepted domains. The next series of pages will ask you to specify the antispam policy properties, antispam actions, and finally, the allow and block lists unique to this policy, and by extension, the users scoped to the policy.

When building allow lists, remember that Microsoft's [Secure by Default](#) approach will disregard customer-defined entries if the message is deemed to contain malware or is classified as high confidence phishing. Instead, these messages will be delivered to the quarantine regardless of their presence in an allow list or modification by a transport rule (e.g., changing the message header to SCL -1 to bypass filtering)

**Real-world:** If migrating from another filtering solution to EOP, I recommend not migrating your allow and blocklists. First, based on the age of these lists, it is questionable which entries are still valid. Second, there is no direct apples-to-apples comparison regarding filtering solutions. Each solution is different. An action taken by one solution may not have been taken by another resulting in redundant entries in your allow/blocklists. When migrating to EOP, this is a good time to start with a clean slate.

## Connection Filter Policy

The connection filter policy allows administrators to maintain a list of allowed and blocked IP addresses. You can access this IP list from the **Defender portal** and navigating to **Policies and rules > Threat Policies > Anti-spam policies > Connection filter policy (Default)**.

It is not possible to create a custom connection filter policy. This means the IPs in the connection filter policy impact all users. Therefore, it is impossible to scope IPs to specific users, groups, or domains like you can with other policies.

To add or remove an IP address, select the **Connection filter policy (Default)**. From the pop-out window, select **Edit connection filter policy** and add the individual IP addresses or address ranges.

Selecting the **Turn on safe list** checkbox includes a list of safe IPs vetted by Microsoft to your connection filter rules.

## Anti-spam Outbound Policy

To stay on top of situations where an internal recipient sends spam messages, an administrator can change the default outbound policy to generate notifications when EOP deems a message suspicious or when a user is blocked from sending messages. In the **Defender portal**, go to **Policies and rules > Threat Policies > Anti-spam policies**. Here you will find any custom policies that have been created, as well as the default policy set. Note that you cannot disable any default policies; however, you can modify their settings.

In addition to the default outbound spam filter policy, you can create custom outbound spam policies to set different notification addresses and sending limits and apply them to specific senders. To create a new outbound policy, click **Create policy** and select **Outbound** from the drop-down menu.

Another reason for custom outbound policies is to define which users or groups can automatically forward messages outside the tenant. This policy does not impact automatic forwarding to internal recipients. Automatic forwarding is any messages automatically forwarded via an inbox rule, mail flow rule, or mailbox forwarding. This is an instrumental setting for organizations concerned with data exfiltration and needs to block some or all users from automatically forwarding messages outside the organization. You can set this policy to **On – Forwarding is enabled** or **Off – Forwarding is disabled**.

The recipient limits in the outbound spam filter policies should not be confused with tenant-wide customizable recipient limits. Recipient limits are the maximum number of recipients you can add per message, with a default of 500 and a range of 1 to 1000 that can be set per mailbox. The outbound spam filter policy allows you to set a maximum number of recipients per hour or per day across all the emails sent by that user in the time interval. If users exceed this number, they can be blocked from sending mail until the following hour or day, or an alert can be generated. The maximum number of recipients you can set per interval of an hour or day is 10,000. When using the default value ("0"), Microsoft determines the number of messages that will trigger the action based on activity rather than a fixed number.

When Exchange Online blocks a user that continuously sends emails that it classifies as spam, the user will receive NDRs for outgoing messages that provide specific information on what they need to do to unblock their account. Blocked accounts show up in the Defender portal under **Review > Restricted Entities**. From there, an administrator can review blocked accounts and unblock them. You can also configure the outbound malware settings in the malware policy to notify administrators when internal users are blocked.

You can also unblock an account by running the *Remove-BlockedSenderAddress* cmdlet. For example:

```
[PS] C:\> Remove-BlockedSenderAddress -SenderAddress Joe.Bowers@office365itpros.com -Reason "Account OK"
```

Note that an administrator can only unblock an account so many times. If the limit for an account has been reached, even the administrator will receive an error, and a support ticket must be opened for investigation.

## Anti-Spam Message Headers

EOP updates the Spam Confidence Level (SCL) header to show whether it considers a given message as spam. The SCL header consists of a single digit that does not explain why a message is considered spam. To provide administrators with additional information, EOP also inserts [two extra headers](#) into each message:

- **X-Forefront-Antispam-Report** is used to provide information on the anti-spam processing of the message; and
- **X-Microsoft-Antispam** provides information specific about bulk mail and phishing results.

The *X-Forefront-Antispam-Report* header consists of many different values, each revealing more information about the message, such as where it was sent from and the result of individual anti-spam tests. The following shows what this header typically looks like:

```
X-Forefront-Antispam-Report:
CIP:199.59.150.72;IPV:NLI;CTRY:US;EFV:NLI;SFV:NSPM;SFS:(8156002)(31570200002)(2980300002)(438002)(286005)(199004)(189003)(64016003)(53416004)(2616005)(19627405001)(110436001)(956004)(476003)(126002)(733005)(18926415007)(118246002)(85226003)(53386004)(606006)(336012)(6486002)(58536013)(106002)(106466001)(2160300002)(36756003)(246002)(16586007)(36736006)(486006)(33656002)(356003)(8676002)(7696005)(26005)(1096003)(620700001)(59450400001)(966005)(551544002)(33964004)(7636002)(84326002)(6306002)(25786009)(53946003)(236005)(4290100001)(146002)(7596002)(16003)(6916009)(270700001)(579004);DIR:INB;SFP:;SCL:1;SRVR:VI1PR0301MB2318;H:spruce-geese-ac.twitter.com;FPR:;SPF:Pass;LANG:en;PTR:spruce-geese-ac.twitter.com;A:0;MX:1;
```

At first sight, it may seem hard to make some sense of the information in the header. However, if you take a closer look, you will see that the header consists of several different fields of which the following give more insight into the anti-spam decision-making process:

- **CIP** has the IP address of the server that delivered the message to EOP. This IP address should be specified when using the IP allow or block list, and the IP address should be listed in the senders' SPF record.
- **IPV** field has two values and is used to decide whether the connecting IP was found on an IP allow list or not.
  - **IPV:CAL**. The message passed anti-spam filtering because the connecting IP address was found on an IP allow list.
  - **IPV:NLI**. The IP address was not found on any IP reputation list.
- **CTRY** specifies the likely country from where the message was received. EOP uses the connecting IP address to determine the country, which may or may not be the same as the original message.
- **SFV** field specifies why a message was marked as spam or not as spam and has several values:
  - **SFV:SFE**. Filtering was skipped, and the message was let through because it was sent from an address on an individual's safe sender list.
  - **SFV:BLK**. Filtering was skipped, and the message was blocked because it was sent from an address on an individual's blocked sender list.
  - **SFV:SPM**. The message was marked as spam by the content filter.
  - **SFV:SKS**. The message was marked as spam before the content filter processed the message. This includes the scenario when a mail flow rule marks a message as spam.
  - **SFV:SKA**. The message skipped filtering and was delivered to the inbox because it matched a safe sender or safe domain list in an anti-spam policy.
  - **SFV:SKB**. The message was marked as spam because it matched an anti-spam policy's blocked sender or blocked domain list.
  - **SFV:SKN**. Before the content filter processed the message, the message was marked as not spam. This includes the scenario when a mail flow rule is used to mark a message as safe.
  - **SFV:SKI**. The content filter skipped processing the message because it was received by the on-premises environment and thus marked as *internal*.
  - **SFV:SKQ**. The message was released from quarantine and was sent to the intended recipients.
  - **SFV:NSPM**. The message was not spam and was sent to the intended recipients.
- **SCL** gives the Spam Confidence Level of the message and denotes the likelihood of the message being spam. The higher the number, the more likely the message is spam.
- **H** specifies the HELO or EHLO string of the connecting mail server.
- **SPF** specifies information about the SPF record lookup result for the message.
- **LANG** specifies the language the message was written in.
- **PTR** identifies the PTR record of the sending IP address (also known as the reverse DNS address).
- **ARC** shows the Authenticated Received Chain headers.
- **CAT** shows the category of the protection policy applied to the message. Multiple policies process a given message, but only the value corresponding to the highest priority one will be stamped in this header, as detailed in the [documentation](#).
- **SRV:BULK** specifies that the message was marked as a bulk email message.
- **SFTY** identifies if the message is a phishing message. The "safety" header value indicates the type of phishing, such as a URL, internal phishing, domain impersonation, etc. SFTY will also indicate if the setting was overridden with a safe sender or domain setting so you can determine if an email was allowed through because of your existing overrides.

The **X-Microsoft-Antispam** header looks like this:

```
UriScan:;BCL:1;PCL:0;RULEID:(7020095)(5600026)(4605076)(4608076)(1401150)(8001031)(1402068)(71702078);SRVR:VI1PR0301MB2318;
```

And this header shows the following components:



- **BCL** specifies the Bulk Complaint Level of the message on a scale of 0-9. If the value is 8 or 9, the email comes from a sender that generates many complaints.
- **PCL** specifies the Phishing Confidence Level of the message on a scale of 0-9. If the value is 0-3, the email is unlikely to be phishing. An email with a value of 4-9 is likely to be a phishing email.

## Message Tracing and Anti-Spam Headers

Often message tracing is used to understand why a message was marked as spam. You must request an Extended Trace report to receive details on why the message was marked as spam. The returned CSV file holds some extra data, including the anti-spam headers explained previously.

Once the search completes, open the CSV file and look at the *custom\_data* field. Note that the example below was trimmed at the end for brevity.

```
S:AMA=SUM|v=0|action=|error=|atch=0;S:AMA=EV|engine=M|v=0|sig=1.193.3192.0|name=|file=;S:AMA=EV|engine=A|v=0|sig=201503191928|name=|file=;S:AMA=EV|engine=K|v=0|sig=19.3.201518:28:0|name=|file=;S:CFA=AS|sfv=NotSpam|rsk=Low|sc1=0|bc1=0|score=|sfs=(601004)|sfp=0|fprx=|mlc=|mlv=|list=1|di=|rd=mail-db3on0089.outbound.protection.outlook.com|h=emea01-db3-...
```

The field contains helpful information regarding the various filters that processed the message. For example, the data from the anti-malware agent is displayed after **S:AMA** (Anti-Malware Agent). Similarly, the anti-spam information is displayed after the **S:SFA** (Spam-Filtering Agent). Finally, information on mail flow rules is shown after **S:TRA** (Transport Rule Agent). Once you have extracted the data, you can use the information explained in the anti-spam message headers topic to understand the decisions of the anti-spam agent.

## Anti-Malware Protection

EOP automatically scans inbound messages for malware. The term malware covers a variety of malicious items such as viruses and spyware. EOP uses a multi-layered approach to detect malware using multiple anti-malware engines to scan messages for malicious code in the message body or attachments.

Unlike anti-spam configuration options, you cannot alter how malware is processed. However, you can control what type of attachments are denied automatically or how users are notified when an email contains malware or a blocked attachment type. The former option is also referred to as **Common Attachment Filtering**. By default, EOP maintains a list of common attachment file types often associated with malware. Among other file types, the default attachment filter will block all .exe, .vbs, and .reg files. By editing the default anti-malware policy or creating a new one, you can add or remove file types to control whether emails containing such attachments are quarantined automatically.

**Note:** EOP blocks attachments in other places than the common attachment filtering settings. For example, OWA extends the list of file types blocked. This block does not affect how Outlook or Outlook Mobile works. Adding additional extensions to the common attachment filter stops these attachments from reaching the user, regardless of their client.

### Custom Malware Policy

The default malware policy is sufficient to suppress malware for most tenants. However, if you need a different policy for a subset of your users, you can create a new malware policy. To create a new policy, navigate to the **Defender portal > Policies and rules > Threat Policies > Anti-malware** and click **Create**.

You can scope the custom policy to a subset of users based on the following conditions:

- **The recipient is** (select a single- or multiple recipients).
- **The recipient domain is** (all recipients within the selected domain(s)).
- **The recipient is a member of** (select group).

Custom policies always precede the default policy and are processed in ascending order. Therefore, the lower the policy priority, the sooner it will be processed.

Like the default policy, the custom policy lets you define common attachment types to block. This is useful if you want to have different common attachment blocks for different groups of users. For example, you may block fewer attachment types for your IT department than your general user population. By default, a new custom policy will have 10 executable types preselected. But you can select up to 96 common attachment types to block. If you need to block additional attachment types not included in the common attachment type filter, you will need to define a transport rule.

Apart from configuring a common attachment filter, you can also configure administrator notifications and toggle the ZAP feature (although it is best to keep this enabled).

**Note:** To notify users of quarantined malware, you must deploy a quarantine policy. A quarantine policy allows you to define user notifications and can be attached to the anti-malware policy.

## Anti-Phishing

Phishing is an attempt to steal something of value from a company. This could range from the exfiltration of sensitive information (such as credit card numbers or personally identifiable information) to encouraging someone to take action that will result in some loss for them. Phishing typically involves impersonation where a bad actor uses an email address that mimics a vendor, partner, customer, or employee to create trust. A typical example is a bad actor using an email address that impersonates a C-level executive requesting money be wired to what will invariably be the attacker's bank account.

The Anti-Phishing policy has several measures to protect companies from these attacks. Like other EOP policies, you manage the anti-phishing policy through the Defender portal. Navigate to **Policies and rules > Threat Policies**, then select **Anti-phishing**. From there, you can view and edit the default policy or click **Create** to make a new policy. The Default policy will always apply and cannot be deleted.

When configuring a policy, you have the following options:

- **Name and Description.** As you can create multiple policies for your organization, it is a clever idea to use a clear name or at least describe the intent of the policy.
- **The Users, Groups, and Domains.** This page specifies the users the policy will be run against. This can be set to all your users or a subset of your users, like all recipients in a specific domain or members of a distribution group. You can also exclude specific users from the policy.

Once these general settings are configured, you can configure the phishing threshold and protection settings for the policy:

- **Phishing email threshold** determines how aggressive machine learning should be when determining what a phish is. Values are from 1 – Standard to 4 – Most Aggressive. The Standard preset security policy (covered earlier) uses a value of 2. However, Microsoft deems 1 to be an appropriate default value. Therefore, it is best to start the threshold at 1 and gradually increase it to meet your organization's needs.
- **Enable users to protect** and **Enable domains to protect** define what email addresses or domains are scrutinized while scanning for incoming phishing attacks. Typically, it would be best to protect high-value users like your C-level executives or people with high privileges within your organization. You can protect up to 350 users (depending on license). To protect all the domains assigned to the tenant, check **Include domains I own**. You can also **Include custom domains** that allow you to add a domain not included in your tenant or that of a partner, vendor, or customer.
- To ensure trusted domains or senders, such as a partner, vendor, or customer, are never treated as impersonators, you can optionally add them to the **Manage trusted senders and domains** allow list.

- Enable/disable **Mailbox Intelligence (Recommended)** to allow the policy to use Microsoft Graph for user email signals. This helps ensure that people you have communicated with before are not treated as spoofed senders.
- Enable/disable the use of **Intelligence for impersonation protection (Recommended)** to enable enhanced impersonation. This also allows actions to be taken on emails impersonating a user.
- Enable/disable **Spoof Intelligence (Recommended)**, which allows you to define who can send mail on behalf of your domains (for example, a bulk mailing service). While the option to enable this setting is part of the anti-phish policy, the allow list is maintained through **Policies & rules > Threat policies > Tenant Allow/Block Lists > Spoofing**.

**Real-world:** The terminology in the wizard is a little ambiguous. “Enable users to protect” and “Enable domains to protect” do not define to whom a policy applies. Instead, these are the email addresses and domains for which anti-phishing will compare the P2 header information of incoming messages. So, for example, you configure ceo@office365itpros.com as a user to protect and office365itpros.com as a domain to protect and apply the policy to all users in your organization. Then, whenever a message is received from ceo@office365itpros.com (note the look-a-like address), that message will be flagged as a phishing attempt if the person receives an email from that address for the first time.

With the phishing thresholds defined, we can now turn to the actions in our policy. An action is performed whenever a phishing threshold has been met.

**Note:** Some actions may be greyed out if you did not enable the corresponding phishing threshold on the prior page of the wizard. For example, “If message is detected as impersonated user” is greyed out, click the **Back** button, select **Enable users to protect**, and define a list of users to protect.

The following actions are available:

- **If message is detected as an impersonated user, impersonated domain, or if mailbox intelligence detects an impersonated user**, you can perform the following actions from the respective drop-down.
  - **Redirect message to other email addresses.** This is useful if you need to send the message to an internal security response team for analysis. The redirection will not deliver the message to the original recipient.
  - **Move message to the recipients’ Junk Email folders** does precisely as the action describes. With proper email security training, delivering the message to the junk email folder is a great way to have users pause and scrutinize the email. This ensures mail is still delivered in the case of a false positive, which could be reported by the user while instilling caution that the message ended up in junk.
  - **Quarantine the message** takes the junk folder concept one step further by keeping the message out of the user’s mailbox and off their devices. The added benefit of the quarantine is that it is out of sight and therefore out of mind. Generally, users only review the quarantine to look for a solicited message they have not received. As phishing emails should be unsolicited, the user will not miss the email they did not expect to receive. Depending on the quarantine policy, proper email security training is still needed as users can release the message to their inbox.
  - **Deliver the message and add other addresses to the Bcc line** sends the suspect message as a blind carbon copy to an internal security response team and delivers it to the intended recipient. This is useful when you have an internal security team to review (mildly) suspicious messages but do not want to delay delivery to the original recipient. Note that this action delivers to the original recipient’s inbox.

- **Delete the message before it's delivered** deletes the message during transport. This is a definitive action, and the message is irrecoverable. Any analysis must be done using the information in message trace logs.
- **Don't apply any action** means that no action applies to the message.
- **If the message is detected as spoof**, only has the actions **Move message to the recipients' Junk Email folders** and **quarantine the message** from the list above. This means you must take corrective action on a spoofed message.
- **Safety Tips and indicators** determine what the user may see in their Outlook client. Safety Tips are available across all Outlook clients, including Outlook for Windows, Outlook for Mac, Outlook on the Web, and Outlook Mobile.
  - **Show first contact safety tip (Recommended)** identifies when you receive an email from someone you know but are using a different email address than what EOP has seen before.
  - **Show user impersonation safety tip, and Show domain impersonation safety tip** informs the user whenever a user or a domain is being impersonated by the sender.
  - **Show user impersonation unusual character safety tip** informs a user when an unusual character is detected in the sender's address. An example of this is a [homoglyph attack](#). A homoglyph attack is when a bad actor swaps a character from one alphabet for a similar or identical-looking character from another. For example, a bad actor may swap out the letter "e" from the Latin alphabet for the letter "e" from the Cyrillic alphabet (Unicode 435). To the human eye, these look identical but are very different.
  - **Show (?) for unauthenticated senders** changes the picture (or initials in the absence of a picture) in the sender's profile card to a question mark. This is when Exchange Online cannot authenticate the sender via sender authentication (e.g., SPF, DKIM, & DMARC). This is explained in greater detail in the *Safety Tips* section below.
  - **Show "via" tag** shows a safety tip if the from address (displayed to the user in their email client) does not match the mailfrom attribute in the message header. An example of this safety tip could be [john.smith@contoso.com](mailto:john.smith@contoso.com) on behalf of [jane.doe@office365itpros.com](mailto:jane.doe@office365itpros.com).

**Note:** The first contact safety tip previously required the presence of a transport rule to apply *X-MS-Exchange-EnableFirstContactSafetyTip* to the email header. While Exchange Online still supports this header, it is no longer required to display a first contact safety tip.

Like other policies, you can also manipulate the policy with PowerShell using the *Get-AntiPhishPolicy* and *Set-AntiPhishPolicy* cmdlets. For example, run the following command to retrieve a list of all anti-phishing policies.

```
[PS] C:\> Get-AntiPhishPolicy
```

To enable mailbox intelligence for the default policy, use the following command.

```
[PS] C:\> Set-AntiPhishPolicy -EnableMailboxIntelligence $True -Identity "Office365 AntiPhish Default"
```

If you enabled an anti-phishing policy when Microsoft first released the feature, a default policy called *"Office365 AntiPhish Default"* is in your tenant. In these circumstances, it is probably best to update the default policy and remove your policy and let the default policy do the work.

## Anti-Spoofing

Spoofing is when someone sends messages using your email domain through an email system other than your primary (internal) messaging system. Often a malicious attacker tries to impersonate a legitimate person's email address to convince the recipient of the message to perform actions helpful to the attacker. The concept of spoofing is not new, and over the last few years, another form called "insider spoofing" has emerged and quickly gained popularity with attackers. Insider spoofing sometimes referred to as a spear-

phishing attack, is no different from regular spoofing, except that the message appears to be coming from an internal recipient, making it much harder for the recipient to figure out if a message is invalid or not. Typically, insider-spoofing impersonates highly-ranked employees (such as a C-level executive) and targets other employees to convince them to take action, such as sending a wire transfer.

Fighting spam and phishing is an ongoing task. A newer form of insider spoofing is where user accounts are compromised (maybe due to a phishing attack), and the spoof emails are sent from compromised accounts. The malicious actor logs in as the user and sends an email from a compromised user's real mailbox. Protection against this attack includes using multi-factor authentication with appropriate password policies and protections.

Although features such as SPF records, DKIM signing, and DMARC help to detect and repel spoofing attacks, only a minority of organizations have successfully adopted some or all these features, so gaps are left for attackers to exploit.

To detect spoofed messages, Microsoft processes each incoming message, inspects the various TO and FROM headers in the message, and compares the values. For example, several checks are performed if a message is sent to someone inside your organization and the FROM field matches an internal domain. If the message was sent from inside the organization, or the message was received through a host or service which may send messages on your behalf (for example, because the host is present in the SPF record), the message is deemed to be legitimate. However, if the message does not pass any of these tests, or if it was received by a host with a bad reputation, the message could be considered spoofed.

When the anti-spoofing feature intervenes to mark a message as spam, the value **SFTY:9.5** is added to the *X-Microsoft-Antispam* header, which also has information about the results of other scanning engines:

```
X-Microsoft-Antispam: UriScan:;BCL:0;PCL:0;RULEID: (71701004) (71702002);SRVR:BY2PR12MB0565;SFTY:9.5
```

For more information on the X-Microsoft-Antispam header, check [this article](#).

Besides the anti-spoofing protection described above, Microsoft offers the Spoof intelligence feature.

## Spoof Intelligence

There can be legitimate reasons that justify the spoofing of your email domain(s). For instance, if you have hired an external marketing company to send electronic surveys to your customers or employees, or if you have a business application that sends notification emails to your internal users. However, differentiating between legitimate and malicious attempts to spoof your domain is not easy, especially because features like SPF, DKIM, or DMARC are often not used or poorly implemented. It is also true that none of these features protect your tenant when accounts are compromised.

Spoof intelligence controls which senders can send messages on behalf of your organization. The Spoof Intelligence dashboard is located at <https://security.microsoft.com/spoofintelligence>. Alternatively, you can access the dashboard by opening the **Defender portal** and navigating to **Policies & rules > Threat Policies > Tenant Allow/Block Lists > Spoofing > View Spoofing Activity**. Using the Spoof Intelligence dashboard, you can allow or revoke a specific sender's permission to spoof your domain. Additionally, when you open the dashboard, you can see a list of senders known to have sent messages on behalf of others (including your domains) in the past 7 days.

Additional details help you determine if a sender should be allowed to spoof messages. For instance, the policy tries to show who the actual sender is next to the information about which users were spoofed. EOP obtains the specific sender information by looking at the reverse DNS (PTR) record of the sending server's IP address. If no PTR record is found, the IP address is displayed in the report.

**Note:** Legitimate senders often have a PTR record that enables Exchange Online to look up and display the sender's hostname or domain name information. Extra caution is advised if no PTR record is found, as

it is highly likely that the sender should not be allowed to spoof your domain. Unfortunately, even malicious senders can have a PTR record!

Although you cannot stop a (malicious) sender from trying to spoof your domain, the Spoof Intelligence feature controls how EOP handles incoming emails. If a sender is not explicitly allowed to spoof one or more users within your organization (domain), messages from that sender will be marked as spoofing attempts, and users will be notified of the fact.

From the Spoof Intelligence dashboard (Figure 7-7), you can change how EOP handles the spoof by selecting **Allow to spoof** or **Block from spoofing**.

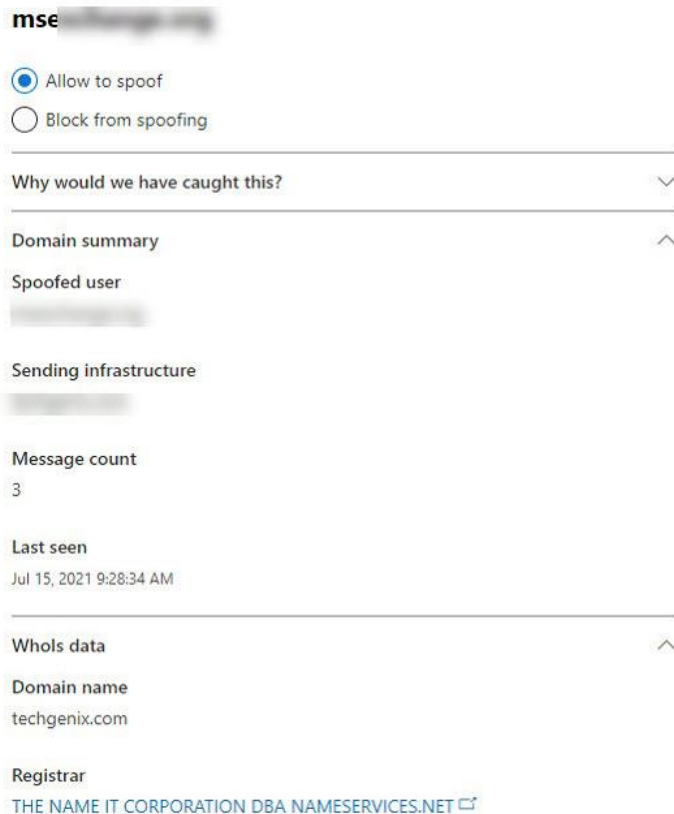


Figure 7-7: Spoof Intelligence Dashboard Action Pane

You can also control the Spoof Intelligence through PowerShell. This is done via the Tenant Allow/Block List (TABL) cmdlets. To get a list of current entries in the Tenant Allow/Block List, we leverage the *Get-TenantAllowBlockListSpoofItems* cmdlet. In the example below, we have two static entries in our list. These two entries identify that *contoso.com* and *fabrikam.com* have been allowed to spoof *office365itpros.com*.

```
[PS] C:\> Get-TenantAllowBlockListSpoofItems | Format-Table SpoofedUser, SendingInfrastructure, SpoofType, Action, Identity
```

SpoofedUser	SendingInfrastructure	SpoofType	Action	Identity
office365itpros.com	contoso.com	External	Allow	7323b926-3334-8d4a-ee09-b94e35f3467
office365itpros.com	fabrikam.com	External	Allow	24f1b37a-766e-0cc1-ca88-7e487732914b

Let's change *contoso.com* to a block rather than an allow. To change an existing entry in the Tenant Allow/Block List, we can leverage the *Set-TenantAllowBlockListSpoofItems* cmdlet. In the example below, we take the ID for *contoso.com* (retrieved from the *Get-TenantAllowBlockListSpoofItems* output above) and specify a block action against the *office365itpros.com* domain.

```
[PS] C:\> Set-TenantAllowBlockListSpoofItems -Identity office365itpros.com\Default -Action Block -
Ids 7323b926-3334-8d4a-ee09-b94e35f3467
```

After making the changes, you can use the *Get-TenantAllowBlockListSpoofItems* command again to verify that the changes were applied correctly. If you need to add a new entry, rather than modify an existing entry, you leverage the *New-TenantAllowBlockListSpoofItems* cmdlet. Similarly, you can remove an entry with the *Remove-TenantAllowBlockListSpoofItems* cmdlet.

In addition to the anti-spoofing features designed to stop spoofed messages from being delivered to the user inboxes, administrators can also request the **Spoof Mail Report**. The report is accessible through the Defender portal or PowerShell. To get the report in PowerShell, use the following command:

```
[PS] C:\> Get-SpoofMailReport | Select Date, Direction, Action, SpoofedSender, TrueSender, SenderIP

Date           : 14/04/2016 0:00:00
Direction      : Inbound
Action         : GoodMail
SpoofedSender  : Boss.Man@office365itpros.com
TrueSender     : hubspot.com
SenderIp       : 50.31.57.0/24

Date           : 11/04/2016 0:00:00
Direction      : Inbound
Action         : CaughtAsSpam
SpoofedSender  : CEO-of-the-company@office365itpros.com
TrueSender     : someonlinemarketingservice.com
SenderIp       : 1.2.3.4/24
```

The report displays information about spoofed messages, like the information in the Spoof Intelligence policy. In addition, it can show you what spoofed messages were received, whom they were impersonating, and who sent the message. In the above example, the top message was sent by an authorized service, as that message was classified as **GoodMail**.

The report serves multiple purposes. First, it helps you understand how many messages inside your organization are spoofed. More importantly, it tells you which account is being spoofed most and can reveal spear-phishing attempts. Secondly, the report enables you to generate a list of hosts sending messages on your behalf. You can then use this list to verify if authorized senders are correctly configured and whether you have represented them on your SPF records, lowering the likelihood of them being marked as spam.

## Message Quarantine

By default, Exchange Online protection routes low confidence spam and phish to the users' junk email folder and high confidence spam and phish, plus messages containing malware, to the quarantine. The quarantine is a vault where marked messages are held instead of delivered to a user's mailbox. Messages can remain in the quarantine for up to 30 days, after which Exchange removes the messages permanently (the default is 15 days). Depending on your quarantine policy, end users can perform any number of actions against quarantined messages, or possibly, no actions at all, which require administrator intervention.

### Quarantine Policy

Quarantine policies in Exchange Online Protection allow you to define what end users can and cannot do in the message quarantine portal. For example, you could define a policy that allows some users to release messages to their inbox and block other users from performing this action. Similarly, you could define different quarantine policies for different threat policies. For example, you could send users quarantine notifications on messages containing spam but not send notifications on messages containing malware.

To view the default policies or create a custom policy, open the **Defender portal** and navigate to **Policies and rules > Threat Policies > Quarantine Policy**.

The default policies *DefaultFullAccessPolicy* and *AdminOnlyAccessPolicy* cannot be modified. The *DefaultFullAccessPolicy* allows users to release the message to their inbox, block the sender, delete the message, and preview the message. In contrast, the *AdminOnlyAccessPolicy* does not allow any user actions. The intention of assigning *AdminOnlyAccessPolicy* to a threat policy is to only allow administrators to act on messages. This is useful when you do not want users to release harmful messages, such as high confidence phishing messages.

To create a custom policy, select **Add custom policy** from the **Quarantine Policy** screen. Give the policy a name that describes its purpose and click **Next**. On the *Recipient Message Access* page, choose either **Limited Access** (which includes all permissions listed below except allowing recipients to release a message) or **Set Specific Access (Advanced)** to specify actions users can perform in the portal. These actions include:

- **Allow recipients to release a message from quarantine** allows the user to release the message to their inbox. This box controls whether the user sees this button in both the quarantine portal and the quarantine notification.
- **Allow recipients to request a message to be released from quarantine;** allow a user to request a message released to their inbox by an administrator. An administrator is notified based on the settings defined in the alert policy (Defender portal > Policies and rules > Alert Policy). By default, this alert goes to quarantine administrators, security administrators, and members of the organization's management role. This box controls whether the user sees this button in both the quarantine portal and the quarantine notification.
- **Delete** allows the user to delete the message from the quarantine portal. The message is not delivered to the user's inbox. This box only controls the button in the quarantine portal.
- **Block sender** allows the user to block future messages from the sender. This box controls whether the user sees this button in both the quarantine portal and the quarantine notification.
- **Preview** allows the user to view the message from the quarantine portal without needing to release it to their inbox. This is useful if the user needs to review the message before acting.

**Note:** The *Allow recipients to release* 5are not honored if applied to an antimalware policy. Users will never be able to release messages with malware regardless of the quarantine policy settings.

With your actions selected, click **Next**. The quarantine notification page allows you to **Enable** whether users should receive quarantine notifications. Using our prior example where we have a different policy for spam versus malware, we may want users to receive notifications for spam in quarantine but not for messages containing malware. Once you have picked your notification setting, select **Submit** to create the policy.

## Quarantine Settings in Protection Policies

To change the quarantine policy assigned to a threat policy, navigate to the **Defender portal > Policies and rules > Threat Policies** and click on the threat policy you wish to configure. You can assign a quarantine action (and quarantine policy) to spam, high confidence spam, phishing, high confidence phishing, bulk mail, user impersonation, domain impersonation, spoofing, malware, and safe attachments. For this example, let's modify the default anti-spam policy. To do this, select **Anti-spam** from the threat policies screen and select **Anti-spam inbound policy (Default) policy**. From the pop-out screen, select **Edit actions**. Under the **Actions** section. Select **Quarantine message** from the **Spam message action** dropdown. When we pick the quarantine action, another drop-down is displayed to allow the selection of the quarantine policy. Let's select **DefaultFullAccessPolicy** and click **Save**. This selection will allow users to perform the following actions on low confidence spam: release the message to their inbox, block the sender, delete the message, and preview the message



## Customizing Quarantine Notifications

The quarantine policies allow us to define whether users receive quarantine messages or not. If you want to change the frequency of those messages or customize the look and feel of the quarantine notification, navigate to **Policies and rules > Threat Policies > Quarantine Policy > Global Settings**.

From the *Quarantine Notification Settings* screen, you can make the following changes.

- **Display Name** allows you to customize the sender shown in the message notification.
- **Disclaimer** allows you to add your custom disclaimer to the bottom of the quarantine notification.
- **Choose language** allows you to customize display names and disclaimers in multiple languages. If you select multiple languages from the drop-down, select each language (not the X) to toggle the display name and disclaimer fields to that language. You can also select Default from the drop-down if you do not wish to specify language-specific display names and disclaimers.
- **Use my company logo** replaces the Microsoft logo on the quarantine notification with your company logo. You can upload your company logo via the **Microsoft 365 Admin Center** by navigating to **Settings > Org Settings > Organization Profile > Custom themes**.
- **Send end-user spam notifications (days)** provides a drop-down where you can specify the frequency of days a notification is sent. It is worth noting that this notification only occurs when new messages end up in the quarantine. Therefore, I recommend a frequency of every day.

## Releasing Messages from Quarantine

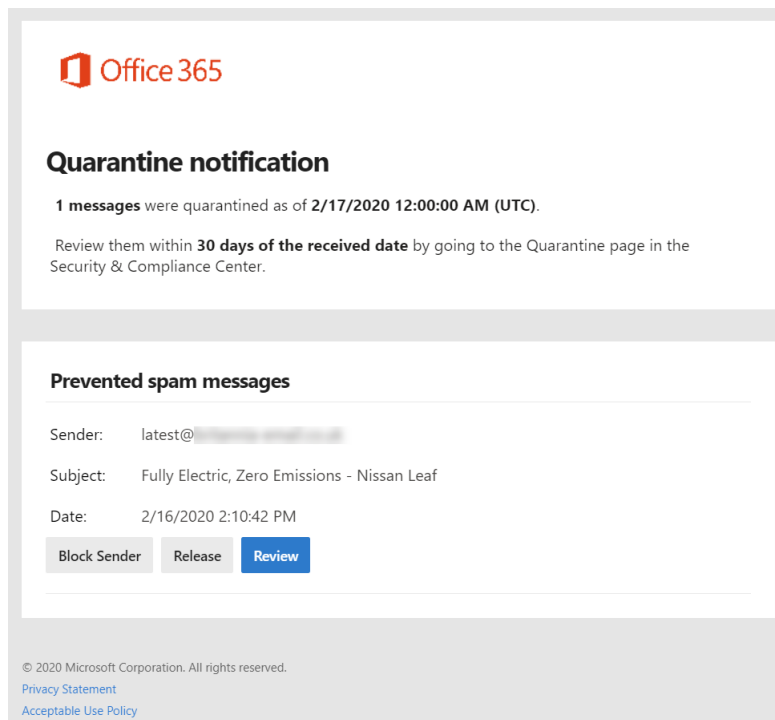


Figure 7-8: A notification that a user has potential spam to handle

Several ways exist to release a message from the quarantine. The fastest and easiest way for a user is to click the **Release** link for the message in the notification digest email, as shown in Figure 7-8. End users with full access to a shared mailbox will also perform quarantine actions for messages sent to the shared mailbox. If the user does not have full access permissions, they will not be able to perform quarantine actions. A user can also click **Review** to preview the message from the quarantine portal before releasing it. From the quarantine portal, the user can also click **Release**.

**Note:** The buttons in the notification email (and quarantine portal) will differ depending on the quarantine policy you have assigned to each quarantine action in your threat policies. If the quarantine policy is set to **Allow recipients to release a message from quarantine**, the **Release** button is available to end-users.

## Quarantine Portal

Users can also log in to the quarantine portal instead of waiting for email notifications to arrive. To access messages held in quarantine, the user must have a valid Microsoft 365 account and an Exchange Online or EOP license.

After signing into the [quarantine portal](#), the user sees a list of their quarantined messages. This list also includes quarantined messages for shared mailboxes where the user is a delegate. The user can search the list based on sender, subject, or, less likely, the message ID. In addition, filtering options such as the date and time the message was received can be applied to focus on specific messages. Depending on the quarantine policy, the user can take the following actions on any message:

- **Release:** Exchange delivers the message to the user's inbox, and they have the option to report it as a false positive to Microsoft.
- **Request release:** Sends a message to an administrator requesting the message be released on behalf of the user.
- **View message header:** Lists the routing information and other message properties with a link to the [Microsoft Message Header Analyzer tool](#).
- **Preview message:** Displays a safe copy of the message.
- **Remove from quarantine:** Permanently deletes the message from the user's quarantine. The message is not released to the mailbox.
- **Block sender:** The sender is added to the user's blocked sender list.

Administrators can work with quarantined messages and files for the entire organization. To grant access to the admin quarantine, use the Quarantine role or the Quarantine Administrator role group within the Defender portal. In addition, members of the Security Administrators and Organization Management role groups also get access to the admin quarantine.

To access the admin quarantine, open the Defender portal and navigate to **Review > Quarantine** or use the direct [link to the quarantine](#). Administrators see a rollup view of all quarantined messages. The interface is like the end-user interface. The difference is that an administrator can see quarantined emails or files for the entire organization and has more options to manage a message. For example, an administrator can release the message to different/additional recipients, download a local copy of the message and attachment, or submit the message to Microsoft for further analysis.

Sort results by

Message ID  Refresh Filter Modify Columns

<input type="checkbox"/>	Received (UTC +01:00)...	Sender	Subject	Quarantine reason	Released?	Policy type	Expires (UTC +01:00) ...
<input type="checkbox"/>	07/08/2020 16:44	ecommerce@cvprovence.c...	LIVRAISON OFFERTE DES 1...	Phish	No	HostedContentFilterPolicy	22/08/2020 01:00
<input type="checkbox"/>	07/08/2020 00:35	pat@caseynet	tony.redmond@	High Confidence Phish	No	HostedContentFilterPolicy	21/08/2020 01:00
<input type="checkbox"/>	05/08/2020 06:42	benjamin@worlddigitalag...	Re: Request for Quote	Spam	No	HostedContentFilterPolicy	20/08/2020 01:00
<input type="checkbox"/>	05/08/2020 06:42	benjamin@worlddigitalag...	Re: Request for Quote	Spam	No	HostedContentFilterPolicy	20/08/2020 01:00
<input type="checkbox"/>	05/08/2020 00:15	usmail@expediamail.com	Longing for a summer get...	Phish	No	HostedContentFilterPolicy	19/08/2020 01:00
<input type="checkbox"/>	04/08/2020 09:45	mundobabyplaza@hotmail...	Request For Quotation 80...	High Confidence Phish	No	HostedContentFilterPolicy	19/08/2020 01:00
<input type="checkbox"/>	02/08/2020 19:51	office227@9pz.org	#External: *G.B.7-31* Plasti...	Spam	No	HostedContentFilterPolicy	17/08/2020 01:00

Figure 7-9: Viewing quarantined messages in the Microsoft 365 Defender portal

In larger environments, the message quarantine might hold thousands of messages. Working with large quantities of messages through the Defender portal might prove challenging. PowerShell can be a better fit for such scenarios, especially when performing bulk actions on messages in the quarantine. For instance, using the *Get-QuarantineMessage* cmdlet, an administrator can look for specific messages in the quarantine. For example, this query looks for high confidence phishing messages:

```
[PS] C:\> Get-QuarantineMessage -QuarantineTypes HighConfPhish | Select ReceivedTime, SenderAddress, Subject, Expires
```

ReceivedTime	SenderAddress	Subject
07/08/2020 00:35:43	pat@casey.net	tony.redmond@ You have 3 messages
04/08/2020 09:45:38	mundobabyp1aza@hotmail.com	Request For Quotation 800014
29/07/2020 18:12:09	account-update@amazon.com	#External: Amazon security alert: Sign-in from new

To release a message from the quarantine, the *Release-QuarantineMessage* cmdlet can be used. This example looks for messages addressed to a certain user that are marked as spam and releases them for delivery. It's unwise to release high confidence phishing messages unless you are certain that the messages are OK.

```
[PS] C:\> Get-QuarantineMessage -RecipientAddress James.Ryan@Office365itpros.com -QuarantineTypes Spam | Release-QuarantineMessage -ReleaseToAll
```

## Configuration Analyzer

One excellent feature of Exchange Online Protection is the configuration analyzer. The analyzer makes recommendations on how to improve a tenant's email security posture, further protecting an organization from bad actors, malware, and spam. The analyzer can be accessed by navigating to **Policies and rules > Threat Policies > Configuration analyzer**.

The **Standard recommendations** tab lets you know how your custom policies match the *Standard Protection* preset security policy. Similarly, the **Strict recommendations** tab measures your custom policies against the *Strict Protection* preset security policy.

In Figure 7-10, the analyzer recommends changing the bulk email threshold in our *Default* policy. It shows our current value of 7, the date the value was set, and the recommended change to a value of 6. To apply the recommendation, we would select the checkbox on that row and click the **Apply recommendation** button. Alternatively, we could click the **View Policy** button to be taken to the properties of the *Default* policy.

### Configuration analyzer

The configuration analyzer can help identify issues in your current configuration, and help improve your policies for better security. [Learn more.](#)

Standard recommendations		Strict recommendations	Configuration drift analysis and history		
Anti-spam	Anti-phishing	Anti-malware	Safe Attachments	Safe Links	
3	4	1	2	3	
Apply recommendation		View policy	Refresh	1 of 13 selected	Search Filter
Recommendations	Policy	Policy group/setting name	Policy type	Current configuration	Last modified
<input checked="" type="checkbox"/> Change 7 to 6	Default	Bulk email threshold	Anti-spam	7	Sep 8, 2021 9:40 PM
<input type="checkbox"/> Change 15 to 30	Default	Quarantine retention period	Anti-spam	15	Sep 8, 2021 9:40 PM
<input type="checkbox"/> Change 9 Entries to 0 Entries	Default	Allowed Senders	Anti-spam	9 Entries	Sep 8, 2021 9:40 PM
<input type="checkbox"/> Change False to True	Office365 AntiPhish Default	Include custom domains	Anti-phishing	False	Jan 29, 2022 5:55 PM
<input type="checkbox"/> Quarantine message	Office365 AntiPhish Default	If email is sent by an impersonated user	Anti-phishing	Move to Junk Email folder	Jan 29, 2022 5:55 PM
<input type="checkbox"/> Quarantine message	Office365 AntiPhish Default	If email is sent by an impersonated domain	Anti-phishing	Move to Junk Email folder	Jan 29, 2022 5:55 PM
<input type="checkbox"/> Change 1 to 2	Office365 AntiPhish Default	Advanced phishing thresholds	Anti-phishing	1	Jan 29, 2022 5:55 PM

Figure 7-10: Configuration Analyzer Recommended Changes

The **Configuration drift analysis and history** tab is an audit log of all changes to each threat management policy, whether modified via the policy or by clicking the adopt button through the recommendations screen.

## Tenant Allow/Block Lists (TABL)

While you can generally trust Exchange Online Protection to make the right decision about things like spoofing, sometimes you need to explicitly allow (or block) specific senders from spoofing mail recipients or domains in your tenant. With Tenant Allow/Block Lists (TABL), you can configure up to 1,024 spoof pairs. A pair is a spoofed domain or user and the actual sending domain, user, or IP address.

To manage the spoofing pairs, go to **Policies and rules > Threat Policies > Tenant Allow/Block Lists** in the **Defender** portal. Click the **Spoofing** tab and click **Add**. From the *Add new domain pairs pane*, add a unique spoof pair to each row, where the first entry on the row is the entity to be spoofed, and the second entry identifies the sender, domain, or IP doing the spoofing. You can then specify whether this is an **Internal** or **External** spoof and whether to **Allow** or **Block** the spoof. Click the **Add** button when you have your desired configuration.

Similarly, you may have a situation where you need to block a specific file, URL, or sender. Tenant Allow/Block Lists lets you configure blocks for up to 500 files, up to 500 URLs (or URL patterns), and up to 500 senders (depending on license). To add one of these blocks, select the relevant tab from the *Tenant Allow/Block Lists* page and click the **Block** button.

To block a sender, you must add each email address or domain as a comma-separated list (or add one entry per row). To block URLs, you need to list a URL or a pattern. For example, you could include just office365itpros.com, or if you needed to include all the subdomains, you could specify ~office365itpros.com.

Managing file exceptions is a little bit harder. You will need to provide the SHA256 hash of the file to add it to the list. There are lots of tools that can compute the hash. Microsoft's [documentation](#) for this feature shows you an easy way to compute the hash value using certutil.exe.

Once you have entered all the URLs, files, or senders to block, you can optionally add an expiration date. This will automatically remove the block entries at the defined time, or you can specify that the entries have no expiration. You can also add a note, which is helpful if you need to describe why the block was implemented.

## Identifying and reporting messages

The following sections outline how users can manage their safe sender and block lists, how users and administrators report junk and phish messages to Microsoft, and how safety tips help users identify suspicious emails.

### Informing users with Safety Tips

To increase visibility to end-users on messages containing potential threats such as spam, malware, spoofing, or phishing, Microsoft includes visual cues in messages to warn users when they meet potentially dangerous messages. Safety tips work in much the same manner as mail tips. These cues are tags inserted into messages by EOP as it processes email before delivery to end-users. For example, whenever a message is considered suspicious, Exchange inserts a safety tip. Similarly, if a message is received from a trusted sender, the notification will make this clear. EOP supports four safety tips (red, green, yellow, and grey). Table 7-4 lists the types of safety tips supported by EOP.

<b>Color</b>	<b>Safety Tip</b>	<b>Condition</b>
<i>Red</i>	<b>Suspicious.</b> These messages are likely to contain phishing scams and should not be opened.	EOP detects the presence of a known phishing message, or the message characteristics are such that EOP considers the message likely to be a scam. See <a href="#">this article</a> for an explanation of why messages are assigned red safety tips.

<i>Yellow</i>	<b>Unknown.</b> The message is spam.	EOP has scanned the message, and it has failed the standard anti-spam tests.
<i>Green</i>	<b>Trusted.</b> The message is from a trusted source.	Microsoft has a list of domains owned by trusted sources (such as itself). Messages from these domains are considered trusted.
<i>Grey</i>	<b>Safe.</b> An informational tab that indicates the message has been marked safe by the tenant or user.	The message is from a domain considered to be safe by the tenant (for example, the IP address for the server is in the <a href="#">IP allowed list</a> ), or on the user's safe senders list, or it was put in the user's Junk Email folder and subsequently moved back to the Inbox to indicate its safe status.

Table 7-4: Types of safety tips

Safety tips can be enabled or disabled for a tenant by updating the policies used by EOP. For example, the `Set-HostedContentFilterPolicy <name> -InlineSafetyTipsEnabled` command is used to configure safety tips via PowerShell. We do not recommend that you disable Safety Tips.

#### Unauthenticated Sender Safety Tip

The unauthenticated senders feature works in conjunction with the analytics gathered by Exchange Online to figure out if the sender is who they say they are or if the sender is spoofing someone else. This is known as sender authentication. If the message sender fails the authentication tests (SPF, DKIM, or DMARC), Exchange replaces the initials/photo shown next to the sender's display name with a question mark. However, a failure does not automatically result in an icon replacement because Microsoft uses extra technologies to help determine if a message is safe even though it failed authentication.

Not every message that fails to authenticate is malicious. However, you should be careful about interacting with messages that cannot authenticate if you do not recognize the sender. Or, if you recognize a sender that normally does not have a '?' in the sender image but suddenly starts seeing it, that could signify that the sender is being spoofed.

The sender authentication feature is enabled via an option in the anti-phish policy. Inside the actions section of the policy, you can disable this feature not to show "?" icons. While this feature is enabled by default, we do not recommend that you disable it.

Exchange Online will never mark senders in your safe sender list as spoofed – even if they are spoofed – and this also applies for messages marked as safe due to mail flow rules, anti-spam policies, and the safe senders specified in those policies as well. Administrators can deal with false positives caused by sender authentication by adding the sender and sending infrastructure to the anti-phishing Spoof Intelligence insight.

For more information on authenticated senders, check the following [Microsoft article](#).

#### How users mark mail as junk, phish, or safe

Exchange Online allows users to control specific junk email settings themselves. Within Outlook and Outlook Web App, a user can do the following:

- Mark a sender as safe
- Mark a sender as blocked
- Add a recipient to a safe-recipient list

In Outlook Web App, when a user right-clicks a message and selects **Mark as Junk**, the message is automatically moved to the Junk Email folder, and the sender is also added to the blocked sender list. More specifically, the information is stored as part of the mailbox configuration in the following AD attributes:

- msExchSafeSenderHash
- msExchSafeRecipientHash
- msExchBlockedSenderHash

This information is accessed by EOP and taken into consideration when processing messages. Messages from safe senders are automatically marked as “not spam.” In contrast, messages from blocked senders will be forwarded to the user's junk email folder or the quarantine, depending on the policy configuration. It is important to note that EOP will not honor the safe sender list if EOP considers a message to be high-confidence phishing. This is because a domain or sender that has been safe-listed might become compromised, leading to a possible compromise of the recipient's mailbox. If you need to have a sender or domain bypass this test, you can still do so with a mail flow rule that adjusts the SCL header on the message.

In a hybrid deployment, these attributes are also written back into the on-premises organization. It is also crucial in a hybrid deployment scenario where the MX records point to the on-premises organization instead of EOP. The write-back capability ensures that Edge Transport servers correctly process messages for cloud-based mailboxes. Note that when an on-premises mailbox is protected by EOP (or a cloud-based mailbox by Edge Transport servers), there can be a delay of at least thirty minutes between a user adding an address to one of the lists and the information being available to EOP or the Edge Transport servers. This is due to the directory synchronization interval and the worker process in Exchange Online that reads the safe and blocked sender lists and writes them to the Exchange Online directory.

In addition to the safe and blocked sender lists, users can configure client-specific settings. In Outlook, the **Junk Email Options** allow you to:

- View or edit the **Safe Senders, Safe Recipients, and Blocked Senders** list.
- Toggle the setting to automatically **trust email from my contacts**.
- Adjust in-client junk email filtering through the general **Options** tab.
- Block messages from specific **countries or regions**. For instance, if a user selects to block messages from Canada (CA), all messages from a domain ending with .ca will be marked as spam.
- Block messages written in an **unfamiliar language**.

Under the **Options** tab, you will find settings that control Outlook's built-in junk email filter. The selection there should always be set to **No Automatic Filtering** for mailboxes leveraging EOP. Otherwise, you will find Outlook making decisions based on its older and less reliable junk email filter. You can also configure Outlook to remove junk email rather than move it to the Junk Email folder.

**Note:** EOP uses only the safe senders, safe recipients, and blocked sender lists. All other features are client-specific and will only affect messages after being delivered to the user's mailbox. However, it should be noted that EOP also provides country- and language-based protection, which can be configured separately in the anti-spam filter policy.

To access the corresponding settings in OWA, click **Settings** (cog wheel) and then **View all Outlook settings**. From there, select **Mail > Junk email**. Apart from the aforementioned lists, you will find settings to restrict message delivery to only people from your trusted lists, to consider all contacts as trusted, and to control whether messages you mark as junk are automatically reported to Microsoft.

**Note:** Outlook has separate tabs for Safe Senders and Safe Recipients, but OWA only shows a single list called **Safe Senders and domains**. If a user makes changes via OWA options, their Safe Recipients list is merged with the Safe Senders list, and that list is copied into both Safe Senders and Safe Recipients in Outlook. Also, any Safe Senders and Safe Recipients from an internal domain are removed automatically from these lists shortly after being added.

Lastly, administrators can review or configure junk email settings on a user's behalf via *the Get-MailboxJunkEmailConfiguration* and *Set-MailboxJunkEmailConfiguration* cmdlets. It is important to note that if you disable the junk email rule using *Set-MailboxJunkEmailConfiguration*, this only disables the user's individually defined junk email lists and junk email settings stored in their mailbox. Junk emails identified by EOP will still be processed and delivered to the user's mailbox as dictated by the anti-spam policies defined in the tenant.

## How users report junk and phishing email

Microsoft develops many of the features to combat spam and malware based on data harvested from inbound emails delivered to Exchange Online and Outlook.com. Using machine learning techniques, Microsoft distills valuable information from the data to improve the efficiency of its protection features. However, because spammers continually introduce new techniques to bypass checks, it is impossible to guarantee that every threat will be detected. As a result, sometimes messages slip past EOP. Users can report these messages (also known as false negatives) by sending them to Microsoft for analysis. If Microsoft's Spam Analysis Team confirms that the message meets the spam classification criteria, they update the EOP filtering systems to block similar messages in the future. Similarly, users can also report false positives. These are messages that EOP detected as spam but are not.

Tenant administrators can decide whether to allow users to report messages to Microsoft. When users believe they have received a phishing email, they can use the **Report Message** drop-down option () in Outlook or OWA to report the message to Microsoft. Messages marked as phishing are automatically moved to the Deleted Items folder.

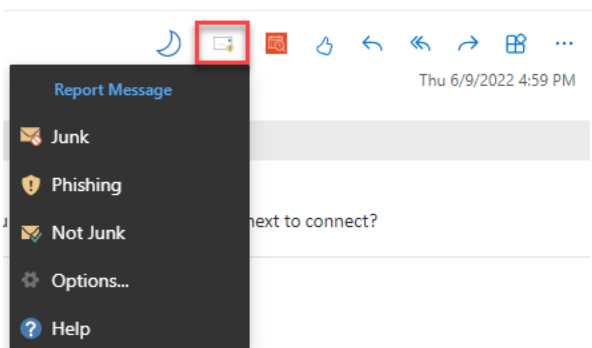


Figure 7-11: Marking a Phishing message with OWA

[The Report message add-in client plugin](#) (Figure 7-11) is available for Outlook and OWA. The add-in [can be enabled by administrators](#) to allow users to submit suspicious messages to Microsoft.

## How administrators report junk and phishing email

The **Submissions** feature, found under the **Actions & Submissions** section of the Defender portal, gives administrators the ability to upload information to Microsoft about messages that should have been blocked or those that are false positives. The information can be the message (in .eml or .msg format), Message ID, a URL found inside the message, or an attachment. Figure 7-12 shows an example of an administrator submitting a false positive using the Message ID to Microsoft

After you submit your findings to Microsoft, you can view the outcome of the analysis. In addition, the report will provide you with information about any policies that have acted on the message, as well as an examination of any URLs and attachments found within. For example, a common scenario is that the message was allowed by a personal exemption by the end-user or via a mail flow rule.

The page will also allow you to review user-submitted messages under the **User reported messages** tab. The experience also allows you to start an investigation (if your tenant has the necessary Microsoft Defender for Office 365 Plan 2 license). Additional information about the Submissions feature can be found in the [online documentation](#).

### Submit to Microsoft for analysis

We will review the information and use what we've learned to improve detection. We will let you know our findings. [Learn more](#)

Select the submission type

Add the network message ID or upload the email file

Add the email network message ID ⓘ

Upload the email file (.msg or .eml)

Choose a recipient who had an issue ⓘ

Select a reason for submitting to Microsoft

Should not have been blocked (False positive) ⓘ

Should have been blocked (False negative) ⓘ

This email should have been categorized as

Phish

Malware

Spam

Figure 7-12: Creating a new submission for a false positive email

## Delisting Your IP Address from the Office 365 Block List

Although Microsoft does not add an IP address to the block list for no reason, sometimes this happens erroneously. Common causes for an IP address to be listed is because it was repeatedly found as the source of spam, malware, spoofed, or phishing messages. In addition, you will find yourself on the list if you have an open relay or do not pay enough attention to outbound messages.

When your IP address appears in the block list, messages sent to Exchange Online recipients fail, and a Delivery Status Notification (DSN) goes to the sender with the following information:

*1.7.1 Service unavailable; Client host [1.2.3.4] blocked using Blocklist 1; To request removal from this list, please forward this message to delist@messaging.microsoft.com*

Like most third-party block list providers, you can use a web form to request removal from the Block List. The form is located at <https://sender.office.com> and takes you through three simple steps:

1. You must enter your email address and the IP address to be delisted.
2. A verification email is sent to the email address.
3. If the verification was successful, you can now continue to delist the IP address.

Once the delist request has been received, the IP address is typically removed within 30 minutes.

## Using a Third-Party Filtering Solution Before EOP

Organizations often consider that they can gain added protection by deploying other email filtering solutions alongside EOP. Although this approach is practical from a technical perspective, it could adversely affect some anti-spam features, including IP throttling, IP blocklists, bulk mail filtering, or anti-spoofing in EOP. This is also why Microsoft recommends against using an added filtering solution with EOP.

For instance, EOP relies on IP throttling to prevent the delivery of spam messages: when a message is received from a new IP address, EOP throttles the incoming connection by issuing an SMTP 450 error. The error indicates a transient error, which means the sending server should retry later. Legitimate servers usually retry



the connection, while most spammers do not. The reason for not retrying is quite simple: the added effort to interpret the error message and then try again later is often too much trouble for spammers. Even though the feature does not stop a spammer from sending messages, it greatly reduces spam delivered to Exchange Online mailboxes. However, when a second filtering solution is used in front of EOP, all messages delivered to your organization originate from a single set of IP addresses. This renders the IP throttling feature useless.

When Exchange Online detects that your MX record does not resolve to EOP, it automatically disables some EOP anti-spam features. An example of one feature they disable is IP throttling. EOP also cannot perform IP reputation blocks, sender authentication checks such as SPF and DMARC, spam filter rules, and more. Therefore, it is imperative to ensure that the third-party solution you choose performs all the same tests as EOP. For example, if you have a security device as your first hop that does not do anti-spam filtering, Microsoft will also not perform all the anti-spam filtering they could otherwise do.

**Filtering hybrid mail:** It is unsupported to use a third-party filtering solution in a hybrid deployment for messages sent between the on-premises organization and Exchange Online. This is because a third-party filtering solution can remove important message headers, like the *X-MS-Exchange-Organization-AuthAs* header, from the message.

When the *X-MS-Exchange-Organization-AuthAs* headers are no longer present, problems can occur. For example, messages might not be recognized as originating from within the same organization. This could lead to internal messages being treated as spam or an external out-of-office message being returned rather than its internal counterpart. For this reason, messages must flow uninterrupted between Exchange Online and the on-premises Exchange servers.

If additional filtering is desired, or you have the requirement to terminate all inbound connections in the perimeter network, you must use a Microsoft Exchange Edge Transport server. The Edge Transport server is the only additional filtering solution that authenticates like an Exchange server in the same manner as the internal Exchange servers. You can continue using a third-party filtering solution to process all external messages entering and leaving your organization to and from the internet.

## Enhanced Filtering for Connectors

If you use an email filtering solution or email gateway other than EOP, you must configure Enhanced Filtering for Connectors. The Enhanced Filtering for Connectors option allows tenants whose MX record does not resolve to Exchange Online to determine the true source of an email. It is used when messages are routed via on-premises servers or a third-party cloud filtering service, both of which make the email seem to come from that source instead of its true origin. Knowing the actual email source, you can fine-tune your spam, phishing, and mail flow rules.

Enhanced Filtering, or skip listing, is a complex routing configuration. In a hybrid scenario, the recommended solution is to route the email directly to EOP by setting the domain MX record to Exchange Online. This is not always possible, however. For example, if you are migrating to Exchange Online, your MX record will resolve to your on-premises Exchange organization until later in the migration project. During this period, or if your MX record remains configured to a third-party email filtering solution, you must ensure that EOP is aware of your internal IP addresses (if routing via on-premises) and/or the IP address range(s) belonging to the third-party filtering service. This includes any analysis or intermediary SMTP servers in your routing pipeline after spam, phishing, and malware filtering services.

When EOP is aware of the IP ranges belonging to other services you trust, it skips those IP address ranges in the SMTP headers to determine the true source of the email. Microsoft's Intelligent Security Graph has information about the trustworthiness and reliability of these sources based on their previous history. This information can be applied to inbound emails. When Enhanced Filtering is not configured, the information in

the previous step in the received SMTP header is all that Exchange Online knows about the previous sender, which is likely to be your on-premises infrastructure or the last server in the third-party cloud infrastructure.

To enable Enhanced Filtering, open the **Defender portal**, go to **Policies and Rules > Threat Policies > Enhanced Filtering**, or use <https://security.microsoft.com/skiplisting>. Enhanced Filtering can be enabled for a subset of users for test purposes. Any email addressed to recipients outside your test user pool will not be processed through Enhanced Filtering. Therefore, if you want to test the feature and some of your pilot users receive email using multiple proxy addresses, you should include all the addresses in the Enhanced Filtering configuration. Once your pilot is complete, you can update the configuration to apply to the entire organization.

When you configure Enhanced Filtering by associating the IP addresses of trusted network ranges with the relevant connector, email domain authentication (aka DKIM, DMARC, and SPF) will improve. Now that EOP can determine the source of the email, it can check this information against anti-spoof technologies and explicitly allow, quarantine, or reject email based on the sender infrastructure.

Two SMTP headers are added to messages after Enhanced Filtering is enabled. These are:

- **X-MS-Exchange-ExternalOriginalInternetSender** shows the true source of the message. This should not be in the IP range of the on-premises servers or your third-party filter. If it is, you have not configured skip listing correctly.
- **X-MS-Exchange-SkipListedInternetSender** shows the true source of the message and is used for reporting purposes.

Enhanced Filtering replaces custom mail flow rules often used to prevent double filtering. These mail flow rules set the spam confidence level (SCL) to -1 so that messages skip EOP filtering.

## Microsoft Defender for Office 365 (MDO)

Microsoft Defender for Office 365 (MDO) is a set of features designed to answer zero-day exploits or new methods built by attackers to bypass and fool protection systems such as EOP. In the current threat landscape of email, having a zero-day malware and link protection service that operates on both internal and externally sourced emails should be considered by all tenant administrators as a required purchase. By providing features such as safe attachments, safe links, safe documents, and advanced anti-phishing controls, MDO can significantly increase your organization's security posture.

Microsoft Defender for Office 365 is included in Microsoft 365 E5 and Office 365 E5 and is also available as an add-on. Like most other services available in Microsoft 365, you can enable the added protection for your entire organization or just a select group of users. However, until you buy at least one license, Microsoft Defender for Office 365 is unavailable through the Defender portal.

In hybrid deployments, messages sent between on-premises Exchange servers and Exchange Online bypass MDO scanning. This is because the connectors created in the hybrid configuration set the Spam Confidence Level (SCL) header to -1 to tell EOP to bypass any additional anti-malware or anti-spam filtering.

### Built-in Protection

Built-in protection is a baseline default policy defined by Microsoft for all Microsoft Defender customers. This policy enables a base configuration for Safe Links and Safe Attachments and is applied to every user in your tenant.

This policy cannot be disabled, but users, groups, or domains can be excluded from the policy by navigating to **Policies and rules > Threat Policies > Preset security policies** and selecting **Add exclusions** under the *Built-in protection* section.

It is not recommended to exclude users from the built-in policy. Instead, if you wish to give your users a different configuration, either create custom Safe Links and Safe Attachment policies and assign those custom policies to those users or assign those users to the strict or standard preset security policies. These policies will always take precedence over the built-in protection policy, so there should be no need to exclude users from the built-in protection.

If a user is assigned to multiple policies, policy precedence occurs in the following order.

1. Strict protection preset security policy
2. Standard protection preset security policy
3. Custom security policies (by order of priority)
4. Built-in protection policy

The built-in protection policy ensures every user has Safe Link and Safe Attachment protection in the event they are accidentally removed from all custom or preset security policies.

For a comprehensive list of built-in protection policy settings and how they differ from the standard, strict, and default custom policy settings, check the [Microsoft documentation](#).

## Safe Attachments

*Safe Attachments* deliver extra protection against zero-day malware. EOP uses multiple anti-virus scanning engines to process incoming (and outgoing) messages. However, these engines use “signature” files to detect known viruses and malware. You can compare it to a database that has the essential characteristics of a virus to enable the anti-virus engine to recognize an instance of the virus when it occurs in email. The problem with this approach is that when a new virus is created, time is needed for security researchers to decipher the virus and construct its signature and then update the database. During this period, messages holding the virus will likely penetrate past standard scanning and arrive in user mailboxes. When MDO is enabled, messages with attachments are rerouted to a virtual sandbox environment where the content is subjected to a “behavioral analysis” based on machine learning. During this process, the attachment is run or opened, scanned, and observed to determine whether it is malicious or not. If no suspicious activity is detected, the message is released for delivery to the user’s mailbox.

Each of the MDO features works as an ‘add-on’ to EOP. This means that messages are only subject to additional scanning or processing if none of the other EOP features have found anything suspicious about the message. In the case of Safe Attachments, a message is only rerouted when EOP’s anti-malware engine successfully processes the message and does not detect a threat. The rerouting of the message itself is fully transparent to the end-user.

In addition to protecting attachments in inbound emails from external sources, Safe Attachments file protection is also available for Microsoft Teams, SharePoint Online, and OneDrive for Business. The option to enable MDO for mailboxes is driven by a policy that applies to whomever the policy is configured for. The option to enable MDO protection for OneDrive for Business, SharePoint Online, and Microsoft Teams is a global setting.

### Configuring a Safe Attachments Policy

Even when you assign an MDO license to an account, the Safe Attachments feature will not process any messages until you have created a policy. To create a policy, navigate to **Policies and rules > Threat Policies > Safe Attachments** in the Defender portal and click **Create**. A policy consists of several elements:

- **Safe Attachments unknown malware response** defines what action Safe Attachments should take when a potentially dangerous attachment is detected. The available options are *Off*, *Monitor*, *Block*, *Replace*, or *Dynamic Delivery*.

- **Monitor** can be helpful when you are first enabling the feature because it allows you to assess how well Safe Attachments are performing. Although unsafe attachments are still delivered to the recipient, all the reporting features are available.
- **Block** will quarantine both the attachment and the message.
- **Replace** mode will quarantine the unsafe attachment and deliver the original message.
- **Dynamic Delivery** will ensure that the message body is delivered instantly while the attachment is being analyzed. Once the attachment has been processed, it is reattached to the original message. Should the attachment be considered malware, it will quarantine the attachment.
- **Off** disables any action taken by safe attachments.
- **Quarantine Policy** defines which quarantine policy to use when the Block, Replace, or Dynamic Delivery actions place an attachment and message into quarantine.
- Whether to **redirect malicious attachments** to other recipients. This might be interesting if you do not want to lose the original attachment, even if it is considered unsafe. Caution is advised if you do decide to open the delivered attachment. Some organizations use this feature to send a copy of a malicious attachment to a security operations center (SOC) for further review.
- **Apply the Safe Attachments detection response if scanning can't complete** determines the action to take when message processing times out or a message could not be processed due to an underlying error. With this checkbox deselected, the message is delivered normally to the user. With this checkbox selected, any attachments with processing errors or timeouts are automatically treated as malware. The message then follows the same action identified in the unknown malware response (e.g., off, monitor, block, replace, or dynamic delivery).

In addition to these options, you must also define the scope of the policy to define which users the policy applies.

- **The Recipient is.** Allows you to select one or more individuals from your organization. This can be helpful if you need to single out a handful of users. However, it is better for larger groups of people to use one of the other conditions.
- **The Recipient domain is.** Allows you to enable a policy for all recipients that use a specific domain. You can select both custom and default domains here. The latter makes it ideal for testing the feature if your tenant does not have a custom domain.
- **The Recipient is a member of.** Enables you to select one or more groups to which the user must belong. If the user does not belong to one of the selected groups, the safe attachments feature will not be executed.

To enable Safe Attachments protection for SharePoint, OneDrive, and Microsoft Teams, you need to click **Global settings** at the top of the page and toggle the switch **Turn on Microsoft Defender for Office 365 for SharePoint, OneDrive, and Microsoft Teams**.

**Note:** It can take up to 30 minutes before a new policy, rule, or setting is fully active.

## Safe Documents

Dubbed **Safe Documents for Office clients**, the feature leverages Microsoft Defender for Office 365 to scan documents before allowing users to open them outside of protected view mode in Microsoft 365 desktop applications (Version 2004, build 12730 or greater). This feature is technically not part of Microsoft Defender for Office 365 and requires a Microsoft 365 E5 or Microsoft 365 E5 Security license.

Safe Documents prevents users from opening a document received from an external source until the document can be checked against Microsoft's threat defense cloud. Since there are certain privacy considerations at play, the feature is not enabled by default. Still, it can be toggled on from the **Safe**

**attachments** > **Global settings** page by selecting **Turn on Safe Documents for Office Clients** or via the `Set-AtpPolicyForO365` cmdlet.

## Safe Links

Phishing messages often include malicious links that redirect users away from legitimate sites to sites under the attacker's control. When EOP processes these messages, a high probability exists that these links are already flagged as malicious, and the message will be caught by the anti-malware or anti-spam engine. Unfortunately, this is not always the case. Sometimes, malicious links are not known yet, or perhaps the link was not yet activated when the message was processed. The weakness in this approach to securing messages is that it only protects messages at delivery. If an attacker updates a malicious web page after it has been delivered, the traditional anti-malware protection can do nothing. Usually, this is where computer-based security software kicks in, but with more users responsible for their own devices under BYOD policies, organizations do not always control what endpoints are used.

The *Safe Links* feature analyzes links at the time they are clicked. To do so, it checks the link in real-time in selected clients or rewrites any hyperlinks in email messages when EOP receives them. Safe Links will, by default, only rewrite links if the sender of the message is from an anonymous source, like external senders, and if the link is in the FQDN format. For example, a link with only a server name (no dots) is not rewritten. Rewriting links for internal senders is disabled by default but can be turned on in the Safe Links policy. This protects against a scenario where a mailbox's login details are compromised. A malicious actor uses the compromised mailbox to send bad links that appear to come from a trusted sender.

For Safe Links in email, the link in the message body is rewritten during delivery in EOP. Once the email is delivered to the recipient, the underlying link looks like this:

```
https://eur02.safelinks.protection.outlook.com/?url=http%3A%2F%2Fwww.spamlink.contoso.com%2F&data=02%7C01%7Cbrian%40nbconsult.co%7C321a4fc34c294a14231908d754b5f730%7C69daf98459d5415ba9b1aa3b30b4a677%7C0%7C637071012816511154&sdata=eYvY4TOvBEPJIOZvilGeNclnLIX0q0nQeFfLF9tjAc%3D&reserved=0
```

**Note:** The base URL (\*.safelinks.protection.outlook.com) will always be the same, the information that follows varies on the message and the link(s) within that message.

Outlook hides the link and shows a popup with the target URL so that users do not see the complicated nature of the rewritten link. Opening the message in a non-Microsoft client or an older Microsoft client will still show the rewritten link. When Safe Links is enabled for Office applications, the user does not see the full link. Instead, they only see the destination the link takes the user to if the target is not flagged as malicious.

Outlook desktop also protects URLs in message subjects, URLs in nested email attachments, and URLs in S/MIME signed messages. This is achieved by processing the link at the time of the click, even if EOP did not rewrite the URL. For example, links within S/MIME protected messages are not rewritten; therefore, the message is not rendered invalid while the link is still protected. Indeed, any text that looks like a hyperlink that is made clickable by Outlook is now protected, even though it is not seen as a link when being processed by EOP. No URL rewriting is done; the protection exists solely within the Outlook client.

When a user clicks a link in a message, the Safe Links feature checks the reputation of the remote host (website) against a set of spam lists. This usually happens within a matter of seconds and does not produce any significant delays for the user. When a website is not on one of the spam lists, the user is redirected to the requested website. However, when a match is found, the user will be stopped from accessing the remote website and receive a warning that they should not proceed.

Depending on how you configure the Safe Links policy, the user can be presented with the choice to continue to the website, despite the warning. By default, users are prevented from clicking through to malicious

websites. If you do allow blocked links to be clicked through, the click is registered and will show up in the *URL threat protection* report in the Defender portal. We do not recommend that you allow users to access links deemed to be unsafe.

## URL Detonation

The Safe Links feature uses a list of known malicious targets to understand if a dangerous link is included in a message. This is great if a URL is known to be malicious. However, whenever an attacker creates a new malicious web page, some time will elapse between the web page being accessible and when it is first reported as malicious and, after that, blocked. The URL Detonation feature attempts to intervene during that time by actively scanning URLs in messages before allowing users to access the targets. When a link is considered malicious or points directly to executable content, and the user attempts to access the link, a warning pop-up window (colored yellow) is shown. The user is prevented from opening the target location until scanning is finished. Unlike the Safe Links feature, there is nothing to configure for the URL Detonation feature apart from turning the option on.

It is also possible to delay the email's delivery to the end-user until the links in the email are scanned (just like scanning the attachments). The option to control this behavior is turned off by default. Check the **Wait for URL scanning to complete before delivering the message** option in the Safe Links policy to turn it on. With this option set, any message containing a URL considered malicious is redirected to the user's Junk Email folder.

## Configuring a Safe Links Policy

To create a Safe Links policy, open the **Defender portal** and navigate to **Policies and rules > Threat Policies > Safe Links**.

When creating a new policy scoped to a subset of your users, you have the following options:

- Scope the policy to users, groups, or domains.
- Select whether URLs should be rewritten in emails from external senders.
- Select whether URLs should be rewritten in emails from internal senders.
- Control the behavior of the URL detonation feature (detailed above).
- Provide a set of URLs that will not be rewritten in emails.
- Select whether Safe Links should check URLs in Microsoft Teams (URLs are not rewritten).
- Select whether Safe Links should check URLs in Office apps (URLs are not rewritten).
- Select whether to track user clicks.
- Select whether to allow users to click through to the original URL.
- Display organization branding and warning pages.

If you need to create a custom list of URLs to block, check the earlier section in this chapter on Tenant Allow/Block Lists (TABL).

**Real-world:** The choice to fully block users from navigating to websites that Safe Links blocks could be very invasive. This is especially true when the protected link is accidentally blocked. To block or not block the final link is a security question, and in most cases, as the blocked site is malicious, it is best to ensure that the user cannot click through to the target site. Note that Safe Links does not stop a user from manually copying and pasting the source link into a browser to open the website from there. While the *Safe Links* feature does increase security, it is by no means a replacement for end-user training and added protection on the endpoint itself.

## Testing Microsoft Defender for Office 365 features

Testing Safe Links is simple. All you need to do is send a message or write a document that includes the following URL: <http://www.spamlink.contoso.com>.

To verify whether MDO scans incoming attachments, it is enough to send a message that includes a regular (safe) attachment of a type that could include malicious code. This can be anything from an executable to a Word or PowerPoint file or a PDF file. Because the regular anti-malware feature will not detect anything in the file, the message will be rerouted through MDO for additional scanning. Therefore, ensure that the attachment type is not blocked by the common attachment type filter.

If you are considering migrating to MDO from another email filtering system, you can enable MDO in evaluation mode. Evaluation mode will provide a log of threats that MDO would have protected you from that your existing email filtering system missed. You can enable evaluation mode in the Defender portal by navigating to **Policies and rules > Threat Policies** and clicking **Evaluation Mode**. It is important to note that the evaluation mode will not alert you to threats already mitigated by your existing email filtering service since those messages or attachments will never reach MDO. You can learn more about evaluation mode and how it works with third-party email gateways in Microsoft's [documentation](#).

## Reporting on Microsoft Defender for Office 365

MDO leaves some information within the message headers. For example, using a tool, such as Microsoft's [Message Header Analyzer](#), you can check the message headers for the following information:

1. The *Received headers* section identifies any delays between EOP servers. A delay, however, does not prove that MDO processed the message – for example, internal routing errors could result in delays that have nothing to do with MDO.
2. In the *Other headers* section, a new header, "X-MS-Exchange-Organization-SafeAttachmentProcessing," is present. The header indicates that the Safe Attachment feature processed the message. Note that this header has no value.

Both elements together can provide you with a first impression of whether MDO processed the message. However, the best way to discover how MDO processed a specific message is by running a message trace. Figure 7-13 shows a series of MDO events as they appear in message trace results. The fourth line from the bottom is one of the MDO information rows and is shown in the popup that appears when hovering the mouse over the row.

Date (UTC)	Event	Detail
01/06/2018 08:06:41	Advanced Threa...	Allowed. Advanced Threat Protection Malware...
01/06/2018 08:06:42	Receive	Message received by: HE1PR03MB1196 using...
01/06/2018 08:06:45	Journal	Message was journaled. Journal report was se...
01/06/2018 08:06:45	Journal	Message was journaled. Journal report was se...
01/06/2018 08:06:45	Send	Message sent to [redacted].mail.protection...
01/06/2018 08:06:45	Receive	Message received by: HE1EUR02TH003 using T...
01/06/2018 08:06:46		Advanced Threat Protection Malware: 109_2919.pdf, Type: PDF, Environment: , Verdict: Good, Evidence: Unknown
01/06/2018 08:06:47		
01/06/2018 08:07:31	Advanced Threa...	Advanced Threat Protection Malware: su11aa_...
01/06/2018 08:07:32	Receive	Message received by: AM3PR03MB1185 using...
01/06/2018 08:07:32	Send external	Message sent to eur02.map.protection.outlook...
01/06/2018 08:06:45	Spam Diagnostics	

Figure 7-13: Message Trace results

The first DEFER event (hidden by the popup in Figure 7-13) shows that the message is rerouted to the MDO scanning environment. Next, the event called Advanced Threat Protection gives information about the outcome of the scanning process. In this case, the attachment was scanned, and the message was redirected back to EOP for delivery to the recipient within six seconds.

## Reporting on Message Latency

With the extra layers of security provided by Microsoft Defender, including the ability to detonate attachments in a sandbox, administrators may be concerned about how much latency this adds to message

delivery. In the previous section, we provided an example message trace that indicated 6 seconds were added by Microsoft Defender before the message was delivered. However, to enable administrators to better visualize this for their entire tenant Microsoft aggregates all message trace data into the *Message Latency* report.

## Mail latency report

50th percentiles   90th percentiles   99th percentiles

This report shows all latency within the mail filtering and delivery pipeline. It does not include client or network latency. [Learn more about this report](#)

Filters: Date (UTC): 4/14/2022-4/26/2022   Message view: Inline detonation




Figure 7-14: Mail Latency Report

You can access this report from the **Defender portal** under **Reports > Email & collaboration > Email & collaboration reports**. The **Mail Latency** report (Figure 7-14) contains a bar chart and a table. The percentile tabs across the top allow you to switch the chart between the 50<sup>th</sup>, 90<sup>th</sup>, and 99<sup>th</sup> percentiles. Each bar identifies the average latency of all messages in a given day. Filters allow you to change the date range or narrow down the scope to only messages that were detonated in a sandbox.

## Reporting on Safe Attachments

In addition to the information provided by message traces, Microsoft includes several reports that give statistical information on processed messages. You can access these reports in **Defender portal** under **Reports > Email & collaboration > Email & collaboration reports**. The **Threat Protection Status** report shown in Figure 7-15 is an example.

The reports available in the GUI are interactive and allow you to adjust the filters on the fly. For example, you can change which file types are shown or alter the date range. However, the information that fuels these reports comes from underlying message trace information. For example, suppose you specify a custom date range that goes beyond 30 days. In that case, Exchange Online generates a historical search automatically, and you will have the ability to download the CSV report later instead of viewing it interactively on screen. More information on message traces and historical searches is in the Message Tracing section.

The reports are also available as tables showing the underlying data. To see this information and explore further (such as message ID, to and from information), click **View data by** drop-down and select something other than **Overview** (for example, *View data by Email > Phish*).



## Threat protection status

The Threat protection status report provides information about threats found prior to email delivery, covering relevant detection technologies, policy types, and delivery actions. [Learn more about this report](#)




Figure 7-15: Safe Attachments reports

### Reporting on Safe Links

Reporting on Safe Links is available in the **Defender portal** under **Reports > Email & collaboration > Email & collaboration reports > URL protection report** widget > **View Details** button.

The **URL threat protection** report has two views. The default view shows the resulting action of each click. By default, this includes blocked clicks, blocked but clicked through, block by tenant admin, blocked by tenant admin but clicked through, click through during scan, and pending scan. Clicking the filter button, you can include allowed clicks, modify the date range for the report, and filter to specific domains or recipients. If your Safe Links policy has the **Do not track user clicks** option selected, this report will contain incomplete data.

### URL threat protection

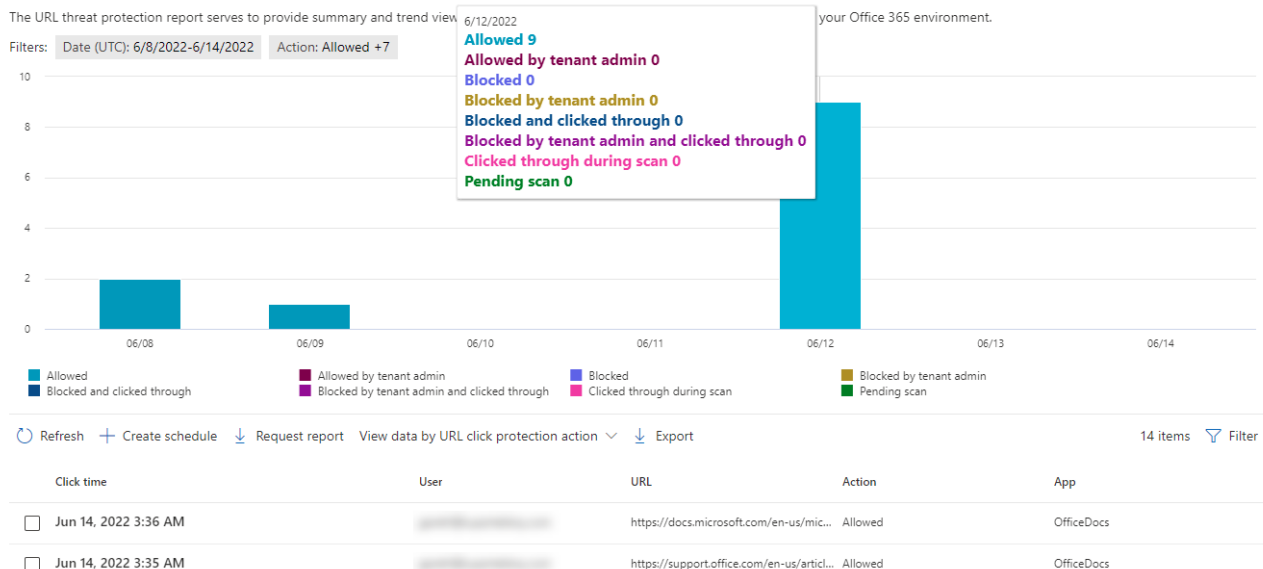


Figure 7-16: URL threat protection report, highlighting the URL click protection action view

Figure 7-16 illustrates the URL threat protection report, highlighting the URL click protection action view. The report shows URLs that resulted in a block page, those block pages where the user clicked through to the

actual destination (if enabled in the policy), and those URLs that were the result of scanning a direct attachment as part of the URL. The actual URLs are displayed below the chart in a table.

The second view is the **URL click by application**, which shows a breakdown across the different applications that support Safe Links. This can give insight into whether users are clicking links in email or other Microsoft 365 Apps (Word, Excel, Teams, etc.).

The basic function of the Safe Links report is to show links that were received. If enabled, the report can also detect and show what links were clicked and by whom. This makes a compelling argument for using the report because it allows you to understand which of your users might need extra security training.

If a URL appears in your Safe Links report that you consider safe, you can add the URL or the domain to your Safe Links allow list. Always do due diligence on URLs. A reason exists for the URL to be blocked. Just because your organization might use the target website does not mean the content is safe.

## Attack Simulation Training

Despite the best efforts of Microsoft Defender for Office 365 (or any other email security solution), it is still possible for malicious emails to reach end-user mailboxes. Therefore, a comprehensive security program will include end-user security awareness training and periodic testing of the program's efficacy. One of the tests is whether end users fall for phishing messages or open attachments to suspicious emails. If you have Microsoft Defender for Office 365 Plan 2, you can automate this testing with Attack Simulation Training.

Attack simulation training can simulate a spear-phishing attack that attempts to harvest valid credentials through phishing messages. A spear-phishing campaign will send sample phishing emails to your users that link them to a realistic webpage that attempts to collect their credentials. If a user clicks the link and enters their credentials, they will be flagged for follow-up. Attack Simulator can also send suspicious-looking emails with attachments that will be tracked if they are opened, or in the case of an attachment with a malicious link, the user will be taken to a page to try and collect their credentials.

You can also use the Attack Simulator to perform password spray attacks and password testing on users in your tenant. Attackers use these techniques to test common, weak passwords against a broad set of users in a tenant. Once the attacker identifies the password for an account, they can use the account for malicious purposes. Multi-factor authentication is a crucial defense against these attacks. [Azure AD Password Protection](#) also works to prevent the use of weak passwords in hybrid environments.

To try out [Attack Simulation Training](#), access the **Defender portal** and navigate to **Email & collaboration > Attack simulation training**.

The **Overview** tab compiles the information of your attack simulation strategy into the following widgets:

- The **Recent Simulations** widget shows any recent simulations. Clicking the **View all simulations** button takes you to the *Simulations* tab. Clicking the **Launch a simulation** button launches the attack simulation wizard (covered in the next section). Note that this widget does not show any of the automated attack simulations.
- The **Recommendations** widget identifies simulated attacks you should launch or schedule for your organization.
- The **Simulation coverage** widget identifies the percentage of users who have and have not received a simulated attack. Clicking the **Launch simulation for non-simulated users** button allows you to create a simulated attack to remedy that gap.
- The **Training completion** widget identifies the percentage of users who have completed the required security awareness training. Clicking the **View training completion report** button shows how each user is progressing with their security awareness training. This report can be exported as a CSV.

- The **Repeat offenders** widget identifies users who have consecutively failed the simulated attacks. Clicking the **View repeat offenders report** button gives you a graph identifying the number of users who have failed each social engineering technique. This includes providing credentials to a fake portal, opening attachments that contain malware, and clicking malicious links. The chart below the graph identifies each user, how often they have repeated the offense, and the simulation types they have failed.
- The **Behavior impact on compromise rate** widget tries to determine the efficacy of the security awareness training and how susceptible the organization's user base is to compromise as a percentage. Clicking **View simulations and training efficacy report** expands the data surfaced by the widget into a more detailed report. This report identifies each simulation, the predicted compromise rate, the actual compromise rate (how well the organization did), the total target users, and how many of those total users fell victim to the simulated attack.

The **Simulations** tab identifies all currently configured simulations and their state (we will cover launching an attack simulation in the next section). States include simulations in draft (saved but not yet submitted), scheduled, in progress, completed, failed, and canceled.

If you select a simulation, you can review the report of that simulated attack. This includes how many users were compromised by the simulation, how many users reported the simulation as a phishing attack, which payloads were used in the simulation, recommended actions to protect users from this type of attack, and the progress of any user awareness training attached to this simulation. Other details include when the simulation was launched, when the simulation will end, how many users were targeted, and the current status of the simulation. To cancel a simulation, select the three dots on the far right of the simulation entry, and click **Cancel Simulation**.

The **Automations** tab allows you to automate one or more attack simulations (we will cover setting up an automated attack simulation in a later section).

The **Simulation content library** tab allows you to create your own custom payloads. Customizing the payload is helpful since you can create a phishing message or fake logon portal that is a more realistic representation of your organization using your logo, color scheme, and other convincing content.

Microsoft offers many different email payloads, including emails with the purpose of credential harvesting, delivering malware as either an attachment or a link, inserting links to websites with malicious code, or convincing the user to grant OAuth consent. Depending on the payload selected, you can configure the sender's display name and email, the subject, phishing links, a malware attachment, language, and theme. You can also import or copy an email into the text editor. A code editor is also available for advanced configuration of the email.

Microsoft offers several templates to create your own fake login portal, including fake portals for LinkedIn, GitHub, and Microsoft. These templates can then be customized either in the text or code editors.

The simulation content library tab also contains the end-user notifications. End-user notifications can be customized, offering either positive reinforcement if the user submitted the simulation as phish or a failure notification if the user fell victim to the simulated attack. This could include links to additional security training and reminders for the user to complete that training.

The **Settings** tab allows you to define what is considered a repeat offender. A repeat offender is someone who has fallen victim to consecutive simulations. The default number is two but can be set higher.

## Launching a simulated attack

To launch a simulation, click the **Launch a simulation** button. The *Select Technique* page (Figure 7-17) identifies all social engineering types, including:

- **Credential harvest** includes an email that links to a fake portal designed to trick users into inputting their credentials on a malicious website.
- **Malware attachment** includes an email with an attachment designed to run malicious code or a macro to compromise or gain access to the user's device, allowing the bad actor to launch additional attacks.
- **Link in attachment** includes an email with an attachment that has an embedded URL. The embedded URL takes the user to a fake portal to trick the user into entering credentials.
- **Link to malware** includes an email with a URL to a malicious file (for example, mimicking a well-known sharing site like SharePoint Online or Dropbox) which could contain malicious code or macros.
- **Drive-by URL** includes an email with a URL to a website that runs malicious code designed to compromise or gain access to the user's device, allowing the bad actor to launch additional attacks.
- **OAuth consent grant** uses an app created by a bad actor, prompting the user to grant permissions to the app. This could allow the app to harvest data or perform other attacks on the user.

Select a technique and click **Next**. Give the simulation a **Name** and **Description** and click **Next**.

From the *Select payloads and login page*, select the payloads sent to users. A payload can contain both an email and a login portal. You can also test each payload from this page by clicking the **Send a test** button. This test is sent to the currently logged-in user. This test is not included in the simulation reporting. When you have selected your desired payload, click **Next**.

On the *Target Users* page, you can include either specific users or groups or the entire organization. Selecting specific users is a great way to test an attack simulation on a small scale before targeting the entire organization. Scope the users for the automated simulation and click **Next**.

On the *Assign training* page, you can assign training to users.

- Select **No training** if you don't want to add any follow-up training to the attack simulation.
- If you want to add a custom training URL, which could be directed to an internal learning management system or security training portal, select **Redirect to custom URL** and enter the appropriate URLs.
- If you want to use security training curated by Microsoft, select **Microsoft training experience (Recommended)**. The Microsoft training experience includes a dozen videos between 3 and 7 minutes. When you select the Microsoft training experience, you can either **Assign training for me (Recommended)** or **Select training courses and modules myself**.

With the custom URL or Microsoft training experience, you can set a training due date of 7, 15, or 30 days.

On the *Select end user notification* page, pick whether to notify users about the results of their participation in the attack simulation. This is a great way to provide positive reinforcement if a user correctly identified and reported a message as phishing or an opportunity to provide security training and awareness to users who fell victim to the simulation. Options include:

- **Do not deliver notifications** if you do not wish to have user involvement with the simulation.
- **Microsoft default notification (recommended)** if you want to use the default Microsoft notifications. This selection allows you to define the delivery preferences for the Microsoft notifications (such as when to deliver the notification) and allows you to review the notification layout before sending.
- **Customized end user notification** if you want to provide your own notifications. We discuss notifications in greater detail while covering the *Simulation content library* tab.

On the *Launch Details* page, select whether to execute the simulation as soon as the wizard is complete or schedule the simulation to start later. With either selection, you must also define (in days) when the simulation should end. The default is two days, ranging from 2 to 30 days. With the launch schedule defined, click **Next**.

Before submitting the simulation, you can test it. Click the **Send a test** button to send the simulation to the currently logged-on user. If the test looks good, click the **Submit** button.

## Automating a simulated attack

To create a new automated attack simulation, click the **Create Automation** button. Next, enter a **Name** and **Description**, then click **Next**.

The *Select Social Engineering Technique(s)* page (Figure 7-17) allows you to select one more attack simulation. Select the simulations you want to automate and click **Next**. From the *Select payloads and login page*, you can manually select or randomize the payloads sent to users. A payload can contain both an email and a login portal. Click **Next**.

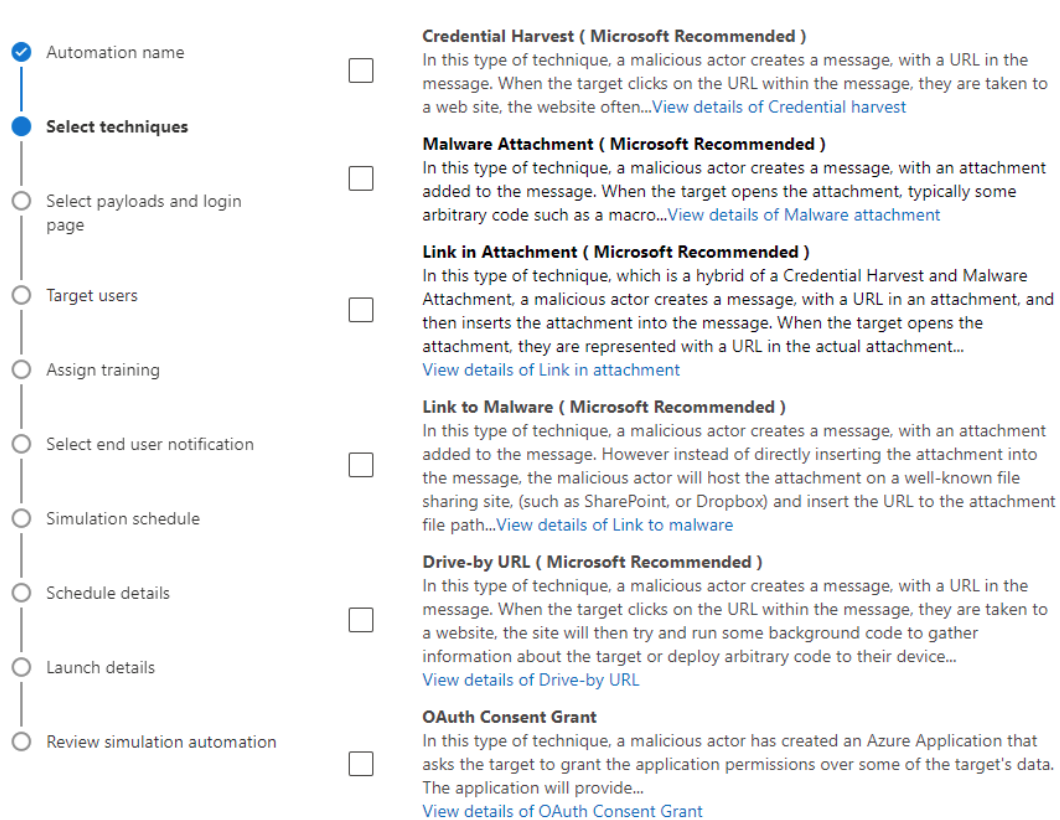


Figure 7-17 Automated Attack Simulator Setup Wizard showing Social Engineering Techniques

The *Target Users*, *Assign Training*, and *Select end user notification* pages are identical to the launch simulation wizard covered in the previous section. For more details on these pages, see the previous section.

On the *Simulation Schedule* page, you can choose whether to create a **Randomized** or **Fixed** schedule to launch the attack simulation. Click **Next**.

- If you selected a fixed schedule, the *Schedule Details* page will ask for **Simulation Start** and **Simulation End** dates. Select a date from each date picker. You can then define how often the simulation reoccurs and whether the reoccurrence is monthly or weekly. If you select weekly, you can choose the day of the week for the occurrence. If you selected monthly, you can choose the day of the month for the occurrence.
- If you selected a randomized schedule, the *Schedule Details* page will ask for **Simulation Start** and **Simulation End** dates. Select a date from each date picker. Then, from the *Simulation Scoping* section, pick which days of the week the randomizer can send simulated attacks. For example, if your organization only works Monday through Friday, you may only wish to select weekdays. The needs of the business should also dictate this schedule. For example, if client invoicing or payroll is every

Friday, you may not wish to launch simulated attacks on that day. You can also choose to randomize the delivery times of the simulation emails.

With the schedule and reoccurrence settings defined, click **Next**.

On the *Launch details* page, you can select additional options such as whether to include the same users on each simulation, whether to target repeat offenders who had fallen victim to the simulation, whether to enable region-aware delivery (so phishing messages are sent during working hours), and whether to use unique payloads for different users. Once you have your options defined, click **Next**.

Review the simulation settings and click **Submit**.

When you submit a simulation, it is saved in a disabled state. To enable the simulation, select it from the **Automations** tab and click the **Turn on** button. When enabled, the *Status* column will switch to *Active* and show the next launch time. If you wish to disable the simulation, select it from the *Automations* tab and click the **Turn off** button. The *Status* column will report *Inactive*, and the *Next launch time* column will be blank. You can also delete a simulation from this tab.

## Using Third-Party Attack Simulators

Some organizations might prefer to use third-party attack simulators to send phish or impersonated emails to their end-users. Traditionally, using third-party simulators requires adding those organizations to various allow lists under several policies, so Microsoft would not act on the simulated attack.

With the advent of Microsoft's [Secure by Default](#) initiative, customer-defined entries in IP allow lists, sender allow lists, domain allow lists, and transport rules are not honored if the message contains malware or is classified as high confidence phish. Therefore, if you previously added your third-party vendor for email attack simulation to any of these lists, their simulations will be blocked. Instead, you should transfer these configurations to Advanced Delivery.

### Configuring Advanced Delivery

To add a third-party attack simulator, open the **Defender portal** navigate to **Policies and rules > Threat Policies > Advanced Delivery**, and select the **Phishing Simulation** tab. From this tab, click **Add**. The *Add Third-Party Phishing Simulation* dialog (shown in Figure 7-18) lets you add all **Sending Domains**, **Sending IPs**, and **Simulation URLs** associated with the third-party attack simulator.

### Add Third Party Phishing Simulations

Phishing simulations are attacks orchestrated by your security team and used for training and learning. Simulations can help identify vulnerable users and lessen the impact of malicious attacks on your organization.

Third party phishing simulations require at least 1 entry for **Sending domain** and at least 1 entry for **Sending IP** categories below. Simulations URLs to allow is an optional field. Specify URLs here to not block or detonate on for your phishing simulation.

The screenshot shows a configuration interface for adding third-party phishing simulations. It is divided into three sections, each with a title and a count of items, and a search input field below it. The first section is 'Sending Domain (1 items)' with a search box containing 'psm.knowbe4.com'. The second section is 'Sending IP (1 items)' with a search box containing '147.160.167.15'. The third section is 'Simulation URLs to allow (0 items)' and is currently empty.

Figure 7-18: Adding third-party phishing simulations

## Investigations

Investigations, or Automated investigation and response (AIR), are part of Microsoft Defender for Office 365 Plan 2. AIR runs a security playbook after an alert is triggered within your organization, either as an automatic response to well-known phish/malware events or manually by an administrator. As part of the playbook, additional information about the event is gathered, and a set of actions is returned as recommendations. However, AIR does not automatically perform remediation actions - an administrator must manually approve recommended remediations.

To access the ongoing and recently completed investigations list, open the **Defender portal**, and navigate to **Investigations**.

By default, this dashboard is filtered to show investigations from the last 24 hours. You can expand this date range by modifying the filter. The filter also lets you scope the dashboard to a specific status or investigation type. For example, you could filter the results to focus on investigations that uncovered compromised users.

To get the properties on a specific investigation, click the **Open in new window** icon. This will launch a new page with the following tabs.

- **Investigation graph** gives you a visual overview of the selected item. You can click on the individual graph components to be transported to one of the other tabs. For example, the investigation graph in Figure 7-19 shows that 7 entities were analyzed (1 email, 1 file, and 5 clusters). If you click the entities analyzed icon, it transports you to the entities tab.
- **Alerts** list the alert that triggered the investigation. You can select each alert to get the details on what caused the alert. Using our example from Figure 7-19, this alert was caused due to ZAP retroactively detecting malicious content in an email. In our case, the alert identified the email (noting the recipient, subject line, sender address, sender IP, and date sent.)
- **Mailboxes** lists impacted users. You can select each user and click the **More details about user** link if you want to see all investigations that the user has been a part of in the last 7 days. This is useful to see if the user is a victim of repeated digital attacks.

- **Evidence** identifies all email items related to the investigation, including the originally reported item and any items found as part of the hunting phase. This could include the message body, attachments, and subject line.
- **Entities** include all the individual objects part of the investigation, such as the user, message, files, IP addresses, etc.
- **Log** gives details on the execution of different steps of the investigation.
- **Pending Actions** lists the actions Exchange Online is recommending taking. For example, the action could be to soft or hard delete the message. From here, the admin can either **Approve** or **Reject** the action. If an admin rejects the action, they must provide a reason. From the action screen, the admin can view the message headers, download the email for further analysis, or open the item in Explorer (covered in the next section).

Figure 7-19 shows the investigation graph of an investigation.

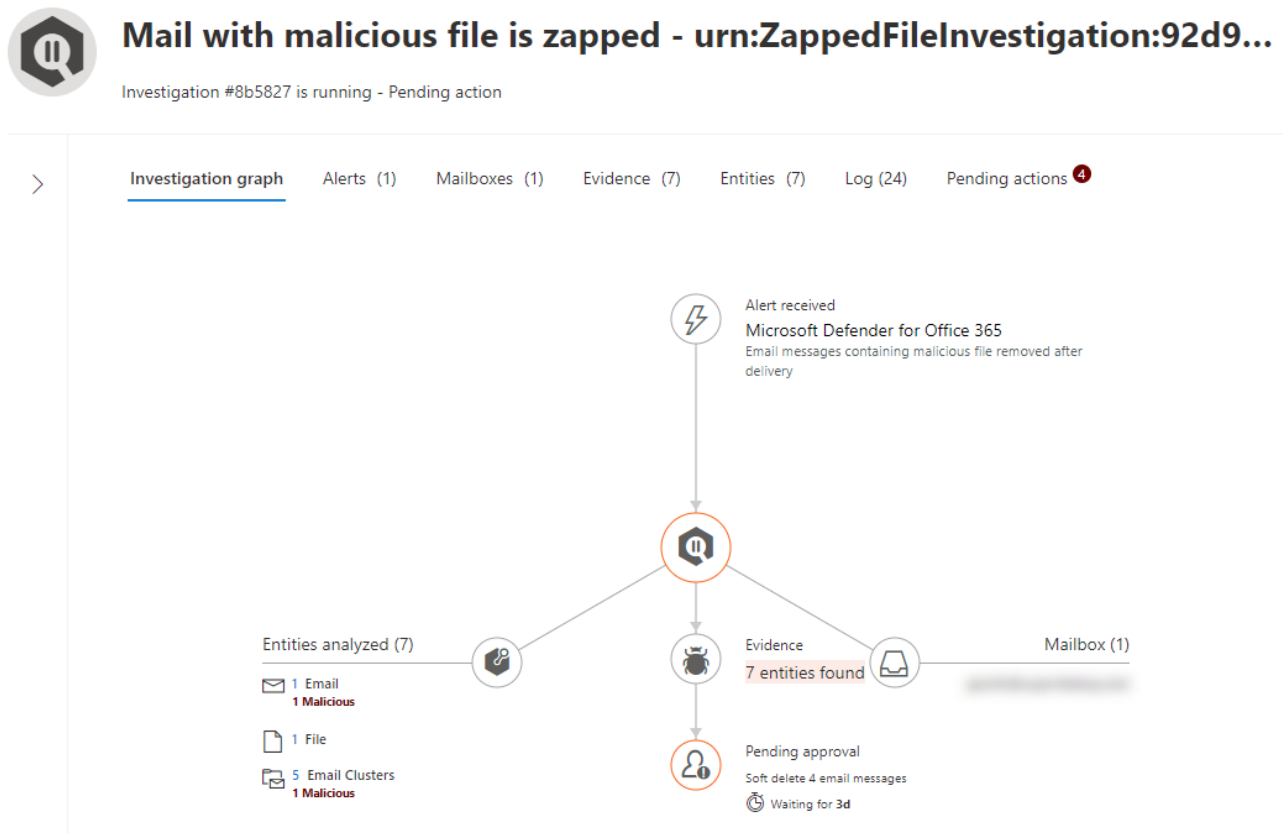


Figure 7-19: Investigation graph of an investigation

## Explorer

Explorer is a great way to see what actions Exchange Online takes on email. The top half of the dashboard displays a stack graph (Figure 7-20) identifying the quantity of email that has been delivered normally, delivered to junk, or blocked. The bottom half of the dashboard provides the individual data making up the graph.

By default, the graph only shows the last two days, but this date range can be increased to the last 30 days using the date and time pickers at the top of the screen. Once you pick a new start or end date and time, click the **Refresh** button.

You can also filter results. For example, to filter all actions taken on messages from a specific sender, select **Sender** from the drop-down, type the sender's email address and click the **Refresh** button. Note that you can



filter multiple values by providing a comma-separated list in the text box. In our example, we could type multiple sender email addresses separated by a comma. There are dozens of other filters, such as looking at messages coming to or from a specific connector, actions taken, impersonated users and domains, and more.

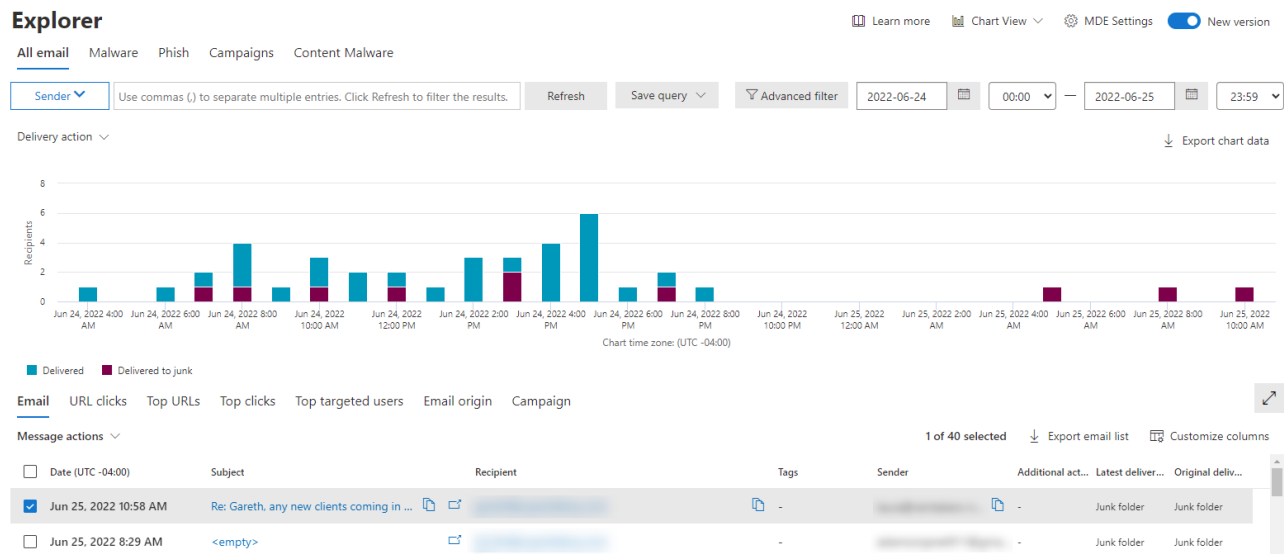


Figure 7-20: Explorer showing mail delivered, delivered to junk, or blocked

If you want to apply multiple conditions to your filter, click the **Advanced filter** button. This filter pane allows you to select one or more conditions and the operator to connect these conditions (for example, an OR operator).

The table in the lower half of the dashboard has multiple tabs allowing you to scope the table to specific results. This includes email, URL clicks, top URLs, top clicks, top targeted users, email origin, and campaigns. Each tab returns different data.

For example, the default **Email** tab identifies each email's delivery date, subject, recipient, sender, and delivery action. Selecting an email and clicking the **Open in new window** button brings up additional detail about that email. This additional detail includes a timeline through Exchange Online, a detailed analysis of the email (which includes message headers, policy actions, and more), if the email contained any attachments or URLs (and any malicious verdicts), and if the emails were related or similar to any other emails. For additional analysis, an admin can preview or download the email from this view.

In contrast, the **Top URLs** tab identifies how many emails a URL appeared. The count columns can also be sorted, allowing you to identify the top URLs found in blocked and junked emails. Selecting a URL allows you to see which emails included the link, as well as WHOIS information on who owns the domain attached to the link. You can view more about the email that contained a URL by clicking the **Open in new window** button.

The **Top Targeted Users** is a great way to identify who is targeted by these digital attacks. Each user identified will have a corresponding number of attempts by their email address. Selecting the user will filter explorer to just that user.

You can also save your queries using the **Save Query** button. This is useful if you need to track a specific event over time. Saved Queries appear in the **Defender portal** on the **Threat Tracker** page.

## Threat Tracker

Threat tracker allows you to access all your previously saved queries from Explorer. Select a query and click the **Explore** link to use a saved query. This will launch that query on the **Explorer** page. For more information, check the previous section on Explorer.

Threat Explorer allows you to modify an existing query by clicking the **Edit** button. You can also delete saved queries by clicking **Delete**.

# Monitoring and Troubleshooting Mail Flow

Microsoft provides a variety of tools and reports to help you keep track of mail flow in your organization. For example, when something goes wrong, tools like message tracing let you explore precisely what happened to a message so that you can troubleshoot a problem. In addition, if you look deep into email messages, there are a plethora of diagnostic message headers that EOP and Exchange Online add to help explain precisely what happened to the message when Microsoft processed the item.

## Mail Flow Dashboard

The mail flow dashboard (accessed through the Home tab of the Exchange Admin Center) is a collection of widgets and reports with associated alerts and recommendations to help you review your mail flow at a glance and identify potential issues.

Figure 7-21 shows the mail flow dashboard. As with most other dashboards, you can customize the arrangement of the widgets to suit your needs. Most of these widgets link back to a mail flow report under the Reports > Mail Flow section in the Exchange Admin Center.

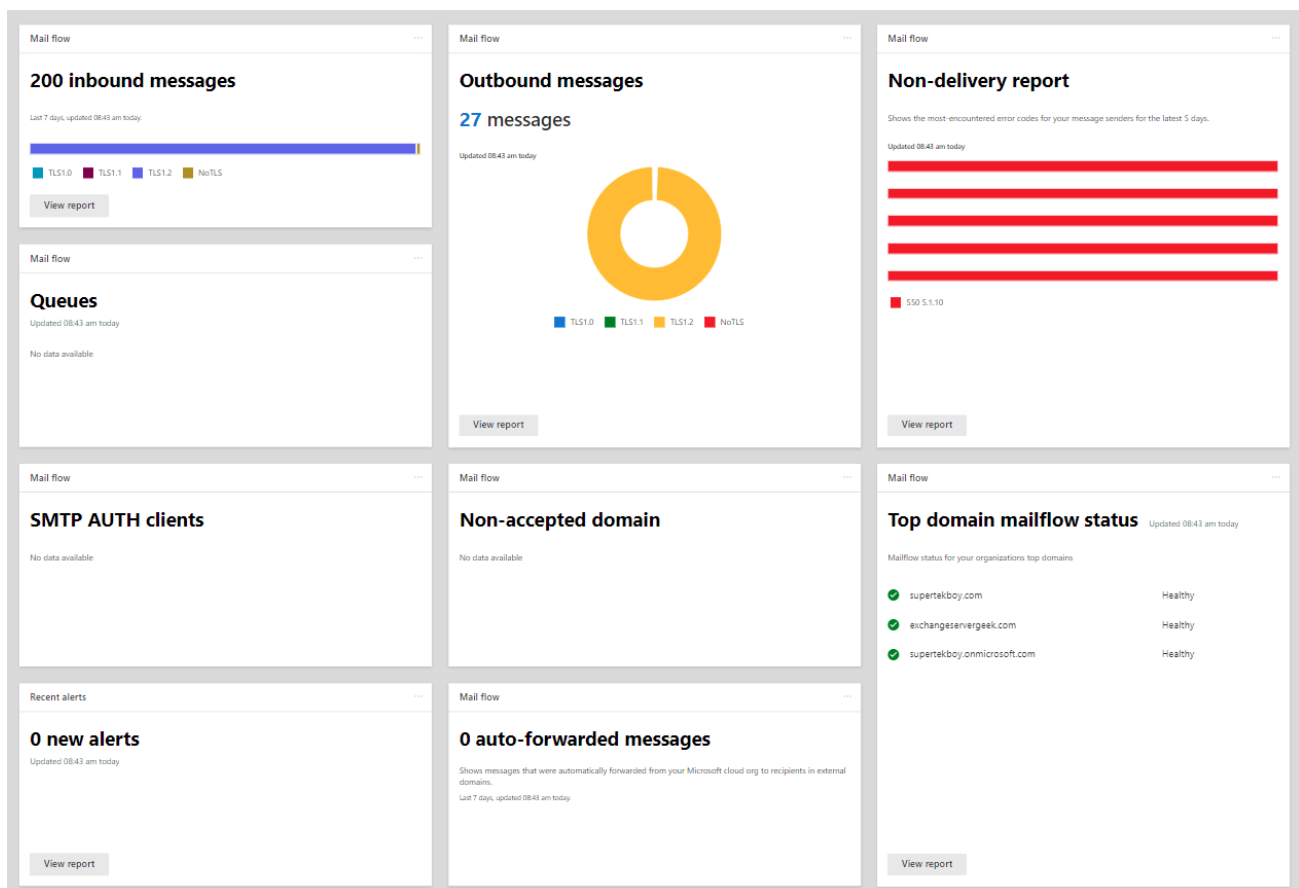


Figure 7-21: Mail Flow Dashboard

## Inbound Messages

The inbound messages widget (Figure 7-21) displays the number of inbound messages received in the last 7 days. Below the graph, a legend identifies each TLS protocol (including messages received without TLS). You

can select each TLS protocol in the legend to focus the graph on just that protocol. Selecting the **View Report** button redirects you to the inbound messages report (under **Reports > Mail Flow**).

The inbound messages report (Figure 7-22) provides greater detail than the widget. This dedicated screen displays daily mail volume, percentage of mail received by TLS protocol, and a detailed line item report breaking down daily mail by each connector.

The various drop-downs allow you to change the date scope between 7, 30, and 90 days (including a custom date range) or select a specific inbound connector (or all inbound connectors). You also have the option to filter the results with the **Filter > New Filter** button. To export the report to CSV, click the **Export** button. Alternatively, click the **Request Report** button to send the report directly to your email.

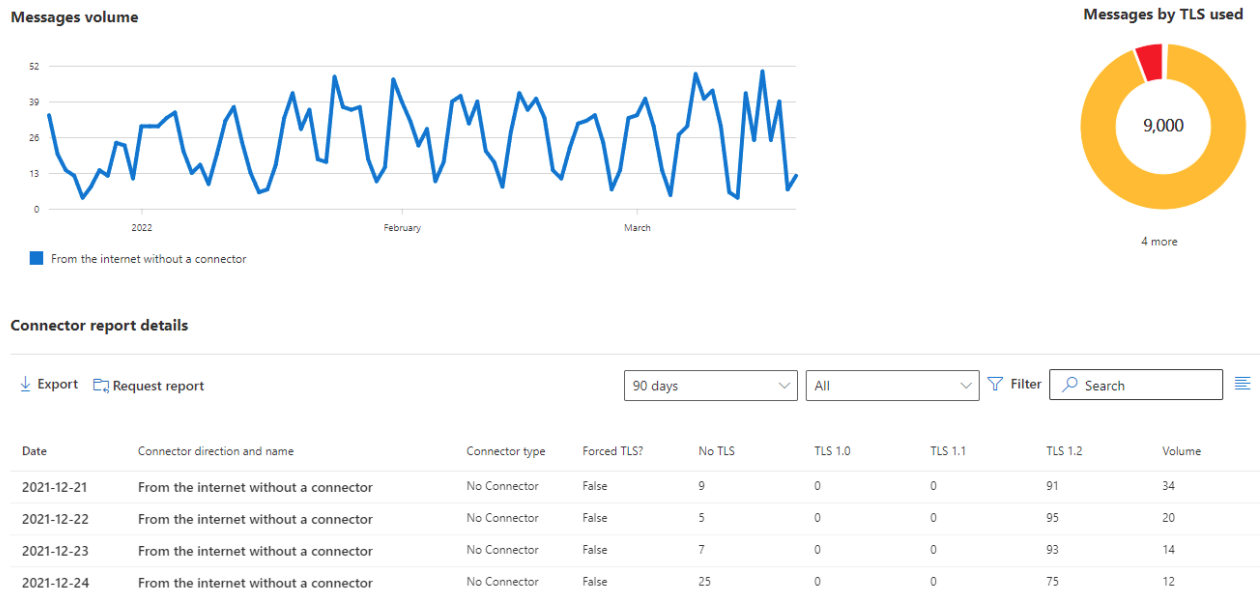


Figure 7-22: Inbound messages report

## Outbound Messages

The outbound messages widget (Figure 7-21) displays the number of outbound messages received in the last 7 days as a donut chart. Below the chart, a legend identifies each TLS protocol (including messages received without TLS). You can select each TLS protocol in the legend to focus the chart on just that protocol. Selecting the **View Report** button redirects you to the outbound messages report (under **Reports > Mail Flow**).

The outbound messages report is identical to the inbound report shown in Figure 7-22. Like the inbound report, you can see outbound mail volume by day, percentage of mail sent by TLS protocol, and a detailed line item report breaking down daily mail by outbound connector.

The drop-downs also operate similarly to the inbound report, allowing you to change the date scope, select a specific outbound connector, or filter the results. To export the report to CSV, click the **Export** button. Alternatively, click the **Request Report** button to send the report directly to your email.

## Non-Delivery Report

This widget shows the most common error codes over the last 5 days, including the number of times that an error was recorded on a given day. For example, Figure 7-21 shows that our only error code in the last 5 days is 550 5.1.10, which equates to "Recipient not found." A fairly innocuous error that is likely due to a mistyped email address.

The **View Report** button redirects you to the non-delivery report (located under **Reports > Mail Flow**). This report gives us a detailed overview of the volume of error codes we have received each day; selecting a key

on the legend filters the graph to that specific error code. This is useful when we want to identify patterns in our errors. For example, in Figure 7-23, we primarily receive error 550 5.1.10 as identified by the widget. However, on 3/10/22 and 3/7/22, we also received error code 550 5.7.133, which means we do not have permission to send it to a distribution list. We can then take the message IDs under the Sample Messages column and perform a message trace to investigate the failure further.

## Non-delivery report

Monitor messages that aren't getting delivered to the intended recipients. When a message can't be delivered, the sender gets an emailed non-delivery report (NDR) with an error code that indicates why the message wasn't delivered. This page shows the details of the NDRs and helps you troubleshoot the issues. [Learn more](#)

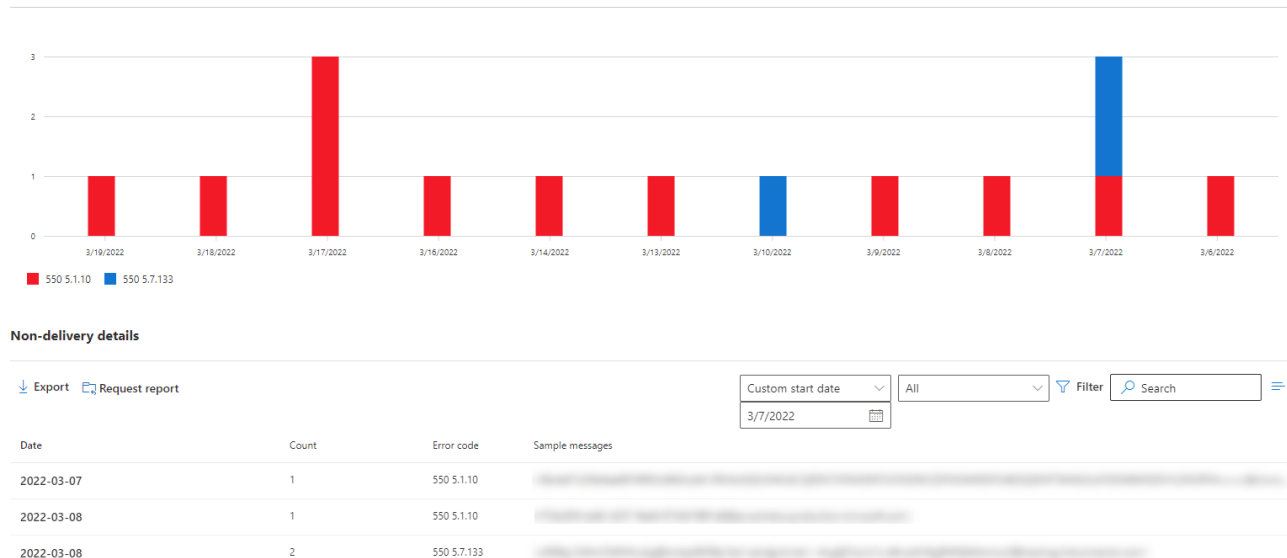


Figure 7-23: Non-delivery report

## Queues

The **Queues** widget is another default part of the dashboard. The Queues widget shows you the number of messages stuck in a mail flow connector for more than one hour. If the number shown is greater than zero, you can click it to go through a series of flyout reports showing information such as the number of queued messages and the connector name (linked through to EAC, where you can fix the connector in error).

The number of queued messages can also be clicked through to show a message trace report for the items stuck in the queue, showing further information on the problem. For example, a queue containing the wrong smart host might show the target host refused to allow an SMTP connection, otherwise known as a socket error.

A companion report is available for this widget by navigating to **Reports > Mail Flow > Queued messages report**.

## SMTP AUTH Clients

Clients that connect directly to SMTP servers are known as SMTP AUTH clients. An example of an SMTP AUTH client could be a multi-function printer performing scan-to-email that relays directly to Exchange Online via smtp.office365.com, rather than an on-premises mail server. The TLS version used by these clients is visible in this widget. This widget will be empty if you do not leverage SMTP AUTH clients.

Understanding the TLS version used for SMTP AUTH is vital as Microsoft is deprecating the use of TLS 1.0 and TLS 1.1 from their primary smtp.office365.com endpoint. Any client needing TLS 1.0 and TLS 1.1 must use the legacy-smtp.office365.com endpoint instead (see the Device and app mail relay to Exchange Online section).

A companion report is available for this widget by navigating to **Reports > Mail Flow > SMTP AUTH clients report**.

**Note:** To reduce your attack surface, Microsoft recommends that organizations enable SMTP AUTH only on mailboxes that require access to the protocol. See the Exchange Online chapter for more information on this topic.

## Non-Accepted Domain

The non-accepted domain widget identifies email relayed from your on-premises environment using a domain not registered in Office 365. For example, if an on-premises application sends mail as "app@office365itpros.co.uk" and the only domain in your tenant is "office365itpros.com," this widget will identify that issue.

Making sure to only send messages from domains registered to your tenant is incredibly important. Microsoft will reject any mail where the *From* address does not match one of the domains in your tenant. This is particularly important in tenant-to-tenant migrations where a domain is removed from one tenant and added to another. During this transition, any on-premises relays must be updated to the new tenant.

A companion report is available for this widget by navigating to **Reports > Mail Flow > Non-accepted domains report**.

## Top Domain Mail Flow

You must have the correct MX record published in DNS to receive emails for a domain. A change in the MX record could result in an email outage. This widget gives you a high-level overview of potential health issues for your registered domains (including the tenant domain). A green checkmark and a status of *Healthy* indicate everything is good.

Selecting the **View Report** button (Figure 7-24) redirects you to the *Top domain mail flow status report* (located under **Reports > Mail Flow**). This report gives us a detailed overview of each domain, its health, the previous and current MX records (if it has recently changed), and if you have received an email to that domain in the past 6 hours.

### Top domain mailflow status report

Domain	Domain status	Previous Mx Record	Current Mx Record	Emails received (past 6 hours)
<input type="checkbox"/> supertekboy.com	Healthy	supertekboy-com.mail.protection.outlook.com	supertekboy-com.mail.protection.outlook.com	Yes
<input type="checkbox"/> exchangeservergeek.com	Healthy	exchangeservergeek-com.mail.protection.outlook.c...	exchangeservergeek-com.mail.protection.outlook.com	No
<input type="checkbox"/> supertekboy.onmicrosoft.com	Healthy	supertekboy.mail.protection.outlook.com	supertekboy.mail.protection.outlook.com	No

Figure 7-24: Top Domain Mailflow Status Report

This report will also identify when domains have expired, which would result in the domain no longer resolving and causing mail flow and other issues.

## Recent Alerts

Security is always something to review on an ongoing and consistent basis. The **Recent Alerts** widget identifies if any of your Exchange Online mailboxes are potentially compromised. The most common item displayed here is alerts on forwarding rules (see the next widget report). Click the **View Report** button to see a report of all alerts in the last 90 days. This report is also accessible by navigating to **Mail Flow > Alerts**.

## Auto-Forwarded Messages

The auto-forwarded messages widget is designed to help you combat a common approach to data exfiltration. For example, this often happens during a business email compromise attack when the attacker

wishes to understand a victim's mail traffic patterns before they send a phishing email to try and lure someone into taking action, such as making a payment. It is recommended that you limit this behavior so that email is only forwarded outside the organization when justified by business reasons.

The widget identifies the number of forwarded items in the past week, including messages forwarded by a user via an inbox rule, their OWA forwarding options, or an administrator who has configured a transport rule or SMTP forwarding.

A companion report is available for this widget by navigating to **Reports > Mail Flow > Auto-forwarded messages report**. The **Summary** tab identifies all forwarded messages, the method used to forward those messages (inbox rules, mail flow rules, or SMTP forwarding), and the forwarding users and domains. The **New Activity** tab shows any new forwarding in the past 7 days.

It is prudent to frequently review this information and block or remove any rules that forward emails as appropriate. It is also worth noting that you can create alert policies to signal an alert when a user forwards emails from their account. Also, Microsoft Secure Score awards a higher score to tenants who do not allow email forwarding.

## Other Widgets

Other widgets (not shown in Figure 7-21) include:

- **Migration Batch** identifies the success or failure of recent migration batches to or from Office 365. Selecting the View Report button takes you to the Reports > Migration tab, where you can dive into the nature of the success or failures of each batch.
- **Mailboxes** provide quick links to everyday administrative actions such as adding a shared mailbox, managing email forwarding, or hiding a mailbox from address lists.
- **Training and Guide** provide links to Exchange Online documentation and training for administrators.

## Reports

Most widgets described in the previous section link to a report under **Reports > Mail Flow**. However, the reports tab also contains additional reports that do not have a corresponding widget. These reports are discussed below.

### Mail flow map report

The *Mail flow map* report provides a visualization of mail flow in Office 365. This is useful to identify trends or anomalies in your mail routing.

- The center of the visualization is Office 365. This shows the count of internal messages in the tenant.
- Messages coming into the tenant, including mail from the internet and any custom connectors, are shown to the left of Office 365.
- Messages leaving the tenant, including mail to the internet and any custom connectors, are shown to the right of Office 365.

The table below the visualization identifies each domain exchanging mail with your tenant in the last 90 days. You can also refine this report by clicking the **Filter** button. For example, to only see messages received by a specific domain, select **Send or receive domain** option from the *Field* drop-down, pick an operator and specify the domain. You can export these results for further analysis to CSV by clicking the **Export** button.

### Mailboxes exceeding receiving limits

The *Mailboxes exceeding receiving limits* report identifies mailboxes that could be throttled due to exceeding transport limits. When mailboxes are throttled, they will experience delays in sending and receiving mail. Throttling has a clear impact on users and the business, so keeping on top of any potential throttling is vital.

This report comprises a heatmap that shows the top 10 mailboxes impacted in the last 24 hours and a table of all mailboxes that exceed transport thresholds.

For a complete list of transport limits, see the Transport Limits section.

## Reply-all storm protection report

The *Reply-all storm protection report* identifies when a reply-all storm was triggered based on the thresholds configured by the tenant admin. By default, reply-all protection blocks replies to an email thread for 6 hours if it had detected more than 10 reply-all messages within 60 minutes to a thread with over 2,500 recipients. These values (aside from the 60-minute detection window) are configurable from the **Exchange Admin Center** by navigating to **Settings > Mail Flow**.

The goal of reply-all storm protection is to stop chain emails impacting the organization. The *Reply-all storm protection report* identifies each message that triggered reply-all protection. This can be useful if you need to provide guidance or training to the message originator to prevent a reply-all storm in the future. The report attributes could also be used if you need to perform a content search to remove the offending messages.

The configurable ranges for the reply-all storm protection can be found in the Transport Limits section.

## Exchange Transport Rule Report

The *Exchange transport rule* report identifies triggered transport rules in a given period. This data is represented by a graph (Figure 7-25), which identifies the daily volume of each triggered rule, and a donut chart, which breaks down rule volume by direction and severity. You can select a key from each chart legend to focus on a specific component. For example, you can select the outbound key to focus solely on the volume of rules triggered on outbound mail flow.

### Exchange Transport Rule report

View matches for the mail flow rules that are set up for your organization. You can manage these rules in the Exchange Admin Center. [Learn more](#)




Figure 7-25: Exchange Transport Rule report

By default, this report shows the volume for the last 7 days. However, like all other reports, it is possible to change this timespan to 7, 30, or 90 days (including a custom date range). In addition, it is possible to focus the graph on a specific rule, mail flow direction, or severity using the drop-downs. This is useful to see whether a rule is triggering more in one direction than the other (e.g., inbound versus outbound) or if you are looking for rules categorized as high severity.

## Non-Delivery Reports

Whenever a mail server cannot deliver a message because of a hard failure, it generates an NDR and sends the NDR to the message's originator. This happens so that the sender knows the message failed. Compare it to snail mail being "returned to sender" because the post office could not deliver the letter to its destination.

NDRs hold information about why the message delivery failed. The following example illustrates the typical information found in an NDR, saying that the recipient does not exist in the remote email system:

### 550 5.1.10 RESOLVER.ADR.RecipientNotFound; Recipient not found by SMTP address lookup

The receiving server generates the information in an NDR. Historically, the recipient's messaging system generates the NDR and sends it back to the originator. However, most modern messaging systems now generate NDRs based on the receiving server's information at the originating system, which significantly enhances the user experience, as outlined below.

Regular NDR messages are not very helpful for most users and often leave them wondering what to do next. On the other hand, NDRs typically also contain other details which are helpful for an administrator trying to figure out what happened.

Your message to [UserX@domain.com](mailto:UserX@domain.com) couldn't be delivered.

UserX wasn't found at [domain.com](http://domain.com)

michael Office 365  
**Action Required** Recipient  
 Unknown To address

#### How to Fix It

The address may be misspelled or may not exist. Try one or more of the following:

- Send the message again following these steps: In Outlook, open this non-delivery report (NDR) and choose **Send Again** from the Report ribbon. In Outlook on the web, select this NDR, then select the link "**To send this message again, click here.**" Then delete and retype the entire recipient address. If prompted with an Auto-Complete List suggestion don't select it. After typing the complete address, click **Send**.
- Contact the recipient (by phone, for example) to check that the address exists and is correct.
- The recipient may have set up email forwarding to an incorrect address. Ask them to check that any forwarding they've set up is working correctly.
- Clear the recipient Auto-Complete List in Outlook or Outlook on the web by following the steps in this article: [Fix email delivery issues for error code 5.1.10 in Office 365](#), and then send the message again. Retype the entire recipient address before selecting **Send**.

If the problem continues, forward this message to your email admin. If you're an email admin, refer to the **More Info for Email Admins** section below.

Figure 7-26: Layout for NDR message(s)

When a user reports an NDR, it is normal to investigate to try to rectify the problem. However, when a remote system responds that a user is not known or that the remote mailbox is full, there is nothing that you can do except wait for the remote server to fix the problem. Microsoft has invested in the layout and text in the NDRs generated by Exchange Online to reduce support calls. These messages are more descriptive and helpful than is the norm with other SMTP-compliant email systems. In some cases, the advice given helps the user to resolve the problem themselves. Figure 7-26 is a good example. The problem is clear, and the steps necessary to deliver the message to the intended recipient are there for the user to take. In addition, the NDR still has



information that could be helpful to the administrator, like the original message headers and SMTP error code.

## Backscatter Filtering

Spammers often send messages to random recipients on the internet and use someone else's email address to spoof the "From:" header. When one of these messages is sent to an invalid recipient, the receiving email server can generate an NDR, which is sent back to the spoofed email address instead of the spammer. These NDR messages are referred to as backscatter.

NDR messages are very useful. They can tell someone who sent a message that it was not successfully delivered, for instance, because the recipient's mailbox is full or maybe because the email address is invalid. However, false NDR messages are a nuisance for end-users, as anyone who has ever received one for a message they know they did not send will understand.

EOP applies a unique signature to each outgoing message to reduce backscatter. The signature is stored in an SMTP header called *x-microsoft-antispam-prvs*, as illustrated in the following example:

```
x-microsoft-antispam-prvs:  
<AM3PR04MB4028F8BB1D18D36001E6876D9E60@AM3PR04MB402.eurprd04.prod.outlook.com>
```

If an NDR message is generated for a legitimate message, it will include the signature header. When EOP processes the message, the header is recognized, and the NDR message is considered valid. However, if the received NDR message does not have the header, Exchange Online considers the message to be backscatter.

If you use EOP to protect on-premises mailboxes in a standalone configuration or a hybrid deployment, but EOP does not handle outbound mail flow, you should enable the NDR Backscatter feature. Go to the **Defender portal > Policies & rules > Threat policies > Anti-spam**, select the **Anti-Spam inbound policy (Default)** and click **Edit spam threshold and properties**. Under **Mark as spam**, toggle **NDR backscatter** to **On**. In on-premises configurations where EOP protects both inbound and outbound mail flow, you do not need to enable the NDR Backscatter feature because the *x-microsoft-antispam-prvs* header is added to all outbound messages automatically.

## Message Tracing

Most Exchange administrators know the following problem: a user calls the helpdesk to report, "Person X sent me a message a few hours ago, but I still haven't received it." Sometimes messages are misaddressed, the sender has not sent the email, or they are not delivered when expected, or the message was sent but never arrived in the recipient's inbox. For instance, a mail flow rule might forward the message to a different mailbox, or the message may go into the user's junk email folder instead of the inbox.

Message tracing proves particularly helpful to troubleshoot scenarios like the ones mentioned earlier as it allows you to figure out what happened to a message, like when it was delivered or why it was rejected, quarantined, or perhaps deferred. Exchange Online keeps message tracking data for 90 days. Today, you can trace message in several ways:

- Through the **Mail flow > Message trace** section of the EAC.
- Using PowerShell with the *Get-MessageTrace* and *Get-HistoricalSearch* cmdlets.

Exchange Online separates traces into recent messages (up to ten days old) and those over a longer timeframe (from ten to ninety days old). Recent message traces are done interactively, and you see the results on-screen. However, if you need to track messages older than ten days or if you need to create a report with extended information, you use a *historical search* instead. Unlike Exchange on-premises servers, message tracking information only exists online for a limited period before it moves into the reporting data warehouse. Historical searches run in the background, and the results are emailed to you as a CSV file once the search is complete. It may take several hours for a historical search to complete.

## Tracing Messages

To use message tracing in the Exchange admin center, go to the **Mail flow > Message trace** section. You can create new custom traces, choose from one of the pre-defined or previously saved trace reports and view the results of finished/pending traces. The pre-defined trace options are a great way to quickly start a trace you regularly use with the same options. Most of the time, however, you will create a new (custom) trace using one of the following parameters:

- **Sender and Recipient** allow you to specify who sent or received the message. If you need to specify an external sender's address in the recipient picker, type it in manually and click **OK**.
- **Time range**. By default, a trace is scoped to the last 48 hours. However, you can edit the scope to span up to the last 90 days. You can also provide a custom scope with specific start and end dates and times (within the last 90 days). Searches over ten days old are performed as background searches, and the results are emailed to you once they are complete.
- **Delivery status** allows the administrator to target messages based on what action was performed. Among other things, you can specify to search for quarantined messages, filtered as spam, failed, or messages that have yet to be cataloged with getting status. If you are unsure what happened to the message, the default selection of **All** works best.
- **Message ID**. Each message has a unique message ID. You can find the ID from the message headers to scope the message trace to a single message.
- **Direction** specifies if the message was sent (outbound) or received (inbound).
- **Original client IP address** is the IP address of the sender's messaging service/server.

When executing a historical or **extended report** search, you must also specify the following:

- **Report Title**. The subject of the email that is sent to the notification email address.
- **Notification email address** of the person to whom the report should be sent.

**Real-world:** Exchange Online automatically decides when to start a historical search. You can quickly distinguish between a regular search from a historical search when creating a new trace. When the text on the **Search** button changes to **Next**, a historical search is created instead, and you will have to wait for the results to be emailed to you. This is true for all traces for which you request an extended report, regardless of the date range.

## Viewing Trace Results

There are two ways to view trace results: interactively (when running a regular search) or through a CSV report (when running a historical or extended report search). The information exposed through the CSV files can be great for reporting purposes. The format of the file makes it easy to sort through the information and create custom views quickly. However, extended reports are restricted to including a sender or recipient or message ID. Standard searches can be time-based only with no other restrictions. On the other hand, the interactive results are generally more interesting for troubleshooting purposes for various reasons: they show you more information on the various events that apply to the message while in transit. In addition, you can quickly find related messages, making troubleshooting easy.

Each trace result details the status of the message and the actual message events. The events include information about the message as it made its way through the various filtering mechanisms in EOP. Although you can derive the same information from the raw message events, the way results are displayed on screen makes it easy to find information about whether a message was delivered successfully and, in case it was not, why not.

In addition to the message status, each trace has extra information in a more human-readable format. For example, it might tell you how to fix a problem when a message was not delivered successfully, or it will tell

you if a message was forwarded to another recipient. Figure 7-27 shows an example where a message goes to a mailbox that cannot accept the message. Note how the GUI makes it easy to find the status and cause of the problem and suggest a solution.

## Message Center Major Change Update Notification

Copy report text below Prepare and email extended report

Sender	Recipient
o365mc@microsoft.com	[REDACTED]

Received Processed Not delivered

**Status**  
Office 365 received the message that you specified, but couldn't deliver it to the recipient due to the following error:

**Error:** 550 5.1.10 RESOLVER.ADR.RecipientNotFound; Recipient [REDACTED] not found by SMTP address lookup

A non-delivery report (NDR) message was sent to o365mc@microsoft.com. The NDR might provide more details about why the email message wasn't delivered and how to fix the issue.

**How to fix it**  
Ask the sender (o365mc@microsoft.com) to follow the instructions in the NDR to fix this issue. The NDR might also include specific information for email admins. If the sender is unable to fix the issue, ask them to forward you the NDR and then follow the guidance for email admins.

**Message events**

Date (UTC)	Event	Detail
6/25/2021, 8:17 AM	Receive	Message received by: BY5... Message received by: BY5PR11MB4465.namprd11.prod.outlook.com using TLS1.2 with AES256
6/25/2021, 8:17 AM	Fail	Reason: [[LED=550 5.1.10 ... Reason: [[LED=550 5.1.10 RESOLVER.ADR.RecipientNotFound; Recipient [REDACTED] not found by SMTP address lookup];[MSG=];[FQDN=]; [IP=];[LRT=]]

Figure 7-27: Viewing trace results in the Exchange Admin Center

## Tracing Messages with PowerShell

Tracing messages in PowerShell is easy. The sole difference from the approach taken by the Exchange admin center is that you use different cmdlets to start a regular or historical search. A regular search, which can trace messages for up to ten days, is started using the *Get-MessageTrace* cmdlet. For instance, to view all messages from the past 48 hours, run the following command:

```
[PS] C:\> Get-MessageTrace -StartDate (Get-Date).AddDays(-2) -EndDate (Get-Date) | Select Received, SenderAddress, RecipientAddress, Subject
```

Received	SenderAddress	RecipientAddress
9 Oct 2018 08:16:39	michel@eightwone.com	tony.redmond@office365itpros.com
9 Oct 2018 08:14:43	newsletter@bloomandwild.com	deirdre.smith@office365itpros.com
9 Oct 2018 08:13:20	michel@eightwone.com	ben.owens@office365itpros.com...

Additional parameters are available for *Get-MessageTrace* to scope trace results, such as looking for messages sent by specific mailboxes.

In this example, we pipe the results to the *Out-GridView* cmdlet, which makes it easier to peruse the message trace information:

```
[PS] C:\> Get-MessageTrace -StartDate (Get-Date).AddDays(-2) -EndDate (Get-Date) | Select Received, SenderAddress, RecipientAddress, Subject | Out-GridView
```

An example script showing how to use the *Get-MessageTrace* cmdlet to report mail sent externally for selected mailboxes is [described in this article](#).

## Extra Detail in Message Traces

When you run a message trace in the Exchange admin center, Exchange Online automatically returns extra routing information with more details about what happened to the message while it was in transit. When using PowerShell traces, to get the same results, pipe the results from the message trace to the *Get-MessageTraceDetails* cmdlet as shown below:

```
[PS] C:\> Get-MessageTrace -StartDate (Get-Date).AddDays(-2) -EndDate (Get-Date) | Get-MessageTraceDetail | Select MessageID, Date, Event, Detail, Data | Out-GridView
```

## Historical Message Traces

The *Start-HistoricalSearch* cmdlet launches a historical search and is used similarly to *Get-MessageTrace*. For instance, to start a search for messages from a certain sender in the past 75 days and send the report to `admin@office365itpros.com`, you would use the following PowerShell cmdlet:

```
[PS] C:\> Start-HistoricalSearch -ReportTitle "Historical Search" -NotifyAddress "admin@office365itpros.com" -StartDate (get-date).AddDays(-75) -EndDate (Get-Date) -ReportType MessageTraceDetail -Sender Kim.Akers@office365itpros.com
```

JobId	SubmitDate	ReportTitle	Status	Rows
ef59-4710-81f3-f511e0bc2222	8 Oct 2018 16:14:04	Historical Search	NotStarted	0

The *ReportType* parameter controls how much detail is returned. For instance, to include the full routing details, the *ReportType* is changed to *MessageTraceDetail*. The results from the historical search are sent by email rather than displayed on a screen, so piping the output to the *Out-GridView* cmdlet will have no effect. The *Get-HistoricalSearch* cmdlet can be used to check the status of searches.

```
[PS] C:\> Get-HistoricalSearch | ? {$_.Status -eq "InProgress"} | Select SubmitDate, ReportTitle, Status | Format-Table -AutoSize
```

SubmitDate	ReportTitle	Status
8 Oct 2018 16:14:04	Historical Search	InProgress

With *Get-HistoricalSearch*, it is not possible to search for all messages in one search. You need to include a sender, recipient, or message ID.

# Transport Limits

The Radicati Group [Email Statistics Report for 2018-2022](#) estimates that over 3.8 billion people used email in 2018, with the number expected to grow to over 4.2 billion email users by the end of 2022. In total, these accounts send and receive roughly 281 billion messages per day. With so many people using email, it is no wonder that email is the most popular medium for sending unsolicited messages (spam).

As a result, Microsoft imposes several transport-related limits on its service to protect its service from those who would like to abuse Exchange Online by using it to send spam and to ensure that every tenant is treated equally. As you plan to use Exchange Online and EOP for your message routing, you should be familiar with the limits Microsoft imposes.

Without limits, a single user could, in theory, send hundreds of thousands of messages daily. As a result, the transport subsystem would consume enormous resources to process those messages and reduce its capacity

to process legitimate emails. To address this, transport limits are imposed at various levels and are the same for all tenants, regardless of the license plan you subscribe to. Table 7-5 lists the various limits imposed in Exchange Online. See [this page for the complete list of limits](#) imposed by Microsoft for other plans.

<b>Feature</b>	<b>Limit</b>	<b>Additional Information</b>
<i>Recipient rate limit</i>	10,000 recipients per day	To discourage the delivery of unsolicited bulk messages, Exchange Online has recipient limits that prevent users and applications from sending large volumes of email.
<i>Recipient limit</i>	500 recipients	The number of recipients you can add to a single message defaults to 500, but a value from 1 to 1000 can be set per mailbox or as a tenant-wide default.
<i>Message rate limit</i>	30 messages per minute	The number of messages the transport subsystem will process per minute, sometimes referred to as the throttling threshold, over SMTP client submission.
<i>Message size limit (Outlook)</i>	150 MB	The maximum size of a message when sent through Outlook.
<i>Message size limit (OWA)</i>	112 MB	The maximum size of a message when sent through Outlook Web App.
<i>Message size limit (Outlook for Mac)</i>	150 MB	The maximum size of a message when sent through Outlook for Mac.
<i>File attachments limit</i>	250 attachments	The maximum amount of attachments that can be added to a single message.
<i>Subject length limit</i>	255 characters	The maximum length of a message subject.
<i>File attachment size limit (Outlook)</i>	150 MB	The maximum size of a single attachment when sent through Outlook.
<i>File attachment size limit (OWA)</i>	112 MB	The maximum size of a single attachment when sent through Outlook.
<i>File attachment size limit (Outlook for Mac)</i>	150 MB	The maximum size of a single attachment when sent through Outlook for Mac.
<i>SMTP Authenticated Submission limit</i>	3 concurrent connections	Three concurrent connections are allowed to send email messages at the same time from the same mailbox. This limit typically impacts third-party email clients and devices such as multi-function printers configured to log in to a single mailbox. The mailbox the client or device authenticates into allows three concurrent connections as well as storing the items it sends in the Sent Items folder.
<i>Reply All Storm Protection</i>	10 reply-all messages to over 2,500 recipients in 60 minutes (default settings)	Further replies to the conversation are blocked for six hours. An NDR is sent to people who try to respond during this time, saying that the "conversation is too busy with too many people" and educating them not to use Reply-All. The number of recipients, responses, and block duration is <a href="#">customizable</a> . The 60-minute detection window is not.
Scanning limits for the content of attachments	1MB	The mail flow rule conditions enable you to examine the content of message attachments, but only the first 1 MB of the text extracted from an attachment is inspected. This 1 MB limit refers to the text extracted from the attachment, not the attachment's file size. For example, a 2 MB file may contain less than 1 MB of text, so all of the text would be inspected

Table 7-5: Transport limits for Exchange Online

In general, Microsoft does not grant exceptions to any of these limits. Although you can raise a support request to increase the limits, the chances of succeeding are slim. Instead, many of these limits can be worked around. A good example is the 10,000 recipients per day limit. It is common for large organizations to broadcast corporate communication messages to all users. If you have more than 10,000 users, you will not be able to send the message to everyone at once. Instead of adding individual recipients to a message, you should use a distribution group. Each distribution group can hold up to 100,000 recipients and is counted only as a single recipient.

Another example is an organization where the marketing department must send customers messages. At a rate of 30 messages per minute, 'only' 43,200 messages can be sent per day. Microsoft recommends using a third-party mass-mailing service for this purpose, such as SendGrid or MailChimp.

**Real-world:** The actual message size can vary from one client to another and is also different for messages routed outside Exchange Online. In the latter scenario, messages could grow roughly 33% because of extra encoding, effectively lowering the maximum size from 150MB to approximately 112 MB. The above limits also used to be much lower. Over time, Microsoft gradually increased the maximum supported message sizes, but they did not change the default maximum message size. As such, even a new tenant imposes a 35 MB message size limit for every user by default. An [excellent article](#) outlines how to determine a user's current configured limit and how you can modify that limit to a value within the boundaries of the supported maximum sizes described in Table 7-5.

# Chapter 8: Managing SharePoint and OneDrive for Business

**Juan Carlos González**

SharePoint Online and OneDrive for Business are core workloads giving tenants the ability to store documents and information and build modern workplace solutions and services on top of both platforms. This chapter focuses on describing how to manage both services (collectively referred to by Microsoft as **ODSP**) using their admin centers, PowerShell, and available APIs. We also describe some basic SharePoint concepts and points to consider when migrating from SharePoint on-premises and on-premises file servers to SharePoint Online and OneDrive for Business. Finally, we cover how applications and services such as Microsoft Lists, Microsoft Search and Microsoft Viva use SharePoint as their platform.

## SharePoint Online

[SharePoint Online](#) is a core Microsoft 365 service used by more than 200,000 organizations worldwide for Intranets, Team Sites, and Content Management among other scenarios. It is also a key building block for other workloads and services such as Microsoft Teams and Planner. In April 2022, Microsoft said that organizations create 8 million new active sites and generate 100 petabytes of SharePoint Online content monthly. Much of SharePoint's success is driven by a thriving developer community and the many applications they create to run on top of SharePoint. Some of those applications are available for SharePoint Online and more will come, which in turn creates further opportunities for tenants to exploit the platform.

Organizations can use SharePoint Online as a standalone service to build Intranet, Extranet, or other types of portals, or to deliver content management services to other applications. Microsoft 365 Groups and Teams are great examples where applications include SharePoint Online in collaborative experiences:

- **Teams:** Each team has its own SharePoint site. Every regular channel inside the team has its own folder to store messages and files shared between team members. Private and shared channels each have a separate SharePoint Online site restricted to the membership of the channel.
- **Planner:** Stores attachments for tasks in the SharePoint Online site belonging to the underlying Group.

Managing documents and other files is the common thread running through SharePoint Online. In fact, SharePoint has returned to its roots and moved away from an attempt to be an application platform to concentrate on what it excels at. The Microsoft 365 substrate together with the Graph and its growing collection of REST-based APIs is the platform; SharePoint Online contributes the ability to manage large collections of documents and metadata in that platform and offers native extensibility through building blocks such as the SharePoint Framework (SPFx), PnP extensibility model, Site Scripts and Site Templates, Microsoft Lists, Power Apps or Power Automate. SharePoint Online is a critical piece of the overall story. Any deployment that seeks to move workload into the cloud must incorporate SharePoint Online into the plans to extract full value from its cloud investment.

Although some feel that the importance of SharePoint Online is less than Exchange Online sometimes appears to be within the Microsoft 365 ecosystem, they are wrong. SharePoint Online is a critical part of Microsoft 365 and Microsoft 365 could not function without SharePoint Online. It is as simple as that.

**Sites and Site Collections:** In SharePoint on-premises and the early days of SharePoint Online, it was common to discuss information architectures based on creating several site collections or several subsites under a single site collection. Today, thanks to the influence of applications like Teams and Groups, most site collections consist of one site and Microsoft has moved away from talking about site collections in favor of sites. Although a site collection is the more correct term and still appears in the documentation, a site is equally valid and is, in most cases, what you deal with.

## Basic SharePoint Online Concepts

Those who have not used the on-premises version of SharePoint are probably unaware of the basic concepts that underpin the operation of the application. Here is a quick guide to the major SharePoint terms and components.

### Sites

[Since 2016](#), Microsoft has modernized SharePoint and delivered change over several waves of enhancements. Modern team sites and communication sites, Hub sites, modern site pages, new SharePoint start page, or Power Apps and Power Automate integration are some examples of the enhancements introduced.

A **Site** is the basic building block for SharePoint Online and is the place where content is kept and managed. To enable easier administration and to ensure that sites that have a common purpose are grouped together, sites are organized into **Site Collections**. Within Microsoft 365, the rule is to use single-site collections. For example, the sites used by Groups (including those used by Teams and Yammer) are single-site collections. This book refers to site collections as sites.

Sites are the entry point to the SharePoint Online start page. SharePoint Online decides what sites appear on a user's start page based on the user's recent activity. Signals stored in the Microsoft Graph gathered from user interaction with SharePoint Online team sites and group document libraries (including those used by Teams) are the basis for suggested sites. From the home page, you can create new sites (modern team sites or communication sites) and news items posted to any site the current user can access.

Each site can have its own security, layout, theme, navigation, regional settings, and custom search settings. A site can also have workflows built-in to ensure that the information held on the site is processed in a specific manner. The latest form of sites is known as modern SharePoint Online sites, including "group-enabled" team sites, modern team sites, and communication sites. A "Group-enabled" site is associated with a Microsoft 365 Group, which is responsible for the management of the site membership. [Communication sites](#) are mobile-friendly and configured to communicate and dynamically display information. Their role is to serve up information to users drawn from different sources available to SharePoint. Finally modern team sites are similar to "group-enabled" team sites but without an underlying Microsoft 365 Group.

A Home site is a particular kind of Communication site bringing together news, events, embedded video and conversations, and other resources to deliver an engaging experience that reflects your organization's voice, priorities, and brand. Currently, there is a limitation of having a single Home site per SharePoint tenant.

### Hub Sites

Communication sites present news and visually rich content, making them the best choice for sites around a single topic. The list of web parts available for Communication sites and the pages are mobile-friendly. They are a great solution for building real-world Intranets. Hub Sites provide the glue to group modern sites with a common purpose which share elements such as branding or navigation.



[Hub sites](#) integrate multiple Communication, Modern Team sites, and even classic sites together to build a real Intranet or general-purpose portals such as employee portals, division portals, or department portals. For example, it is possible to have several enterprise departments, each with its Team site for collaboration, together with multiple Communication sites with news and updates, interconnected to become an information gateway, without having to create the sites as a single hierarchical collection.

**Hub Site Limits:** The [number of Hub Sites that can be created by](#) a tenant is currently limited to 2,000.

A Hub site can aggregate news, events, and highlighted contents from all its associated sites and fulfill some other valuable functions: it adds consistent branding across all sites and a common menu, making it easier to navigate between the sites. Each hub site includes cumulative analytics for its connected sites through the site usage page. Users can search within the complete Hub site components and can use the digest option to organize selected news stories to go out as an email summary including images and links.

## Tenant Home Site

You can designate a Communication site as the Home site for the tenant by running the *Set-SPOHomeSite* cmdlet or through the [Home site setting](#) in the SharePoint Online Admin center. The site can be registered as a hub site but can't be associated with a hub. After executing the cmdlet, access to the Home site is available in the SharePoint mobile App. In addition, the Home site is automatically set to be an organizational news site and configured for organization-wide search. You can run the cmdlet again to switch the Home site to a different site.

## Modern Information Architectures for SharePoint Online

Before the arrival of modern SharePoint Online sites, the typical SharePoint deployment followed a classic pattern based on a hierarchical structure of sites and subsites with inherited permissions, navigation, and common branding. This architecture pattern is used extensively in the on-premises world and many organizations brought the approach to SharePoint Online. While popular, architectures based on this pattern can be difficult to maintain, inflexible, and sometimes suffered performance problems., Microsoft introduced a different information architecture based on single-site collections that can be associated with a Hub site (see the information about Hub sites later in this chapter). The new approach has a flat structure that makes it easy to share navigation, branding, and other elements. It is known as a *Flat Information Architecture* and is the [de facto standard promoted by Microsoft](#) for the design of SharePoint Online deployments.

## SharePoint Search

SharePoint Online uses Microsoft Search Service to create content indexes from the files kept on sites. The Search box on the SharePoint Online home page allows users to find content using different scopes (Sites, Files, People, and News). In the case of the Files and Sites scopes, Microsoft Search returns all the sites (including OneDrive for Business sites) and files and/or folders the user has access to that contain the search term. News and people scopes provide results of any news containing the search term and any user whose user profile includes that term. The way that Microsoft Search displays results is like the approach taken by Delve, with the difference being that the results of SharePoint searches can include sites, documents, people, and news while the results from Delve searches provide references to "boards" and "people." Each site also has a search capability, but in this case, the scope of the search is limited to the site.

## Pages

Modern pages are one of the key building blocks of any SharePoint Online site. They are fast, easy to author, and support rich multimedia content to make sure that they look great on any device, in a browser, or within the SharePoint mobile app. Examples of how to use modern Sites and pages are showing status and trip reports, guides, and frequently asked questions. SharePoint Modern pages are built with modern Web Parts that can be customized based on business needs, supporting the addition of documents, videos, images, site

activities, Yammer feeds and increasingly more items as Microsoft releases new Web Parts. In addition, the user experience of adding content to modern pages is different: just click the plus sign and pick a Web Part from the toolbox to start customizing a modern page. Using SPFx, developers can build custom Web Parts that show up directly in the toolbox. SPFx also supports the creation of SPFx extensions to customize specific areas in SharePoint modern pages such as the page footer and page header, or modern list and document libraries, by placing new actions in the list/document library command bar or the item/document menu or customizing Microsoft Search in SharePoint Online. Microsoft Viva Connections uses SPFx through the deployment of out-of-the-box Adaptive Cards Extensions (ACE) or custom ones to the Viva Connections dashboard.

## Apps

Sites can include **Apps** (or Add-ins) and Web Parts to expand the functionality available through the site. If enabled, site administrators or users with full control permission can add custom apps developed by the organization or by third parties to a site using the SharePoint store.

**Document libraries** and **Microsoft Lists** are very common examples of Add-ins used with SharePoint sites. Document libraries manage documents belonging to the site. Microsoft Lists are a great way to [create business Apps](#) to share, store and track information of interest either for individual (personal Lists) or workgroup purposes.

On an ongoing basis, Microsoft adds new features to improve the overall user experience when working with Microsoft Lists and Document libraries. Examples include:

- Allowing users to create a new list from scratch, create a new list from an existing list or template, or create a list from an Excel Table.
- Adding a new rich-text editor for text fields in lists and document libraries.
- Conditionally show or hide columns in a list or a document library form based on a value in another column.
- Being able to customize list and document library forms using JSON code.

## Metadata

Document management systems depend on metadata to organize and find information. Although users do not need to add metadata to documents (apart from a title and file name) when they add them to a SharePoint document library, it is a good idea to coach people to pay attention to the titles, tags, and comments that they can add to documents, plus any other metadata considered necessary for certain types of documents. The big pay-off is that documents with good metadata are easier to find with SharePoint Search, Delve, and eDiscovery searches.

## File Versioning

Versioning is a useful feature that does not exist for documents stored on local drives or file shares. File versioning allows you to keep multiple versions of documents and revert to an earlier version if necessary. For example, a file might become corrupt because of a hardware or software problem or the files in a document library might be encrypted by a cyberattack. File versioning supports many SharePoint features, including:

- Users can restore a previous version of a document to become the current version with the Version history option in the SharePoint Online and OneDrive for Business browser clients. Restoration of a previous version is also possible for any file synchronized with the OneDrive sync client through the Version history option in File Explorer, macOS Finder, and OneDrive Activity Feed.
- Microsoft 365 Enterprise Apps (desktop) and Online applications (Word, Excel, and PowerPoint) use file versioning for document AutoSave. These apps can open a previous version of a document and compare its content against the current version. Apps can also restore a previous version.

- The continual synchronization of changes made to documents permits real-time updates for co-authoring of Office documents.
- SharePoint Online and OneDrive for Business can restore document libraries to a point in time within the past 30 days.

[SharePoint Online supports the concept of major and minor versions](#) while OneDrive for Business only supports major versions. However, there is no technical difference between a major and minor version of a document. Both represent updates to files and the difference between the two exists in how a user regards the number and importance of changes they made in a version. By default, SharePoint Online and OneDrive for Business libraries support [a minimum of one hundred major versions](#). You can't disable versioning in library settings through the versioning settings page for a given Document library and you must choose a value between 100 and 50,000 for the number of versions kept. To use versioning on lists, you must enable it first.

Keeping hundreds of versions of documents might seem excessive, but if those versions are not present, it might not be possible to recover important documents to a specific point in time. Organizations that don't want to enable minimum versioning must run *Set-SPOTenant* to disable the feature. To disable the feature, download the latest version of the PowerShell module for SharePoint Online and run the command:

```
[PS] C:\> Set-SPOTenant -EnableMinimumVersionRequirement $False
```

## AutoSave

The AutoSave feature in Microsoft 365 Apps (Word, PowerPoint, and Excel) automatically captures changes made to documents stored in SharePoint and OneDrive for Business libraries. The idea is that a user can work on a document without having to worry about saving it because any change they make is saved to the server on an ongoing basis. The feature isn't supported for files stored on local or shared drives. AutoSave uses versioning to capture edits made using the Office desktop applications and online apps. AutoSave is unsupported for non-Office files because only the Office apps can synchronize small differential changes back to SharePoint to create new versions on the server.

New versions aren't generated for each change made to a document. Sets of changes are gathered by the server and are periodically committed to create a version. During a session, a new version might be generated every ten minutes or so, depending on the volume and type of edits in a document. For instance, pasting a large amount of text into a document usually forces Word to create a new version. If the app loses connection to the server, the server creates a new version to make sure that no data is lost.

A long editing session can easily generate 20-30 versions of a document. In some cases, SharePoint Online keeps many more versions. For instance, the chapter files for this book typically generate several hundred versions over the course of a year. The storage requirement for versions is mitigated using shredded storage in SharePoint Online to capture and store versions efficiently. The BLOBs used for document storage are broken up into "shreds" (roughly 64KB). As users update documents and create versions, SharePoint only needs to update the affected shreds. When documents are opened, SharePoint combines the shreds and delivers the complete file.

Continuously saving documents as changes occur makes the Office co-authoring feature more seamless. Updates made to a file by one user are synchronized to the server and back to the copies open by other users within seconds. Co-authoring is possible without AutoSave, but users must manually refresh their copies of documents to see changes.

You can [change AutoSave options in the Microsoft 365 Apps or update group policy settings](#), but you cannot affect how the cloud Office apps save document edits.

## SharePoint Online Quotas and Limits

Many of the [SharePoint on-premises boundaries and limits](#) are present in its cloud counterpart, but there are also [some specific limits](#) that only apply to SharePoint Online. For instance, a tenant can create up to two million sites and a single site can store up to 25 TB of information.

On-premises SharePoint servers use SQL databases to store site data. SharePoint Online also uses SQL (or rather Azure SQL), but the management of storage is much simplified in that Microsoft takes care of this for tenants. The storage limit for a tenant is calculated based on the Microsoft 365 licenses owned by the tenant plus any added storage the tenant buys from Microsoft. Briefly, the storage quota for a tenant is:

*1 TB plus 10 GB per licensed user plus purchased storage*

Thus, a tenant with 200 licensed users has a SharePoint Online storage base of 3 TB. If this is insufficient, the tenant can buy an unlimited (in practice) [amount of extra storage](#). The added monthly charge is based on the amount of added storage. If your storage requirements decrease, you can change the amount of added storage or cut it altogether. In addition, users with a frontline worker license don't add any storage to the available pool for a tenant.

The Export option in the Active Sites section of the SharePoint admin center downloads a CSV file which includes storage information for each site. If you prefer, you can write your own version with PowerShell, perhaps using [this sample script from GitHub](#).

## Managing SharePoint Online

The management of SharePoint Online follows the same approach as other workloads, through a dedicated portal and PowerShell. This section describes how to manage SharePoint using both approaches.

### The SharePoint Online Admin Center

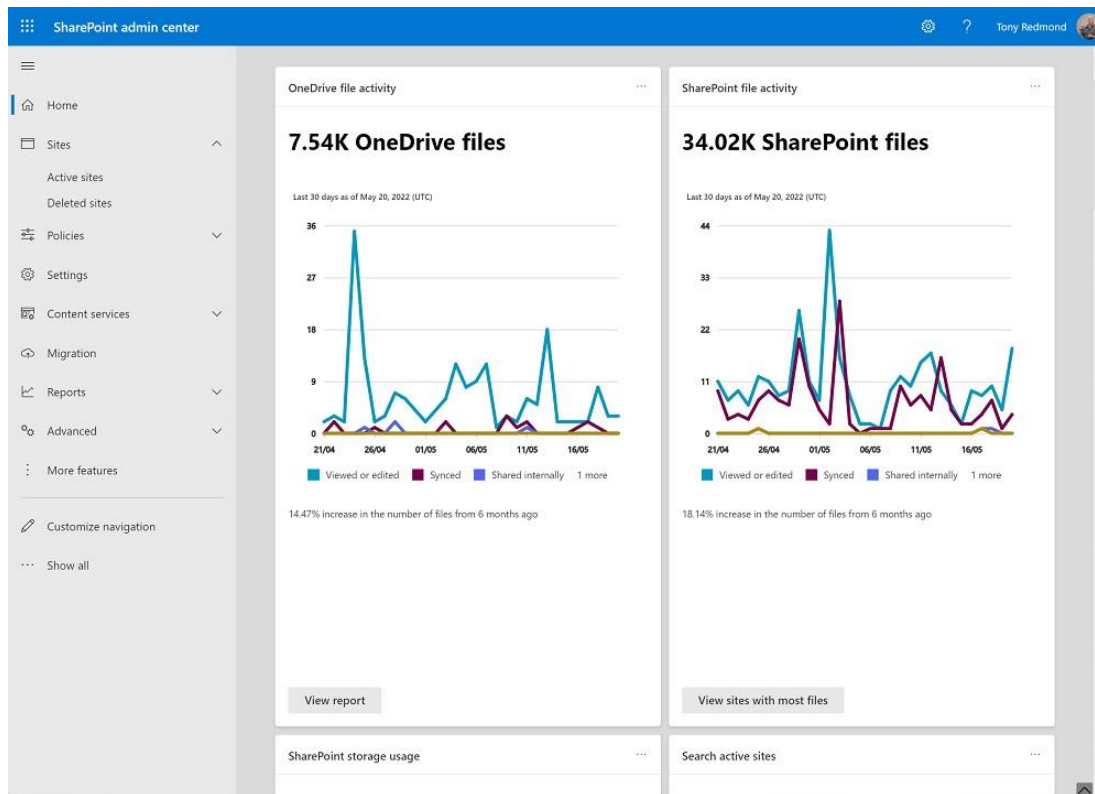


Figure 8-1: The SharePoint Online admin center

The SharePoint Online admin center (Figure 8-1) is the portal for administrative access to SharePoint Online and OneDrive for Business. To access the portal, select **SharePoint** from the **Admin Centers** section in the

Microsoft 365 admin center, or type the URL *https://[tenant]-admin.sharepoint.com* into a browser. Note that the "-admin" part in the SharePoint admin center is critical as this tells SharePoint that you want to access the administrative portal instead of sites.

The SharePoint Online admin center displays site information in a list view that allows the administrator to apply filters, add extra columns to expose information such as the sensitivity label assigned to sites, or find sites connected to a Microsoft 365 Group. Site management in the SharePoint admin center includes an option to export site data to a CSV file. The SharePoint admin center shares the same look and feel with other Microsoft 365 admin centers. Along with site management, the SharePoint admin center includes features such as displaying a summary of the SharePoint activities in the tenant and relevant messages from the Message Center and the Service Health status.

The integration of Groups and SharePoint team sites means that when Microsoft 365 Groups are available in a tenant, when a new team site is created, SharePoint creates a new Group to manage site membership. During group creation, the provisioning process creates a mailbox to hold a shared inbox and the Group calendar, a OneNote notebook, a Planner Plan for task management, and a SharePoint team site. Each Group gets a Modern home page, document library, lists, and business apps. Management of group membership is possible using a variety of interfaces, including PowerShell and the Microsoft 365 admin center. A group-enabled site can also be linked to Teams. If you create a team for a site, you can also add site resources such as pages or lists as tabs in the General channel of the new team. Additionally, managing modern team sites lines up to a great degree with the administration of Microsoft 365 Groups by following usage guidelines, naming conventions, and classification, meaning that there are some site administration tasks that team sites rely on that are controlled within the SharePoint admin center (like storage quota or the ability to adhere to a custom site provisioning experience) and some others through the Microsoft 365 admin center; this divergence might be addressed in the SharePoint admin center.

**Classic SharePoint admin center features availability:** Some classic SharePoint admin center features are still available in SharePoint Online, but Microsoft has started retiring some of its components. Initially, the following classic pages are going to be retired: Site collections, Sharing, Access Control, and Geolocation.

## Admin roles to Manage SharePoint Online

To administer SharePoint Online, several administrator roles can be assigned to users:

- The **Global Administrator** role. When Microsoft 365 provisions a new tenant, SharePoint Online creates several default sites and adds the Global admin as the primary administrator for each site.
- The **SharePoint Online Administrator** role. Accounts holding this role have access to the SharePoint Online admin center and can manage settings for SharePoint Online and OneDrive for Business.
- The **Global Reader** role gives read-only access to all the information and settings available in the SharePoint admin center.
- The **Site administrator** is assigned to users by a SharePoint Online administrator or a global administrator to give the assigned users permissions to manage a specific site. A single site can have several administrators, but only one primary administrator. Group owners are automatically added as Site administrators for sites that are group-enabled, including those used by Teams. The SharePoint Online administrator can assign permissions to the primary site administrator when creating a site and can add more administrators for the site afterward. Site administrators do not have access to the SharePoint Online Admin Center. A site administrator can be added to a site in three different ways: from the Active site list in the SharePoint admin center, PowerShell, and the site settings page.

## Managing Active Sites and Deleted Sites

The Active sites list in the SharePoint admin center manages active (not deleted) sites. To see the details of a specific site, select it in the list and click the “i” button or the site name to display the site details panel with site usage statistics and site properties (Domain, URL, Template, etc.). Administrators can change the default All Sites list view at any point to use one of the default views (like Microsoft 365 Groups sites) or a custom view. For a selected site, the following management actions are available:

- **Edit:** Allow administrators access to the site details panel where they can update different site settings.
- **Permissions:** Depending on the site type, you can change/modify Group owners (Group sites) for the selected site or change/modify primary and secondary site admin (Classic sites and Communication sites).
- **Hub:** You can promote the selected site to become a [Hub site](#) (“Register as hub site” option) or to join the site to an existing Hub site (“Associate with a hub site” option). If the site is already a Hub site, a SharePoint admin can change the Hub site settings (Display name and the users that can associate new sites to the Hub) or unregister the site as Hub site. Finally, if the site is already joined to a Hub site, the Hub site association can be changed using the “Change hub association” option.
- **Sharing:** This option configures the external sharing capability for the site (Sharing options for SharePoint sites are detailed later). From the sharing setting panel, it is also possible to access the global sharing settings page.
- **Delete:** It deletes the selected site by moving it to the site recycle bin.
- **Storage:** If manual quota management is enabled, you can modify the storage quota for the site selected.

If the site selected is the root site in the tenant, a “Replace site” option is available. The organization can replace the root site with any other site in the tenant if that site is a team site or a communication site. The replacement site cannot be a hub site or connected to a Microsoft 365 group.

Logically, to create SharePoint sites, use the **Create** button. Currently, it’s possible to create modern team sites (with or without associates Microsoft 365 groups), communication sites, a content center site for SharePoint Syntex, and (but not recommended) classic sites. In the Active Sites section, you can execute the following bulk edit options:

- **Sharing:** Configure the same sharing option for the selected sites.
- **Hub associations:** Associates the selected sites with a Hub site in the tenant.
- **Delete:** Deletes the selected sites and move them to the Sites recycle bin.

When you select a site from the Active sites view, SharePoint opens the site details panel to display information about the site in four sections:

- **General:** Displays several modifiable site properties (Site name, Site Url, Hub association, and Storage limits) along with some that are not (site template, if the site connects to a Microsoft 365 Group, domain, site description, creation date, and who created the site). If the site is a Hub site, you cannot edit the Hub association here. For team-connected sites, a banner flags the site as “Connected to Microsoft Teams.” If the team has any channel sites (used for private and/or shared channels), the panel shows the number of channel sites together with a link to access additional details of those sites. SharePoint administrators cannot access the content of a channel site unless they are a member of the private or shared channel owning the site.
- **Activity:** Gives information about recent site activity such as Last site activity, files stored for the site, the number of views in the last 30 days, the number of files viewed/edited in the last 30 days, and storage usage. The statistics displayed here are usually two or three days old.

- **Permissions:** Gives the same information about site permissions as shown in the “Permissions” action in the Active sites list together with lists of the membership of the default SharePoint Groups in the site (Owners, Members, and Visitors).
- **Policies:** Change the Sharing settings of the site or assign a sensitivity label to the site. Before you can assign a sensitivity label with container management settings to SharePoint sites, the labels must exist and be published to users to allow them to apply the labels to sites.

Hovering over the name of a user (on the Active sites list or in the details panel) lets you see details about the user. An edit link is available in some sections to allow admins to change site properties.

Deleted sites in the tenant are listed in the Deleted sites section in the SharePoint admin center. An administrator can select any deleted site and restore it using the “Restore” option. To permanently remove a site, just click the “Permanently delete” option.

**Renaming Sites:** The SharePoint admin center includes the ability to partially change the URL of a site by renaming the site. This is done by selecting the site and editing the name in the site details panel. A detailed overview of the process to rename a site [is explained here](#). You can also rename site URLs with PowerShell using the *Start-SPOSiteRename* cmdlet. Be aware that some consequences of site renaming exist, as [explained in Microsoft’s documentation](#).

### Controlling Site Creation

It is possible to allow users to create new Sites Collections or restrict creation to only administrators. Because of the tight connection that Team Sites have with Microsoft 365 Groups (including Teams), this is not a simple matter as you must consider whether you want users to be able to create sites, group-enabled sites, or both.

The basic control over Site creation is achieved by showing or hiding the Create site command on the SharePoint home page of the root site. This is the button that invokes the Create Site wizard. A SharePoint tenant setting controls whether the button appears on the page for all sites. To access the setting, open the SharePoint admin center and then **Settings**. On the Settings page, click on **Site creation** and then simply switch on/off the **Let users create new sites** option.

### Creating a Group-Connected Team Site

If the SharePoint settings for a tenant allow the creation of group-connected sites and the user can create Microsoft 365 Groups, the team site creation process creates a group. The new group holds the membership for the site and the Microsoft 365 Groups membership service is used to allow site members access to group resources.

Unlike the creation process used by OWA, Outlook, PowerShell, or an Admin console, SharePoint creates the new team site at once rather than waiting for the first user to access the site. In effect, this is the opposite of what happens when OWA or another email client creates a new group, where Exchange Online creates the group mailbox first to allow conversations to occur and notifications to go to new group members. After SharePoint creates the new group object in Azure AD, a process of forwarding directory synchronization makes the existence of the new group known to Exchange Online and forces the creation of the group mailbox. While this process proceeds, the user can build out the contents of the SharePoint team site with new lists, document libraries, a customized home page, and so on. The same creation process to set up a new team site and group occurs when a user invokes this choice from OneDrive for Business.

**Connect existing SharePoint Team Sites to new Microsoft 365 Groups:** Existing SharePoint sites, including classic sites, can be connected to a new group (or “groupified”) using *the Set-SPOSiteOffice365Group* cmdlet, CSOM API, or using the **Connect to new Microsoft 365 Group** feature available in site settings options. When this happens, the existing content, hierarchy, and permissions for a site stay intact, and SharePoint Online connects the site to a new group and populates with the group membership using the existing site membership. The owner can adjust the group membership after

creation. A tenant or SharePoint Online Administrator can disable this option by selecting the option **Prevent site collection administrators from connecting sites to new Microsoft 365 groups** in the **Settings** page in the classic SharePoint admin center. Microsoft [gives detailed information](#) about what SharePoint sites can be connected to a new Microsoft 365 Group, a [scanner tool](#) to analyze if existing SharePoint sites are suitable to be connected to Groups, and PowerShell scripts to enable modern features in the Sites and configure group membership.

### Settings and Permissions for Group-enabled Sites

When a new group-enabled team site is created, SharePoint creates a team site to host the document library and shared notebook for the group. The gear (options) menu for a site reveals the Site Information panel used to customize the details published to users about the site as well as its classification, or label defining the importance of the information held in the site and the Hub site association in case the site is not a Hub itself. Also, the owner can access the Site Permissions panel and set the desired level of access for different member types.

The default configuration for a group-enabled site is that group owners have SharePoint full control permissions over the site, while group members have SharePoint Edit permissions. If the site is public, every account in the tenant excluding external users has SharePoint Edit permissions. Amend the permissions for the site to restrict access as needed but be aware that because you are working with Microsoft 365 Groups rather than individual user accounts (as is the case for on-premises SharePoint or classic sites), you should take care to ensure you do not interfere with the ability of group members to access the resources available to the group, or for group owners to manage the group. It is also possible to select an individual user account and give them full control of the site. This permission is a SharePoint permission over the team site and does not make the user an owner of the group.

One of the permissions you can assign to a site is "Share Site Only" access, which means that the owner can give someone access only to the SharePoint resources instead of all the resources belonging to the group. Also, it is possible to assign a user read-only access to the site, which removes their ability to edit or remove data from the site (they might still be able to share files from the site). Assigning read-only access to users is a good approach to take for sites holding information (like HR documents) and it is necessary to make the site content generally available.

Apart from the ability to set access to site control through these settings, it is unwise to try and use traditional SharePoint access control over team sites used with groups because these controls expect to deal with individual users rather than when an identity is shared by group members.

### Hub Sites

A Hub site is a modern team site, or a communication site registered as a Hub using PowerShell, or a site converted to be a Hub site using the "Register as hub site" option in the SharePoint admin center. [Registering a hub site with PowerShell](#) can be done with cmdlets from the SharePoint Online module or PnP cmdlets (we discuss how to manage SharePoint Online with PowerShell later). To use the following example, you must use the latest version of the PnP module. The code opens a connection to the tenant and then creates a team site, elevates it to a Hub site, and associates other sites with the Hub.

```
[PS] C:\> Connect-SPOService -Url https://<TenantName>-admin.sharepoint.com
Connect-PnPOnline -SPOManagementShell # Will ask for the site URL
New-PnPSite -Type TeamSite -title "New SiteColl" -alias "NewSite" -Description "New SiteColl"
# Create a new Site
Register-PnPHubSite -Site https://<TenantName>.sharepoint.com/sites/NewSite # Register the SiteColl
as a Hub site
New-PnPSite -Type TeamSite -title "Sub-SiteColl" -alias "SubSite" -Description "Sub-SiteColl"
# Create a sub-sitecollection
Add-PnPHubSiteAssociation -Site https://<TenantName>.sharepoint.com/sites/SubSite -HubSite
https://<TenantName>.sharepoint.com/sites/NewSite # Associate the subsite to the Hub site
```



Associating a site with a Hub site is a task that only site administrators and site owners can do. Hub owners can manage visitor permissions for the Hub site itself and sites joined to the Hub. This feature, off by default, can be enabled by a Hub owner through the Site permissions panel for the Hub site. When enabled, a new SharePoint Group with read permissions is created in the Hub site and Hub owners can add the users, Groups and security groups that later can optionally be synchronized to any Site joined to the Hub. The synchronization of visitor permissions from the Hub site to a site joined to the Hub must be explicitly done for each site.

A [Hub site can be associated with another Hub site](#) as a way to expand search results across multiple Hubs in an organization allowing in fact to define information architectures with more than two levels deep. Hub to Hub association is a very simple process that can be done in the SharePoint admin center:

- In the active sites list, select an existing Hub site.
- Edit the Hub settings through the Hub settings panel.
- In the **Parent hub association** field, select the Hub to associate with.
- Save the changes.

Some limitations exist for Hub sites: there's no workflow around the publishing process, anyone who has edit permissions for a site can change any element of that site, and any new story pushes older ones down the stack, making it impossible to create a "Most important" news that stays always at the top. Additionally, Hub sites cannot be nested, limiting their potential as well, and they cannot deal with multiple languages. It is foreseeable that Hub sites will mature in the future because they are ongoing work.

### SharePoint Spaces

SharePoint Spaces creates mixed reality experiences on SharePoint sites. When enabled on a site (go to the Site settings page and enable the Spaces feature on the Sites features page), users can create spaces by defining a structure, background, and theme. The next step is to add Web Parts to visualize 3D objects, 360-degree images, videos, 2D images, and text. Once a space is ready, it can be viewed in a web browser or with a mixed reality headset.

### Modernization of Classic Sites

Although Microsoft introduced modern sites, it's still common to find classic sites in use. Microsoft recommends that tenants upgrade classic sites to take advantage of modern site technology. To help this modernization, Microsoft and the SharePoint community have released tools to speed up the upgrade process:

- Microsoft automatically updates any classic site with a home page that isn't customized with a modern page.
- The [SharePoint Modernization Scanner](#) is a tool originally built by the PnP Community to identify the work that must be done to modernize sites.
- The [Enable-SPOCommSite](#) cmdlet (in the SharePoint Online module) converts a classic team site (STS#0 template) that meets the requirements described in [this link](#) into a Communication site.

### Managing Performance for SharePoint Sites

A key factor in the success of any SharePoint site is achieving good performance when navigating across pages on a given site. To optimize performance in SharePoint sites, two out-of-the-box tools are available:

- [Site performance page](#): is accessible by Site owners and editors through the site settings menu. This page provides access to the Page Diagnostics for SharePoint tool that identifies any performance issue present on modern SharePoint pages and suggests actions to address issues found. Page Diagnostics for SharePoint tool uses a browser extension available for Microsoft Edge and Google Chrome.

- [SharePoint Portal Launch Scheduler](#): a feature to launch specific SharePoint sites that are expected to receive high volumes of traffic using a phased deployment. The tool also includes, if needed, an automatic redirect for existing sites. The SharePoint Portal Launch Scheduler is available through the `New-SPOPortalLaunchWaves` cmdlet.

## Managing Access Control

The SharePoint admin center includes settings through the Access control policies page that can be used to control access to SharePoint Online sites and OneDrive for Business accounts:

- **Unmanaged devices**: Settings to control how unmanaged devices can access SharePoint sites, OneDrive for Business (ODFB), and information stored there. Access control to SharePoint Online and ODFB is done through the creation of conditional access policies, which require an [Azure AD Premium Subscription](#). Among other settings, these policies can exert control over access from devices that are not compliant or joined to a domain (unmanaged devices). By default, SharePoint Online and OneDrive contents are accessible from unmanaged devices but it's possible to configure limited access or even block access through Azure AD conditional access policies, including the use of authentication contexts with sensitivity labels. You can also apply conditional access policies at the site level as explained in [this link](#). The Identities chapter contains examples of how to process unmanaged devices in conditional access policies.
- **Idle session sign-out**: Disabled by default, [this setting](#) is a mechanism to first warn and then sign out users on unmanaged devices if they have been inactive for a period and they don't opt to stay signed in when they sign into SharePoint Online, OneDrive for Business, or Microsoft 365. Two settings can be configured when this setting is enabled: the idle session period (15 minutes minimum, 24 hours maximum) and the warning notice (one minute minimum, 30 minutes maximum). Idle session timeout can also be configured running the `Set-SPOBrowserIdleSignOut` cmdlet. For example, this command enables a session timeout and sets a warning period of 45 minutes (2700 seconds) and a forced sign-out after an hour (3600 seconds). The timeout values apply to both SharePoint Online and OneDrive for Business sessions.

```
[PS] C:\> Set-SPOBrowserIdleSignOut -Enabled $True -WarnAfter (New-TimeSpan -Seconds 2700) -SignOutAfter (New-TimeSpan -Seconds 3600)
```

- **Network location**: When enabled, it lists the IP addresses from which access to SharePoint and ODFB is allowed.
- **Apps that don't use modern authentication**: Enabled by default, the setting blocks access to SPO from Apps that don't use modern authentication. Office 2013 or earlier clients are examples of Apps that don't use modern authentication and can pose a potential security problem for the information stored in SharePoint Online. In general, it's recommended to use this setting to block access for older clients.
- **Limit OneDrive access**: This setting limits the use and access of OneDrive for Business to users in specific (up to 10) security/Microsoft 365 groups.

## Managing Global SharePoint and OneDrive Settings

These Global Settings are managed through the Settings section in the SharePoint admin center:

- **Default admin center**: Toggles the default admin center for SharePoint Online. If the setting is on (it's off by default), the SharePoint admin center is opened.
- **Home site**: Allows to select an existing site to be set as Home site in the tenant. The site chosen can be removed at any point and replaced with a different site.
- **Notifications**: Controls if users of the SharePoint mobile app receive notifications about site activity.
- **Pages**: This section allows administrators to change global settings for creating and commenting on modern SharePoint pages. By default, both settings are enabled.

- **Site creation:** In this section, configure the managed path and time zone to be used when a new SharePoint site is created and if end users can create sites from the SharePoint home page.
- **Site storage limits:** Defines how the available SharePoint storage is managed. Two options are available: Automatic and Manual.
- **Link to the Classic settings page:** Gives access to several other settings from the classic SharePoint admin center.
- **Stream:** Defines the default destination for the Stream app launcher tile. There are three possible destinations to choose from:
  - Automatically switch to Stream (on SharePoint) when recommended: This is the setting configured by default and it means that Microsoft controls what version of Stream users access through the app launcher
  - Stream (on SharePoint).
  - Stream (Classic).

OneDrive settings that can be configured in this section are explained later.

## Managing Content Services

The content services section in the SharePoint Admin Center manages and publishes Content Types and managed metadata for a SharePoint Online tenant.

### Term Store

The Term Store is a central location to [create and store taxonomies, which](#) can be used in three ways:

- Create manage metadata fields that allow to classify a document or a list of items through a term that is part of a specific taxonomy.
- Configure the global navigation of an Intranet solution deployed in the tenant.
- Improve end user experience when using search capabilities utilizing the search dictionaries that contain specific Term sets to provide features such as including or excluding words for query spelling correction.

A custom taxonomy consists of a Group, at least a term set in the Group, and terms in a term set. The maximum number of levels of nested terms in a custom taxonomy is seven. The Term Store also contains enterprise keywords and hashtags as a mechanism to allow users to tag content with existing or new keywords created by end users. As part of the modernization of the Term Store, an increase in the number of terms available has increased from 200,000 to 1,000,000. Finally, the filters panel in modern Document libraries has been enhanced to support managed metadata.

### Content Type Gallery

The Content type gallery is a central location to manage and publish custom types to any site in the tenant. To create a new content type:

- Access the Content type gallery in the SharePoint Admin center and click **Create content type**.
- Type the information required to create a content type (Name, Description, Parent content type, and Category) and click on Save.

After the content type is created, click its name in the list of content types to access to the content type details page where you can take the following actions:

- Modify the details of the content type.
- Add site columns to the content type. You can choose between adding existing site columns or creating new columns. To create a new site column, provide the name, description, category, and column type. Depending on the column type, you might need to configure additional settings.
- Publish the content type so it can be used on any site in the tenant.

- Access to advanced content type settings such as the document template to use (only for document types), permissions to enable/disable the modification of the content type, and enabling/disabling the update of the content type on sites and lists when it's updated.
- Access to the content type policy settings.
- Delete the content type.

To use the content type in an existing document library:

- Browse the document library where you want to use the content type, click **+ Add column**, and then **Content type** to open the Add content type panel.
- In the panel, select the content type you want to add to the document library and click **Apply**. Note that before applying the content type you can optionally indicate if a specific view based on that content type is added to the library or if the content type columns should be added to the current view.

Any published content type in the Content Type Gallery can also be added to a site by using the [Add-PnpContentTypesFromContentTypeHub](#) cmdlet. If the content type already exists in the site, the published version is synchronized to the site.

## Managing Reports

The Reports section provides access to Content services and Data access governance reports. Content services reports are available only for tenants with a SharePoint Syntex subscription. The reports contain insights about the usage of terms and terms sets in the global Terms store across the tenant:

- **Terms store operations:** A graph showing details of the Terms store operation over the last 15 days.
- **Terms store composition:** Displays an overview of the number of terms in the Term store identifying the terms distribution (regular, hashtag, and keywords). From this report, it's possible to promote keywords to become terms in the Term store.
- **Open and closed term sets:** Indicates the total number of term sets and the distribution between open and closed term sets.
- **Terms without synonyms:** It reflects the percentage of terms in the Terms without synonyms or abbreviations specified.

Two types of [Data access governance reports](#) are available:

- **Sharing links** reports monitor sharing activity across the tenant and identify potential oversharing situations. Three kinds of sharing reports identify the sites where users create the highest number of "Anyone", "People in your organization" and "Specific people" links. It's possible to run all reports together or one by one. It can take a few hours to generate these reports.
- **Sensitivity labels applied to files** reports the sites including files with a specific sensitivity label. To get a report for a specific label, you need to add the report and then run it. Note that you cannot create more than ten sensitivity labels reports.

## Managing Advanced SPO Settings

The SharePoint admin center also includes the ability to manage some advanced settings in the platform such as API Management or Geo locations when multi-geo is enabled for the tenant:

- **API access:** From this section Admins can approve/reject permissions requested by third-party APIs that can be called by SharePoint Framework (SPFx) solutions and custom scripts. An administrator can easily approve or reject a request from the list using the Approve or reject panel.
- **Geo locations:** Lists available geo locations for the tenant. From here an administrator can add new satellite locations.

## Managing SharePoint Online Services and Settings

As in SharePoint on-premises, you can manage some of the core services as well as key settings. This section provides an overview of the services and settings that can be managed through the shortcuts provided in the **More features** section in the SharePoint admin center.

### User Profiles

Identity management for Microsoft 365 is centralized in Azure AD, but each core workload also has its own directory (SharePoint Directory Store or SPODS in the case of SharePoint Online) that is synchronized bidirectionally with Azure AD. Although SharePoint Online does not allow administrators to manage Azure AD users directly, it includes the [User Profiles](#) service as a mechanism to modify existing user properties or add new properties. Behind the scenes, a SharePoint Timer Job imports account information from Azure AD into the User Profiles service.

In the User profiles service page, Manage User Properties and Manage User Profiles links are most used by a SharePoint Online Admin. The first link takes us to a page that lists all the user profile properties and from where we can create new properties. The second link redirects to a central location from which we can look for a specific user profile to modify some of his/her properties and their visibility (only for some of the properties), delete the user profile, grant/control access to the user's OneDrive or [add additional administrators to users' OneDrive accounts](#). Depending on the subscription purchased, some of the properties needed to take full advantage of SharePoint Online features might be blank, so they must be manually or automatically filled (via PowerShell). An example of such properties is the Work email property that is empty in standalone SharePoint Online plans used for sharing functionality and alerts.

### Records Management

Despite its name, the Records management section in the SharePoint Online Admin Center simply allows administrators to [create new Send To connections](#) that can be used by a [Content Organizer configured on a SharePoint site](#) to route documents to a specified location (a document repository or a records center). It is also possible to configure Send To connections to be used by end users, so they can manually copy, move or move and leave a link when a document is selected in a document library.

### Search

Search is one of the core features in SharePoint Online that administrators [can customize](#) to ensure that end users have a great experience when searching documents by content and metadata through any site in the tenant. The search administration in SharePoint Online includes the following items:

- **Search Schema:** This option provides access to the list of managed properties and crawled properties in the tenant. An Administrator can update existing managed properties or create new ones. Each managed property can be mapped to one or more crawled properties. A crawled property is just content and metadata that is extracted during a crawl from an item stored in SharePoint, such as a document, a SharePoint page, or a list item. A good overview of crawled and managed properties can be found in [this Microsoft article](#).
- **Search Dictionaries:** Allows the creation of [a list of terms in the Term store](#) to include or exclude company names extracted from the contents indexed by the search engine or to include or exclude words for query spelling correction.
- **Query Suggestions:** Using this option we can upload query suggestions to the search engine so when the user enters some search criteria in the search box, search suggestions for that search criteria are shown.
- **Result Sources:** The ability to [limit searches to a subset of search results](#) is achieved through Result Sources. From the Result sources page, a SharePoint Online administrator can review existing Result sources or create new ones. The Result sources managed at this level can be used for all sites. A Site administrator or a site owner can manage result sources for a site.

- **Query Rules:** Search results can be improved through [Query rules](#) that give a mechanism to specify conditions and actions to promote the relevance of the search results that meet those conditions. As an example, we could create a query rule that boosts technology news in our Intranet that are tagged with a specific category so when someone searches for that news category, all the news in that category is displayed on the top of search results. Over the time Query rules in SharePoint online search will be replaced using answers in Microsoft Search.
- **Remove Search Results:** Lists the URLs removed from search results.
- **Usage Reports:** To evaluate how users are performing searches in SharePoint Online, a set of predefined reports is available. The reports provide insights about the number of queries performed by users, top search queries, queries with no results, abandoned queries, or query rules usage. While these reports are still available, the recommendation for search analytics across Microsoft 365 and SharePoint online is to use Microsoft Search insights.
- **Search Center Settings:** This page allows administrators to define the URL of the Global Search Center in the tenant and the search results loading strategy used in the tenant.
- **Export Search and Import Search Configuration:** Those options provide a simple mechanism to export to a text file any customization done in search configuration (query rules, result sources, result types, ranking models, and site search settings) or to import a search configuration file.
- **Crawl Log Permissions:** This page lists the users with read access to crawl log information in the tenant.

Changes done in the search administration page are applied to the whole tenant, but it's also possible to customize search at the site level.

## Secure Store

The [Secure Store service](#) is designed to manage and store the credentials needed to connect to an external data source. Credentials are just one of the settings of a Secure Store Target Application that also defines the mapping between them and a group of users in SharePoint Online, who can use the credentials to access the external data source. A Secure Store Target Application can be used by services such as Business Connectivity Services (BCS) to connect to an external system and integrate data in SharePoint Online sites in the form of SharePoint external lists, external data columns, default BCS Web Parts or a custom solution.

## Apps

SharePoint Online can be extended through [Apps or Add-ins](#). In general, two types of Apps can be added to a SharePoint site:

- **SharePoint Add-ins**, [defined by Microsoft as self-contained extensions](#) that may include logic and data deployed in the cloud, SharePoint components (such as Content types, Lists, Document libraries, etc.), and client-side scripts.
- **Web Parts and Extensions created with the SharePoint Framework.** SPFx Web Parts can be added to both classic and modern pages. SPFx extensions can only be used in the modern SharePoint experience.

SharePoint Apps are only installable from the SharePoint Store or the [global App Site](#). A tenant can have only one global App Site, a special type of site that needs to be created by a SharePoint Online Administrator. SPFx solutions (Web Parts, Extensions, and Adaptive Cards Extensions) for SharePoint sites, Microsoft Teams, and Viva Connections can also be distributed from the global App Site. From the global App Site, an Admin can distribute custom Apps or browse and install Apps that meet the business requirements directly from the SharePoint Store. End users can request new Apps to be added to a site through the My Apps page by browsing the SharePoint Store to select and request an app. Tenant administrators can then approve or deny the app requests. Finally, site owners can delete requested apps, even if administrators approve their installation, from the My requests section on the My Apps page.

For more granular Apps deployment and isolation, Apps catalogs can be created in specific sites. To create an App catalog, connect to SharePoint Online with PowerShell and run the following command:

```
[PS] C:\> $sSiteCollectionUrl = "https://<Site_Collection_Url>"  
[PS] C:\> Add-SPOSiteCollectionAppCatalog -Site $sSiteCollectionUrl
```

You can only install Web Parts and Extensions deployed to a site catalog in the root site and any site created in the collection, but not in any other site. To create a site catalog, the user executing the *Add-SPOSiteCollectionAppCatalog* must be a site admin for the global App catalog.

## Hybrid Picker

SharePoint Online and SharePoint On-Premises can be integrated and connected in a hybrid deployment. Assuming the [pre-requisites required for a hybrid deployment](#) are met, the list of hybrid features currently available is the following:

- **OneDrive redirection:** If enabled, on-premises users are redirected to OneDrive for Business.
- **Hybrid search:** provides the ability to combine on-premises and cloud search results in the same search index to make it possible to search content from both cloud and on-premises sources in the on-premises farm. On the other hand, hybrid federated search allows to search content from on-premises and cloud indexes in a single search center.
- **Hybrid App Launcher:** This feature allows access to services such as Teams or Stream in the on-premises App launcher.
- **Business to Business (B2B) Extranets:** By using this feature, it is possible to create an Extranet in SharePoint Online to collaborate with partners, so they don't have access to any on-premises data.
- **Hybrid Taxonomy:** Delivers a mechanism to replicate taxonomies created in SharePoint Online to SharePoint on-premises. In a nutshell, hybrid taxonomy allows an organization to have a shared taxonomy between SharePoint Online and SharePoint on-premises.
- **Hybrid Self-Service site creation:** If enabled, any user that creates a new site is redirected to the self-service site creation page in SharePoint Online.

## InfoPath Forms Services and Business Connectivity Services

[InfoPath Forms Services](#) and [BCS Services](#) are two legacy services supported by SharePoint Online. InfoPath enables tenants to easily design electronic forms that can be deployed to any SharePoint site in the tenant using InfoPath Designer 2013 as an authoring tool. The latter provides a simple mechanism to connect a SharePoint site to external business data sources such as SQL Azure Databases or any other one exposed by a Windows Communication Foundation (WCF) service.

Microsoft has committed to support [InfoPath Forms Services 2013 and InfoPath 2013 clients](#) until 2026 and has positioned Power Apps as its natural replacement. The general recommendation about InfoPath usage in SharePoint Online is to avoid the creation of new InfoPath Forms and carefully plan the migration of existing ones to Power Apps or custom solutions built using any of the frameworks and platforms provided by Microsoft (including SPFx).

BCS supports the integration of data stored in external sources in SharePoint sites through SharePoint external lists, BCS Web Parts or external data columns can be added to existing lists and/or document libraries. While Microsoft has not made an official statement about the future of BCS in SharePoint, they have not made any kind of investment in BCS since the release of SharePoint 2013. In SharePoint Online, the integration of external data in SharePoint sites can be achieved by using a combination of modern cloud technologies and platforms such as Power Apps, Power Automate, Azure Logic Apps, or cloud development patterns.

## Managing SharePoint Online with PowerShell

The SharePoint Online Management Shell is a Windows PowerShell module designed for command-line operations and inclusion in PowerShell scripts. The module enables batch processing for tasks like reports and is the only way to achieve some management tasks in SharePoint Online and OneDrive for Business.

Microsoft updates the SharePoint Online Management Shell regularly. The updates include new cmdlets, new parameters for cmdlets, and other tweaks. If you use PowerShell to work with SharePoint Online, it's important that you use the latest module. See the PowerShell chapter for advice about how to update PowerShell modules.

### Connecting to SharePoint Online with PowerShell

The *Connect-SPOService* cmdlet is used to connect to the SharePoint administration endpoint for a tenant (the same used by the SharePoint admin center). To build the *endpoint*, take the normal SharePoint root URI for your tenant (like <https://office365itpros.sharepoint.com/>) and insert an "-admin" after the tenant name. For example:

```
[PS] C:\> Connect-SPOService -URL "https://office365itpros-admin.sharepoint.com"
```

The SharePoint Online module is designed for administrative tasks, so you should always connect with an account that has Global Administrator or SharePoint Administrator rights for the tenant.

### Basics of SharePoint Online Cmdlets

To see the available SharePoint Online cmdlets, run:

```
[PS] C:\> Get-Command -Module "Microsoft.Online.SharePoint.PowerShell"
```

The current cmdlets available in the module can be divided into several types:

- Tenant-Level cmdlets like *Get-SPOTenant* and *Set-SPOTenant*.
- Site-Level cmdlets like *Get-SPOSite* and *Set-SPOSite*.
- Cmdlets for specific operations like *Start-SPOSiteRename*, *Invoke-SPOSiteSwap*, and *Test-SPOSite*.

It's very common to want to retrieve information about the sites in a tenant. To do this, run the *Get-SPOSite* cmdlet. The *Limit* parameter specifies that all sites are to be returned.

```
[PS] C:\> Get-SPOSite -Limit All
```

This command returns all types of sites found in the tenant, including redirect sites (created because of site URL renames), hub sites, the global App site, Group sites, and sites used by Teams private and shared channels. In most cases, it is best to be more precise when using *Get-SPOSite* to find sites by specifying the template for the type of sites you want to process. For instance, this command returns the sites connected to Teams:

```
[PS] C:\> Get-SPOSite -Limit All | ? {$_.IsTeamConnected -eq $True}
```

While this command returns the set of sites connected to Teams private channels:

```
[PS] C:\> Get-SPOSite -Limit All -Template "TEAMCHANNEL#0"
```

Alternatively, you can run this command to list the channel-connected sites together with the type of channel:

```
[PS] C:\> Get-SPOSite -Limit All | ? {$_.IsTeamsChannelConnected -eq $True} | Format-Table Title, TeamsChannelType
```



For Modern Pages, SharePoint Online displays usage information (the number of views) and allows users to like and enter comments for pages. It is possible to disable this feature via PowerShell. At the tenant level it can be done by running the following command:

```
[PS] C:\> Set-SPOTenant -SocialBarOnSitePagesDisabled $True
```

Alternatively, you can disable the feature for specific sites by running the *Set-SPOSite* cmdlet and passing the URL for the site. For example:

```
[PS] C:\> Set-SPOSite -Identity https://Office365itpros.sharepoint.com/Projects/ -
SocialBarOnSitePagesDisabled $True
```

## Combining SharePoint Online and Other Objects

Sometimes you need to retrieve information about a SharePoint Online site for use elsewhere. For example, if you want to include a document library belonging to a group, team, or team private channel on an eDiscovery case or content search, you need to specify the site's URL as a search location.

If you use PowerShell to examine a group's properties, you will see three SharePoint Online URLs returned for the site, the document library, and the shared OneNote notebook. The value returned by the *Get-UnifiedGroup* cmdlet (part of the Exchange Online module) in the *SharePointSiteUrl* property is the one needed when you wish to add a site to content searches, found using the *Get-UnifiedGroup* PowerShell cmdlet:

```
[PS] C:\> Get-UnifiedGroup -Identity "Office 365 for IT Pros" | Format-List Share*Url

SharePointSiteUrl      : https://Office365ITPros.sharepoint.com/sites/0365ITPros
SharePointDocumentsUrl : https://Office365ITPros.sharepoint.com/sites/0365ITPros/Shared Documents
SharePointNotebookUrl  : https://Office365ITPros.sharepoint.com/sites/0365ITPros/SiteAssets/Office
365 for IT Pros Notebook
```

The URL retrieved from the group can be used with *Get-SPOSite* to find further information about the site belonging to the group.

It's also possible to discover what group a site belongs to by using the *GroupId* property stored for the site. For example:

```
[PS] C:\> Get-UnifiedGroup -Identity (Get-SPOSite -id https://office365itpros.sharepoint.com/sites/
0365ITPros).GroupId.Guid -Detailed | Format-Table DisplayName, SharePointSiteURL

DisplayName      SharePointSiteUrl
-----
Office 365 for IT Pros https://office365itpros.sharepoint.com/sites/0365ITPros
```

Many SharePoint Online sites are connected to a Microsoft 365 group (with or without Teams). To find the group information for a team site, run the *Get-SPOSite* cmdlet with the *Detailed* parameter to return the group identifier in the *GroupId* property. You can then use the identifier with the following cmdlets:

- *Get-MgGroup*: to return details of the Azure AD group.
- *Get-UnifiedGroup*: to return details about a Microsoft 365 group.
- *Get-Team*: to return details of a team if the group is team-enabled.

For example, this code returns the display name of the Azure AD group that holds the membership of each group-connected site. Not all sites created with the *Group#0* template might have corresponding Microsoft 365 groups. This can happen after the deletion of a group if the site is retained by a retention policy.

```
[PS] C:\> [array]$Sites = Get-SPOSite -Template "GROUP#0" -IncludePersonalSite:$False -Limit All
Write-Host "Processing" $Sites.Count "sites"
ForEach ($Site in $Sites) {
    $GroupId = (Get-SPOSite $Site.Url -Detailed).GroupId.Guid
    $Group = Get-MgGroup -GroupId $GroupId -ErrorAction SilentlyContinue
```

```
If ($Group) {  
  Write-Host ("SharePoint Online site {0} is connected to the Azure AD group {1}" -f $Site.Title,  
$Group.DisplayName) }  
}
```

## Enabling or Disabling Web Parts

Modern pages in SharePoint sites can be customized with modern Web Parts to show information from other Microsoft 365 services and non-Microsoft services such as Twitter, Kindle, or YouTube. SharePoint Administrators [can use PowerShell to hide specific web parts from end users](#) by running the *Set-SPOTenant* cmdlet to set the *DisableWebPartIds* property. For example, to hide the Web Part for Twitter, you run this command:

```
[PS] C:\> Set-SPOTenant -DisableWebPartIds f6fdf4f8-4a24-437b-a127-32e66a5dd9b4
```

Multiple Web Parts GUIDs can be specified in a comma-separated list. To view a list of disabled Web Parts, you can use the *Get-SPOTenant* cmdlet. To re-enable some disabled web parts, run *Set-SPOTenant* and specify only the GUIDs for the Web Parts that you want to be disabled in *DisabledWebPartIds*. To re-enable all Web Parts, run:

```
[PS] C:\> Set-SPOTenant -DisabledWebPartIds @()
```

## Enabling or Disabling the Weekly SharePoint Auto Digest e-mail

The SharePoint Auto-News Digest feature allows tenants to send a weekly automated email to users containing the latest News posts that they have not yet read that might be relevant to them. This feature is enabled by default and can only be disabled in two ways:

- Users can stop receiving the Auto Digest e-mail by clicking on the unsubscribe link present at the bottom of the digest e-mail.
- SharePoint Admins can completely disable it by running the *Set-SPOTenant* cmdlet to set *EnableAutoNewsDigest* to *\$true*.

```
[PS] C:\> Set-SPOTenant -EnableAutoNewsDigest $True
```

## Enabling or Disabling the SharePoint App bar

The [SharePoint App bar](#) and the SharePoint start page give users direct access to relevant sites, news feeds, lists and files. To help an organization use the App bar or mitigate situations where it might impact site customizations built with SPFx extensions, SharePoint Admins can disable the App bar using the *Set-SPOTemporarilyDisableAppBar* cmdlet. For example:

```
[PS] C:\> Set-SPOTemporarilyDisableAppBar $True
```

To reenable the SharePoint App bar, run:

```
[PS] C:\> Set-SPOTemporarilyDisableAppBar $False
```

## Enabling or Disabling the Link to Return to Classic SharePoint

Modern SharePoint lists and document libraries can include a link to "Return to classic SharePoint" in the lower left corner of the page to revert the current list or library back to classic mode. On the contrary, the classic mode for lists and document libraries has a hint to "Exit classic experience". To control the visibility of those links, a SharePoint Admin can hide both by executing the *Set-SPOTenant* cmdlet to set *DisableBackToClassic* to *\$true*:

```
[PS] C:\> Set-SPOTenant -DisableBackToClassic $True
```

## Disable Creation of New SharePoint 2013 Workflows

Although Microsoft [has deprecated](#) SharePoint 2013 workflows in SharePoint Online, it's still possible to create new SharePoint 2013 workflows in a tenant. To disable the creation of new workflows, a SharePoint Admin can use the *StopNew2013Workflows* parameter in the *Set-SPOTenant* cmdlet:

```
[PS] C:\> Set-SPOTenant -StopNew2013Workflows $True
```

Existing SharePoint 2013 workflows won't be affected when this setting is true and will continue running. If necessary, the workflows can be modified. Microsoft strongly recommends that organizations move 2013 workflows to Power Automate. To help, an open-source [SharePoint 2013 Workflow assessment tool](#) is available. The tool provides usage data of SharePoint 2013 Workflows and generates a Power BI report to help plan the migration to Power Automate.

## PowerShell and SharePoint Sharing

As discussed earlier, you can limit the ability of users to share content with others in different ways. The SharePoint PowerShell module has several cmdlets that work with sharing settings. The most basic example is how to limit the type of sharing for a tenant using the *Set-SPOTenant* cmdlet. In this example, we allow sharing with authenticated external users. To disable external sharing, set the value to *Disabled*, or use *ExternalUserAndGuestSharing* if you want to allow the generation of anonymous guest links (or to allow these links to be used in team sites).

```
[PS] C:\> Set-SPOTenant -SharingCapability ExternalUserSharingOnly
```

In the next example, we create a list of external domains that we allow users to share content with and add an extra level of security for sharing links by setting *RequireAcceptingAccountMatchInvitedAccount* to *\$True* to ensure that only the account (email address) that receives a sharing invitation can redeem it to access the shared content. If you leave this setting at the default (*\$False*), then anyone who has a link can use it to access the information.

```
[PS] C:\> Set-SPOTenant -SharingAllowedDomainList "locklan.com.au Microsoft.com"  
-SharingDomainRestrictionMode AllowList -RequireAcceptingAccountMatchInvitedAccount $True
```

To revert to the default configuration, reverse the settings:

```
[PS] C:\> Set-SPOTenant -SharingAllowedDomainList $Null -SharingDomainRestrictionMode None  
-RequireAcceptingAccountMatchInvitedAccount $False
```

SharePoint Online supports anonymous guest links that site owners (or users with full control permission for a site) can send to external people to allow them to access site content. The trouble with guest links is that anyone who has the link can use it, so if you decide to allow these links, it is best to force the links to expire after a period (expressed in days). As shown below, a tenant administrator can run the *Set-SPOTenant* cmdlet to assign a default expiry period in the *RequireAnonymousLinksExpireInDays* setting. Users can send guest links that expire sooner than the default, but they cannot exceed it. This setting applies to links sent from both SharePoint Online and OneDrive for Business.

```
[PS] C:\> Set-SPOTenant -RequireAnonymousLinksExpireInDays 10
```

## Site Swap

Site Swap allows SharePoint Admins to replace the location of a SharePoint root site with another site using *Invoke-SPOSiteSwap*. For example, this [PowerShell command performs a Site swap](#):

```
[PS] C:\> Invoke-SPOSiteSwap -SourceURL  
https://office365itpros.sharepoint.com/sites/NewMarketingComms -TargetURL  
https://office365itpros.sharepoint.com -ArchiveURL  
https://office365itpros.sharepoint.com/sites/OldMarketingComms
```

*Invoke-SPOSiteSwap* requires the [following attributes](#) to be configured:

- *SourceURL* parameter indicates the site that is going to replace the root site.
- *TargetURL* parameter indicates the root site to be swapped.
- *ArchiveURL* parameter indicates the new URL for the former root site.

Once the cmdlet is invoked, a background job performs the site replacement and after some minutes and some occasional 404 errors, the new root site will be available.

There are also some restrictions in this feature that should be considered before making a Site swap. The source or target sites can't be associated with a modern Team Site or a Hub Site.

## SharePoint PnP PowerShell Module

The functionality available through the SharePoint Online PowerShell module is limited and restricted to basic administration tasks performed by a SharePoint Online administrator, such as managing sites and tenant settings. To get extra functionality, use the cmdlets available in the [SharePoint PnP PowerShell cmdlets project on GitHub](#), part of the Patterns & Practices community initiative.

To install the PnP PowerShell module from [the PowerShell Gallery](#), run this command:

```
[PS] C:\> Install-Module -Name PnP.PowerShell -Force
```

Many good examples of using the SharePoint PnP cmdlets are available on the web.

## Sharing

As the application name implies, sharing documents, folders, and sites with internal and external users is an important SharePoint feature. Microsoft's terminology can be confusing because of the way SharePoint documentation refers to "external users." For SharePoint, an external user is someone outside the tenant. That user might have a guest user account in the tenant, created by an administrator, or because they were invited to join a Group or a Team. On the other hand, an external user can also be an ad-hoc recipient of a sharing invitation, in which case they authenticate with an access code, or they can receive an anonymous link, which allows anyone with the link to access the content.

The current advice is that SharePoint Online sharing should use the following approach:

- Use Groups and Teams to control ongoing access for external users to content in SharePoint team sites. External people who join Groups and Teams automatically get a guest user account as part of the invitation process and use those accounts to access the document library belonging to the team or group they join.
- Use guest user accounts to control access for external users for a sustained period to content in traditional SharePoint sites. You can create these accounts in the Users section of the Azure AD blade of the Azure portal. SharePoint creates guest accounts when someone shares a site with an external user.
- Use one-time passcodes (OTP) for one-time access for external users. The one-time passcodes are one-time eight-digit numbers sent to an email address contained in a sharing invitation to allow them to open content. Codes are good for 15 minutes.

## SharePoint and Azure B2B Collaboration

The long-term direction for SharePoint Online is to use Azure B2B Collaboration to control sharing. This means that when someone creates a sharing link, SharePoint creates a guest account in the tenant. The sharing link contains an invitation for that person, much like adding someone to a team or group creates a sharing invitation that the external recipient must accept before they can access resources. When the person uses the sharing link and goes through an OTP challenge to validate their credentials, the invitation is accepted, and the guest account becomes fully active and thereafter can use their guest account to access

tenant resources. The resources include individual documents, folders, and sites shared with them and where their account is added to the membership of Groups (including Teams). When the [integration with Azure B2B](#) collaboration is enabled, SharePoint Online creates guest user accounts for all types of sharing recipients as described in Table 8-1.

Target Sharing Recipient	Sharing Steps	Guest account created
Microsoft (MSA) account.	Sharing link sent. User goes through the account validation process. Sharing happens.	Yes
Account in another directory (for example, Zoho Mail). Gmail can federate with Azure AD, in which case users sign-in directly.	Sharing link sent. User redeems OTP. Sharing happens. User must redeem a new OTP for each session.	Yes
Account in another Azure AD tenant, including those without Office 365 (like Yandex.com).	Sharing link sent. User redeems OTP. Sharing happens. User doesn't need OTP for future access.	Yes

Table 8-1: Guest accounts created by sharing activities in SharePoint Online

More information about how to manage guest accounts is in the Managing Users chapter.

**Site People View:** The SharePoint browser interface keeps track of external people with whom users share resources. Sometimes, errors creep into the list of external people, such as when people enter incorrect email addresses into sharing links. To remove these errors and prevent them from showing up as suggested users in sharing dialogs, a site administrator can access the list by adding `/_layouts/people.aspx?MembershipGroupId=0` to the site URL. This exposes the All People view listing all the internal and external users with access to the site. You can remove users from this list whose entry is incorrect for some reason (like a bad email address) or who no longer have access to site resources.

## Sharing Controls

SharePoint Online supports up to 50,000 external shares per securable object (a site, a list/document library, a folder, or a(n) item/document), which should be enough for even the most dedicated sharer. For performance reasons, Microsoft recommends keeping the number of shares on an object to 5,000 or less. Somewhat confusingly, the settings that control how users can share SharePoint content are in two administrative consoles:

- **Microsoft 365 admin center:** Select Settings, then Org settings, and then sharing. A link to the sharing controls is in the Sharing side panel.
- **SharePoint admin center:** Select Policies and then Sharing.

The same basic settings for sharing controls are made available through any of these consoles and the same values appear in each. However, you will find some differences in presentation and emphasis. On the other hand, the Microsoft 365 admin center presents a somewhat more simplified view of sharing controls than either of the other consoles. And of course, you can use PowerShell to adjust settings too.

The **Sharing** section of the SharePoint admin center (exposes the settings to control how sharing with external people occurs within a tenant. Sharing occurs through the exchange of sharing links between the person sharing a file or folder with some other internal or external person. The sharing link is a revocable secret to allow access to designated content. The link can be transferrable or confined to a set of recipients.

SharePoint Online and OneDrive for Business support four types of sharing links, listed here in order from the most permissive to the most restrictive:

- **Anyone:** Also called anonymous sharing, this is the loosest permission. It allows users to share files and folders with anyone who has an email address. SharePoint sends an email with a link to allow the recipient access to the content, but anyone who subsequently has access to the link can use it to access the associated content. If you allow anonymous sharing, you can limit the lifetime of a link (to say, 7 days) and restrict access to view-only rather than view and edit.
- **New and existing guests:** Users can share with external users if those users can authenticate themselves by signing into the tenant using a guest account or by using a verification code.
- **Existing guests:** Users can share with external people, but only if the target users already have a guest account in the tenant directory. It is a good idea to create guest accounts when users need to share information with known external people over a sustained period. Guest users use their credentials to access content shared with them.

**Only people in your organization.** In other words, users cannot share files or folders with external people. Microsoft 365 Apps such as Teams and Lists use the same controls to create sharing links for content within and outside the tenant.

**Shared with information in the sharing control.** The sharing control used across Microsoft 365 [lists the set of people](#) whom a file, a folder, or a list item is already shared with so that document owners know how many people already have access to a file and who they are.

If you do not change the default sharing setting, users can create Anyone links. Changing the sharing setting for a tenant can be a slow process and can take several hours before a change is effective. The sharing settings also allow you to control the type of link created when a user shares a document with someone (choose from Specific people; Only people in your organization; Anyone with this link) as well as the [default permission](#) (View, View and edit for Files; View, View, edit and upload for Folders). Because users usually accept the default link type, you should set the default permission to view-only. We can also restrict users so that they can only share files with authenticated external users who belong to specific domains. The sharing settings apply to every SharePoint Online site and OneDrive for Business account in a tenant.

You can also [customize sharing for a site including the](#) per-site Anyone Link expiration policy, but you cannot make sharing more permissive at a site level than it is for the tenant. [Per-site Anyone Link expiration](#) can be configured through the Sharing option in the SharePoint Admin Center or using the *Set-SPOSite* cmdlet to set the *AnonymousLinkExpirationInDays* parameter that overrides the tenant policy and set a more or less restrictive expiration policy for target sites. This parameter controls how all anonymous/anyone links that have been created (or will be created) will expire after the set number of days. It only applies if the parameter *OverrideTenantAnonymousLinkExpirationPolicy* is set to true. For example, the following command sets a 10-day Anyone link expiration period for the <https://Office365itpros.sharepoint.com/sites/Confidential> site:

```
[PS] C:\> Set-SPOSite -Identity https://Office365itpros.sharepoint.com/sites/Confidential -AnonymousLinkExpirationInDays 10 -OverrideTenantAnonymousLinkExpirationPolicy $True
```

Administrators can also override the Anyone Link expiration period for OneDrive for Business sites by running the *Set-SPOSite* cmdlet and passing the desired expiration period in the *AnonymousLinkExpirationInDays* parameter.

## Advanced Settings for External Sharing

The advanced settings for external sharing include:

- Create an allow list for external domains with which you are happy for users to share files. Conversely, you can create a block list to prohibit sharing with specified domains. You can have either a block or

an allow list; you cannot have both. You can enter a maximum of 120 domain names separated by spaces.

- Restrict external sharing only to users that are members of a specific security group in the tenant.
- Control whether external users must accept sharing invitations from the same email address used in the invitation. In other words, an external user cannot sign in from *DomainA* to accept an invitation sent to them in *DomainB*.
- Control whether external users can share items they do not own (a bad idea, normally).
- Define the number of days after which access for guest users to a Site and OneDrive for Business expires.
- Establish how long a verification code sent to external users is valid before they must re-authenticate. The default value for this setting is 30 days.

## Expiring Access Policy

The expiring access policy allows organizations to control how long (between 30 and 730 days) external users can access SharePoint Online and OneDrive documents. The mechanism works by placing an expiration date on the sharing links created for external users. When the expiration date approaches, site owners receive prompts in the SharePoint Online and OneDrive for Business browser GUIs and via email to manage expiration. They can either [extend access or let the access expire](#).

To enable the Expiry access policy for a tenant, access the sharing policy in the SharePoint Online admin center and edit the *Guest access to a site or OneDrive will expire automatically after this many days* setting. Alternatively, run the *Set-SPOTenant* cmdlet. For example:

```
[PS] C:\> Set-SPOTenant -ExternalUserExpireInDays 60 -ExternalUserExpirationRequired $True
```

Administrators can override the tenant guest expiration setting for individual sites by running the *Set-SPOSite* cmdlet. For example:

```
[PS] C:\> Set-SPOSite -Identity "https://office365itpros.sharepoint.com/sites/CriticalSite1" -OverrideTenantExternalUserExpirationPolicy $True -ExternalUserExpirationInDays 730
```

It's important to underline that this setting applies only to sharing links, direct permission changes, and SharePoint group membership and does not apply to guest access to SharePoint Online sites granted through the membership of Microsoft 365 Groups. Guests who are members of Microsoft 365 Groups can continue to access the content of the group-connected sites for as long as they are members of those groups. Setting an expiring access policy for a tenant or site applies only to sharing granted after the policy becomes effective and only for the sharing links created for guest accounts.

## Other Settings for External Sharing

There are other Sharing settings available in the admin center Sharing section:

- Enable (default option) or disable the setting to display the names of people who viewed their files.
- Provide site owners with the ability to display the names of the people who viewed files or pages in SharePoint.
- Shorten links or change their default permission.

## Per-Site Sharing controls

Admins can configure Sharing settings at the site level. Select a site from the Active sites list, select an existing site and click **Sharing**. In the Sharing settings panel, you can then select the following configurations:

- Change the current **External sharing** setting to another value such as Anyone or Existing guests only. The choice of permitted values depends on the global sharing setting at the tenant level. The external

sharing capability for a site can also be set by assigning a sensitivity label with container management to the site. See the Information Protection chapter for more information.

- Create a list of **external domains** which site members are allowed to share site contents.
- Set the **default sharing link type**. Note that to choose the *Anyone with the link* setting, the site must first be configured to allow anonymous sharing. This link type also allows users to override the default *link expiration*.
- Override the **default link permission** (view or edit) for sharing links.

If external sharing is disabled on a site, an administrator can configure a custom sharing help link to explain why external sharing is disabled or how to request a policy change. This feature can be enabled by running the PowerShell command `Set-SPOTenant -CustomizedExternalSharingServiceUrl <url-address>`.

**Per-Site Sharing Links default to people with existing access.** A SharePoint Admin can also set the default sharing link for a site to “People with existing access” by running the `Set-SPOSite cmdlet` with the `-DefaultLinkToExistingAccess` parameter. Once the setting is applied, any user sharing a file or a folder from the site will get an existing access link that does not change permissions on the file or folder being shared.

## Link Settings to Control Access to Documents

Sharing links created by SharePoint Online and OneDrive for Business can [block recipients from downloading copies](#) when an Office document is shared anonymously, to all users in the organization, or to specific people. To block downloads, a user edits the sharing link settings and toggles the “Block download” setting (Figure 8-2). When a sharing link blocks downloads, recipients can use Office Online to view the content.

As the name implies, an *Anyone* link allows anyone with a link to open the associated content. When a password is set, people must enter the password before SharePoint Online displays the content. The password is a custom word set by the sharer when they create the link: the sharer must transmit the password to authorized recipients separately (it is a bad idea to include the password as a comment sent with the link). The presence of the password means that the content is safe if the link is forwarded or reshared with others (unless the password is also shared). The ability to block the download of other file types (images, 3D, PDF, and more) is also supported and it's [the default setting for some audio and video files](#). For added protection, you can also set a password when sharing a file or folder using *Anyone* links. As the name implies, an *Anyone* link allows anyone with the link to open the content the link points to.



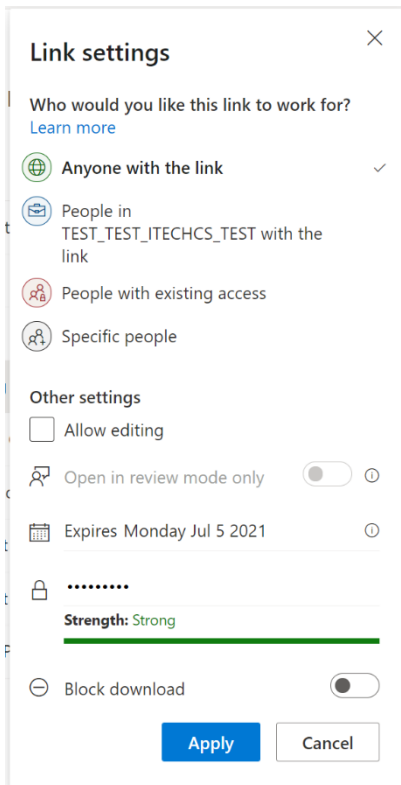


Figure 8-2: Editing a sharing link to block downloads and set a password

**Open in review mode in Word Online.** Users can choose to share Word documents with the "Open in Review mode" option. When enabled, recipients can only open the Word file in the online app with the "Reviewing" mode enabled. While in Reviewing mode, users can suggest changes (using the track changes feature) and make comments, but they cannot make edits that can't be tracked. Owners of the document have the option to accept or reject the suggestions made by the recipient.

## User Experience When Accessing a Shared File

When a user shares a file, recipients receive a link to the file previewer. The link does not automatically download the file. If they want to download the file using a download link, users must open the file details pane, click on More detail at the bottom of the pane, and then use the Copy action next to the Path section.

## Sharing with Microsoft 365 Groups or Teams

Every group has a SharePoint site with a document library. This applies whether the group supports conversations through Outlook, Yammer, or Teams. SharePoint sets the sharing capability for the sites belonging to new groups to *ExternalUserSharingOnly*, which is a level of access enough to allow guest accounts to work with content in the document libraries.

During the creation of a new group or team, SharePoint Online provisions a site to hold the document library and other resources. The tenant sharing settings control what sharing group users can do for files in the document library, but a group owner or tenant administrator can make the sharing more restrictive by changing the *SharingCapability* setting for the group's site. The values that you can set are equivalent to the four settings available for the SharePoint tenant as explained above. Again, we list the permissions below from most restrictive to most permissive.

- **Disabled:** External sharing is not possible.
- **ExistingExternalUserSharingOnly:** Sharing is only possible with external people who already have a guest account in the tenant directory.

- **ExternalUserSharingOnly:** Users can share documents freely with external people with guest accounts and can invite authenticated external users to share documents using one-off time-limited codes. This is the default setting applied to sites created for new groups if a tenant supports guest access.
- **ExternalUserAndGuestSharing:** Users share documents freely with guest users and authenticated external users and can also create links that can be used by anyone who can access the link (also called anonymous or guest link sharing). Be careful about using this sharing capability because it means that any group member can share documents from the document library with anyone else by emailing a sharing link to that person. While acceptable for libraries that do not hold confidential material, this setting should never be present for libraries that hold anything other than documents intended for public access.

The tenant sharing setting prevails over the setting for a site. If you want to use a setting like *ExternalUserAndGuestSharing* for the site belonging to a group, you must first make sure that the organization allows anonymous sharing. To check the tenant setting, use the SharePoint admin center or connect to SharePoint Online with PowerShell and run this command:

```
[PS] C:\> Get-SPOTenant | Select SharingCapability
SharingCapability          : ExistingExternalUserSharingOnly
```

To check the sharing capability for a group, run the command (A connection to both SharePoint Online and Exchange Online is required):

```
[PS] C:\> $sGroupName="Office365forITPros"
Get-SPOSite -Identity (Get-UnifiedGroup -Identity $sGroupName).SharePointSiteUrl | Select
SharingCapability
SharingCapability          : ExistingExternalUserSharingOnly
```

In this instance, the sharing capability for the site matches that of the tenant and is appropriate to support guest access to the site's document library through Groups or Teams.

Other sites might implement a different sharing model to that used by Groups. If you want to align the two sharing models, you can restrict sharing to external users who already have guest accounts in your tenant directory. When this happens, administrators must create guest accounts in the directory before users can invite the holders of those accounts to access their documents. You should also allow SharePoint users to search for guest accounts in the people picker used to select with whom to share a document. By default, guests do not show up in the people picker, so to enable them to appear, you need to run the *following cmdlet*:

```
[PS] C:\> Set-SPOTenant -ShowPeoplePickerSuggestionsForGuestUsers $True
```

Not all team sites are linked to Groups as some are still of the classic variety to serve specific purposes. To list all the team-enabled sites in a tenant, use this cmdlet, which includes the template for group-enabled sites to isolate the set we want to see:

```
[PS] C:\> Get-SPOSite -Template "GROUP#0" -IncludePersonalSite:$False
```

Taking this a little further, here is how to loop down through the set of group-enabled team sites and report some information about the sharing activity for each site.

```
[PS] C:\> $SiteNumber = 0
ForEach ($Site in Get-SPOSite -Template GROUP#0 -IncludePersonalSite:$False)
{
    $SiteNumber++
    Write-Host "Site number: " $SiteNumber " " $Site.Url
    Write-Host "Owned by: " $Site.Owner
```

```
Write-Host "Sharing: " $Site.SharingCapability
Write-Host "-----"
Get-SPOExternalUser -SiteUrl $Site.Url
}
```

**Tracking Document Sharing Through Audit Records:** Each time someone shares a document or folder, SharePoint records the event in an audit record. You can search the audit log to retrieve and analyze these records to understand what sharing occurs within a tenant. See Reporting and Auditing chapter for details.

## Move Files and Keep Sharing

When users move files from a SharePoint site to another site, from a SharePoint site to OneDrive, or from OneDrive to a SharePoint site, they have the option to retain sharing of the file with the same people at the new destination. When this option is used, those with access to the file receive an email to tell them that the file has been moved and that they have a new link or direct permissions to match those at the source location.

## Sharing with LinkedIn Contacts

If a tenant is configured to connect to LinkedIn, Microsoft 365 users can connect their accounts to their LinkedIn accounts (see Managing Users chapter for details). When an account is connected to LinkedIn, first-degree contacts are downloaded from LinkedIn and included in the “suggested people” list used by browser interfaces such as the SharePoint Online and OneDrive for Business clients. The suggested people list also includes tenant users (including guest accounts) and email addresses from Outlook’s auto-complete list.

When the user next shares a document from SharePoint Online or OneDrive for Business or the online Office apps, the name they enter is checked against the suggested people list. If a match is found against a LinkedIn contact and the user goes ahead and shares the document, the sharing invitation is sent to the contact’s email address.

## At a Glance Summaries in Sharing Emails for Word Documents

When a user shares a Word document, the notification to inform the recipient about their access to the document includes a list of key points in the text and the time estimated to read the content. Files identified as sensitive by Data Loss Prevention policies will not include this information.

Administrators can disable the feature for the tenant by running the *Set-SPOTenant* cmdlet:

```
[PS] C:\> Set-SPOSite -IncludeAtAGlanceInShareEmails $False
```

## Sharing from Microsoft Lists

Sharing capability for Microsoft Lists is achieved in two ways:

- By granting access to the entire list with specific permission (Full control, Edit, View) to specific users, mail enabled security groups, Microsoft 365 Groups, or security groups. In practice, granting access in this way implies the creation of unique permissions for the list.
- Through the sharing link generated when a specific list item is shared through the universal sharing dialog.

## Managing Sharing and Access Request Settings

Site admins and site owners can manage **Sharing and access request settings** directly from the site permissions panel (through the **Change sharing settings** link). This simplifies the previous experience which involved first going to the Advanced permissions settings page and then selecting the “Access Request Settings” option in the ribbon. Configuring settings through the new experience uses a simple configuration panel where a site admin or a site owner can:

- Modify sharing permissions by choosing one of the following options: “Site owners and members can share files, folders, and the site,” “Site members can share files and folders, but only site owners can share the site” and “Only site owners can share files, folders and the site.”
- Allow or deny access requests to the site. This setting, enabled by default, also includes the choice to direct site access requests to site owners or are sent to a specific e-mail address. A custom message for the access request page can be also added here.

## Managing Access to Files and Folders in SharePoint Online

The **Manage access** pane makes it easy for people to manage access permissions, remove individual recipients from shared links, and stop sharing overall. This pane, available for both SharePoint Online and OneDrive for Business, shows up when the user clicks the **Manage access** link under the **Has access** link in the files/folder details pane. From here, the user can easily grant access to the selected file/folder or manage existing sharing permissions.

## Sharing Reports for SharePoint Online and OneDrive for Business

Site owners can generate a CSV file to report site content shared with any user. The option is available on the Site usage page. To create a Sharing report, a site owner browses the site usage page and clicks the **Run report link** in the **Shared with external users** section. When prompted, the user selects a folder to save the CSV report and then clicks **Save**. Once the report is generated, SharePoint sends a notification to say the CSV file is available in the chosen location. The report contains a row for every file or folder shared by the user with [the following information](#): File/Folder path, item type (Web, Folder, Document), permission level applied to the item (Total Control, Collaborate, Read), user name, user e-mail, user or group type (Internal, External, SharePoint Group, Security Group or Microsoft 365 Group), link identifier, link type (Anyone, People in the organization with the link, People with existing access or Specific people) and the link ID used to access the item.

A similar sharing report feature is available for OneDrive for Business, the difference being that instead of focusing on the sharing activities for everyone on a site basis, the OneDrive variant only looks at sharing of a specific user’s personal files. The **Run sharing report** link used to generate a sharing report is available on the modern OneDrive settings page. An example of how to use the Microsoft Graph to report sharing within OneDrive is [posted here](#).

## Tracking Shared Files

The *Set-SPOTenant* cmdlet also controls a setting often used by administrators to keep track of files shared with external users. The example shown below sets the *BccExternalSharingInvitations* setting to *\$True* to force SharePoint Online to generate a BCC message to a nominated list of email addresses (in a comma-separated list with no spaces) each time a user shares a file stored in a SharePoint Online or OneDrive for Business document library with an external person. The users specified in the *BccExternalSharingInvitationsList* parameter will receive a message saying that the user wishes to share a file with the addressee of the message.

```
[PS] C:\> Set-SPOTenant -BccExternalSharingInvitations $True -BccExternalSharingInvitationsList administrator@Office365ITPros.com
```

**OneDrive for Business Notifications:** SharePoint Online controls many settings for OneDrive for Business, among which are notifications for events connected with sharing files. The Sharing settings available in the SharePoint admin center allow you to opt for users to receive notifications via email when:

- Users invite external users to access shared files.
- External users accept invitations and open shared files.
- Users create or update an anonymous link used to share files.

## SharePoint Online Extensibility Options

Traditionally, one of the strongest characteristics of SharePoint has been its natural extensibility: an organization can change the system to meet business requirements in several ways. Unfortunately, the introduction of SharePoint Online initially featured a strong reduction of the extensibility options when compared to SharePoint on-premises. This was due to the intrinsic multitenant design of the platform and because Microsoft wanted to reduce the fragility of custom development.

SharePoint Online inherited from its on-premises counterpart the Add-ins model as a customization mechanism. This model supports the creation of integrated applications in the platform while running totally outside the main structure. Furthermore, it was (and still it is) impossible to install and run server-side code in SharePoint Online, something that removes the possibility to create traditional SharePoint artifacts such as classic Web Parts, Timer Jobs, Event Handlers, etc. Slowly, as the SharePoint Online platform matures and evolves, substitutes for older integration mechanisms have been added to SharePoint Online through the natural integration with other platforms (Azure, Power Automate), the work done by the SharePoint Team with SPFx and specific SharePoint community initiatives such as the PnP project.

### SharePoint Add-ins and the Global App Site

SharePoint Add-ins are self-contained pieces of functionality that extend the capabilities of SharePoint to solve business problems. Add-ins don't have custom code that runs within SharePoint Online: all custom logic moves to servers outside the SharePoint platform. Keeping custom code out of SharePoint guarantees that the Add-in can't harm SharePoint or reduce the performance of the system.

The Global App Site makes internal custom Apps available for users to install when they search apps using the "From my organization" filter on the Site Contents page. Site owners can add these apps to customize sites with specific functionality or to display information. After a Global App site has been created, it is possible to upload any custom App that the organization has developed by copying the manifest document to a document library and setting some properties. The global App Site also supports the management of app requests from users, checking how Add-ins are used, and the maintenance of license information.

The Global App Site is not created by default in a new tenant but can be added from the SharePoint admin center. In the More features section, click the "Open" button under Apps, select App Catalog, and then follow the steps to create a new App Site in the tenant. This creates a Global App Site, meaning that every Add-in uploaded to the App Site will be available for all sites in the tenant. But tenant administrators can choose to enable "Site App Catalogs" for specific sites so that solutions deployed to the Site App Catalogs can only be installed on that specific site.

Site App Catalogs are configured and managed using PowerShell and the Global App Site must exist already. Use the *Add-SPOSiteCollectionAppCatalog* cmdlet to create a Site App Catalog, indicating the site where the app catalog should be created with the *-Site* parameter. An "Apps for SharePoint" document library will be added to the site to deploy SharePoint Add-ins and SPFx solutions. To disable the Catalog, use the *Remove-SPOSiteCollectionAppCatalog* cmdlet indicating the site in the *-Site* parameter; this prevents new components from being added and any previously installed components from executing their code. It is not possible to enumerate with PowerShell all the site collections in the tenant that have the Site App Catalog enabled, but it's possible to see those Site App Catalogs through the "Site App Catalogs" list in the Global App Site. This is a hidden list that stores a record for every Site App Catalog created in the tenant. In addition, be aware that although solutions installed in Site App Catalogs can only be used in these specific sites, they can theoretically use resources from other sites in the tenant.

### SharePoint Framework (SPFx)

SPFx is a Page and Web Part development model that empowers client-side development, integration with SharePoint Online and on-premises, and integration with the Microsoft Graph. It is based on open source

tooling and JavaScript technologies. SPFx is used in SharePoint as the extensibility paradigm for the client-side development framework to allow developers to implement new functionality in SharePoint.

SPFx is used extensively by Microsoft to build the modern user experiences seen in SharePoint Online, but external developers can use the same technology, tools, and techniques to build more productive experiences and apps that are responsive and mobile-ready.

From the infrastructure point of view, the files required for SPFx are just JavaScript archives plus the accompanying style sheets, graphics, and other required files that should be deployed to an accessible place for the client computers. This can be implemented using the SharePoint Content Delivery Network (CDN), a SharePoint document library, the Azure CDN, or any other available private or public CDN. Additionally, a manifest file must be deployed to the SharePoint Catalog, so that SharePoint is aware of the existence of the components. Normally a distribution package is created by the developers and delivered to the administrators that install them in the SharePoint Catalog, in a similar way as for SharePoint Add-ins.

## Microsoft 365 Patterns and Practices (PnP)

The [Microsoft 365 Patterns and Practices community](#) initiative comprises components aimed at enhancing the functionality of SharePoint on-premises and Online in several ways, filling the development and infrastructure gaps left by Microsoft. The components include:

- [PnP Framework](#), a .NET standard 2.0/.NET 5.0 library targeting Microsoft 365 containing useful extensions to extend SharePoint and Microsoft 365 APIs.
- [PnP Transformation Framework](#), a generic solution designed to transform any web-based/content-based solution into a modern SharePoint Online Portal.
- [PnP PowerShell cmdlets](#).
- [PnP Scripts Samples](#), a web site with several samples demonstrating how to use PowerShell in different use cases about getting information and managing SharePoint Online and other Microsoft 365 workloads.
- [PnP Core SDK](#), a unified object model to work with SharePoint Online and Teams that is agnostic to the underlying APIs. Over the time, the developers plan to extend the SDK to support other Microsoft 365 workloads.
- **A complete [PnP Provisioning Engine](#)** is included in the PnP Framework.
- [PnPJS](#), a JavaScript library to help consume SharePoint and Office 365 APIs securely.
- PnP tools, such as the PnP SPFx Yeoman Generator, the CLI for Microsoft 365, or the Microsoft Graph PowerShell SDK among others.
- [PnP guidance](#).
- **Ready to use solutions** such as the [PnP Starter Kit](#).
- **The SharePoint look book**, a [free service](#) to deploy customized sites in a SharePoint tenant based on a selected site template. The service includes several ready-to-use templates. Each template consists of a customized home page, lists, document libraries, and custom Web Parts.

Using PnP reduces the development effort needed to work with SharePoint. For example, to create a new list or document library using PowerShell requires the use of the SharePoint CSOM API (create a context, recall the web and site objects, create a new List object, execute the operation, etc.), as well as some knowledge about programming and the internal structure of the service. With PnP PowerShell cmdlets, you only need to connect to SharePoint Online for a tenant and use *New-PnPList* as follows:

```
[PS] C:\> Connect-PnPOnline -Url "https://tenantname.sharepoint.com/sites/namesitecollection" -
SPManagementShell
[PS] C:\> New-PnPList -Title "PnPList" -Template GenericList -Url "Lists/PnPList"
```

The *Connect-PnPOnline* cmdlet needs valid user credentials to make the initial connection to the tenant. Then a new list called "PnPList" is created using the *New-PnPList* cmdlet and the *Template* parameter to specify the

list template to use. Be aware that the connection is only valid for the given URL endpoint; if you want to use another site, you need to make a new connection.

Provisioning is probably the most used and most powerful PnP feature. The recommended way to provision in SharePoint Online is to create a new "standard" object in SharePoint, and then apply all the customizations using scripting (thus not creating new templates as traditionally used). The PnP provisioning engine is powerful enough to create and modify almost any element in a site: the creation of subsites, Lists, Libraries, Custom Fields, Content Types, views, attach users, etc. There are several cmdlets and ways to provision them using PnP:

- Make a template (just an XML file) that contains all the elements to be created, deleted, or modified, and run the *Invoke-PnPSiteTemplate* cmdlet pointing to the XML file and URL endpoint. The creation of fully customized SharePoint objects using this way is possible (and recommended).
- Manually create all the customization required in a dummy site and generate the XML template using the *Get-PnPSiteTemplate* cmdlet. Then, the template can be used to provision as many clones as necessary using *Invoke-PnPSiteTemplate*.
- PnP Provisioning Engine and XML templates can easily be used in .NET code, so the same process described to get and apply a provisioning template with PowerShell can be done programmatically.
- There are also specific PnP cmdlets to manage some provisioning objects. For example, to create a new view for the early created List, use:

```
[PS] C:\> Connect-PnPOnline -Url "https://tenantname.sharepoint.com/sites/namesitecollection"
[PS] C:\> $myList = Get-PnPList -Identity "Lists/PnPList"
[PS] C:\> Add-PnPView -List $myList -Title "myView" -SetAsDefault -Fields ID,Title,Created
```

## Site Scripts and Site Templates

Site scripts and site templates are [extensibility mechanisms for SharePoint Online](#) that allows administrators to create, deploy and apply templates to [modern and classic sites based on a specific site template](#). Site scripts and site templates can be applied during site creation, as part of a hub association, or to existing sites. End users with the required permissions can also apply [scenario-based templates](#) or organization templates to existing SharePoint sites by using the *Apply a site template* option in the site settings menu. This option is also based on the site scripts and templates extensibility model.

The following site types support site scripts and site templates:

- Modern Team Sites linked to a Microsoft 365 Group (GROUP#0).
- Modern Team Sites not linked to a Microsoft 365 Group (STS#3).
- Communication sites (SITEPAGEPUBLISHING#0).
- Channel site templates (TEAMCHANNEL#0, TEAMCHANNEL#1). The latter template is now generally used for sites belonging to both private and shared channels.

**Site templates and Hub Sites:** Site templates can be [applied to Hub Sites to apply to the associated sites too](#). This option is an optional setting for a given hub site and requires the site template and site script to already be available in the tenant.

Site templates are always based on an out-of-the-box template, and can be deployed through PowerShell as follows:

```
[PS] C:\> Add-SPOSiteDesign -Title "HR Site" -WebTemplate "64" -SiteScripts "<ID>" -Description "HR Department Site"
```

Site templates always reference one or more site scripts (by their identifiers) to apply the customizations defined by the script. *WebTemplate 64* is the identifier for the Modern Team sites while 68 is for Communication sites and 69 for Teams channel sites. It's important to note that there is a limit of 100 site scripts and 100 site templates per tenant.

A site script is a JSON file defining the customizations to be added to a new site (where the site template is the actual template to define the site). Site scripts are uploaded to a gallery at the tenant level, making them available for all sites. Site scripts are deployed using the *Add-SPOSiteScript* cmdlet, which returns the ID to be used by *Add-SPOSiteDesign*:

```
[PS] C:\> Add-SPOSiteScript -Title "Create HR lists" -Content $mySiteScript -Description "Creates lists for HR site"
```

In the above example, the *\$mySiteScript* variable contains the JSON string with the actions to implement.

**Site template JSON Schema:** Microsoft updates the [JSON Schema](#) regularly. The schema defines all the actions and sub-actions that can be included in a site template. To simplify the creation of a Site template, there is a [free online tool](#) to visually define the JSON structure for a specific Site Design.

## Using PowerShell to Apply Site Templates

A site template can be applied to an existing site by using the Apply a site template option in the settings menu for a site or by running the *Invoke-SPOSiteDesign* or *Add-SPOSiteDesignTask* cmdlets. The first option allows the site administrator to view all the site templates available and apply them to the site while updating site templates with PowerShell is a good way to automate updates across a set of sites. The two cmdlets take a different approach in how they implement site templates:

- *Invoke-SPOSiteDesign* applies the site template to the site immediately.
- *Add-SPOSiteDesignTask* adds the site template application to a schedule to be run as a background job. In addition, this cmdlet extends the 30-action limit in a site script that is applied synchronously to 300 actions (or 100,000 characters). More information about the limits applying to site scripts and site templates [can be found in Microsoft's documentation](#).

Site scripts can also be generated from existing sites, lists, and document libraries, which means that existing site settings, lists, and document libraries can be easily ported to another site. The *Get-SPOSiteScriptFromList* cmdlet can be used to autogenerate a site script from a list. Alternatively, the *Get-SPOSiteScriptFromWeb* cmdlet can autogenerate a site script that includes the following site settings: branding, theme, regional settings, sharing capability, and the lists and document libraries specified in the *IncludedLists* parameter:

```
[PS] C:\> $sSPOSiteUrl=" https://redmondassociates.sharepoint.com/sites/0365ExchPro"
[PS] C:\> Get-SPOSiteScriptFromWeb -WebUrl $sSPOSiteUrl -IncludeBranding -IncludeTheme
-IncludeRegionalSettings -IncludeSiteExternalSharingCapability
-IncludedLists("Lists/ListName1", "Lists/ListName2")
```

**Discover which site templates are applied to a site:** The *Get-SPOSiteDesignRun* cmdlet shows all the site templates that have been applied to an existing site. The *Get-SPOSiteDesignRunStatus* cmdlet returns the result of each action from every site script in a site design. Site owners can also [view templates](#) through the View template history link available in the Site Information panel.

## Configuring Site Regional Settings with Site scripts

Site scripts can set default regional settings for new sites. This is important if you want to ensure that sites display the correct document creation and modification times in local format when accessed through a browser. By default, SharePoint Online sets the default time zone to Pacific Daylight Time (UTC -8), which is fine for people in Redmond but not so good for users elsewhere in the world. You can always update these settings by accessing the Regional Settings of a site. However, given that Teams and Groups create many of the new SharePoint Online sites and their owners are usually people with little knowledge of SharePoint internals, it is best if the settings are correct from the start.

An example of [how to configure regional settings](#) is available on GitHub. The steps are straightforward. First, create a JSON-formatted input file holding the settings that you want to use. In this example, I select GMT as



the default time zone, that we use the 24-hour format instead of the 12-hour format, and that Ireland is the default country (locale id or LCID). Examples of other locale ids are 3082 (Spain), 5129 (New Zealand), and 1035 (Finland). A [full list of locale ids](#) is available online.

```
{
"$schema": "schema.json",
"actions": [
  {
    "verb": "setRegionalSettings",
    "timeZone": 2, /* Greenwich Mean Time */
    "locale": 6153, /* Ireland */
    "sortOrder": 25, /* Default */
    "hourFormat": "24"
  }
],
"bindata": { },
"version": 1
}
```

Save the JSON data to a file. Then connect to SharePoint Online with PowerShell and execute the following command to retrieve the settings from the file and load them into a variable (*\$SiteScript*). Make sure you install the latest version of the SharePoint Online PowerShell module before trying to use the site template cmdlets.

```
[PS] $SiteScript = (Get-Content "C:\Temp\RegionalSettings.json" -Raw | Add-SPOSiteScript -Title "Set Ireland Regional Values" -Description "Sets locale, time zone, and hour setting for Ireland")
```

Now run the *Add-SPOSiteDesign* cmdlet to load the settings as default regional values.

```
[PS] C:\> Add-SPOSiteDesign -Title "Set Ireland Regional Values" -WebTemplate "64" -SiteScripts $SiteScript -Description "Applies Ireland regional settings" -IsDefault
```

To test that the change is effective, create a new team or group and examine the regional settings for the site that SharePoint Online provisions. Open the site with SharePoint, select Site Contents, then Site Settings, and then Regional settings. The locale and time zone settings should be the values set in the script. Alternatively, you can apply the site template to an existing site through the Site templates dialog available in the settings menu, just select the desired site template and click on the "Apply to site" button.

If you make a mistake, you can update the site design with the *Set-SPOSiteDesign* cmdlet or remove the site template and start over. Run the *Get-SPOSiteDesign* cmdlet to return a list of site templates in the tenant and note the identifier (Id) for the design you want to remove. Then run the *Remove-SPOSiteDesign* cmdlet to remove it.

```
[PS] C:\> Get-SPOSiteDesign

Id                : f466a1da-e2a2-4107-b558-35954ad199de
Title             : Default
WebTemplate       : 64
SiteScriptIds    : {9e287dd9-b2b1-4ceb-8649-998f43b24c1d}
Description      : Applies Ireland regional settings
PreviewImageUrl  :
PreviewImageAltText :
IsDefault        : True
Version          : 1

[PS] C:\> Remove-SPOSiteDesign -Identity f466a1da-e2a2-4107-b558-35954ad199de
```

Setting default regional settings for new sites with a site script does not update regional settings for existing sites. If you want those sites to use the correct time zone and locale, you can either update the settings manually or use a script. Here's an example that looks for Microsoft 365 Groups that are provisioned for SharePoint and then applies a site template (identified with a GUID) to each site.

```
[PS] C:\> $SitesUpdated = 0
$DesignID = " f466a1da-e2a2-4107-b558-35954ad199de"
$Groups = (Get-UnifiedGroup | ? {$_SharePointSiteUrl -ne $Null} | Select SharePointSiteUrl,
DisplayName, Alias)
ForEach ($G in $Groups) {
    Try {
        Write-Host "Processing" $G.SharePointSiteUrl "for group" $G.DisplayName
        Invoke-SPOSiteDesign -Identity $DesignID -WebUrl $G.SharePointSiteURL -ErrorAction Stop
        $SitesUpdated++
        Set-UnifiedGroup -Identity $G.Alias -CustomAttribute13 "Site Design Updated" }
    Catch {
        Write-Host "Problem Processing" $G.SharePointSiteURL}
}
Write-Host $SitesUpdated "sites updated successfully. You need to check the following and update
them manually"
Get-UnifiedGroup -Filter {CustomAttribute13 -eq $Null} | Sort DisplayName | Format-Table
DisplayName, SharePointSiteURL
```

The list of groups generated at the end of the script includes all groups that we cannot update. In some cases, this is because SharePoint is not fully provisioned for the group (a more common problem for older groups) and no one has ever tried to access the document library. If the site does not exist, we cannot update its regional settings. In other instances, the account used to run the script might not have the necessary permissions to update site settings. Administrators can update these sites manually as time allows.

## Managing Built-In Site Templates

SharePoint administrators can control which of the built-in site templates are available to site creators. To retrieve the current state of specific built-in site templates, run the [Get-SPOBuiltInSiteTemplateSettings](#) cmdlet. Then, to show or hide a template, run the [Set-SPOBuiltInSiteTemplateSettings](#) cmdlet. For example, to hide the "Event planning" site template, run the following command:

```
[PS] C:\> Set-SPOBuiltInSiteTemplateSettings -Identity '9522236e-6802-4972-a10d-e98dc74b3344' -
isHidden $true
```

## Additional Features in SharePoint Online

Apart from the day-to-day business of site and service management, several operations exist that only need to be performed on an on-demand basis. We cover these actions here.

### Tenant Rename

Currently available in preview, Tenant rename allows organizations to [change the SharePoint domain name](#) using the following steps:

- Add and verify the new domain name in the custom domain names blades in the Azure AD portal.
- Initiate the SharePoint tenant rename using the *Start-SPOTenantRename* cmdlet with the attributes *DomainName* and *ScheduledDateTime* configured with the target domain name and the date and time to start the rename process. For example, this command instructs SharePoint Online to begin the domain rename process on July 31, 2022, at 10:25 UTC.

```
[PS] C:\> Start-SPOTenantRename -DomainName "NewDomainName" -ScheduledDateTime "2022-07-31T10:25:00"
```

- Give the tenant rename process some time (it can take several hours or even days depending on the number of SharePoint sites and OneDrive accounts to process) before starting to review if everything works as expected under the new domain.

To check the status of the rename operation, use the *Get-SPOTenantRenameStatus* cmdlet. Similarly, the *Get-SPOSiteRenameState* cmdlet checks the state of a specific site.

## Microsoft Defender for Office 365

[Microsoft Defender for Office 365](#) can be enabled on E5 tenants or E3 tenants to run check files in document libraries to make sure that they're not infected with malware. Microsoft Defender for Office 365 P1 is also available in Microsoft 365 Business Plans.

Microsoft Defender for Office 365 works across all SharePoint Online and OneDrive for Business sites in a tenant. It uses a mixture of technologies, including advanced threat heuristics generated from intelligence gathered by Microsoft about malware to find suspect files. Instead of scanning all files, which would consume enormous server resources, Microsoft Defender for Office 365 focuses on blocking the spread of malware. When a user shares a document, it checks whether sharing should go ahead. If all is well, sharing proceeds as normal. If not, SharePoint locks the file and disables the ability to download, open, or share it. The user can remove the file to solve the problem. See the Mail Flow chapter for additional information about Microsoft Defender for Office 365.

## Restore This Library for SharePoint Online

To help recover documents from accidental deletions, a malware attack, or other data loss, the Restore this library feature (available to administrators and site owners in the library settings panel) is a self-service recovery mechanism to allow site administrators and site owners to roll back changes made to a document library to any point in time during the last 30 days. You can select from several out-of-the-box restore points (Yesterday, One week ago, Three weeks ago, Custom date and time), or use a slider to move through the set of changes for the library to select any point up to the 30-day limit (Figure 8-3). The changes that can be rolled back include updates to document properties such as title or assigning a retention label.

After you select a restore point, SharePoint prompts you to confirm to go ahead with the restore, and if confirmed, SharePoint rolls back all the changes from the library to the chosen point in time to restore the library to its state at that time.

Restore HR Department - Documents

If something went wrong, you can restore this library to a previous time. Select a date preset or use the slider to find a date with unusual activity in the chart. Then select the changes that you want to undo.

Select a date

Three weeks ago All changes after 3/21/2020, 12:00:00 AM will be rolled back

Restore Cancel

Move the slider to quickly scroll the list to a day.

Days ago

Select a change in the list below to highlight it and all the changes before it. Then select the Restore button to undo all the highlighted changes.

Change	File name
2 days ago - 4/9/2020 (3)	
Renamed by Juan Carlos González Martín 4:44:23 PM	HR Document #31.docx > HR Document #6.docx
Renamed by Juan Carlos González Martín 4:44:09 PM	HR Document #21.docx > HR Document #5.docx
Renamed by Juan Carlos González Martín 4:43:55 PM	HR Document #11.docx > HR Document #4.docx
6 days ago - 4/5/2020 (12)	
Updated by Juan Carlos González Martín 11:46:33 PM	HR Document #2.docx

Figure 8-3: Using the Files this library feature for SharePoint Document libraries

## Multi-Geo Support for SharePoint Online

[Multi-Geo support for SharePoint Online and OneDrive for Business](#) is a mechanism to control where data from both services are stored at rest to meet data residency requirements when users work in multiple geographies. [Multi-Geo capabilities are available as an add-on to specific Office 365 Plans](#) for Enterprise Agreement (EA) customers with a minimum of 250 paid licenses. For more information about multi-geo, including pricing, contact your Microsoft account team.

## Encryption for SharePoint Online and OneDrive for Business Storage

Files in SharePoint Online and OneDrive for Business are protected by unique, per-file keys that are exclusive to a tenant. The keys (AES 256-bit), which are created and managed by the SharePoint Online service itself or by customers (when customer-managed keys are used), are used to encrypt files stored in SharePoint Online and OneDrive for Business.

## Comments on non-Office files in SharePoint Online and OneDrive for Business

The Office 365 file viewer supports more than 350 file types and includes the ability for users to collaborate and communicate about non-Office files with comments and replies. Anyone who opens a non-Office file can view the comments through the details pane and leave additional comments. File owners will receive notifications when someone comments on their files. Notifications are also sent to users when someone leaves a reply to their comments. Users may unsubscribe from these e-mail notifications via the "Unsubscribe" link in the notification e-mail.

## Intelligent File card in SharePoint Online and OneDrive for Business

The File card in SharePoint Online and OneDrive for Business incorporates several [intelligent features](#) to quickly get deeper information about relevant activities happening around the document (edits, mentions, and comments), document statistics (number of views, number of viewers and details about who viewed the file), suggestions about to whom share the document based other people the current user frequently works with, and [conversations](#) happening in e-mail or Teams about the document.

## Enabling Loop Components in Applications

Microsoft's Fluid framework is a technology ([now available as open source software](#)) designed to help developers build better collaborative applications using "live" components. The big selling point for the Fluid framework is its synchronization capabilities, which allow components to coordinate updates made by multiple people to the same content and coherently present the information. Apart from the new Microsoft Loop application, scheduled for availability sometime in 2022, Microsoft 365 supports loop components in Teams chat messages, email in OWA and Outlook for Windows, and Microsoft Whiteboard. See the Teams architecture chapter for more information about the implementation of loop components in Teams.

Microsoft controls the use of loop components in Microsoft 365 applications using a setting in the SharePoint Online tenants configuration. ***IsLoopEnabled*** controls the availability of Loop components in applications like Teams, OWA, and Outlook desktop. This setting replaces the older ***IsFluidEnabled*** setting, which is due for deprecation in November 2022. A related ***IsWBFluidEnabled*** setting controls if Whiteboard uses OneDrive for Business storage. Microsoft will retire this setting after Whiteboard completes its transition to OneDrive for Business in all tenants.

You can update the settings with the *Set-SPOTenant* cmdlet. For example:

```
[PS] C:\> Set-SPOTenant -IsLoopEnabled $True -IsWBFluidEnabled $True
```

To disable the setting, update it to `$False`. Due to an issue with offline access to Loop components in items recovered by eDiscovery searches, some organizations do not allow the use of Loop components. Microsoft is working on an offline access capability, but this is not yet available.

# OneDrive for Business

OneDrive for Business (or OD4B as it is known in the Microsoft 365 ecosystem) is a cloud-based file sharing service based on SharePoint Online to provide users with a personal storage space. From an administrative perspective, a OneDrive account behaves like a personal version of a SharePoint document library. Microsoft Search indexes all the information stored in OneDrive, meaning that it is discoverable and available for compliance purposes.

Referring to anything as a personal storage space is a bit of a mouthful. In a practical sense, users can store anything in their OneDrive for Business account, just like they would do with a personal network share on an old-fashioned file server. Indeed, the prime purpose of OneDrive for Business is to help companies to wean themselves off file servers by moving the content held on the servers into the cloud. Users who have personal files stored in network shares can move their information to OneDrive for Business while Teams or SharePoint Online team sites are good homes for organizational or shared information.

## Access on Multiple Devices

The big advantage of using OneDrive for Business from a user's perspective is that once they store files on OneDrive, the files are available from any device. When network connectivity was not as pervasive and capable as it is now, it would not have been possible to contemplate such a movement, but it is now. Although you need to have network access to browse OneDrive for Business sites, the OneDrive sync client allows users to work with files copied to a local cache. Some restrictions exist in terms of the types of files and the names of files that OneDrive for Business supports. It is sensible to check on the [current limits and restrictions](#) before you start a project to help users move their information from file servers. The limits for the current OneDrive sync client are much higher than in earlier versions, so you should confirm whether any limits that stopped a project moving forward still exist. It is also sensible to implement whatever controls are necessary for client synchronization (discussed in the next section) at the start of a migration project instead of trying to retrofit some restrictions afterward.

The ability of OWA, Outlook, and Teams to access files held in OneDrive for Business gives further encouragement to use the service. For example, users can upload files to OneDrive for Business and then send a link to the attachment (a "cloudy" attachment). Users can also save attachments from Outlook or OWA direct to OneDrive for Business or to the document libraries in Groups where they are members. The intention behind making these features available is to help users change their habit of storing email attachments on their local PC and refocus toward cloud-based storage instead. Whether it is possible to break the "always send full copies of attachments by email" habit is debatable. What is for sure is that breaking the habit needs people to change their personal workflow, and that is always hard.

Another advantage gained by keeping files in OneDrive for Business is that users can grant others the right to edit the document in place, which might encourage better collaboration. By using cloudy attachments, recipients get to see the latest version of the information they receive because access is always to the version held in OneDrive for Business rather than a personal copy, thus avoiding the problem that someone might update a copy received in email by the time you read an attachment.

**Install OneDrive as a Progressive Web App:** OneDrive can be [installed as a Progressive Web App](#) in Microsoft Edge, Google Chrome, or any other modern web browser that supports Progressive Web Apps.

## OneDrive for Business Contents and Limits

The first time a user accesses OneDrive for Business, SharePoint Online provisions the site. Once the site is ready, the user can upload files to the site and share those files (or complete folders). Users can attach files

from OneDrive to messages as easily as local files. To ensure that users have enough space to move their personal information from old file servers, [depending on their license type](#), users receive up to 1 TB of online storage, which comes from an allocation that is separate to and not counted against the pooled storage used for SharePoint Online. Tenants that have five or more licenses for E3 or higher enterprise plans (plus government and academic equivalents) can take advantage of “unlimited storage”.

OneDrive assigns a default storage quota of 1 TB to each user for their personal site, which should be enough to move personal files off file shares and even from local PC hard drives. If you want to adjust the OneDrive for Business storage for an account, you can either send a support ticket to Microsoft or ask an administrator to run the *Set-SPOSite* cmdlet to increase the quota. Here is how to increase the storage quota from the default 1 TB to 5 TB for the *Tony.Redmond@Office365ITPros.com* account.

```
[PS] C:\> Set-SPOSite -Identity
https://office365itpros-my.sharepoint.com/personal/tony_redmond_office365itpros_com -StorageQuota
5242880
[PS] C:\> Get-SPOSite -Identity
https://office365itpros-my.sharepoint.com/personal/tony_redmond_office365itpros_com | Format-List
```

Later, if the account approaches the 5 TB storage limit, you can increase the quota to 10 TB (set the value to 10485760). You can keep on increasing the quota in 5 TB chunks as the need arises. An attempt to increase the quota will only succeed if the used quota is approaching a chunk boundary. In other words, you cannot increase the quota assigned to an account from 3 TB or 15 TB but must first increase the quota to 10 TB and then wait until the storage used passes 12 TB or thereabouts before upping the quota to 15 TB.

To check the quota assigned to each user and how much of the quota they use, you can run the following PowerShell snippet.

```
[PS] C:\> Get-SPOSite -IncludePersonalSite $True -Limit All -Filter "Url -like
'my.sharepoint.com/personal/'" | Select Owner, StorageUsageCurrent, StorageQuota
```

Owner	StorageUsageCurrent	StorageQuota
james.gangley@office365itpros.com	1	1048576
brian.weakliam@office365itpros.com	1	1048576
tony.redmond@office365itpros.com	3727	5242880

No error checking is in this simple code, so you will see an error if an attempt is made to check a user's OneDrive site that has not yet been provisioned because the user has never used OneDrive. To fix that problem, you can request OneDrive for Business to provision a site for the user by running the *Request-SPOPersonalSite* cmdlet. The input is a comma-separated list of email addresses for the users for which OneDrive is to create sites. You can specify up to 200 email addresses in the list.

```
[PS] C:\> Request-SPOPersonalSite -UserEmails "Sanjay.Ramaswamy@office365itpros.com,
David.Pelton@office365itpros.com" -NoWait
```

OneDrive uses a timer job to create the sites. The actual time when the sites will exist depends on system load.

## Information Views in OneDrive for Business

OneDrive for Business includes several views to give the user the ability to see their information as well as the information they have shared with other users and the information other users have shared with them. Table 8-2 details the current information views included in OneDrive for Business.

Information View	Information View Details
My files	Shows the files and folders stored in the user's OneDrive for Business account.
Recent	Shows a list of recent files the user has worked on, including the date and time of the last access.

Shared	Two additional information views are available. "Shared with you" shows the files shared with the user including the person sharing the file and the last file activity. This view also includes the "Popular around you" section that surfaces files recommended to the user based on his/her working relationships. This section is intended to help find relevant and trending information faster and to also discover content the user might not have been aware of. "Popular around you" only shows the information the user can access and does not change file permissions. "Shared by you" shows the files the user has shared with other users and the last file activity.
Quick access	Provides access to the different Document libraries the user has access to. Document libraries are in two categories: "Frequent" and "Followed". From the Shared libraries view, users can create new Document libraries and/or access the SharePoint Online landing page.

Table 8-2: Information Views in OneDrive for Business

## Sharing in OneDrive for Business

Sharing is one of the key capabilities provided by OneDrive for Business driven by the same features previously described for SharePoint sites.

### OneDrive for Business Link open receipt

When a user opens a file or folder shared from OneDrive for Business, the service sends an e-mail to the sharing user to notify them that the file or folder has been opened. The notification e-mail also includes a link to file or folder permissions management to review how the file or folder was shared and even remove the sharing link.

### Customer Branding in Sharing Email

If a tenant configures Azure AD with custom branding for the organization, their [logo appears in the sharing e-mail sent](#) when a file is shared from OneDrive for Business or SharePoint Online. Information about how to apply custom branding for Azure AD is in the Identities chapter.

## Request Files

The Request files feature allows users to create a special sharing link to ask other users to upload files to a specific folder. When a user selects a OneDrive for Business folder and clicks "Request files," OneDrive generates a file request link. Anyone can use this link to upload files to the target folder, but they cannot view, edit, or see who else may have uploaded files using the link. When someone uses the Request files link, he/she is redirected to a special page where he/she can indicate the files to be uploaded together with some personal details (First Name and Last Name) to let the requestor know who uploaded them to the folder.

## Add Shortcut to OneDrive

The Add Shortcut to OneDrive feature allows users to add shortcuts to SharePoint Online and OneDrive for Business folders where they commonly access files. The folders can be:

- Folders shared with the user from OneDrive or SharePoint.
- Folders in any SharePoint Online document library the user can access.

Shortcuts appear in OneDrive's My Files view. Folders pointed to by Shortcuts can be accessed from any OneDrive App, synced across all devices, and shared like any other folder owned by the user. Owner information is visible in the Sharing column to differentiate it from the user's own content. The group name appears for content shared from sites owned by Microsoft 365 groups. Shortcuts to shared folders retain all

policy, compliance, and permissions settings from the source. Shortcuts [can be moved from the Files root in OneDrive to a public or shared folder](#).

## OneDrive Client and Synchronization

A key OneDrive for Business benefit is the ability to synchronize content stored in SharePoint Online document libraries and OneDrive for Business accounts to workstations using the OneDrive Sync Client. System requirements for the OneDrive Sync Client, including the list of supported platforms and operating systems, are in [this support article](#). In addition, [Microsoft has previews of a native OneDrive for Business Sync Client](#) for ARM devices for Windows and Mac OS and Apple Silicon devices.

The support lifecycle of the OneDrive Desktop application (sync app) follows the Windows support lifecycle, meaning that support for Windows 7 and Windows 8.1 will end on January 10, 2023.

The synchronization mechanism depends on the host operating system and application. For Windows, OneDrive uses the Windows Push Notification Service (WNS) to synchronize files. The basic approach is:

- A change occurs in a file stored in Microsoft 365. WINS notifies the OneDrive sync client of the change.
- The OneDrive sync client adds the change to its Internal Server Changes Queue. If the change involves file metadata, like renaming or deleting a file, the client processes the change immediately. Otherwise, the sync client starts a download session to update the local copy of the file.
- When a file change happens locally, WINS notifies the OneDrive sync client that it needs to update the copy in Microsoft 365.
- The client starts a session to upload the file to Microsoft 365.

The synchronization mechanism differs depending on the file type. For non-Office files, the client uses Background Intelligent Transfer File (BITS) sessions. The client transfers files smaller than 8 MB in a single HTTPS request. It divides larger files into chunks and transmits them separately to the service. Each chunk has a unique identification and separate key for protection. After the client has transmitted all the chunks, the server can reassemble the file. This mechanism allows the OneDrive sync client to process very large files up to the 250 GB current limit.

Office applications have built-in synchronization capabilities to enable features like autosave and co-authoring. The OneDrive sync client offloads synchronization processing to the relevant Office application, which sends changes to SharePoint Online or OneDrive for Business. This can only happen when the Office application is active. If it is not and it must copy an Office document, the OneDrive sync client processes it like it would any other file.

### OneDrive Sync Client and Files on Demand

The OneDrive for Business sync client (OneDrive.exe) shares a common code base with the OneDrive consumer client. It supports OneDrive for Business and SharePoint sites, including the document libraries used by Groups and Teams, and can perform read/write synchronization of document libraries with more than 20,000 files, including required metadata. The client is integrated with Windows File Explorer to allow users to choose which files have local copies and which are copied on demand (when needed). A feature called Differential Sync copies only the parts of files changed during edit sessions instead of the entire file. This feature is especially important when the OneDrive sync client processes large files up to its 250 GB limit. The sync client can pause synchronization in low-bandwidth environments (such as when connected to Wi-Fi on an airplane). The synchronization problems experienced by some users in the past are largely gone and the synchronization of documents from SharePoint Online and OneDrive for Business to and from local storage is reliable.



**Restrictions and Limitations in OneDrive for Business Synchronization:** OneDrive for Business can sync both OneDrive for Business libraries as well as SharePoint Online document libraries from team sites, but there are some [documented restrictions and limitations](#) that must be considered when synchronizing files to a local workstation. The file size limit for both SharePoint Online and OneDrive for Business is 250 GB. This limit applies to web uploads and synchronization activities.

The **Sync** section of the OneDrive for Business admin center includes the **Show the Sync button on the OneDrive site** setting to control access to the sync client for the tenant. You can hide the button completely if you do not want users to be able to synchronize OneDrive to their workstations.

Files On-Demand is a feature that allows users to control local storage for files held in OneDrive for Business or SharePoint Online sites. The feature is managed through the **Settings** tab of the OneDrive sync client, where users can set the checkbox to enable Files on Demand for their account. Afterward, they can select how to store remote files locally for the folders from OneDrive and any synchronized SharePoint site that they want to see on the device. Synchronized files can be in the following states:

- **Online-only files** are in OneDrive or SharePoint sites but appear in File Explorer to let the user know that they exist. These files occupy no disk space on the local computer.
- If users open online-only files, File Explorer downloads them from the host site and makes the files **locally available**. A locally available copy can be opened at any time, even if a network connection to a document library is unavailable. When a network connection becomes available, the OneDrive sync client synchronizes any changes made on the workstation with the host site.
- Users can also decide to mark files as **Always keep on this device**. This means that the OneDrive sync client copies these files from the host site to make them available locally even if the user never opens them.

File Explorer uses different icons to show users the status of files. Online-only files have a cloud icon, while locally-available copies have a white circle with a green checkmark, and files always present on the device have a green circle with a white checkmark (Figure 8-4).



Figure 8-4: Files On-Demand icons

When the OneDrive sync client is processing a file, the icon changes to two curved arrows in a circle. To change the synchronization status of a file, select it (or several files at one time), and use the right-click menu. You can see the icon showing the synchronization state of the files in the Status column. Files on Demand settings are unique to a device. If you use several devices, you must configure the settings on each device.

**Reset the OneDrive Sync Client:** Sometimes the OneDrive sync client can experience issues that cannot be resolved through normal troubleshooting techniques such as pausing synchronization for a short time or stopping and restarting the sync client. If necessary, you can resolve the problems by resetting OneDrive by typing the following instruction in the Windows command line:

```
%localappdata%\Microsoft\OneDrive\onedrive.exe /reset.
```

This command forces OneDrive to perform a complete synchronization of all connected personal and business sources, including the user's OneDrive for Business account and any SharePoint Online document

libraries they have opted to synchronize. After running the reset, you might need to manually restart OneDrive. As the operation resets all OneDrive settings, if the user has selected to synchronize specific folders, they must [make the folder selection](#) again.

## Files on Demand Prerequisites

Files On-Demand prerequisites for Windows 10 and Windows Server 2019 are detailed in [this support article](#). Prerequisites for Mac OS are explained in [this other article](#). Files On-Demand is part of the [Storage Sense](#) feature, so when disk space runs low, OneDrive frees up space automatically by moving back the oldest files (not marked as “always keep on device”) to a cloud-only state

**Support for UPN changes in the OneDrive sync client:** the OneDrive sync client for both Windows & Mac automatically syncs the correct OneDrive location after a user’s UPN changes. To learn how UPN changes affect OneDrive, see this [Microsoft article](#).

## Disable Windows Permission Inheritance in Folders Synced Read-Only

Through [this setting](#) in the Group Policy for OneDrive for Business, admins can instruct the sync client to manage permission inheritance for folders that are synced as read-only on a Windows Computer. When enabled, the performance of the sync client is better when synchronizing folders for which the user has only read-only permissions.

## Smart Upload Management

[Automatic bandwidth management for upload](#) controls the sync client upload rate based on bandwidth availability. In this mode, OneDrive consumes only unused bandwidth to cause no interference with any other network application. Smart Upload Management uses the Windows LEDBAT (Low Extra Delay Background Transport) protocol and is available on devices running Windows 10/11 or Windows Server 2016 (or later).

## OneDrive Sync Client Mass Delete Prompt

OneDrive monitors the removal of files from a device and prompts for confirmation if the user deletes many synchronized files on their computer at one time. If the user didn’t intend to delete so many files, they can tap **Restore Files** to recover. If not, OneDrive removes the files from the device and the cloud locations. If a user wants to skip the new behavior, they can select the Always remove checkbox to skip the prompt for future mass deletes.

## Error Resolution for Illegal File and Folder Names in OneDrive

There are situations where the OneDrive client can’t sync a file or a folder because its name contains invalid characters. To solve the problem, the OneDrive client includes a **Rename** button to automatically fix these sync problems when file and folder names meet the following criteria:

- Beginning or ending with a space.
- Beginning or ending with a period.
- Containing unsupported [Unicode](#) code points.
- Named with [surrogate](#) pair issues.

In all these situations, OneDrive replaces the illegal character with an underscore. As detailed in [this support article](#), there are still other invalid characters that cannot be renamed by the Sync client. To help users fix invalid characters, Microsoft created the [Incident deflection feature for the Sync Client](#), an in-app resolution to sync errors. For example, when the sync client meets an invalid condition like using an asterisk “\*” symbol or any file name starting with “~\$.” In these situations, the OneDrive sync client renames the file and continues with the upload.

## Request Assistance Feature

A [“Get Help” option](#) in the OneDrive activity center is available for end users so they can initiate a support ticket in case they have problems with the Sync client. Administrators can turn off this setting by running the PowerShell command:

```
[PS] C:\> Set-SPOTenantSyncClientRestriction -DisableReportProblemDialog $True
```

## Sync Machine Specific Files

The OneDrive sync client is designed to ignore machine-specific files used by File Explorer on Windows and Finder on macOS, but situations exist where these files might end up in OneDrive such as during a file migration. To prevent the visible errors displayed when those files are copied to OneDrive, the Sync client will delete the erroneous copy without touching the local copy.

## Per-machine Sync Client Installation

instead of installing the OneDrive Sync client for each user account under the “%localappdata%” folder, you can install a shared version of the sync client under the “Program Files (x86)” directory. This means that all the profiles on the computer use the same OneDrive.exe binary. This installation mode enables the following features to help in the transition from the previous generation sync client:

- Automatic transitioning from the previous OneDrive sync client (Groove.exe).
- Automatic conversion from per-user to per-machine.
- Automatic updates when a new version is available.

## Exclude Specific Files from Upload

[This setting](#) prevents the sync client from uploading specific files to OneDrive or SharePoint sites. Excluded files appear in File Explorer with a “do not enter” icon in the Status column. This setting is supported from version 20.201 of the OneDrive for Business client. It can be set in the admx/adml files. Once enabled, the sync client will not upload new files that match the specified restrictions. There is no impact on existing files uploaded to OneDrive and SharePoint.

## Controlling Sync Client Auto-pause Behavior

OneDrive users can [control the OneDrive sync client’s auto-pause behavior](#) on their PCs when they are connected to a metered network or if the device is in battery saver mode. The options to control this behavior are found in the Settings menu.

## Controlling Client Synchronization

A common issue often raised by management and those responsible for the protection of the intellectual property is the potential undesirability of allowing users to synchronize documents to their PC, possibly including a PC at home. The same synchronization issue occurs for SharePoint Online document libraries too. Users often store confidential information on home PCs (the blurring of home and work life encourages this practice). It’s also true that various blocks can prevent this behavior (such as [restricting synchronization to domain-joined PCs](#)). However, users have been swapping data between PCs ever since floppy diskettes and will likely find new methods of moving information to desired locations as quickly as IT closes off loopholes.

**Automatic Updates:** Microsoft releases updates to the new OneDrive client independently of other updates. The updates are downloaded and installed automatically on client computers unless a [group policy is implemented to control OneDrive updates](#). Like many current Microsoft products, OneDrive updates are released in “rings”. You can update a PC to receive updates faster by configuring a system registry DWORD value called **EnableTeamTier\_Internal** at HKEY\_CURRENT\_USER\Software\Microsoft\OneDrive. Set the value to 1 if you want to receive early updates (fast ring) or 0 (zero) to receive updates as normal (slow ring).

Several approaches can help control the content held in OneDrive for Business sites. First, sensitivity labels (see Information Protection chapter) can protect confidential information so that only the intended recipients can open protected files. Second, Data Loss Prevention policies can stop users from inadvertently sharing sensitive information outside the company. Third, as mentioned above, you can configure SharePoint Online so that the OneDrive sync client will only copy files to workstations that belong to specific on-premises Active Directory domains. The intention here is to stop corporate data from ending up on PCs not managed by the organization, such as home devices.

**Protecting OneDrive for Business:** A [Microsoft script](#) helps administrators to configure rights management for users' OneDrive for Business sites. The script can also reverse the process if necessary (some problems have been reported with the sync client when protection is enabled). To use the script to remove protection from sites, change the line `$list.IrmEnabled = $true` to be `$list.IrmEnabled = $false`.

As shown below, the synchronization block is set by running the `Set-SPOTenantSyncClientRestriction` cmdlet to enable checking against a set of known domains, each of which is identified by a GUID. You can use the `Get-ADForest` PowerShell cmdlet to return the set of domains in an Active Directory forest and then `Get-ADDomain` cmdlet to find the `ObjectGUID` for each domain (see [this page](#) for information how to determine the `ObjectGUID` for a domain).

```
[PS] C:\> Set-SPOTenantSyncClientRestriction -Enable
-DomainGuids 'Domain-GUID1; Domain-GUID2; Domain-GUID3'
```

When the block is in place, users see that they cannot synchronize OneDrive for Business libraries for offline access on an unmanaged PC. Administrators can use information in the audit log (See Managing Reporting and Auditing chapter) to create a report to highlight computers that OneDrive blocks from syncing files to discover whether users have tried to take files offline and then take the necessary action to tell those users why this facility is unsupported.

Although this synchronization block will not stop mobile clients from using OneDrive for Business, it does block Mac clients. For this reason, you should not implement the block if Mac clients are in use in your organization. Another restriction is to constrain users to upload files with a set of known extensions. For example, if you do not want to have users uploading executables, zipped files, and PSTs to OneDrive for Business, you can run the following cmdlets.

```
[PS] C:\> Set-SPOTenantSyncClientRestriction -ExcludedFileExtensions "exe;zip;rar;pst"
```

```
[PS] C:\> Get-SPOTenantSyncClientRestriction
```

TenantRestrictionEnabled	AllowedDomainList	BlockMacSync	ExcludedFileExtensions
False	{}	False	{exe, zip, rar, pst}

Separate each file extension with a semi-colon and do not leave any extra spaces. The `Get-SPOTenantSyncClientRestriction` cmdlet can check that the proper restriction is in place. If necessary, you can clear the set of excluded files by setting the property to `$Null`. The setting only controls the OneDrive.exe Windows client, so you can add excluded files to a OneDrive for Business site with another client before the exclusion kicks in to prevent file synchronization with the PC.

**Blocking OneDrive Consumer:** OneDrive sync client can synchronize both OneDrive for Business and OneDrive consumer accounts to a single PC. However, some organizations do not want to allow users to synchronize their OneDrive consumer accounts to corporate PCs. It's possible to block this through a Group Policy setting called `DisablePersonalSync`. See [this Microsoft page](#) for information on how to control the setting and to learn about other Group Policy settings used with OneDrive.

The last approach that you can take is to remove the temptation for users to synchronize files from certain document libraries by hiding the Sync button for those libraries. If the Sync button is not available for a

document library, users cannot start the synchronization process. To do the job, you need to write some PowerShell code that uses the SharePoint CSOM API. Fortunately, [a script exists](#) to help start the job. You can take that code and change it to meet your needs. For instance, you might use a CSV as the input to specify the target sites to disable synchronization. It is important to realize that removing the Sync button will not interfere with or stop synchronization if the user previously set it up for a document library. However, it stops new synchronization.

It might seem that controlling client synchronization is a big issue for OneDrive, but it is not really. Similar issues to those listed above often happen in how people use (or abuse) personal file shares. The best thing about ODFB is that all the features of Microsoft 365 security and data governance protect the information in user sites. In addition, because OneDrive for Business is an area of focus for Microsoft, a reasonable expectation exists that more features will become available over time.

### Prevent Users Syncing Libraries Shared from Other Organizations

[B2BSync](#) allows users to sync SharePoint and ODFB content shared with them by people from other organizations. This feature works with the external sharing integration previously described and requires recipients to have a guest user account in Azure AD. B2BSync can be disabled through the *BlockExternalSync* setting as described in the `adm\OneDrive.admx` and `OneDrive.adml` files.

### Allow Syncing OneDrive Accounts for Specific Organizations

This [setting in the OneDrive for Business policy](#) prevents users from uploading files to other organizations by creating a list of allowed organization IDs (Tenant IDs). When enabled, users see an error if they try to add an account from a prohibited organization. This setting takes priority over Block syncing OneDrive accounts for specific organizations.

### Configure OneDrive to Automatically Sync SharePoint Team Site Document Libraries

This feature allows administrators to [automatically connect and synchronize specific SharePoint team site document libraries](#) as part of a OneDrive deployment. To help configure a team site to sync automatically, a new group policy (“Configure team site libraries to sync automatically”) is available for admins to deploy.

Once deployed, the OneDrive sync client automatically syncs the contents of the shared library as online-only files the next time the user signs in, within an 8-hour window to help distribute network load. Once configured, the user cannot stop syncing the shared library unless the policy is disabled. Disabling the policy however will not unmount the shared library on behalf of the user, but the option to stop syncing the library will be available to them.

### PC Folder Backup

[PC folder backup](#) simplifies the process to move user contents located in well-known PC folders (Desktop, Documents, and Pictures folders) to OneDrive for Business. The feature was previously known as KFM (Known Folder Move). PC folder backup automatically syncs user content held in several well-known PC file locations to OneDrive for Business. To start using this feature, end users only need to click on a PC folder backup toast notification or click the **Manage backup** button under the Backup tab of OneDrive Client settings. Before the OneDrive sync client starts synchronization, it checks for the presence of any unsupported files and then begins to copy the folders from the workstation to OneDrive. Well-known folders are created in the user’s OneDrive using the same locale present in the user’s workstation. PC folder backup works on Windows 7, 8, 10, and 11 workstations.

PC folder backup also comes with the ability to silently redirect Windows known folders to OneDrive without user interaction. To enable this setting, a group policy must be enabled.

To keep admins aware of the PC folder backup deployment status, Microsoft [provides a PowerShell Script](#) to generate a report with the following information:

- PC folder backup eligibility, indicating if PC folder backup is possible for the current device.
- PC folder backup status, reports if the known folders have been moved to ODFB.
- PC folder backup GPO State: lists the GPOs in use.

The script is intended to run on a single device, but it can be modified to be run across all the devices in an organization.

## OneDrive Sync Health Reports

The [OneDrive sync health reports](#) section in the Microsoft 365 Apps Admin Center is a dashboard to check and identify OneDrive sync app versions, sync status, sync errors on individual devices, and monitor the deployment of OneDrive features such as KFM. To feed that dashboard, Admins can configure the OneDrive sync client to send health and diagnostics data by using the Group Policy Object [EnableSyncAdminReports](#). Note that [sync health reports for macOS](#) are currently in preview.

The dashboard includes a summary of how well OneDrive synchronization works within the organization in terms of how many devices have at least one sync error, the percentage of computers synchronizing known folders (KFM), and the percentage of computers with the current version of the OneDrive sync client. The Details tab provides deeper and more granular information on the OneDrive sync status for each user in the organization. Administrators can use this information to track sync errors, review error details, and do some basic troubleshooting tasks before contacting users experiencing sync problems. The Issues tab shows an aggregated view of the top sync errors happening to end users including the number of affected devices.

## Restore a OneDrive for Business Account

Ransomware has the nasty habit of infecting important documents, no matter what their location. If an attack happens against a OneDrive site, you can use **Restore Your OneDrive** (click the cogwheel to expose the option) to restore files from the document library versions to a point in time. OneDrive comes with some out-of-the-box restore points (such as a week ago), but you can select any point up to a 30-day limit. When you select a restore point, OneDrive prompts you to confirm to go ahead with the restore, and if confirmed, OneDrive restores files as they were at that point.

The consumer version of OneDrive also supports the ability to restore synchronized files to a point in time. However, the feature is only available if you have a paid subscription (Microsoft 365 Personal or Microsoft 365 Family).

# Managing OneDrive for Business

While OneDrive for Business is a core service in Microsoft 365, it no longer has a dedicated administrative console. Instead, all the required OneDrive admin features are part of the SharePoint admin center. Currently, the following OneDrive admin controls are available in the SharePoint admin center (Settings page):

- Notifications: Enable/Disable device notifications about OneDrive activities.
- Retention: Set the default OneDrive retention period for deleted users.
- Storage: Set the default OneDrive storage that applies to all the users in the tenant.
- Sync: Manage the following OneDrive sync controls:
  - Show/Hide the sync button in OneDrive and SharePoint Document libraries.
  - Enable synchronization only on devices joined to specific domains that are identified by a GUID.
  - Block the upload of specific files in OneDrive.

Admins can also manage specific settings of OneDrive sites through the user details panel in the Microsoft 365 Admin center. Specifically, the following OneDrive options can be administered through the panel:

- Storage limit: By default, this setting honors the global OneDrive storage limit value, but it can be overridden to set a specific value for the user.
- Manage external sharing: This setting updates the global OneDrive sharing setting to another sharing option.
- Create an Admin link to the user's OneDrive so the administrator is added as a OneDrive admin.

## Using PowerShell to Manage OneDrive for Business

Although Microsoft offers no specific tools to manage OneDrive for Business with PowerShell, the same principles and guidelines used for SharePoint Online apply. A SharePoint Admin can also use the SharePoint Online Management Shell, PnP Cmdlets, and SharePoint and Microsoft 365 APIs to manage OneDrive. For example, to generate a CSV file with details of the storage quota, consumption, and percentage of storage used for the OneDrive for Business sites, run the following PowerShell code:

```
[PS] C:\> $ODFBSites = Get-SPOSite -IncludePersonalSite $True -Limit All -Filter "Url -like 'my.sharepoint.com/personal/'" | Select Owner, Title, URL, StorageQuota, StorageUsageCurrent | Sort StorageUsageCurrent -Desc
$TotalODFBGBUsed = [Math]::Round(($ODFBSites.StorageUsageCurrent | Measure-Object -Sum).Sum /1024,2)
$Report = [System.Collections.Generic.List[Object]]::new()
ForEach ($Site in $ODFBSites) {
    $ReportLine = [PSCustomObject]@{
        Owner      = $Site.Title
        Email      = $Site.Owner
        URL        = $Site.URL
        QuotaGB    = [Math]::Round($Site.StorageQuota/1024,2)
        UsedGB     = [Math]::Round($Site.StorageUsageCurrent/1024,4)
        PercentUsed = [Math]::Round(($Site.StorageUsageCurrent/$Site.StorageQuota * 100),4) }
    $Report.Add($ReportLine)}
$Report | Export-CSV -NoTypeInformation c:\temp\OneDriveSiteConsumption.CSV
Write-Host "Current OneDrive for Business storage consumption is" $TotalODFBGBUsed "GB. Report is in C:\temp\OneDriveSiteConsumption.CSV"
```

Because many different types of sites exist in SharePoint Online the code needed to create a storage report for SharePoint Online sites is a little more complicated. An example script is [available in GitHub](#).

## Modern OneDrive Settings

The modern OneDrive settings page gives access to the minimum settings required by any user to manage his/her OneDrive for Business site. Compared with the classic OneDrive for Business admin options, the new settings page only includes two sections: Notifications and More settings. Notifications are all about enabling/disabling the different notifications related to OneDrive activities:

- Reminders for missed Sharing emails.
- Notifications sent when OneDrive detects massive files deletion over a short period.
- Notifications sent when others reply to user comments.
- Notifications sent when others comment on the user's documents.
- Notifications sent when the recipient clicks the link in a sharing e-mail.
- Notifications sent when others upload files to file requests.

Additional settings include links to different OneDrive configuration options such as Manage guest expiration, Site administrators, Run sharing report, Regional and language settings, Site features and storage metrics. The page also includes a link to the classic OneDrive settings page.

# Migrating to SharePoint Online and OneDrive for Business

A very common need from any organization that is on Microsoft 365 or thinking to start the Microsoft 365 journey is the migration of corporate assets (On-Premises and/or in other cloud services) to SharePoint Online and ODFB.

## Moving Windows File Servers to SharePoint

Moving individual files from a network file share to a SharePoint document library, perhaps one used by a Group or a Team can be a tiresome process. The OneDrive sync client is often used to move files between different libraries, and you can use it to move files from network shares too. Here's how:

- Synchronize the target SharePoint document library to a workstation with the OneDrive sync client.
- Map the source network file share on the workstation.
- Use File Explorer to copy files from the source network share to the folder holding the synchronized content from the SharePoint document library.
- The OneDrive sync client will then synchronize the copied files to the document library.

This is an effective approach when you have a few hundred files to move. Things become more complicated when you have more files to process. At this point you can consider:

- Run a PowerShell script to move files. Microsoft's SharePoint PowerShell module does not have cmdlets to directly copy files to SharePoint Online or OneDrive, but it provides cmdlets to migrate on-premises contents by creating a migration package that can be submitted to be [imported by the Office 365 Import service](#).
- The SharePoint [Patterns and Practice \(PnP\) module](#) includes cmdlets to copy contents to a SharePoint document library. Here is [an example of a script that uses the PnP module](#) to copy files from a network file share to a SharePoint document library.
- Use Microsoft's free SharePoint migration tool (see the section later).
- Use the Migration manager.
- Use a commercial third-party product. These products vary in terms of capability and cost, so some testing is necessary to figure out which best suits your circumstances.

## SharePoint Migration Tools

Microsoft's free [SharePoint Migration Tool](#) can move files from on-premises SharePoint Server (2010, 2013, and 2016) sites, file servers, and [Azure file shares](#) to SharePoint Online sites. The tool also supports migration to ODFB. The sole requirements are to be able to read data from the sources and write to the target sites. After checking that access is available, the tool extracts documents from the source and creates a set of XML manifests and content. The tool then uploads the content package to a secure location in Azure and invokes a migration job in SharePoint Online to fetch the content and import it into the target locations. [Version 3 of the toolset](#) includes features such as:

- Select SharePoint Online Sites, ODFB, or Teams in Microsoft Teams as possible destinations for a given migration.
- For SharePoint Server to SharePoint Online migrations, choose to migrate the site structure as it is in the source (if using a classic information architecture) or promote all the subsites to site collections and associate them with a hub site.
- Create the destination site for a migration in case it does not exist.



- JSON and CSV support for bulk migrations.
- [Migrate full out of the box SharePoint sites](#) including some site features (those included in [this list](#)), some Web Parts (those included in [this list](#)), Pages, Site Description, Site Icons, Site Navigation, Content Types, and also existing taxonomies defined in the Manage Metadata Term Store.
- Migrate SharePoint lists created using a [list template supported by the tool](#).
- (Public Preview) Migrate SharePoint Server 2010 out of the box list workflows, including Approval and Collect feedback workflows, to Power Automate Flows.
- SharePoint Server subfolder selection.
- Document sets and document templates.
- Select Teams and channels directly from the destination selection page.
- Ability to create sites during file share migration.
- Support for migrating SharePoint on-premises Record libraries.
- Support for migrating basic SharePoint on-premises settings such as site logo, title, description, and some other general settings.
- Support for migrating files up to 250 GB.
- Provide feedback to Microsoft from the tool.

See the [Migration Resources Center](#) and SharePoint Migration Tool [release notes](#) for more information about the SharePoint Migration Tool.

**PowerShell support for the SharePoint Migration Tool:** There are [eight PowerShell cmdlets](#) documented to automate migrations to SharePoint Online using the SharePoint Migration Tool. A sample migration script showing how to use the cmdlets can also be found in the cmdlet documentation

Migration utilities typically work on the basis that they move documents and metadata from a source to a target destination (a SharePoint Online site). Inevitably, some documents end up in the wrong place and need repositioning to the correct destination afterward. Until Microsoft upgraded SharePoint's "Move to" feature to allow the movement of files to a different site, rearranging documents post-migration was difficult and time consuming. Now, you can move documents to any folder in a SharePoint Online or OneDrive for Business site to which you have write access. Moving is different from copying because a move includes all the versions of a file instead of just the final version. In addition, moves preserve metadata between source and target destinations by matching column names (if a column name does not exist for the target, the move drops that metadata). Moves also respect compliance policies.

You cannot beat the zero-cost price point for the Microsoft migration tool. However, the tool offers a limited feature set and if you have more complex requirements to move documents, schema changes, metadata, and customizations, you might need a more sophisticated approach to migration such as those available from Quest, AvePoint, or Sharegate. All SharePoint migration utilities use the same APIs and are subject to being throttled. Third-party migration products usually have their own ways to increase throughput or minimize throttling to move data to SharePoint Online more quickly, and this might be a factor in your planning if you have large quantities of data to move. See [this page](#) for more information about likely throughput for SharePoint and OneDrive for Business migrations.

You should decide early on what migration tool to use for your project. Making that choice can absorb a lot of effort to analyze, test, assess, and eventually decide. Do not underestimate the importance of this work.

**On-Premises inventory/audit:** Before starting a migration to SharePoint Online, it's advisable to make an inventory/review of the contents and structures to be migrated so any issue that surfaces during the migration is found beforehand. Microsoft's [SharePoint Migration Assessment Tool](#) (SMAT) helps identify the impact of migrating from SharePoint 2010 / 2013 to SharePoint Online. As reflected in Microsoft's documentation, SMAT also includes [support for the identity migration](#) from SharePoint On-Premises to SharePoint Online. Of course, as with the SharePoint Migration Tool, there are also commercial inventory

tools that give more detailed information about a SharePoint Farm and issues that can arise when starting a migration project. The SharePoint Documentation Toolkit from SysKit and Metalogix Expert from Quest are examples of third-party inventory/auditing tools.

## Migration Manager

[Migration manager](#) includes a set of migration tools designed to simplify complex migrations where it is necessary to create migration tasks on several migration machines and orchestrate the tasks to a successful conclusion. Migration Manager is integrated into the SharePoint admin center. For the migration of file shares, the process managed by the Migration manager has four steps:

- Install the Migration manager agent in each server from where information is to be migrated to SharePoint Online sites.
- [Scan and assess the file shares](#): The Migration manager performs this step automatically after a data source is added to the service. The scans provide an overview of the contents to be migrated in terms of size, the number of files, or issues that might affect the migration. Scan log files are also available for download to enable deep analysis and troubleshooting on individual files identified.
- Create migration tasks by typing the path of the file share to be migrated and the destination location (a OneDrive for Business site, a Microsoft Teams team, or a SharePoint site).
- Monitor how migration is progressing across the different clients and access reports generated automatically by the Migration manager.

Migration Manager currently supports the migration of contents from file shares, [Box](#), [Google Workspace](#), [Dropbox](#), and [Egnyte](#) to SharePoint Online and/or OneDrive for Business.

## Microsoft Lists

Microsoft Lists is an app that allows end users to create custom lists using a dedicated space hosted in OneDrive for Business (personal lists) or shared lists stored in a SharePoint Online site. [A lightweight version of Microsoft Lists](#) (currently in preview) is also available for small businesses and individuals that don't have a Microsoft 365 but regularly use Microsoft Accounts (MSA).

Users can install the Microsoft Lists app as a Progressive Web App in Microsoft Edge, Google Chrome, or Firefox. A Microsoft List is a modern version of a traditional SharePoint list created from scratch or by using one of eight out of the box templates (including Issue tracker, Employee onboarding, and Event itinerary). Users can also create lists from an existing list or Excel worksheet in their OneDrive for Business account or from lists in SharePoint sites the user can access. Microsoft Lists incorporate features such as:

- Add new list views to show list data in different formats (List, Calendar, Gallery, Board).
- Apply rich formatting at different levels in a List without writing a line of code through column formatting, view formatting, and Forms customization. Additionally, list forms can be customized with PowerApps or the Lists Forms layout designer.
- Create simple business rules to notify when a change in a list record happens (a column changes, a column value changes to something, creation or deletion of an item, and so on). The integration of Power Automate enables the creation of more sophisticated business processes tied to a list. Business rules are supported in Microsoft Lists and regular Document libraries.
- Export list content to an Excel or CSV file.
- Mention someone in a list comment by using "@".
- Ability to easily share a personal list created by the user in OneDrive for Business with other users.
- [Create a Power BI report from list data with a single click](#). This feature requires either a Microsoft 365 E5 or Power BI Pro license.

- Work offline with list data synchronized locally using the [Project Nucleus technology](#) enabled for users through the Nucleus service plan included in all Microsoft 365 SKUs. Synchronization is on by default for eligible lists as explained in [this support article](#). Administrators can manage the sync settings for Lists through [a specific group policy](#).

Lists can be accessed through the Lists app (which presents both personal and shared lists), the SharePoint app bar, the Site contents page in a SharePoint site, the Lists app in Teams, or through the mobile Lists App available for iOS and Android. A Global or SharePoint admin can control the following [Microsoft Lists settings](#):

- Disable the creation of personal lists by running the *Set-SPOTenant* cmdlet to set the *DisablePersonalListCreation* property to *\$True*. This setting stops users from creating personal lists in OneDrive for Business.
- Disable built-in list templates that are not relevant for an organization by running the *Set-SPOTenant* cmdlet to populate the *DisableModernListTemplateIds* with the set of disabled templates. For example, to hide the Issue tracker list, run the following command:

```
[PS] C:\> Set-SPOTenant -DisableModernListTemplateIds 'C147E310-FFB3-0CDF-B9A3-F427EE0FF1CE'
```

To re-enable a built-in list template, write its GUID into the *EnableModernListTemplateIds* property.

- Disable list comments in any Microsoft Lists list by running the *Set-SPOTenant* cmdlet to set the *CommentsOnListItemsDisabled* property to *\$True*. Note that list comments can be disabled per list through the list settings page.

## Custom List Templates

In addition to the default set of list templates available for Microsoft Lists, a SharePoint administrator can add custom list templates using the following steps:

- Extract the site script from the desired list using the *Set-SPOSiteScriptFromList* cmdlet.
- Run the *Add-SPOSiteScript* cmdlet to add the site script extracted to the list of available site scripts in the tenant.
- Run the *Add-SPOListDesign* cmdlet to add the custom list template to the organization lists catalog.
- (Optional) Run the *Grant-SPOSiteDesignRights* cmdlet to indicate who can see and use the custom list template.

The following script is an example of the process:

```
[PS] C:\> $SPOListURL="https://office365itpros.sharepoint.com/Lists/Projects"
$SPOSiteScriptFromList=Get-SPOSiteScriptFromList -ListUrl $SPOListURL
Add-SPOSiteScript -Title "Corporate Projects" -Description "List that contains Corporate Projects"
-Content $SPOSiteScriptFromList
Add-SPOListDesign -Title "CompanyDevices" -SiteScripts fcbd3f82-28bb-4de0-b9bb-3e0cb6e6125e
```

After the script executes, users will see the new custom template in the *From your organization* tab in the *Create a list* dialog. To remove a custom list template, run the *Remove-SPOListDesign* cmdlet. To grant permissions on a custom list to specific people in the organization, administrators can run the [Grant-SPOSiteDesignRights](#) cmdlet to pass the identifier for the custom template and the list of people allowed to view the template.

## Formatting of Columns, List Views, and List Forms

Document library and list columns can be customized with JSON or using the Format column designer through the "Format this column" setting available for any column in modern lists and document libraries:

- Advanced mode formatting allows site designers to change the column look and feel with JSON code in different ways, as described in [this How-To article](#).

- Format column designer allows non-developers to add conditional color coding to SharePoint columns, apply color and rounded styling for each choice (Choice pills template), and Data Bars for numeric columns without needing to create JSON scripts (Figure 8-5). The Format column designer also supports defining rules that must be met to apply a color to the column.

Similarly, designers can customize list views with JSON code and/or use the Format view designer to:

- [Apply alternating row styles to the list without writing any JSON code](#). The Format view designer also supports rules creation to define the conditions to apply alternating row styles on lists and document libraries. Column and List View formatting is also supported in the [new enhanced quick edit view](#) in SharePoint lists and document libraries.
- Visually configure the fields to be displayed in cards and their order in views created with the Gallery or Board layouts. A developer can use the [Gallery/Board](#) formatter to define custom cards layouts or field values inside a card using the same JSON syntax used for column formatting.
- [Customize the list command bar](#): modify basic aspects (icon and/or text), hide and/or reposition existing options.

It's also possible to customize List forms with JSON code. The following areas in a List form can be currently customized with JSON: List form header, body, and footer.

Title	Device Type	Numer of Devic...	Purchase Date
Surface Pro 7	Tablet	30	6/15/2020 12:00 ...
Surface Book 2	Laptop	20	6/15/2020 12:00 ...
iPhone	Smart Phone	50	
iPad	Tablet	8	
Screen 15''	Other	10	
Screen 17''	Other	11	
Count		6	Sum
		6	129

Figure 8-5: Sample of modern list customized using JSON Column formatting

## Delve and Microsoft Search

Delve is an app that is heavily associated with SharePoint Online and OneDrive for Business. Apart from offering users the ability to search for documents across all sites and accounts that they have access to, Delve replaces the User Profile functionality found in SharePoint on-premises. Users can update profile settings such as the projects they have worked on, skills, and contact information (if allowed by the organization).

Delve was the first application to exploit Graph queries to display documents of interest to a user. This capability is available to Microsoft 365 applications through the Graph Insights API. Three types of insights are available: people, meetings, and items (documents). See this article for more information about how to control

[the use of insights within applications](#). Delve chapter in the companion volume contains more information about Delve.

## Microsoft Search

Microsoft Search is available throughout Microsoft 365 to perform searches in clients like OWA, OneDrive for Business, SharePoint Online, and Teams. Microsoft Search is designed to only show results relevant to the user who performs the search. Results are security trimmed, meaning that each user can only see results for documents and other objects that he/she can access. Microsoft Search does not change permissions.

Connecting Microsoft Search to services like SharePoint Online sites, OneDrive for Business, Exchange Online, Teams, and Yammer means that Bing can include content from these services in its searches. This capability known as [Microsoft Search for Bing](#) is another way to expose content stored in SharePoint Online, OneDrive for Business, and Teams to users.

Microsoft Search is enabled by default in Microsoft 365. No setup is necessary, but, as described here, it's possible to streamline and improve the overall Microsoft Search experience for users. [Management of Microsoft Search settings is through the Search and Intelligence section](#) under Settings in the Microsoft 365 admin center.

The first task is to verify that **Microsoft Search in Bing** is enabled for the tenant through the Configurations tab. You can then configure settings such as Acronyms, Bookmarks, Floor plans, Locations, and Q&A to improve the findability of content. See [this page](#) for more information.

Behind the scenes, the services connect to Microsoft Search in Bing (this process might take up to 24 hours to complete), and users see a **Work** link in the result bar to expose results from the organization when they use Bing.com to search or configure Bing as the default search engine for their browser. Clicking the link shows results from the connected sources (Figure 8-6), divided into the following sections:

- **People:** People in the tenant's directory.
- **Groups:** Microsoft 365 Groups and distribution lists related to the search term.
- **Sites:** SharePoint sites in the tenant matching the search term.
- **Files:** Documents found in SharePoint and OneDrive for Business libraries accessible to the user.
- **Messages:** [Messages from Teams chats and channel conversations and Outlook](#) (Exchange Online). Users can only search their personal mailbox. They won't find items in shared mailboxes or other user mailboxes they have delegated rights to.
- **Yammer:** Messages from Yammer communities and personal conversations (the Yammer network must run in Microsoft 365 mode).
- **Power BI:** Datasheets and other Power BI components matching the search.
- **Learning:** Content from both the organization and learning partners such as LinkedIn learning.
- Any custom vertical configured by an Admin.

Brief snippets appear for the most relevant hits in each category. Links are available to bring the user to the underlying applications to view content or perform more exhaustive searches. Users only see information that they can access, and Bing filters the results based on the permissions held by that user.

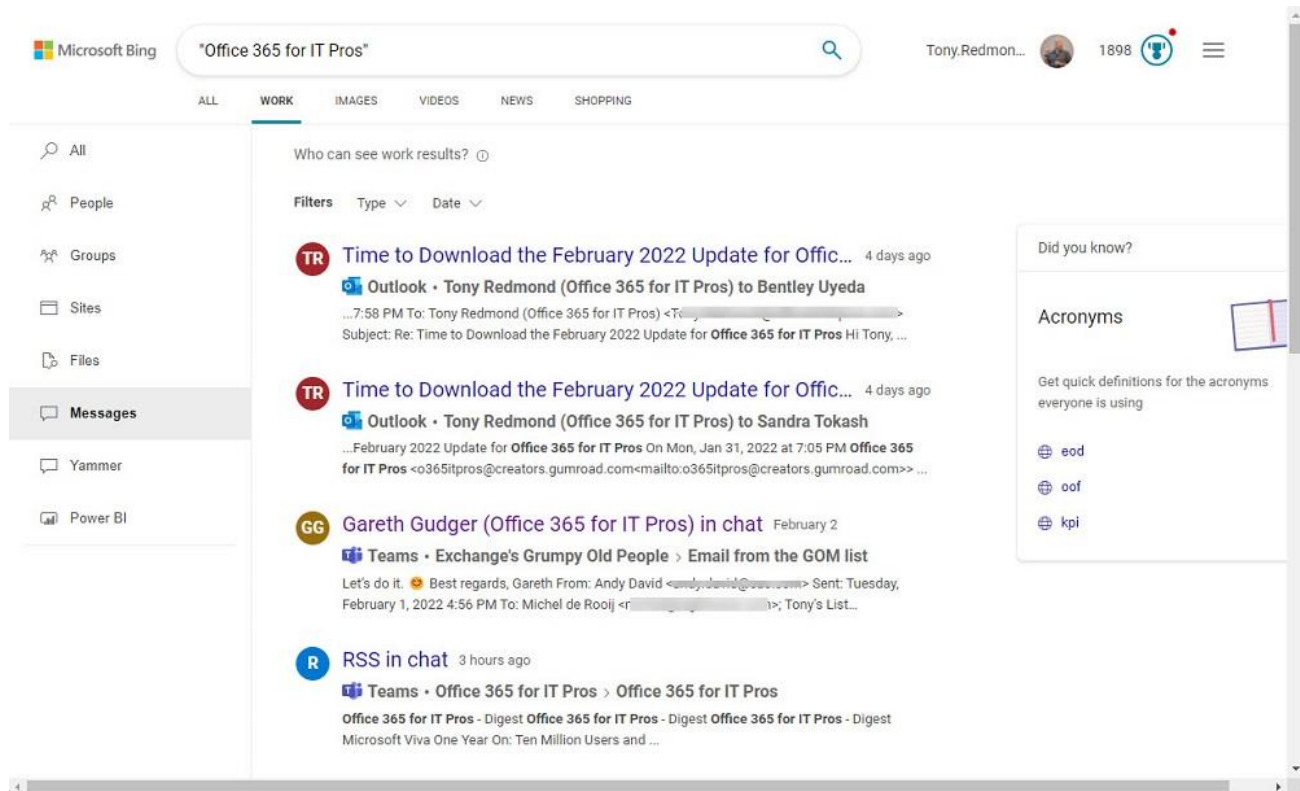


Figure 8-6: Integrating Office 365 sources into Bing search results

Integrating Microsoft 365 sources with Bing is only useful if people use Bing as their preferred search engine and is useless when they use other search engines like Google or DuckDuckGo. However, if your corporate standard is Bing, integrating Microsoft 365 workloads into Bing search results is a surprisingly useful and worthwhile extension.

Note that SharePoint search and the search feature in Office.com also present results drawn from across Microsoft 365 like conversations from Teams and Outlook. One difference is that Microsoft Search in Bing covers Yammer whereas the other searches do not.

## Admin Roles for Microsoft Search

There are two specific Admin roles to manage Microsoft Search in Microsoft 365:

- **Search admin:** This role can manage all the settings that can be configured in Microsoft Search administration including creating and managing search result content, defining query settings, and all the content-management that can be done with the Search editor role.
- **Search editor:** Create, manage, and delete editorial content in Microsoft Search such as frequently asked questions and answers, important places and locations or Floor plans.

## Acronyms

To help users find terms that might have multiple references within an organization terminology, Microsoft search provides Acronyms. An Acronym is a type of searchable term supported by Bing, Microsoft 365, and SharePoint. To get answers to an Acronym search, users must enter a specific search pattern in the Search box. For instance, if users are looking for the term DLP (Data Loss Prevention), they should use queries such as *What is DLP, Define DLP or DLP means*.

There are two types of Acronyms in Microsoft Search:

- **Admin-curated acronyms:** created by users with the Search admin role in the Microsoft Search administration.

- **System-curated acronyms:** mined by Microsoft Search from user's personal email and documents and publicly available data within the organization. Bear in mind that mined acronyms from new e-mails and documents can take up to 7 days to appear in Microsoft Search results.

Acronyms can be added to Microsoft Search in two ways:

- **Manually:** Acronyms are created by Search administrators and defined by the following fields: Acronym (Mandatory), stands for (Mandatory, spelled out abbreviation), Description (Brief description of the acronym), and Source (URL of the page or website where the users should go for more information about the acronym). Once an acronym is saved, it must be published before it can be used in Microsoft Search.
- **Automatically:** Importing a CSV file that contains all the Acronyms to be added to Microsoft Search.

## Bookmarks

Bookmarks are a way to provide shortcuts to most-visited internal or external resources in an organization or corporate applications used to complete tasks such as entering vacation time or reporting expenses. Microsoft Search includes some predefined Bookmarks and additional Bookmarks can be manually added or imported in bulk through a CSV file or from SharePoint Promoted results from sites.

A Bookmark is defined by the following elements: Title (Mandatory), URL (Mandatory, URL to an internal website or page), Bookmark description, Keywords (Mandatory, search terms commonly used to find the Bookmark), Reserved keywords, and Categories (To organize and group Bookmarks). A Bookmark can have several keywords and several bookmarks can share the same keyword. Some optional Bookmark settings can be configured such as the Dates when the result is going to be available, countries where it's going to be available, or the Groups (Security groups or Microsoft 365 Groups) that can use the Bookmark.

Existing and new Bookmarks can be edited when required. If a Bookmark is not ready to use yet, it can be saved as a draft before being published. Several Bookmarks can be imported/edited in bulk by using the Import action.

To simplify the process of creating and curating Bookmarks, Microsoft Search generates [recommended Bookmarks](#) by scanning content across SharePoint sites to identify potential bookmarks. A recommended bookmark includes the following mined information: URL, Title, Description, and Keywords. Administrators make the final decision to publish or reject the recommended bookmarks.

## Floor Plans

Floor Plans are used to find people and meeting rooms within a building. Follow [these steps](#) to enable Floor Plans in Microsoft Search:

- Get the list of codes used to identify the user's Office locations.
- Verify each user has the Office location properly informed. You can do that verification through the users' list in the Microsoft 365 Admin Center.
- Make a search to find a user and verify the Office location is appearing correctly (In case you have just updated the Office location, you may need to wait up to 72 hours for the updates to be visible in the search results).
- Prepare Floor plans files that must be in DWG format. Each DWG file should contain at least 10 rooms marked with labels.
- Add building locations to Microsoft Search through the Locations section in the Microsoft Search Administration.
- Start the process to enable Floor Plans in Microsoft Search. In Floor plans select Get started and wait until the process finishes. Bear in mind this process can take up to 48 hours to be completed.

- Follow the wizard to add Floor plans to Microsoft Search. For each floor uploaded enter the floor number represented in the Floor plan file and then the building code that can be found on user's office location.
- Publish the Floor plans so they can be available in search results when users search for a co-worker's office.

## Locations

To find addresses and locate the organization's buildings, Microsoft Search includes Locations. A Location has the following attributes: Name (Mandatory), Country or region (Mandatory), Address, Keywords (Mandatory), and Reserved keywords. Locations can be added individually or imported in bulk using the import feature. After adding or updating locations, it can take several hours for the new data to appear in search results.

## Q&As

Q&As enable [rich-text answers to user questions](#) instead of just providing links to web pages. As in the case of Bookmarks, once a Q&A is added or changed, it's immediately available for users. In case a Q&A and Bookmark share the same keyword, the Bookmark result appears first.

To manually create a Q&A, navigate to the Q&A section in the Microsoft Search Administration and click on Add. In the Add Q&A panel, you need to add the following properties: Title (Mandatory), URL (URL to an internal website or page), Answer description (Mandatory, a brief description of the answer. You can apply some formatting to the description using a rich text toolbar or using existing HTML content), Keywords (Mandatory, search terms commonly used for the Q&A) and Reserved keywords. A Q&A can have several keywords and several Q&As can share the same keyword. Finally, as happens with Bookmarks, there are some optional Q&A settings that can be configured such as the Dates when the result will be available, countries where it's going to be available or the Groups (Security groups or Microsoft 365 Groups) that can use the Q&A and targeted variations to provide different Q&A messages based on the user's device and location.

Existing and new Q&As can be edited when required. If a Q&A is not ready to use yet, it can be saved as a draft before being published. Several Q&As can be imported/edited in bulk by using the Import action.

## Microsoft Search Usage Reports

[Usage reports](#) covering Microsoft search requests over the last 7, 14, 31 days or 12 months are available to search administrators to help them understand how end users look for content across Microsoft 365 and specifically in SharePoint Online using out of the box search boxes. Microsoft Search Usage Reports also help identify possible issues faced by end users when they search Microsoft 365 sources. The following reports are currently available.

- **Query volume trend graph:** Displays the total number of search queries over the selected time frame. This report includes query volume trends and periods of high and low search activity.
- **Top queries:** Displays the most commonly used search queries. It can help understand the types of information users search for.
- **No results queries:** Shows searches that return no results.
- **Abandoned queries:** Shows popular searches where users do not click on the results or the results have a low number of click throughs. This might indicate that the search results do not include information users consider useful or interesting.

The No result queries and Abandoned queries reports give search administrators insight that can help them improve the discoverability of content and identify search queries that don't return useful information and therefore create user dissatisfaction. To improve matters, search administrators can create new answers in Microsoft Search to improve search results or ingest new content by configuring Microsoft Graph Connectors.



At the Site level, [Site admins have the same reports](#) plus an Impressions by result type report with insights about the number of results by type for the last 31 days. Over time, those reports will replace some of the search usage reports available in the SharePoint Online search service.

## Customizing the Microsoft Search User Experience

The [user experience in Microsoft Search can be customized](#) in the following ways:

- **Modifying existing verticals and creating new ones** to show results of a certain type or from certain content. There are seven default Search verticals that tenants can modify: All, People, Files, Sites, News, Images, and Power BI. Note that's not the case for the [Messages](#) and Viva Learning Search verticals. To create a new vertical, you must provide a name, a valid content source, optionally a KQL (*Keyword Query Language*) query against the content source, and also optionally custom filters. Note that a connector to a data source must be created and the content indexed (it can take up to 48 hours to index a Connector) before it is effective. Custom filters, which are based on search managed properties, can be added to new and existing verticals. Custom Search verticals are visible to users when they make searches in the [Microsoft Office Home](#), SharePoint at different levels (SharePoint Home page or Sites), and Bing. Once you add a new Search vertical to Microsoft Search, it might take some time before users can view it.
- **Adding new Result types** from content sources and customizing the layout for the data indexed using result layouts designed with the [Search Layout Designer](#). A Result type allows to show search results to be displayed in different ways through one or more conditions and result layouts to use for search results that meet the conditions defined. To display results on a Search vertical, at least one Result Type must be created. A Search vertical can have multiple Result types as a mean to distinguish different types of search results. For example, you could have a Search vertical to display ServiceNow support tickets and use different visualizations depending on the severity of the support case opened by end users. To create a new Result type, you must provide a name, a content source to be used to render search results, the rules (this is an optional setting) to be used to distinguish search results depending on the conditions met and a results layout to display the search results.
- **Creating custom result layouts** to customize search results using the Search Layout Designer. A result layout can be created from scratch or any of the existing layouts provided in the tool. Once the layout is ready, the JSON definition must be exported for use in the Result type creation process.

## Microsoft Graph Connectors

[Microsoft Graph Connectors](#) support the indexing of third-party data hosted on-premises or in public/private clouds to integrate data into Microsoft Search results. Currently, there are thirteen default Connectors available in Microsoft Search (some in preview): Azure DevOps, JIRA, ADLS Gen2, Oracle SQL database, ServiceNow, CSV, File share, Microsoft SQL Server, Enterprise websites, MediaWiki, Salesforce, Confluence, and Azure SQL. In the future, Microsoft plans to add more built-in connectors, and it is possible to [create custom connectors and add them to Microsoft Search](#). A [Microsoft Graph connectors gallery](#) is available with more than 100 connectors created by Microsoft partners.

Before creating Search verticals and Results types, a Connector to a supported data source must first be created. A Connector is created through a step-by-step process that depends on the data source being added. For instance, to create a Connector to index data for a web site, the following settings are needed:

- **Name:** the connection by providing a connection name, ID, and Description. In this step, you must acknowledge that Microsoft will index data from the data source in the Microsoft 365 tenant.
- **Connection settings:** URL of the website to be indexed, the authentication type to be used (None, Basic) and the credentials (if required) to connect the website. The URL indicated in this step is the URL that initiates the crawl and is used for authentication.
- **Add URLs to exclude:** this is an optional step that allows admins to exclude URLs from search results.

- **Schema:** Defines the schema for the results that will be indexed. In this step, a list of the source properties indexed is shown. For each property, at least one of the following attributes must be set: Queryable, Searchable, Retrievable.
- **Manage search permissions:** For this data source there is no support for ACLs (Access Control Lists), so any user in the organization will be able to make searches against it.
- **Refresh settings:** It allows to indicate the refresh schedule interval for content crawling. This data source only supports full crawls.

Once a Connector is configured, it can be saved in a Draft state or published for use in Search verticals and Result types. Bear in mind that you must wait for some time until the connector starts indexing contents.

## Microsoft Search in SharePoint Online and OneDrive for Business

[Microsoft Search is integrated with Microsoft 365 and all its core workloads](#), including SharePoint Online and OneDrive for Business. For SharePoint Online, the integration is available in two ways:

- The search box added to the Microsoft 365 navigation bar to collect search queries that lead to results tailored for the signed-in user:
  - When the user clicks on the search box, Microsoft Search lists the last applications, files, contacts, and news the user interacted with.
  - As soon as the user types a search term, Microsoft Search immediately suggests search results that meet the search criteria.
  - If the search results returned are not the ones expected by the user, he/she can always expand the results and be redirected to the search results page.
- A dedicated Microsoft Search settings page is available through the Site settings page on a given site. Here a site admin can:
  - Create custom result types based on any of the Search connectors configured in Microsoft Search administration.
  - Create custom verticals to display search results coming from the custom result types configured in the tenant.

**Limitations when configuring Microsoft Search in a Site:** Some limitations currently exist when Microsoft Search is configured for a specific site such as the inability to modify a custom result type or a search vertical once they are created.

For OneDrive for Business, users can search Files and Folders not only in OneDrive but also in any shared library they can access across the tenant. The search box in OneDrive for Business includes a dropdown to allow users to choose where they want to search.

## How Microsoft Viva and SharePoint Syntex Brings AI to SharePoint

Under the umbrella of Microsoft Viva and SharePoint Syntex, Microsoft has introduced a set of new AI Services in SharePoint Online to recognize content types, extract information and knowledge, and automatically organize it into topics. Those AI services have the following [main features](#):

- Automatically recognize topics across contents and conversations, organize related information, and add generated topics to pages curated by subject matter experts. Topic pages created are available through a knowledge center, a SharePoint site that uses a set of specific Web Parts. For every topic

discovered, the AI generates topic cards with a summary of the topic, connected people, and resources. Topics cards appear automatically across Microsoft 365 in applications like Outlook, Word, SharePoint modern pages, Teams, and results posted by Microsoft Search. This technology is available in **Microsoft Viva Topics**.

- Capture content, extract information and automatically classify content employing image and text recognition, forms processing, and machine teaching. The first technique identifies objects in scanned or uploaded images and extracts text from images and PDFs, while the second permits to establish rules to automatically detect important information in forms and extract that information as metadata. Machine teaching uses AI models to recognize information in unstructured documents such as contracts, proposals, or training materials. This technology is now available in **SharePoint Syntex**.

SharePoint Syntex and Viva Topics licenses each cost \$5/user monthly.

## SharePoint Syntex

[SharePoint Syntex](#) was the first product in the Microsoft Content Services initiative. It is available as an add-on for Office 365 and Microsoft 365 plans. SharePoint Syntex uses machine teaching and AI Models integrated into the platform to automatically capture, recognize, and extract key information to ensure that uploaded documents are properly stored and classified in SharePoint document libraries. If configured in the AI Models, Syntex can [apply retention labels](#) or [sensitivity labels](#) automatically to the documents processed.

Among the SharePoint Syntex capabilities are:

- [Prebuilt models options for invoices and receipts](#) that can be configured to detect and extract a set of standard fields present in these file types. These prebuilt models can be published globally (tenant-level) or locally (site level) so they can be discovered and used by end users.
- [Local mode creation](#): Users can create and train local understanding models in sites outside of a Content Center site. These models can then be applied to libraries in those sites.
- [Content assembly](#): Supports the conversion of a Word document into a template by defining placeholders in the document for dynamic text. This template can be used later to generate new documents by filling values in the placeholders either manually or selecting data from a SharePoint list.
- Two specific SharePoint Syntex site templates accelerators: Content center site template and Contract management accelerator.

To evaluate the capabilities of SharePoint Syntex without the need to acquire licenses or even start a trial, Microsoft provides a content center site template to evaluate how document understanding models are created, trained, and managed. Models created in the center can be trained to classify content and extract information but cannot be applied in a document library to run against uploaded files. To get advice on where SharePoint Syntex might fit in an organization, Microsoft offers a [SharePoint Syntex Assessment](#) service. The output is a Power BI report with recommendations based on factors such as the size and structure of libraries, existing use of metadata or content types, and use of retention labels. SharePoint Syntex Assessment is included as a module in the Microsoft 365 Assessment Tool.

### Advanced Search Options Powered by SharePoint Syntex

As explained in the Microsoft Search section, the default user experience when searching is limited to typing a search term and getting results meeting those criteria. When SharePoint Syntex is added to a tenant, the search box in Document libraries is enhanced with the following search filters: Keywords, File name, People, Modified date, and File type. These search filters are present in Document libraries of all SharePoint sites, but they are not available in other SharePoint search locations including the SharePoint tenant landing page.

## Microsoft Viva Topics

[Viva Topics](#) uses artificial intelligence, Graph APIs, and Microsoft Search to analyze, capture and categorize the information from SharePoint Online sites to identify and generate Topics. Originally a Microsoft Research initiative called [Project Alexandria](#), Viva Topics searches bodies of information to find material to highlight as topics in applications. A topic is a phrase or term that is significant for a given organization, has related resources such as a page or site, and provides knowledge to corporate users. While the potential exists for many types of topics, Viva Topics initially recognizes the following types: Project, Event, Organization, Location, Product, Creative work, and Field of study. Users with edit permissions in the SharePoint Syntex service can [create/merge topics from taxonomy terms](#) (through both the SharePoint admin center and the site-level term stores). This process can take up to 24 hours.

To generate topics, artificial intelligence scans the sites selected by administrators to identify a set of suggested topics. The suggested topics appear in the Viva Topics center, a special SharePoint Online site used to manage the lifecycle of topics from suggestion to publication. To publish a topic, SharePoint creates a page called a topic card. Applications can then display topic cards to users when they detect relevant content in news items, posts, and messages.

A topic page can contain the following information:

- Alternate names and acronyms for the topic.
- A short description of the topic.
- People who might be knowledgeable about the topic.
- Files, pages, and sites that are related to the topic.

Like SharePoint Syntex, Microsoft Viva Topics is available as an add-on. To start using Viva Topics, you need to [follow a step by step process](#) where you have to define the following:

- SharePoint sites are used as sources of information to generate and build topics. By default, the Viva Topics setup wizard recommends including all the sites in the tenant.
- If necessary, topics to be excluded from the Viva Topics mining process. Those topics are identified through a list of names added to a CSV file with a specific structure. For every topic in the file, the following information is required: Name, MatchType (Exact/Partial). Once the CSV file is uploaded into the backend, changes will be effective within 24 hours.
- The users can view discovered topics generated from contents and conversations they have access to. Note that to view discovered topics in the Viva Topics center, user accounts must have a Viva Topics license.
- Users who act as knowledge managers process the topics generated automatically by the AI. Knowledge managers can also create topics from scratch.
- The URL and Name for the SharePoint site to host the Viva Topics center, which is the fulcrum for topic management and publication. The Viva Topics App for Teams can also be used for topic management. Currently [deploying the app](#) requires the creation of an app manifest file with PowerShell.

After configuring Viva Topics, users with the required permissions and a Viva Topics license can start to build a **Knowledge Network** in the Viva Topics center [taking into account considerations such as](#):

- Viva Topics AI requires some time (up to 14 days), as many contributors as possible (between 50 and 100 active users), and at least a minimum of 20,000 input files to identify potential topics and generate topics for review based on evidence (alternate names, acronyms, descriptions, people, pages, or sites related to each topic). Note that Viva Topics only supports English language documents, with more languages due to be added in the coming months.

- Topics generated by the AI must be confirmed, reviewed, curated, connected with other topics, and published by knowledge managers to make the topics [available for use in different Microsoft 365 Apps and Services](#) such as Microsoft Search, SharePoint Online pages, Microsoft Teams, OWA, or Microsoft Word.
- Knowledge managers can track topics discovered by the AI using the [Topic management dashboard](#). This dashboard also provides insights into how Topics are progressing between the different statuses available (Suggested, Confirmed, Published, Removed).
- Knowledge Managers can create additional topics in the Knowledge Network to expand the network and cover information not discoverable in the source locations. From May 2022, knowledge managers can add links to external pages to topics.

**Viva Topics Insights:** This is a dashboard available for Viva Topics administrators that visualizes the count of the files processed and topics that are discovered from those files. These insights are part of the Insights tab under the Search & intelligence section in the Microsoft 365 admin center.

Each topic has its own dedicated page in the Viva Topics Center site. When a knowledge manager publishes a topic page, a background process creates a topic card from the contents of the page and makes it available to applications that support Viva Topics. Users with Viva Topics licenses can then insert topic cards by typing a hash sign (#) followed by all or part of a topic name. Viva Topics shows the set of available topics and the user chooses one to insert in the text. Licensed readers can then see the topic card when they view text in a supported application.

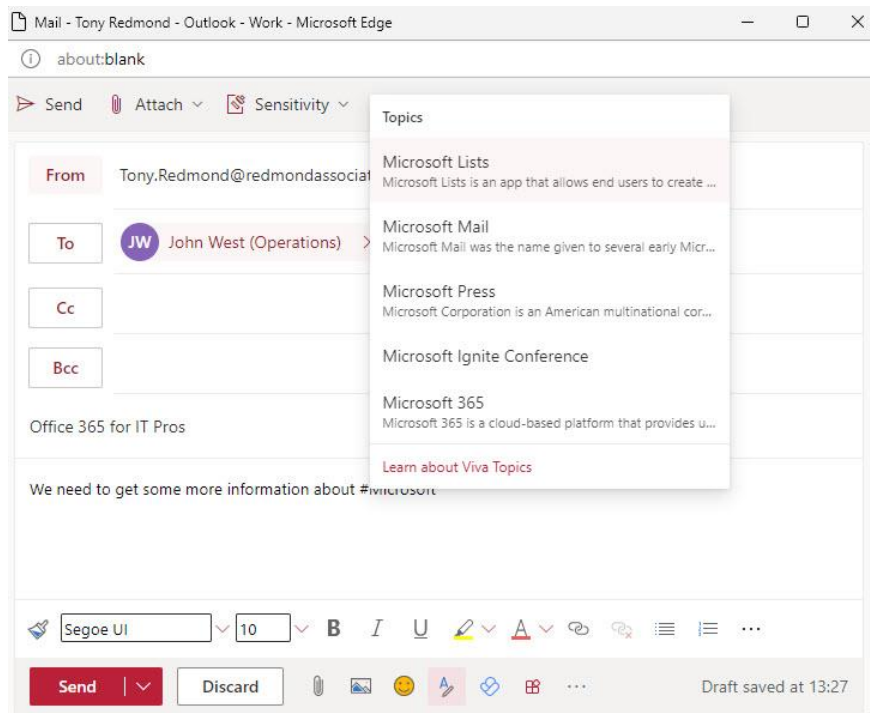


Figure 8-7: Including a topic in an OWA message

The set of applications supporting Viva Topics in text includes:

- SharePoint Online news items
- SharePoint search results.
- The Office Online apps.
- OWA (Figure 8-7).
- Teams chat.

Developing a knowledge network takes time and patience. It is not a one-off operation that can be completed and left to run on its own. To be effective, knowledge managers must monitor and manage the topics in the network and ensure that they are kept up to date.

## Microsoft Viva Connections

[Microsoft Viva Connections](#) provides native integration of a modern SharePoint site as an App in Microsoft Teams. Viva Connections enables end users to directly access and browse corporate resources published on the site directly within Teams, in other words, it brings the modern Intranet into Teams. Deploying Viva Connections requires a SharePoint Admin to [follow a step by step process](#) that involves the following actions:

- (Recommended, but not mandatory) Set up a SharePoint home site by running the *Set-SPOHomeSite* cmdlet or the Home site setting in the SharePoint admin center
- (Recommended, but not mandatory) Enable global navigation in the home site and customize any navigation link required for end users.
- (Recommended, but not mandatory) Set up the Viva Connections dashboard that can be customized with custom dashboard cards.

Once the SharePoint Admin has done his/her job, a Teams Admin is required to perform the following actions:

- Browse the Manage apps section in the Teams admin center. Look for the Viva Connections Teams App, verify it's enabled, and change (if required) the following parameters:
  - The name for the app. This name will appear in the Teams navigation bar once the app is deployed.
  - A short and long description of the app.
  - Privacy policy and terms of use.
  - Company name and company website.
  - Large and small app icons (sized at 192x192 and 28x28 pixels).
- (Optional) Manage and pin the app by default for corporate users by modifying an existing Setup policy or creating a custom one as [described here](#).

Microsoft Viva Connections does not need a specific license. Any user with a Microsoft 365 license can use Viva Connections.

## Other Services Using SharePoint Online and OneDrive for Business

SharePoint Online and OneDrive for Business are not only core workloads but also a key part of other services in combination with other services such as Exchange Online and Azure AD:

- **Microsoft 365 Groups.** Each time a Group is created, a site with a default home page with four Web Parts and a single document library is provisioned for the Group. In the same way, each time a team site is created, a Group is provisioned. Group members can easily store files, create lists, and document libraries and subsites.
- **Planner,** the task management application is deeply integrated into SharePoint Team sites through the Planner Web Part, which allows developers to [incorporate plans in pages](#) in a site.
- **Teams** uses Microsoft Groups as their membership service, which means each time a Team is created, a Group is created behind the scenes together with all its building blocks (an Azure AD Object, an Exchange Online mailbox, and a team site). Again, all the documents generated through team activity are stored in a document library. Teams also uses OneDrive for Business to store files shared in

personal chats. In addition, each time a private channel or a shared is created for a Team, a new site is created in SharePoint Online.

- **Yammer** can be configured to use native mode for Microsoft 365. In this mode, each time a Yammer community is created, a special type of group is provisioned including a team site, a Planner Plan, and a OneNote notebook. All the files uploaded to the community are stored in a document library within the site.

SharePoint Online can also take advantage of other cloud services:

- **Microsoft Power Apps** can customize list and document libraries and forms or create standalone Apps that do not require the author to write any code to [create custom forms that are mobile ready](#). The user only needs to click on the “PowerApps” action in a list to either create a custom app or customize the list forms. [A modern PowerApps Web Part](#), still in preview, supports the rendering of Power Apps in SharePoint pages.
- **Power Automate** enables users to design and deploy business processes that work with SharePoint lists and document libraries (see Power Platform chapter for more information about how to build Power Automate apps).

# Chapter 9: Tasks

**Paul Robichaux**

## Delivering Lightweight Planning

The quality of planning done for a project can spell the difference between success and failure. In the 1960s, project-planning tools were mostly restricted to extremely large engineering-focused projects such as NASA's manned lunar missions. However, as computer power and capability became more democratized, so did project planning and scheduling methodologies. Some of the early successes of third-party PC software included Borland's Sidekick utility and Lotus Agenda, both DOS-based tools including a to-do list manager. Since then, the modern business world has been inundated with tools and strategies to help people better manage their work, from David Allen's seminal book [Getting Things Done](#) to bullet journals and Kanban boards.

Microsoft's portfolio has long included Microsoft Project, a sophisticated planning product. Perhaps the kindest thing that can be said about Project is that it is easier to use than its competitors—which does *not* mean it's easy for a newcomer to manage. Every version of Outlook has included a tasks module to create and track tasks alongside email and calendar items, and since the early 2000s, Microsoft provided mobile device access to synchronize these tasks with Exchange mailboxes. Today, Microsoft 365 has a Task ecosystem where multiple applications generate tasks that share many common elements. The fact that we have Outlook, Microsoft To Do, Microsoft Planner, Microsoft Project, and various features to flag or assign items in other Office 365 applications can be confusing. It helps to reduce confusion once you know that Microsoft now thinks of tasks as falling into three categories:

- **Me tasks** belong to an individual who is usually the person who creates, manages, and views the tasks. For example, my list right now in Microsoft To Do includes "update GPS," "pay mortgage," and "reorder cat food." Usually, people manage tasks of this nature with To Do or Outlook, but you could also use tasks in apps like Planner, To Do, Outlook, or access the task objects programmatically using the Graph API.
- **We tasks** are owned, shared, or worked on by a team. While you could use a shared task list in To Do, it's much more natural to think of these tasks as those associated with plans in Planner.
- **Organizational tasks** represent shared work items defined by an organization. Think of the work a retail company must do each time they open a new store. From site selection to hiring and training to putting out signs on the sidewalk, there's a repeatable process followed each time, but the people who create the process aren't the same ones who carry it out. For example, the local store manager usually hires store staff. Teams includes the ability to publish tasks making up a work plan based on an organizational structure (the team targeting hierarchy). See [this article](#) for more information.

## Multiple Clients for Managing Your Tasks

Based on this division of task ownership, the way Microsoft has organized its task management applications makes more sense. There are a few organizing principles. First is that it's okay—in fact, desirable—for there to be multiple ways to access, edit, and update the same set of tasks. The use of a common task object across Microsoft 365 is why Tasks created in Outlook show up in the To Do app and vice versa. Second, the service should always try to collect all the tasks that "belong to" you in one place. In this case, "belong to" means both tasks that you created and those assigned to you, no matter the type. That's why the Teams Tasks app shows both your personal tasks and tasks assigned to you in Planner plans.



Another organizing principle is that Microsoft distinguishes between task management and project management. As individual people, we all have tasks to manage, but managing larger-scale projects requires different tools.

Table 9-1 shows a summary of some of Microsoft's task management ecosystem. With that summary in mind, we can dig into the individual applications in more detail.

<i>Individual ("Me")</i>	<i>Team ("We")</i>	<i>Company ("Org")</i>
Outlook Tasks	Planner (including Tasks in Teams)	Project Server / Project Online
Microsoft To Do	Microsoft Lists	Teams organizational tasks

Table 9-1: The continuum of task management from to-do items to projects

## Planner

Microsoft Planner is a lightweight task-oriented planning application whose big advantage and unique selling point is its integration with other Microsoft 365 workloads facilitated by using the membership and identity services of Microsoft 365 Groups. Planner is available for tenants with enterprise licenses in both commercial and [U.S. government clouds](#), the Business Premium and Business Essentials SKUs, and most educational and non-profit SKUs. As Microsoft's newest task management solution, Planner competes with Project Online, other Microsoft offerings like SharePoint Online, third-party offerings that plug into Microsoft 365 for task management (like [Tasks In A Box](#)), and third-party products. It's a fair question to ask what Planner's purpose is given this crowded space.

Part of the answer comes from the fact that the same engineering group develops Planner, Project, and Project Online, which means that it's no surprise that Planner and the web version of Project share some code. In the spectrum of Microsoft project management software, Planner is "Project lite," an entry-level task management application. Although Planner lacks many of the features of Project, like dependencies between tasks, its user interface offers better sorting and filtering of tasks and people often consider Planner easier to work with for many projects.

Probably the most accurate way of thinking about Planner is that it is Microsoft's version of the popular and widely used Trello application, which is itself an app that embodies [Kanban-style planning methodology](#). Trello has been around for a while and the two applications share many visual similarities in layout and visual design, albeit with slightly different naming conventions. For instance, Trello lists are equivalent to Planner buckets. See [this page for a comparison](#) between Planner and Trello.

## Microsoft To Do

Most of us try pretty hard to organize our own lives for the better. Whether electronic or paper, to-do lists often help people to prioritize and get things done. Almost every system designed to help people with office-oriented tasks has included some form of to-do list alongside email and a calendar. Microsoft bought the Wunderlist application in 2015, then promptly threw it away and built a new app on top of its cloud infrastructure, called Microsoft To Do. To Do is described as a "simple and intelligent to-do list app that empowers users to keep track of and focus on the things they need to get done."

To Do is supported for both Microsoft 365 and Microsoft consumer accounts. The To Do mobile (for iOS or Android), Windows 10, and browser clients use the Outlook Tasks Graph API to store task info in the Tasks folder of either an Outlook.com or Exchange Online mailbox. If users create a new list with the app, the to-do items for that list go into a sub-folder under Tasks named for the list. Tasks created using Outlook.com or Exchange Online synchronize with the To Do service. If you flag a message in Outlook, it shows up in To Do as a task in a separate (and optional) *Flagged Email* view. You can use To Do to edit details of the task, assign steps (individual checklist items for the task), mark it as important, set reminders, and add the task to the *My Day* view. Although it seems simple, this feature made To Do much more useful than before and convinced

some people who were dedicated users of Outlook Tasks to adopt To Do as their personal task management app. In addition, To Do tasks can be integrated with Microsoft's Cortana personal assistant, allowing you to add tasks by voice from Windows PCs, mobile devices with the Cortana app, Xbox One devices, and a variety of other voice-enabled devices. Being able to shout "Hey, Cortana, add 'buy cat food' to my to-do list" is surprisingly useful (when it works, that is).

By default, Microsoft enables To Do for users with [eligible Microsoft 365 subscriptions](#). There's no longer a way to control To Do on a tenant-wide basis, but you can control access for individual users by editing the options bundled in the Microsoft 365 license assigned to the user to remove or allow access to the application. Like many Microsoft 365 client applications, there's an "Insider" program for To Do that allows you to get early access to new features in the client and service. There are separate insider programs for Windows, [Android](#), and iOS.

## Microsoft Project

At the high end, Project is Microsoft's headline project and task management application. While To Do is intended to manage personal tasks, and Planner helps you manage tasks for your team or workgroup, Project is the kind of tool you'd use if you wanted to launch a space mission, build a hospital, or do something similarly complex. Originally launched for DOS in 1984, the capabilities of the Project application have expanded dramatically since. The latest iteration is composed of Project clients and servers, the latter being available in on-premises and [cloud](#) versions. Project is a powerful and sophisticated time and task management application that is much beloved by professional project managers. It can handle the complexities involved in very large projects such as the design and construction of major infrastructure projects. Like many highly functional software packages, learning to use Project can be overwhelming and it can take a long time to become proficient in its use. Those who have mastered Project will probably consider Planner to be simplistic and deficient in terms of charting, dependencies, and other areas. However, it is worth underlining again that Microsoft does not intend Planner to take on the kind of complex, heavy-duty planning requirements that Project is meant for. In short, Project is the right tool to use when you must coordinate multiple milestones, dependencies, and complicated schedules. If needed, you can [connect tasks managed in Project Online with Planner](#). However, you cannot transfer tasks from Planner to Project or convert a plan into a project or vice versa.

## Outlook Tasks

Outlook's Task functionality dates to its original debut back in the early 1990s. Don't expect much from Outlook as a task manager: you can create and manage lists of tasks, assign due dates, and so on. Perhaps the best task-related feature in Outlook is the ability to treat a flagged message as a task, or maybe the ability to drag a message into the Tasks folder to treat it as a task. When you absorb Microsoft's principle of "many clients, one task," it's easier to let go of the concept of using Outlook for task management instead of the easier-to-use and more-functional To Do and Planner clients.

## Microsoft Lists

At first glance, Microsoft Lists seem very similar to Planner. However, while some crossover exists, the two applications are very different.

- Planner is an out-of-the-box general-purpose task management application for teams. You don't need to customize anything to start creating and managing tasks. It's all done for you in the app. The downside is that you can't customize how Planner works.
- The Lists app is more of a toolbox to create customized applications to track and process different forms of data. You can use Lists to manage tasks, but you'll need to define what a task is and any associated components, like attachments, links, and so on.

Today, no connection exists between Planner and Lists. Microsoft has talked about how the two applications could interact in the future, but no solid plans for this are yet public.

## The Planner Back-End

Because Planner supports shared tasks and plans, there's more to it than just another front-end that reads tasks from users' Exchange Online mailboxes. Before we go any further, it's probably helpful to understand that the fundamental object type in Planner is the *plan*. Think of a plan as a container for tasks that "belong" together because they're related to the same project or owned by the same user or group.

### Planner, Microsoft 365 Groups, and Applications

Planner originally had a dependency on Microsoft 365 Groups. Every Group can have a corresponding plan; if a plan exists in a Group, the plan belongs to that Group. Originally, a Group could only have a single plan, the default plan, but this situation changed to accommodate the need for Teams to support multiple plans through channel tabs. The membership of a team depends on the underlying Microsoft 365 group. However, each channel inside a team can have multiple plans, all of which belong to the team (and the Group). One unique feature of the integration between Planner and Teams is that when you're assigned a task in Planner, the assignment [appears in the activity feed in your Teams client](#) (although thankfully you can turn this integration off).

Group-enabled SharePoint team sites can also create multiple plans, all of which are associated with the team site. You can add these plans (or the default plan for the Group) as links accessible from the home page of the team site or embed one or more plans in site pages using the Planner web part, deciding whether the web part displays boards or charts. The Planner hub does not display the name of the underlying team site that a plan belongs to, so it is wise to include the name of the site when naming a plan.

Ideally, a user needs to have an Exchange Online mailbox to enjoy the full functionality available in Planner, particularly the close connection with the underlying Microsoft 365 Group. If a hybrid user has an on-premises mailbox, they are unable to add a comment to a task using OWA or Outlook. Marking plans as favorites is also a feature that does not work properly when a user does not have a cloud mailbox. However, a hybrid user whose mailbox is on-premises can post comments to tasks through the Planner application.

### Lightweight Plans: Plans without Groups

In early 2021, Microsoft announced plans for "roster containers," a new way of creating plans without needing to use a Microsoft 365 Group to host the plan. Instead, the container has a membership roster (or list). Microsoft hasn't yet said how they will use roster containers, but it's likely to be associated with features like Teams shared channels, which use external federation rather than Azure B2B Collaboration to allow the sharing of information with people from external tenants.

The release of the roster container feature was [delayed for several months](#). Microsoft then announced plans to roll out support for "lightweight Plans" (the new name for the roster containers feature) in September 2021. For now, you'll only be able to create and work with lightweight Plans using the Graph API. Now that [Loop components are available in Teams chat](#), Microsoft will bring support for lightweight Planner plans in Teams chats, as opposed to the simple to-do list component now available. This will be a significant upgrade because the existing fFuid task component stores tasks in the Loop components file, which lives in the originator's OneDrive for Business account—meaning that these tasks break the principle of "see your tasks everywhere." We'll have to see if Microsoft delivers on this promise fully with lightweight Plans; expect to see further updates here as this feature evolves.

### Planner Services

Planner stores the metadata for plans, including information describing the tasks and buckets that make up each plan, in an Azure data service, protected with using Azure Storage Service Encryption and complying

with the Service Organization Controls Report [SOC 1](#), SOC-2, and SOC-3 standards, plus ISO 27001. [Planner is deployed in some, but not all](#), country-level Microsoft 365 data center regions, so your Planner services may come from a regional data center. For example, Planner data for tenants in Australia is hosted in Australia, but tenants in Brazil have Planner data held in the US.

Planner uses the Microsoft 365 Groups service for membership and access control. Each plan is owned by a Group and multiple plans can be created per Group (or team). The links to the Group and its resources used to store conversations and attachments are part of the plan metadata.

Planner data is accessible through the Microsoft Graph API. However, Microsoft hasn't made an API available to backup and restore Planner data, or to move plans to another Office 365 tenant. There are third-party tools available to migrate plans from one tenant to another, but they can't capture 100% of the plan's data because not all the objects and fields are available through the public Graph API endpoints.

## Planner Clients

Like many other Microsoft 365 workloads, Planner's first client was browser-based, followed by mobile apps for [iOS](#) and [Android](#) and integration with Teams as an app and channel tab. One of the chief useful features of these clients is that they can aggregate all tasks assigned to a user in a single view on the "My Tasks" page, so no matter what team or Group the original plan is in, all the tasks an individual is responsible for appear in a single place. As with Outlook, new Planner features usually show up in the web version first. Another useful feature of the iOS client is the ability to select Planner as the target to share an item (like a website, Facebook post, tweet, or another interesting snippet). You can post the item to any plan you can access.

No desktop client currently exists for Planner. Interestingly, the Planner team did start work on one, but deferred work on it in favor of providing integration between Planner and Teams. Microsoft has an iCalendar connector to link tasks assigned to a user to their Outlook calendar (see the later section). In the meantime, a third-party add-in is available in the Office Store from [iGlobe](#) that makes Planner appear as an Outlook resource (just like public folders). You can navigate between tasks and view their details. Another tool from the same company [gives a reporting capability](#) for Planner.

SharePoint Online supports a Planner web part that can be added to a page for a modern site. The web part displays tasks from a plan belonging to the underlying Group. If a plan doesn't exist, you can create it and then add it to the web part.

## Programmatic Access to Planner Data

Planner supports a [Graph API](#) endpoint to support integration with outside solutions. There's also a [PowerShell module](#) but the process for using it is complicated: you have to download the PowerShell module itself, unblock the PowerShell module (.psm1) file and an accompanying DLL, set the execution permissions for the module on the workstation you're using, and import the module. After all this work, you get a limited set of PowerShell cmdlets that only a Global admin role holder can run. That's probably why Microsoft calls the existing PowerShell module the "Planner Tenant Admin" module. However, you can still use PowerShell with Planner to do other useful things. You can:

- Use the PowerShell cmdlets for Groups to manage details of the groups used by Planner.
- Use Graph calls made through PowerShell to access Planner data.

An article describing how to use the Planner API with PowerShell to fetch data about all the plans accessible to a selected user is [posted here](#). The article explains how to create an app to interact with the Graph, assign the needed permissions to the app to access Planner data, and obtain an access token for the app to fetch data through the Graph. These are common steps used to create an app to use with the Graph.

After connecting, the code finds the set of Groups the user belongs to. Not every Group has a plan, and some Groups have multiple plans, so the code loops through each Group to see if plan data is available. Here's the PowerShell code to call the Graph to retrieve the plans for a Group:

```
[PS] C:\> $PlanURI = 'https://graph.microsoft.com/v1.0/groups/' + $Group.GroupId + '/planner/plans'
[array]$Plans = Invoke-WebRequest -Method GET -Uri $PlanURI -ContentType "application/json" -Headers
$Headers
```

If plan data is found, the code loops through each plan to find details about the plan, tasks, and buckets. Here's a snippet of those commands:

```
[PS] C:\> ForEach ($Plan in $Plans.Value) {
    $PlanId = $Plan.Id
    $PlanNumber++
    $PlanCreated = Get-Date($Plan.CreatedDateTime) -format g
    $PlanOwner = $Plan.Owner # Microsoft 365 Group
    $PlanTitle = $Plan.Title
    $BucketURI = 'https://graph.microsoft.com/v1.0/planner/plans/' + $PlanId + '/buckets/'
    $Buckets = Invoke-RestMethod -Method GET -Uri $BucketURI -ContentType "application/json" -
Headers $Headers
    $NumberBuckets = $Buckets.Value.Count
    $TasksURI = 'https://graph.microsoft.com/v1.0/planner/plans/' + $PlanId + '/tasks/'
    $Tasks = Invoke-RestMethod -Method GET -Uri $TasksURI -ContentType "application/json" -
Headers $Headers
```

After fetching the Graph data, normal PowerShell commands process the data to create a report listing the number of tasks (and the status of the tasks) and buckets in each plan, when the plan was created, and the number of days since the last task was posted as an indicator whether the plan is active.

The script isn't perfect, but it is a working example of how to combine PowerShell with the Graph to interact with plan data. An example of how to use PowerShell and the Graph to do real work with Planner is in Alexander Holmset's script to [move Planner data from one tenant to another](#).

## Planner Basics

You can access the Planner application with your normal Microsoft 365 credentials by:

- Clicking the Planner icon in the Microsoft 365 app launcher.
- Navigating to the Planner Hub at [tasks.office.com](https://tasks.office.com).
- Selecting Planner from the [...] menu of a Microsoft 365 Group when working with conversations in OWA.
- Navigating to it from a Yammer community.

The Outlook integration with Groups doesn't support Planner. We'll discuss the Teams integration later.

The Planner web application uses a streamlined interface composed of a navigation pane to the left and a details pane to the right. The user interface accommodates a range of screen sizes and form factors and resizes elements to fit the available space. Figure 9-1 shows how the Planner Hub displays Microsoft 365 Groups as Plans in the navigation pane. Plans marked as favorites show up on top. If you select a plan as a favorite, the groups in the favorites section of Outlook and OWA include the associated Group for the plan.

The "Assigned to me" link in the navigation pane displays tasks assigned to the user from all plans, while the **All** pivot lists all the Groups to which the user belongs, even if these Groups have no plan information associated with them. The **Recent** pivot exposes Groups where recent activity has occurred, but not necessarily planning activity. To work with a plan, click the name of an existing plan or create a new plan. You can force a plan to open in Teams by selecting the "..." menu next to a plan in the **Recent** or **All** pivots and selecting the **Open in Microsoft Teams** command.

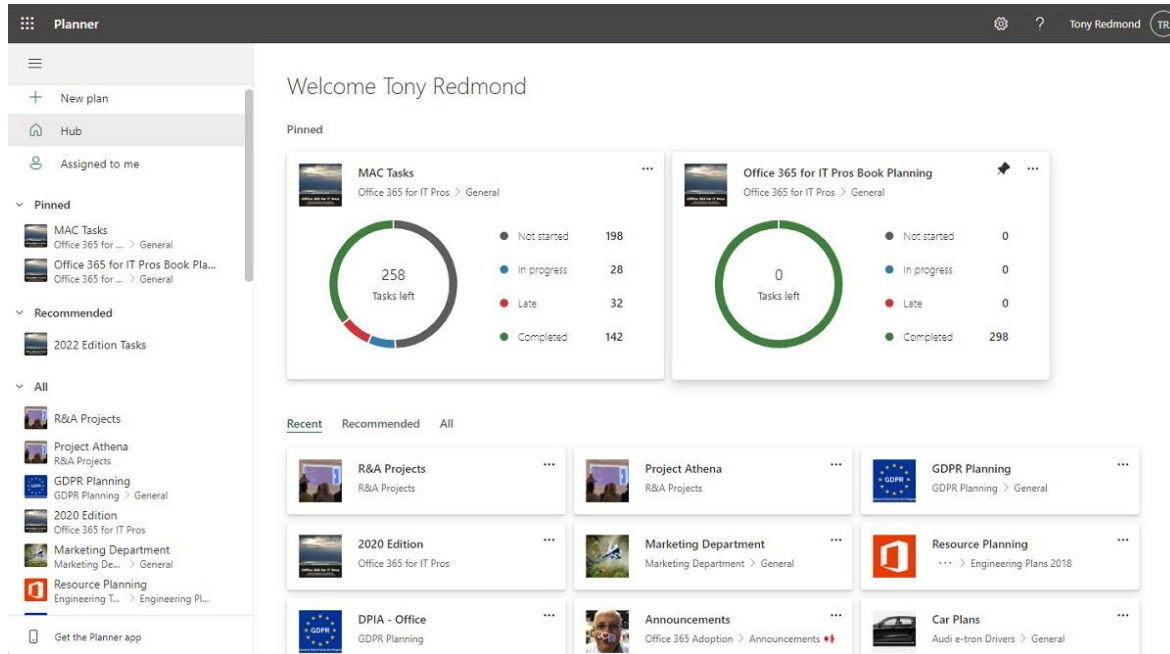


Figure 9-1: The Planner Hub

## Creating New Plans

You can create a new plan through the **New plan** command. There's no separate policy available to control plan creation, so Planner uses the creation settings from the Azure AD policy for Groups. If a user is allowed to create new Groups, she can create a plan and a new Group at the same time. If she's not allowed to create new Group objects, Planner will still allow her to create a new plan, but she must choose an existing Group to contain it. Planner doesn't give you a way to see multiple plans that belong to the same Group together; they'll still be treated separately. Plans created in an existing Group take on the properties of that Group (for example, public or private access) and are immediately available to its members.

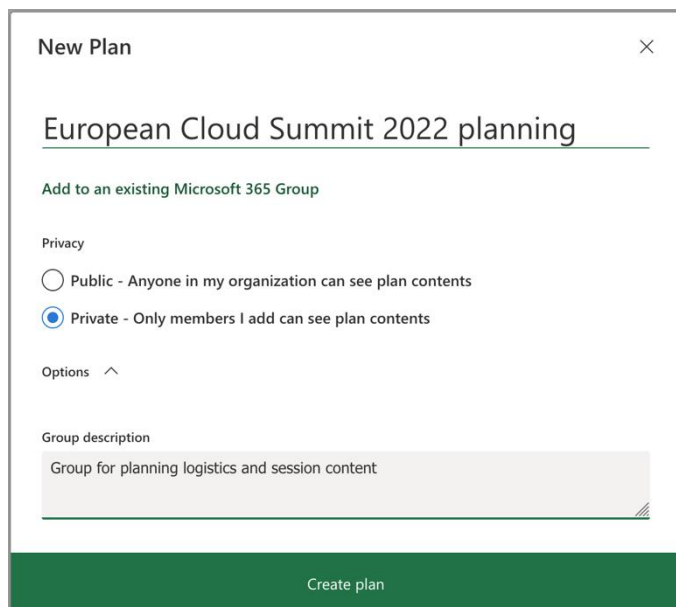


Figure 9-2: Creating a new plan

Figure 9-2 shows the interface to collect information about a new plan for a user who's allowed to create new Groups. The first choice to make is where to create the new plan, which can either be attached to an existing

Microsoft 365 Group (or team) or created from scratch. If you opt to attach the new plan to an existing Group, the next step is to select the target Group or team.

Figure 9-3 shows how the Planner hub displays multiple plans belonging to a Group called GDPR Planning. The default plan for the Group takes the name of the Group while the other plans have the Group name shown under their name.



Figure 9-3: Multiple plans for a Group

## Provisioning New Groups

If you choose to create a new plan, Azure AD creates a brand-new Group to host the plan. A Microsoft 365 Group created by Planner is the same as any other Group with a SharePoint team site and other resources. If a Group naming policy is in force, the names given to new groups or teams must comply with the policy. The exception is for plans created by tenant administrators, who do not come under the control of the naming policy.

After creating a new plan, the next thing to do is to add some members. Click Members in the menu bar to add tenant users and guest users to the membership of the Group. Remember that these people are now full members and therefore have access to the other resources belonging to the Group.

## Planner Tasks, Buckets, and Boards

Tasks form the basic building block for a plan. Each task describes a piece of work, including information such as the person assigned the work, start and end dates, and a priority for the task. Planner uses a card metaphor to display information about tasks. You can customize the cards to display different information about the tasks such as including a graphic, an expanded description, or even a checklist item.

Planner organizes tasks into “buckets” and arranges the buckets on a “board” (the analogy is to pin index cards with details of tasks to a corkboard). Think of a bucket as a collection of tasks that constitute a major section of a plan, or a collection representing tasks in a specific state. Each plan begins with a “To Do” bucket. You can rename this bucket and create as many other buckets as you need. For instance, some teams use buckets to track tasks for recurring meetings like weekly team briefs. They create a new bucket for each meeting and move tasks still outstanding from previous weeks to the new bucket, which then becomes the central focus of discussion for what needs to be done this week.

You can even use buckets to coordinate several plans within a single overarching plan. In this scenario, each plan, or sub-plan, becomes a bucket. If you remove a bucket from a plan, you also remove all the tasks in the bucket.

Tasks belong to an individual plan. You can move a task between plans in a Group and (from October 2021) to a plan owned by another Group.

Tasks are arranged for display in six ways:

- **Buckets:** This grouping allows you to view the tasks assigned to each bucket. Some people use buckets to group tasks by importance, some to divide up a project into logical pieces of work, and some to show task status by having buckets with tasks that are completed, on hold, not started, or whatever other status you would like to use.

- **Assigned to (person):** This grouping displays the tasks assigned to each person. By default, the view only shows tasks that are in progress or not started. You can add the completed tasks to the list by scrolling to the bottom of the list and clicking “Show Completed.”
- **Progress:** This view shows tasks in the three stages supported by Planner – Not Started, In Progress, and Completed. Figure 9-4 shows an example of tasks for a plan grouped by progress with the completed tasks listed in a column.
- **Labels:** This view groups tasks into the labels (or categories) that you can assign to tasks.
- **Priority:** This view creates a virtual bucket for each of the priorities you can assign (Urgent, Important, Medium, and Low).
- **Due date:** Tasks not started or in progress are arranged according to the date set for their completion as:
  - Late.
  - Today.
  - This week.
  - Next week.
  - Future.
  - No date.

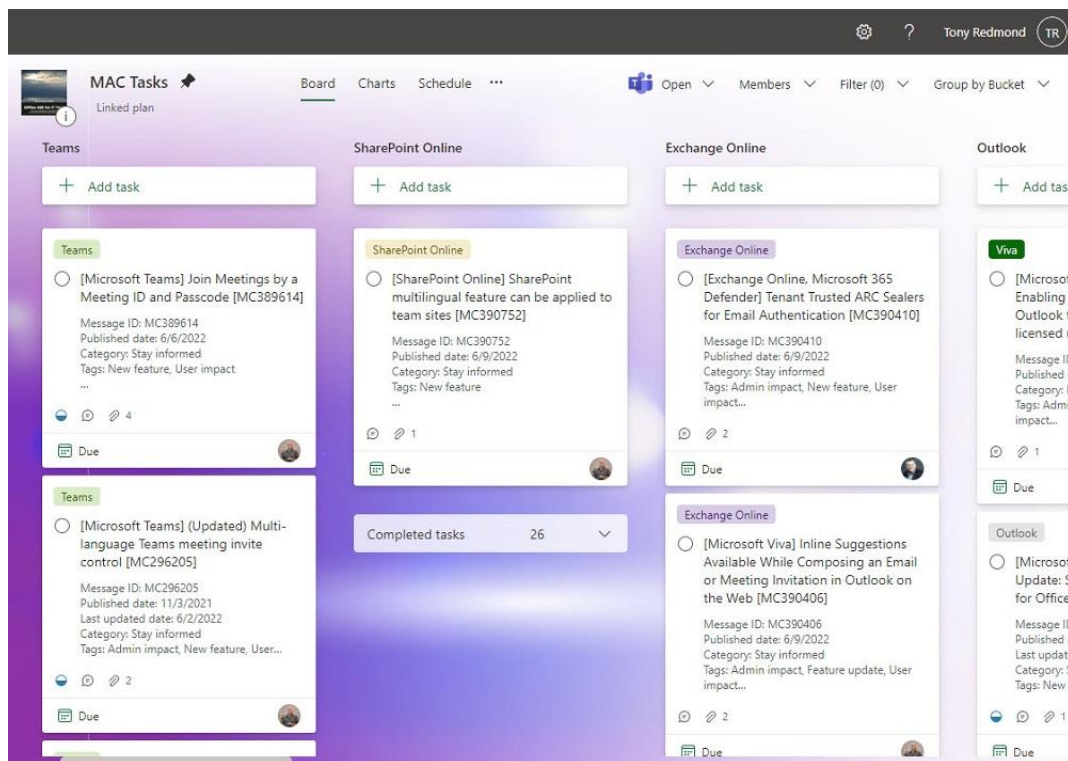


Figure 9-4: Tasks for a plan grouped and viewed by progress

Keep in mind that, when you move a task, some of its items may not move because they remain associated with the original Group—so moving a task will currently cause it to lose its labels and comments.

Planner lists tasks in a bucket in the order of creation. There is no way to select a preferred sort order such as by the due date. If you have assigned priorities to your tasks, you can use the display by priority view to see tasks grouped by the different priorities. A nice touch in the UI is that when you click a member’s icon, Planner highlights the Teams tasks assigned to that person.

Buckets are primary elements of the Planner user interface. Unless you use a very wide screen, though, it is best not to have more than five or six buckets per plan; with too many buckets, the user interface becomes



cluttered, and it can be difficult to find an individual task. If you think you will need to use more than six buckets to organize tasks, perhaps it is best to split them across multiple plans.

**Arranging buckets:** You can display the tasks for a plan in different groupings. If you use buckets to arrange tasks into segments of a plan, you can group the tasks by bucket and then drag and drop the buckets into whatever order makes sense to create an overall view of the plan. Grouping by priority is a useful feature that helps make up for the lack of a sort-by-priority option. You cannot reorder the other grouping (by progress or by assignee) in this manner.

## Task Filters

Because plans can span hundreds of tasks, Planner supports filters for views to allow users to focus on different collections of tasks. Apart from being able to group tasks in buckets, you can filter by:

- **Due date:** Select Late, Today, This week, Next week, Future, or No date.
- **Priority:** Select Urgent, Important, Medium, or Low.
- **Label:** Select one of the 25 labels assignable to tasks.
- **Bucket:** Select one of the buckets in the plan to see only tasks associated with that bucket.
- **Assignment:** Select a plan member to see the tasks assigned to them.
- **Progress:** Select to see all tasks with a specific state of progress.

As shown in Figure 9-5, you can combine one or more values for the different filters to highlight specific tasks. In this case, we want to see any task scheduled for this week in the Teams bucket.

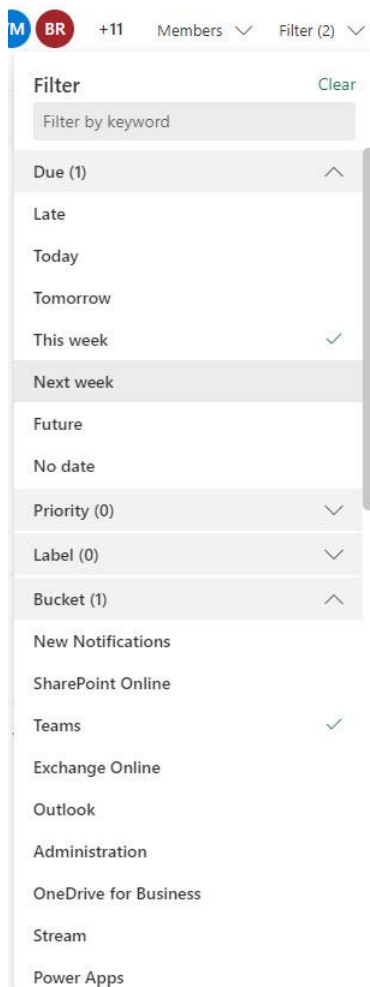


Figure 9-5: Applying filters to tasks

One common complaint is that it's difficult to search Planner items to find specific tasks. Filtering by keyword gives you similar results as being able to search the title field of your tasks, but there is currently no way to search tasks by labels or description.

## Charts

The Charts view presents another way of looking at the tasks in a plan with the screen divided into a chart view of the complete plan and a list of selected tasks on the right-hand side of the screen. In this case, tasks are listed by bucket, but you can apply any of the filters available in Planner to display matching tasks.

With just four charts available (status, bucket, priority, and members), Planner has limited ability to present graphic views of tasks. A simple embellishment would be to allow graphs to chart information for one or more selected buckets rather than a plan. Another obvious enhancement is the addition of a timeline view as people often measure the progress of plans by date. It would be very convenient to be able to drag and drop tasks along a timeline and have their due dates adjusted to match. Microsoft has committed to considering a timeline view in the future but has said that they are unlikely to add more complex or project-oriented charts such as Gantt or burndown charts to Planner as these are more suitable for a product like Project.

## Schedule View

The Schedule view (Figure 9-6) presents open tasks in a calendar, which is a natural way for many people to think about the priority order for their work. The view is like OWA's calendar and lists all open tasks in the plan unless the view is filtered (see below). Clicking a task reveals its full details.

Dates aren't assigned to tasks by default; if you don't assign correct dates, you won't get correct results.

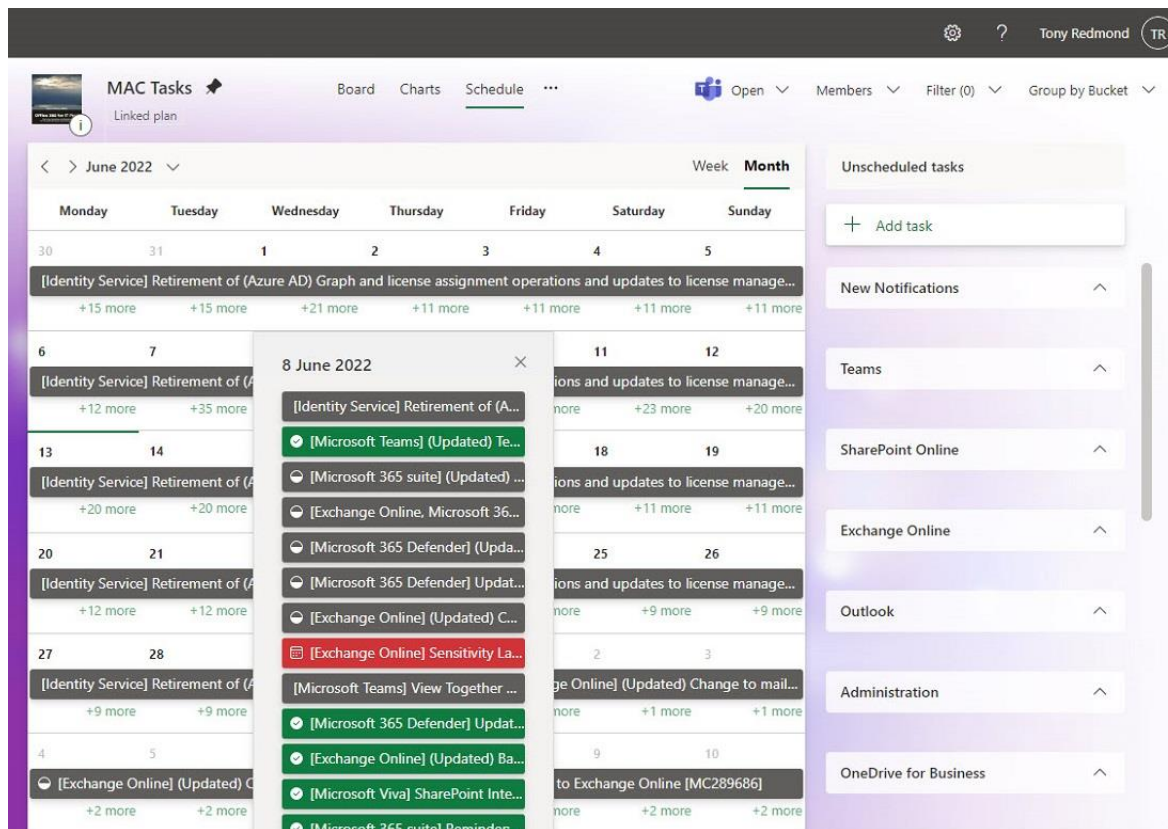


Figure 9-6: Planner's schedule view

You can add new tasks into the calendar or drag and drop existing tasks to assign them new due dates and put them into the proper position in the plan.

## Opening Linked Plans

If a plan has a label that says “Linked plan” under its name, that means that its content will be visible in other applications. For now, that means either Teams, SharePoint, or both. In Teams, if you add a plan as a channel tab, it becomes linked, and you’ll see it in the tab, in the Teams tasks app, and in Planner. For SharePoint, you can add a plan as a web part to any page, at which point it will be considered as linked.

Linked plans can be viewed and edited either in the Planner applications or in the linked applications. You’ll see a button labeled **Open** in the Planner menu bar (just to the left of the **Members** pulldown) that, when selected, lets you choose where to open the linked plan.

## Plan Settings

The options available in the ellipsis menu [...] in the menu bar for a plan are:

- Links to other Group resources. These choices open new browser tabs.
  - **Conversations:** Opens OWA to view email conversations in the Group mailbox.
  - **Members:** Opens OWA to view Group membership. Group owners can add, modify, or remove members here.
  - **Files:** Opens the document library in the SharePoint team site belonging to the Group.
  - **Notebook:** Opens the shared OneNote notebook belonging to the Group.
  - **Sites:** Opens the home page of the SharePoint team site belonging to the Group.
- **Pin:** As with other parts of Microsoft 365, Planner lets you pin items that you want to refer to often. Pinned plans show up in a separate Group in the left navigation rail. If you pin a plan, its Group also appears in the list of favorite groups in OWA or Outlook. If the plan’s already pinned, this command changes to **Unpin**.
- **Copy plan:** Creates a copy of the plan in a new Microsoft 365 Group. See details in the section below.
- **Export plan to Excel:** Creates an Excel worksheet with all the tasks in the plan. Each task occupies a single row. Once you have the plan in Excel, you can print it, combine multiple plans for analysis, or do anything else you can do with an Excel file.
- **Copy link to plan:** Copies a URL for the plan to the Windows clipboard. The link will start with *tasks.office.com*, then include the tenant name and a unique ID for the plan (e.g. <https://tasks.office.com/redmondassociates.org/en-US/Home/PlanViews/knwuDbAxtE2A7iYVAEJ-vpYAFE4L>). When pasted into a browser, it opens the plan.
- **Plan settings** are divided into three tabs:
  - **General:** Change the plan name and the background used by Planner. The backgrounds are generated automatically by the PowerPoint Designer component based on the title of the plan. Any Group member can change the plan background, and Designer does suggest different backgrounds to different people. You can’t force Planner to use a specific background like a corporate logo. Group owners can also delete the plan here (the Group remains intact when a plan is removed, and other plans attached to the Group are unaffected).
  - **Group:** (only visible to Group owners) Change the display name and description of the underlying Microsoft 365 Group. If the tenant uses sensitivity labels for container management, you can select a label to apply to the Group. A label setting dictates the privacy (public or private) of the Group. If sensitivity labels are not used, the Group owner can choose the privacy setting and classification for the Group.
  - **Notifications:** Planner sends two kinds of notifications (Figure 9-7). The first type of notification is via email when tasks are assigned or completed. Only Group owners can change this setting. Individual users can choose to receive notifications when they are

assigned a task and when a task assigned to them is late. Notifications for task assignments are generated when the assignment happens and are sent by email, a notification in the Teams Planner app, and using a push notification to the Planner mobile app. Late task notifications are generated by a background process and arrive only by email. Messages include a link to open Planner and interact with the relevant task.

- **Add plan to Outlook calendar:** Publishes the plan as an iCalendar feed that can be consumed by Outlook. This process is explained later.
- **Leave plan:** A member of a Group can leave a plan at any time. Group owners can't leave a Group if this action would leave the Group without an owner.

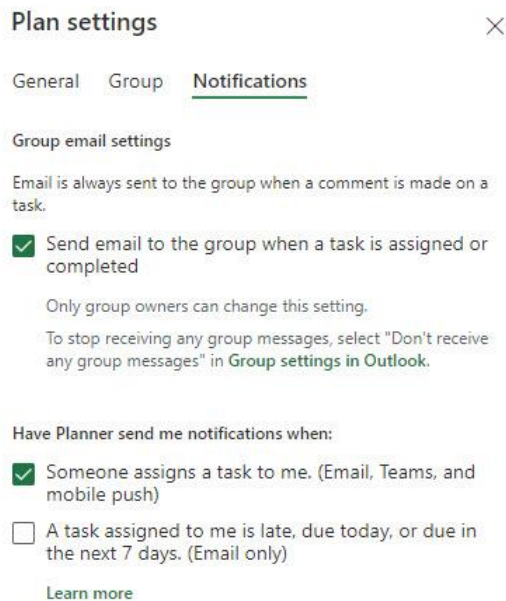


Figure 9-7: Email notification settings for a plan

Plans can span hundreds of tasks. No restrictions exist on how tasks are assigned. The tasks can be left unassigned or assigned to a single user. The largest plan I have used included over 500 tasks (you can assign up to 300 tasks to an individual user within a single plan). A user can have multiples of this number in terms of open tasks spread across different plans.

You may make mistakes when setting up a plan and want to start over. A complete bucket can be selected by clicking its name and then **Delete Bucket** from the ellipsis menu. If you want to remove the complete plan, go to the ellipsis menu for the plan, select **Plan Settings**, and then **Delete Plan** (see below).

## Private or Public

The privacy setting assigned to its underlying Microsoft 365 Group controls whether a plan is public or private. If public, any tenant user can view plan content but cannot change it. If a non-member wants to update tasks, they must join the Group (using OWA or Teams). If the plan is private, only members have access. If you change the plan's access type, you change the access type for the underlying Group and its team (if enabled).

The method used to change the privacy setting for a plan depends on if the tenant uses sensitivity labels for container management. If a tenant uses sensitivity labels for container management (see Chapter 20), choosing a new sensitivity label imposes the privacy setting inherited from the selected label. If sensitivity

labels are not used, Group owners can edit the properties of the plan or run the *Set-UnifiedGroup* cmdlet. For example, to set a Group (and its plan) to be private:

```
[PS] C:\> Set-UnifiedGroup -Identity "My Big Plan" -AccessType Private
```

Because access to a plan depends on its Group, changing a plan's access type might have some unexpected consequences. In the case of changing a plan from private to public, the intention might be to reveal details of a plan to invite comments from outside the project team on the set of tasks defined for the project. This is laudable, but perhaps the project team does not want to expose the other shared resources (documents, conversations, calendar, and notebook) to all and sundry. Some thought is necessary to figure out whether a plan should be public or private when it is created or if you edit a private plan to change its status to public.

Remember that if a sensitivity label is assigned to the Group, the privacy setting in the label controls whether the Group is private or public. In this case, you can't change the access setting without changing the sensitivity label.

## Copying a Plan

Copying an existing plan solves the problem where many similar projects exist in an organization, all of which have the same basic structure and need. The idea is that you can save some time by creating a template of a plan for these projects and then copying the plan to create a new plan as each new project spins up. You can copy a plan to an existing Group or create a new Group to host the copied plan. Although the choice to copy a plan is revealed to all users, only those allowed to create new Microsoft 365 Groups can create new plans (and Groups) in this manner.

The **Copy plan** option is available in the ellipsis menu when a plan is open or when viewed through the Planner hub. In either case, when you copy a plan, Planner does the following:

- Copies the plan data to the selected Group. By default, the display name of the new Group is "Copy of" appended to the display name of the source plan. You can override this and compose a different display name and description before you copy the plan (or afterward). The plan's photo is not copied.
- Copies the bucket structure, label assignments, and the tasks (including the chosen details of the tasks) to the new plan.
- Sets the person who copies the plan to be the owner of the new plan (and if a new Group is created, they become the owner of the new Group). Members of the source Group are not added to the new Group, so you must add new Group members after the new plan is copied.
- Sets the privacy and classification settings to the selected values.

Figure 9-8 shows the dialog used to collect information for a plan before copying. In this case, the copy creates a new Microsoft 365 group. If your tenant uses sensitivity labels for container management, the access to the plan is set by the sensitivity label you choose.

Copying a plan creates a new plan based on the source plan's structure. The following information is not copied from the source plan:

- Task assignments, due dates, and progress (all tasks are marked "Not started").
- Comments (stored as conversations in the Group mailbox).

Copying even a very complex plan is very fast. Once Planner finishes, you can open the copied plan to complete the process of establishing the new plan by adding a description and photo, removing unwanted buckets and tasks, updating task descriptions and due dates, and assigning tasks to plan members.

Copy Plan

### Copy of Microsoft 365 Message Center

Copy to Group:

+ New group  
Public

Group options

Privacy

Public - Anyone in my organization can see plan contents

Private - Only members I add can see plan contents

Classification

Confidential

Usage guidelines

Group description

Optional group description for new members

Include

Attachments

Priority

Dates

Description

Checklist

Labels

Copy plan

Figure 9-8: Defining details to create a copy of a plan

## Closing or Archiving a Plan

Apart from abandoning a plan, Planner doesn't include a way to close or archive a plan. No option exists to mark a plan as "done" or place it into a read-only mode. Planner doesn't support the Microsoft 365 information governance framework so you can't apply retention policies to manage completed tasks. If you want to clear out a plan before starting afresh (for instance, at the start of a new year), you must remove old tasks manually. Another approach is to create a new bucket called "Archive" and move all the completed tasks to it before starting work on the next stage of a plan.

Some organizations use the Export to Excel feature to create a monthly copy of plan data in a worksheet that is then stored in SharePoint Online. Although this is a manual process, it has the benefit of creating a snapshot of the plan when the export occurred.

Archival is not just a matter of simply moving tasks to some other repository because tasks often have associated documents and comments. This data is stored in SharePoint Online and Exchange Online. It is reasonably easy to copy the documents held in the SharePoint site belonging to the Group but more problematic to copy the conversations (comments) from the Group mailbox. Even if you write the necessary code to extract the necessary information from the Group mailbox, the challenge still exists to link everything together in a way that you can reliably rebuild a plan if needed.

# Planner and Microsoft 365 Compliance

Archiving inevitably brings compliance to mind. As discussed in this book, Microsoft 365 includes a comprehensive suite of data governance technology designed to work across the service. In early 2021, Microsoft announced that Planner will create “digital twins” (compliance records, sometimes called secondary copies) of tasks in Exchange Online mailboxes to ensure their availability for eDiscovery and other compliance purposes. The implementation involves:

- Planner creates compliance records for tasks assigned to a single user in the *AllToDoTasks* folder in the hidden part of their Exchange Online mailbox.
- For tasks assigned to multiple users, Planner creates a copy of the compliance record for the task in the *AllToDoTasks* folder in each user’s Exchange Online mailbox.
- Compliance records for tasks assigned to hybrid and guest users are in the special cloud-only hidden mailboxes used to hold compliance items.

It’s important to emphasize that the data held in Exchange Online is for compliance purposes only. Task data in the Planner data store in Azure remain the repository of record.

Unassigned tasks are ignored until they are assigned to a team member. Compliance records are generated when tasks are created or edited. Planner does not go back to generate compliance items for old tasks, meaning that records for these tasks do not exist unless they are updated.

Initially, the Planner compliance records will be used for eDiscovery purposes to allow Planner data to be found when content searches process user mailboxes. Because they are stored in Exchange Online mailboxes, the records are indexed and can be found by content searches. Later, Microsoft is likely to include Planner support in other compliance features such as retention policies and communication compliance policies.

In addition to the compliance records, Microsoft Search indexes the documents and comments belonging to plans stored in SharePoint Online and Exchange Online to make sure that they are discoverable. Finally, although Group creation and updates (like adding a new member to a plan) generate audit records, plan actions such as creating a new bucket or task creation, update, assignment, deletion, and completion do not.

## Deleting Plans and Plan Data

To delete a plan, choose **Plan settings** from the ellipsis menu and then **Delete this plan**. Upon confirmation, Planner removes the plan and all its tasks. Removing a plan from Planner does not remove tabs in Teams channels linked to the plan. You’ll have to do this manually.

To delete the underlying Microsoft 365 Group, choose **Plan settings**, then the **Group** pivot, and then **Delete this Group**. If you go ahead, the entire Microsoft 365 Group and all its provisioned resources are removed and put into a soft-deleted state. The resources might include a team, the Group mailbox, the SharePoint Online site, the shared notebook, and so on. If you remove a Group accidentally, you can recover it using the procedure described in Chapter 11.

When you remove a user from Azure Active Directory, their [Planner data is not deleted](#). This is intentional, as a plan will normally be a collaborative object accessed by multiple people, and it would be rude to just wipe it out when the creator’s account is removed. If you want to remove plans, tasks, or other data from Planner when a user leaves, you will have to log in with appropriate privileges and remove it manually.

# Creating and Managing Tasks

Creating a new task is very straightforward. Click **Add task** and enter details of the task such as the bucket it belongs to, the due date, and the person to whom you assign the task. You can also select an existing task and copy elements of the task (such as the description, checklist, attachments, and labels) to create a new task. To copy a task, click the ellipsis [...] menu when editing the task that you want to copy, select the elements to copy, and give the copied task a new name.

A task can also be copied to another plan (even to plans owned by other groups). The person copying the task can select any plan they belong to and choose a target bucket in that plan. When a task is copied between plans, only the progress, dates, description, and checklist are copied. You'll have to assign the task after it is copied to the target plan and add attachments and labels.

Although each task must have a name, the names do not have to be unique. You can create as many tasks as you want with the same name and assign them all to the same person, which might be a little confusing. Task names can be up to 254 characters long. If you try to use a longer name, Planner trims the name back to the limit. Special characters are fully supported and a task name such as "1"\$%^&\*()\_+{}|;:@'~#><.,.?/\\" is possible, even if it doesn't mean very much to anyone. On the upside, plenty of opportunities exists to dream up some interesting code names for tasks, especially because you can also use emoji in task titles.

## Adding Tasks from Teams

Microsoft lets you quickly add a task from within Teams, which is immensely useful as a way to quickly capture things that you need to do. The "More actions" context menu for any personal or group chat message will contain a "Create task" item; when you select it, it takes the selected message text and uses it as the title of the new task. The notes for the task will contain a link to the thread. You can also create tasks from channel conversations, in which case you can select whether the task should be stored as a personal task for you or in a plan associated with one of the teams that you're a member of. Keep in mind that an individual Team may have more than one plan associated with it.

## Assigning Tasks

You can assign a task to one or more people. Planner populates the drop-down list used for task assignment, with the Group membership divided into those assigned to the task and those who are not (yet).

When you assign a task to someone, Planner can notify them, depending on what notification settings are in effect. By default, Planner will email a notification to the Group to tell the assignee about their new task. If you use Outlook to access the Group mailbox, these notifications serve as the starting point for conversations about the task. Any comments made about the task become part of this conversation. Planner adds the comments to the conversation in the Group mailbox and circulates copies of the comments to members who have previously commented on the task.

You can also control whether Planner will attempt to send notifications of task assignments to individual users. These notifications are sent via email, push notifications to the Planner mobile app, and through Teams; this setting is at the plan level and is accessed through the plan settings interface (click the [...] menu and select **Plan settings**). You can't control which methods are used to send notifications. However, users can still turn task assignment notifications off in the Teams or Planner apps themselves, which will override the plan settings value.

If needed, you can leave tasks unassigned until the right person is available to take on a task. As already noted, if someone leaves a Group, their tasks stay assigned to them until someone reassigns the tasks to another team member. No bulk reassignment function is available to move a set of tasks to someone else



within one plan or across all plans in a tenant, such as when someone leaves the company and the need arises for others to take over this work. For this reason, it is sensible to check whether an individual has any assigned tasks when they leave the company.

Like the other applications which use Microsoft 365 Groups to manage their membership, everyone in a Group shares equal access to the plan. Anyone can create or edit a task. Anyone can assign a task to another member or reassign a task that they have received to someone else. And they can remove a task or update the progress of a task. It's all very self-liberating and empowering if you're prepared for this mode of working. At present, none of these changes are audited or tracked (although you can see which team member completed each task).

If you need a more hierarchical form of management where only defined members can add, change, or assign tasks, then you need to use traditional project management software such as Project Online. Another option is to use [the "Project" template for SharePoint Online](#) to create a site to hold tasks for a project. Out of the box, this will not deliver the same kind of user interface that Planner provides, nor will it have the same close connection to the Group, but because it is based on SharePoint Online, it is possible to customize and organize the site to meet your requirements.

## Adding Users to the Group Via Task Assignment

An owner can add any account from the tenant directory, including guest users, to the Group membership by assigning them to a task, including when a plan is accessed through a channel tab in Teams. It's logical to say that someone can't be assigned a task unless they can access the plan but the result is that the newly added member has access to all the resources available to the Group, including its document library. In other words, you cannot add someone to a plan without also adding them to Group membership. No facility exists to allow someone to have view-only access to a plan.

Planner warns owners when they add new people by assigning tasks to highlight that the user will join the Group membership. However, it's easy to overlook the warning and end up with people having access to Group resources when this is not desired.

## Fleshing Out Task Details

Newly created tasks won't usually contain much information— just enough to be able to populate the card. This might be enough for some plans but not for others. To add more details to a task, select it to reveal the task dialog box (Figure 9-9), which allows you to view and update the task metadata. Because the task shown here comes from the Message Center in the Microsoft 365 admin center, it is already populated with data, including:

- The title is set to the title of the notification posted in the Message Center.
- Progress is set to Not started.
- Priority is set to Medium.
- The start date is the date posted in the Message Center.
- Notes come from the text of the notification.
- Attachments might be available. In this case, a graphic is available. Weblinks are also commonly included.

Examples of changes you might want to make include:

- Set the expected end date for the task. As discussed earlier, the start and end dates are important if you plan to use the Schedule view to manage tasks.
- Assign the task to Group members.
- Add more attachments.
- Add checklist items.

- Use the 25 colored labels to give a visual marker to the task. To assign a label, click Add label. You can then pick a label or give a new name to a label.
- Decide what to show on the task card.
- Mark the task as complete or in progress.
- Set a different priority.

The screenshot shows a task card in Microsoft Planner. At the top, it has a title "[Microsoft Teams] Teams meetings, end of meeting notification [MC22...]" and a subtitle "Last changed moments ago by you". Below the title is the creator's name, "Tony Redmond", with a profile picture. There is an "Add label" button. The task is categorized under "Teams" in the "Bucket" field. The "Progress" is set to "In progress" and the "Priority" is "Medium". The "Start date" is "18/09/2020" and the "Due date" is "30/09/2020". There is a "Notes" section with a "Show on card" checkbox checked. The notes contain a message ID (MC222346), a published date (17/09/2020), and a paragraph of text about alerting Teams meeting participants. Below the notes is a "Checklist 0 / 2" section with a "Show on card" checkbox unchecked. The checklist items are "Blog post", "Update book", and "Add an item". At the bottom, there is an "Attachments" section with an "Add attachment" button.

Figure 9-9: Details of a task

Microsoft has previously announced, twice, that the task notes field will gain support for rich text and images, which is useful both to emphasize or highlight certain parts of a task but also when pasting notes into a task from another source. This capability depends on a new Microsoft Graph API, which is initially supported only in the Tasks app for Teams and the Planner web app. Behind the scenes, two task note fields exist: one for plain text and one for rich text. Rich text notes are converted to plain text and synced to the plain text field, but the reverse is not true; for now, if you edit the plain-text version, its changes will overwrite the contents of the rich-text version. That means if you add rich text to a task note in the Planner web client, then edit that task note on (say) the iOS version, the edits you make in iOS will overwrite the existing content and formatting in the rich text note you see in Planner. Their latest update pegs this feature for a September 2022 delivery... stay tuned.

You cannot customize or otherwise change the metadata available for tasks. In other words, you cannot add new fields for users to complete to describe or categorize a task. This is one key advantage that Microsoft Lists offers as you can customize the fields available for the items in each list.

The task name is the default information shown in the card, but if you prefer, you can display other elements in task cards. For instance, it is often easier to pick out a specific item by scanning a list for an image, so Planner allows you to use graphic files added as attachments to tasks on the task card.

Checklist items (such as those listed in Figure 9-9) allow task owners to note activities and things to do for a task. In one way, you could consider checklist items to be sub-tasks. However, checklist items have no formal connection to the outcome of a task because you can mark a task as complete even if many checklist items are unfinished. The idea here is that a checklist item is not a formal activity to complete before a task can finish. Instead, it is a free-form reminder of something that should happen, and tasks can complete with incomplete checklists. However, to encourage a sense of achievement, Planner displays a progress bar to reflect the completion of checklist items. Also, Planner displays “confetti” (an animation) when the last checklist item is complete. Unlike comments, checklist items are part of the plan metadata and Planner doesn’t generate email updates to Group members when checklists change.

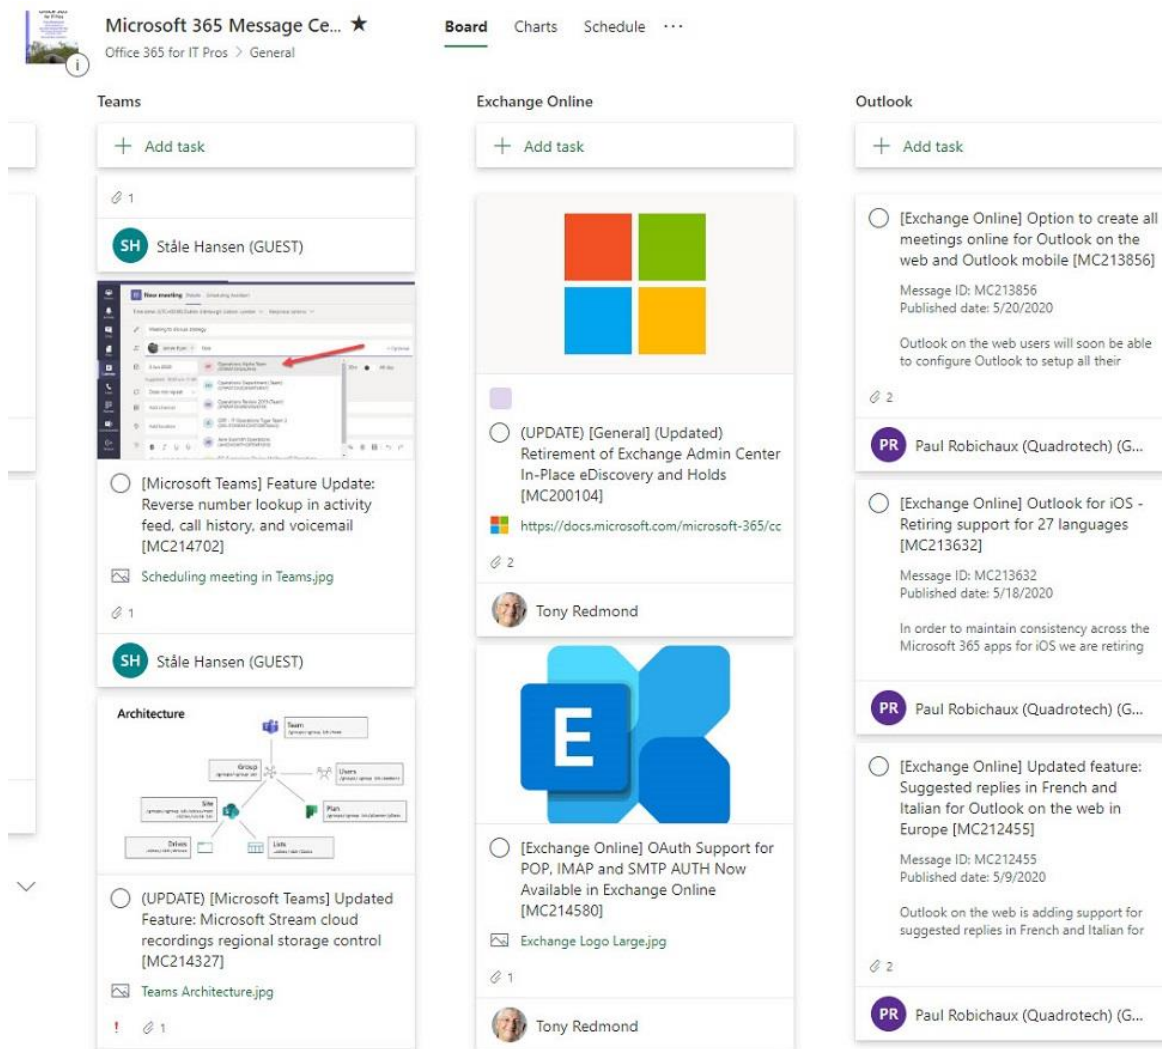


Figure 9-10: Using graphics to highlight tasks

You can add the following attachments to a task:

- A file from your computer. This includes files stored in any network location available to the PC, such as the OneDrive and SharePoint sites synchronized to the PC.

- A file from the SharePoint document library associated with the Group that owns the plan.
- A URL.

Planner stores the files uploaded as attachments for a task in the Documents folder of the Group document library (you cannot select a different folder). The same folder is used for attachments uploaded to tasks for any plan belonging to the Group. You cannot upload the same file twice as it would cause a duplicate file in the folder, but once you upload a file to the library, you can attach it to multiple tasks. If you try to create a link attachment that points to a link already attached to the task, Planner overwrites the original link. You can attach up to nine files or links to a task.

Task cards with graphic previews (Figure 9-10) take up more screen real estate than text-only cards, but sometimes you want to highlight a specific item and a graphic can be an effective way of doing this. The task card can also display checklist items or the free-form text description of the task. You can add up to 20 checklist items to a task.

The different elements of a task show up in the user interface for the task cards too. The visual hints that you can add to cards give a quick insight into the tasks within a plan, bucket, or assigned to a specific user. A big difference exists between a card that displays a simple line of text when compared to one which uses graphics. What you add to a card in terms of comments, attachments, and checklist items help plan members understand the full context of a task, including the task owner and their progress towards completion. A lot of information about a task is therefore available at a glance.

Unlike Project, Planner tasks have no dependencies on each other. You cannot link one task to another in any way except by making sure that they are arranged in display order so that one follows another as needed. This situation is problematic for some, but it reflects the general “let’s make it simple” design philosophy followed throughout Planner.

## Task Status

Managing tasks effectively requires you to update task status as the work unfolds. At any time, you can mark a task status as complete or incomplete, edit the dates associated with the tasks, or add or remove new tasks. Eventually, everything will come together and all the tasks on a board will be complete. And then, when all the tasks across all boards are complete, the plan is complete.

Planner allows a task to be in just three states – not started, in progress, and completed. There is none of the precision and exactitude of measuring a task to be say 87% complete. On the other hand, the simplicity is appealing as it is easy to know whether a task has started, is being worked on, or has finished.

## No Recycle Bin

Once you mark a task as completed, Planner draws a line through it and leaves the task in a completed state. You can delete the task if you want by selecting it and using the Delete option in the [...] menu. But be aware that once you delete a task it’s gone for good and can’t be recovered. Planner doesn’t have a recycle bin and there’s no way, even for Microsoft, to restore a deleted task.

## Restricting Task Deletion

By default, any user who has access to a plan can create, remove, or edit tasks in that plan. It would be nice if Microsoft provided more granular control over this behavior. For now, the best you can do is to block people from removing tasks they didn’t create. Blocked users can still delete tasks they created or edit any task in the plan. To do this, use the *Set-PlannerUserPolicy* PowerShell cmdlet with the *-BlockDeleteTasksNotCreatedBySelf* flag, like this:

```
[PS] C:\> Set-PlannerUserPolicy -UserAadIdOrPrincipalName andy.ruth@contoso.com  
-BlockDeleteTasksNotCreatedBySelf $true
```

Of course, you can use the same cmdlet to unblock a user and allow them to go back to deleting other people's tasks just by using `$false` as the parameter for `BlockDeleteTasksNotCreatedBySelf`.

## Labels

To highlight tasks, you can assign one or more of the 25 labels (sometimes called categories or tags) available for a plan. Colors (Blue, Purple, etc.) corresponding to the label color are used as the default names. Planner displays the assigned labels in the task card. Because you can use labels to group or filter tasks, they are a useful way to flag high-priority tasks or tasks that may need specific attention from someone.

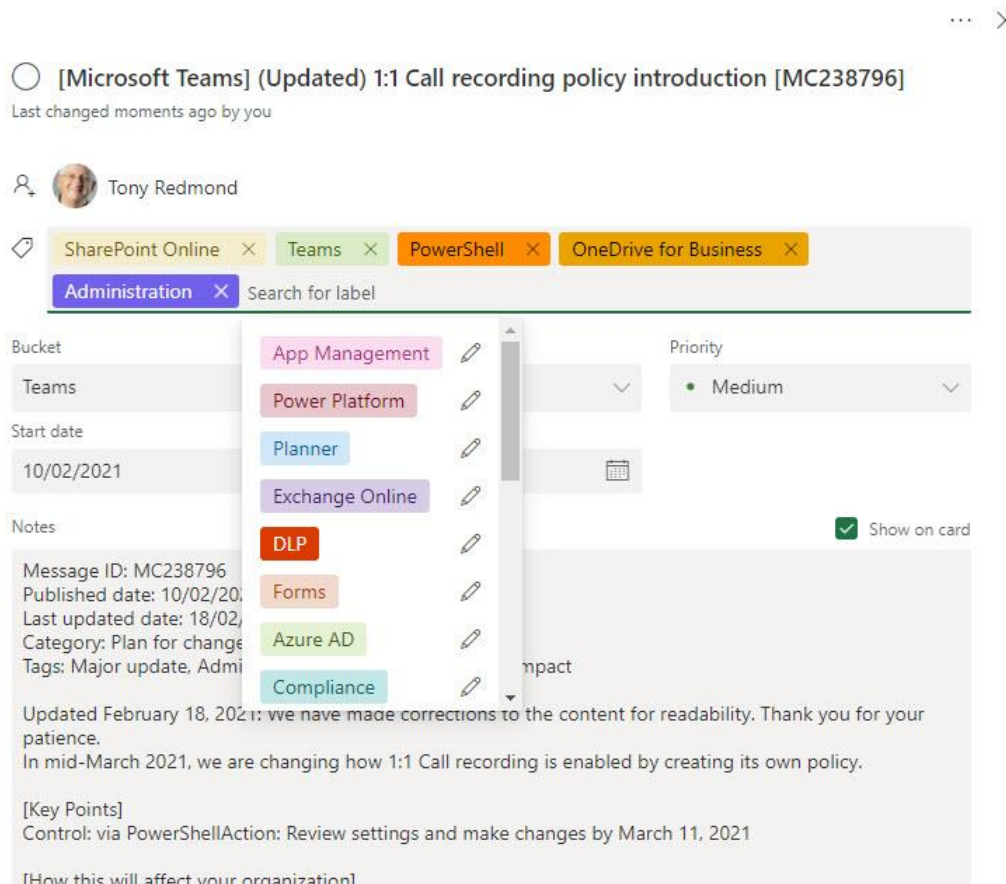


Figure 9-11: Assigning labels to a task

You can't change the colors used for labels, but you can change the names to assign names more appropriate to the plan such as Urgent, Action, or Critical with these steps:

- Access a task.
- Select Add label to expose the current set of labels.
- Select the pencil icon alongside the label you want to change and input the new name (Figure 9-11).

Labels can be renamed using the web or mobile clients by any member of a plan. This aspect of the user interface encourages freedom of operation but seems a little odd in the context of an application designed to help people organize work. For instance, if I name the red label *Important*, I am not sure that I want someone else to rename it to *Can be ignored* without asking why the label name is currently *Important*. These customizations are specific to the individual plan.

## Comments and Conversations

While the work is ongoing to complete tasks, some communication between the Group members is likely, and that's where conversations come in. Or rather, "comments" that Group members make about tasks, for this is how Planner refers to them.

As tasks are created and assigned to people, conversation items are created in the Group mailbox. People who choose to follow the Group Inbox receive these items in their mailbox, meaning that when you create a new plan for an existing Group that has a large set of existing subscribers, a certain amount of message traffic is likely to be generated as tasks are created, assigned, and perhaps reassigned. The email notifications help members keep track of task assignments as well as to know when tasks are complete without having to open the plan.

Figure 9-12 shows that a task has four comments with a new comment being prepared. The text editor used for comments is rudimentary and does not support common text formatting shortcuts such as CTRL+B or bold selected words. If you want to emphasize text in comments, you will have to use another client.

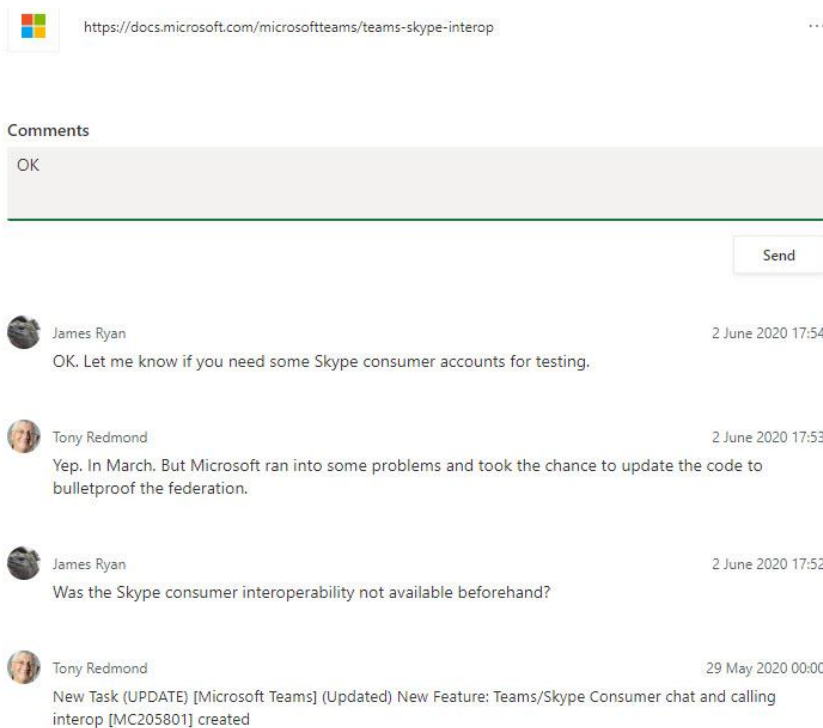


Figure 9-12: Comments for a task

The comment posted when a new task is created is one thread in the Group mailbox; the other comments form a threaded conversation for the task. Unlike task notifications, members don't receive copies of comments in email unless they have commented on a task. Members who receive comments in email can either reply directly to the message, or they can use the link embedded in the message to open Planner and go to the task and add a comment there. Posts added to the conversation through email appear as comments for the task. And, as you'd expect, if you delete a conversation, all the comments for the associated task disappear.

Behind the scenes, Planner uses the [Graph API](#) to post comments to the Group mailbox. When a user contributes to a conversation, the item is submitted through the Exchange Online transport system to apply transport rules before the comment shows up in conversations. For example, a rule might block someone from posting sensitive data such as credit card information. Another might stop the posting of any text holding an offensive term.

## Printing Planner Data

Planner includes no method to print information about a plan such as a list of tasks assigned to a user or a schedule view of tasks due in the coming week. This is a much-requested feature that is partially provided by the ability to send tasks to Outlook (see below), from where users can print the information as needed. Alternatively, you can export tasks to an Excel worksheet and format and print task information from Excel.

## Synchronizing Tasks to Outlook Calendars

Outlook calendar synchronization is automatically enabled for tenants with Planner as part of their subscription (to disable the feature, follow the [instructions in this article](#)). When a user wants to connect Planner to Outlook, they select **Assigned to me** in the navigation pane and click the ellipsis menu to reveal the choice to **Add “Assigned to me” to Outlook calendar**. Click the button and then select **Publish**. Planner generates an iCalendar link to the All Tasks view. The link looks something like this:

```
https://tasks.office.com/b662313f-14fc-43a2-9a7a-d2e27f4f3478/Calendar/User/Qi6exSZyQkC737py0An6HpYAAZ1X?t=0_95716443-a6d2-425b-a936-27fc76a889be_2018-05-10T12%3a29%3a13.4330492%2b00%3a00
```

Now click **Add to Outlook**. Planner launches OWA at the Calendar subscription window, copies the iCalendar link from Planner, and creates a default name for the new calendar (you can overwrite the name if you wish). By default, the calendar is created in the “Other calendars” section, but you can move it to one of the other sections. Click **Import** to continue. OWA creates a new calendar folder in the user’s mailbox and synchronizes details of “Not Started” and “In Progress” tasks assigned to the user. You can’t filter the tasks as the connector is configured to fetch all open tasks assigned to the user.

Synchronization is one-way from Planner to Outlook. New items do not synchronize at once because Outlook refreshes Planner data via the connector every three to four hours. You can’t force synchronization to happen. The information synchronized to the calendar for a task includes:

- **Date:** Planner items are scheduled for all-day calendar slots as Planner bases task assignments on days rather than hours. If a task has a due date, the calendar item is scheduled for that day. If it has both start and due dates, the item is scheduled for that period.
- **Location:** None added as Planner does not capture this data.
- **Progress:** The status of the task in Planner.
- **Checklist:** A note about how many checklist items exist and are complete.

Calendar entries created through the connector do not tell you what plan or bucket within a plan a task belongs to, but each entry has a link to Planner to bring the user back to the original task, where whatever changes are necessary can be made. Changes are then synchronized back to Outlook. Although the user cannot edit details of the planner task, they can add a reminder to a task.

If you prefer not to use the iCalendar connector, you can also link Outlook to Planner using [several standard Power Automate templates](#) published by Microsoft. Power Automate is good at inserting items into a calendar. It is less successful at tracking changes made to tasks such as completing tasks or synchronizing changes made to a task like changing its name, due date, or status. This is probably because the [set of triggers](#) supported by Planner for the connector available to Power Automate only covers task creation, assignment, and completion, and not task updates.

## Working with Planner Tasks in To Do

To Do clients support visibility of assigned tasks (from all plans) in To Do clients where the tasks appear in the *Assigned to Me* list. The bi-directional synchronization supports some actions against Planner tasks performed in To Do. Users can:

- Update the completion date for the task.
- Mark a task as completed or incomplete (but not mark a task as in progress).
- Add checklist items.
- Update the task description.
- Hide completed tasks.

You must enable this integration from the To Do client; there's currently no way to do so from within the Planner interface. Open the settings page for your To Do client (the exact location will depend on whether you're using Windows, the web, or macOS), then find the "Connected apps" section, then toggle the "Assigned to you" control as shown in Figure 9-13.

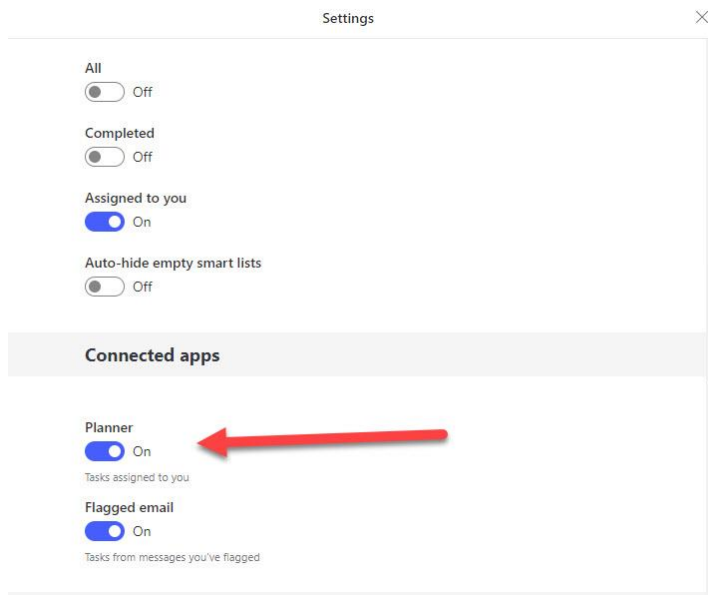


Figure 9-13: Enabling Planner integration with To Do

You can't change assignees, add attachments, or change the tabs for tasks. You also can't create new tasks. However, an Open in Planner link is available to open the task in Planner if a user needs to perform one of these actions.

## Moving Planner Data to a Different Tenant

Many Microsoft 365 workloads have formally defined processes for moving data from one tenant to another. APIs exist to import and export data from some workloads, and both Microsoft and clever third-party vendors have used these to provide a way to move workload data between tenants. Planner data can be moved between tenants, but it requires that you open a support request with Microsoft to accomplish the move. Before they will execute a requested move, you [must authorize it](#) using the rather scary-sounding `AllowTenantMoveWithDataLoss` flag with the `Set-PlannerConfiguration` cmdlet. No plans or task data are moved when Microsoft runs the move operation, so it's not clear why you'd ask Microsoft to do such a move instead of using a tool such as [MVP Sean McAvinue's PowerShell script](#).

## Using Planner Offline

The Planner iOS client can display tasks when offline but cannot update anything unless the client is connected. However, the Planner browser app includes more extensive offline capability. While you can't connect to Planner without a network connection, if the connection becomes unavailable during a session, you can continue working with tasks. When offline, you can:



- Add a new task to a plan.
- Update task properties like progress (not started, completed), start and due dates, notes, comments, and priority.
- Add or remove a web link or file attachment.
- Add or remove checklist items for a task.
- Move tasks between buckets.
- Assign tasks to team members.

You cannot:

- Start Planner when offline.
- Add a SharePoint item to a task.
- Create a new plan for an existing Group or with a new Group.

Using the cached data, Planner charts and the schedule view are both available when offline. When the network link is restored, Planner synchronizes any changes made offline to the server.

## Planner and Guest Users

Planner supports the Azure B2B Collaboration model for guest access. As explained in Chapter 11, to join a Group, external users receive and redeem an invitation. During the redemption process, if the external user doesn't already have a guest account, Azure AD creates one for them in the tenant directory. The same membership as used for Microsoft 365 Groups and Teams controls access to Planner, and a guest can be added to a Group through Outlook or Teams. Once a guest user becomes a member, their membership automatically allows access to the plans associated with the Group, including any plans created in channels belonging to a team. If a guest user already exists in the tenant directory, an owner can add them to a plan by assigning them a task.

To access Planner, guest users must specify the service domain of the tenant to which they want to connect. Planner can't switch guests between tenants, so the connection is always to a specific tenant. For example, to access plans in the `office365itpros.com` tenant, a guest connects to the URL `https://tasks.office.com/office365itpros.onmicrosoft.com`. If you're already signed into Planner in your browser, you should use a private browser session to connect to Planner in another tenant. After successful authentication, the Planner browser client displays the set of plans available to the guest. These plans include those created before Planner supported external access.

Planner displays a restricted user interface to guest users by removing choices they cannot use (Figure 9-14). Guests can create, assign, and edit tasks, change task status, post comments, add checklist items and even update a plan's background. However, guests cannot create new plans because they cannot create new Microsoft 365 Groups. Likewise, they cannot remove a plan. Guests can edit a plan's name, but they cannot change other plan settings, like its privacy level (which a sensitivity label might control). Finally, guests cannot browse the host tenant to look for plans (and Groups) to join.

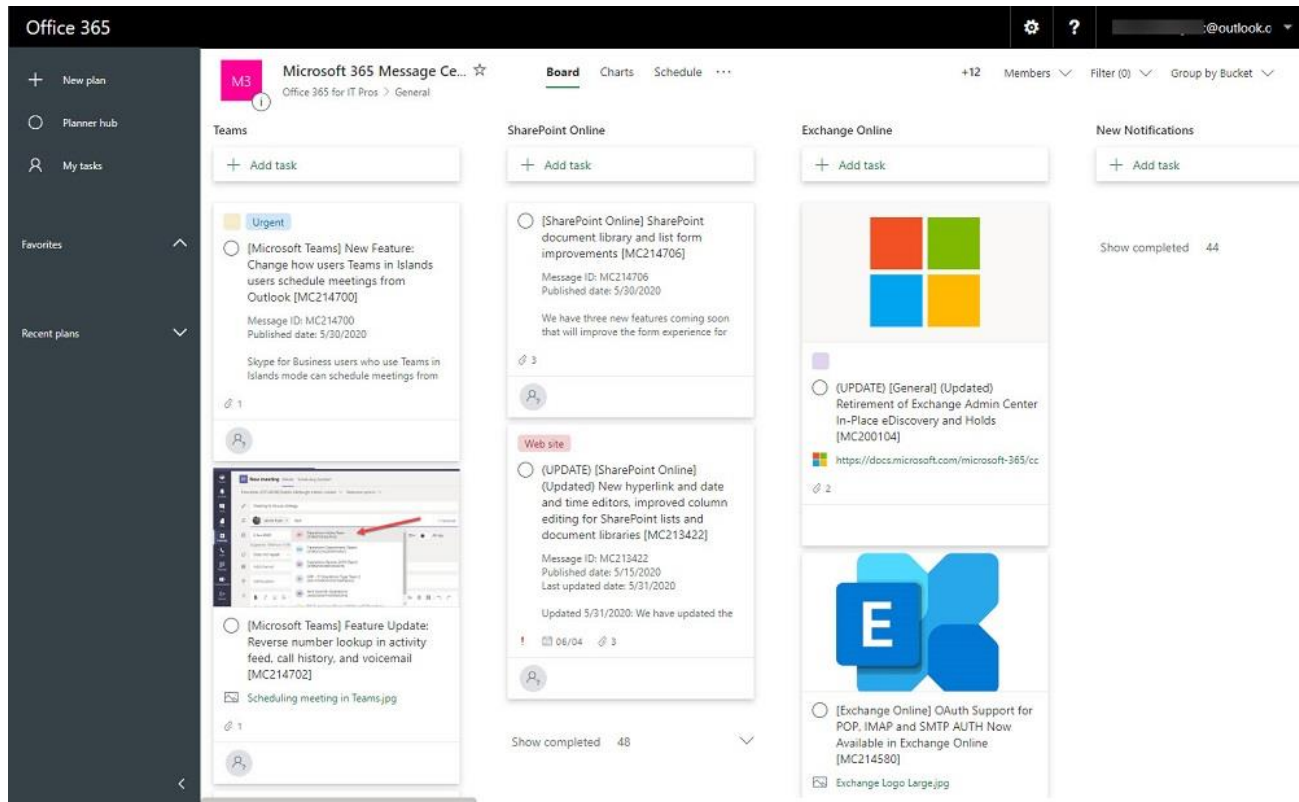


Figure 9-14: The Planner interface for a guest user

## Planner, Tasks, and Teams

Teams can be used with Planner in most enterprise, government, and education plans. One point to consider during any implementation that includes frontline workers is that Teams is available for frontline (F1) licenses while Planner is not.

Teams and Planner share two points of integration. First, a user can open the Tasks app to access all their tasks, including personal tasks created in Outlook or To and tasks assigned to them through their membership of a team or Group. The Tasks app replaces the earlier Planner app for Teams and continues the [overall Microsoft strategy](#) of providing multiple endpoints for creating and managing tasks that all end up working against the same data store.

The second point of integration is when Planner is added through a channel tab to access the tasks in a plan belonging to a team. We'll discuss this method shortly.

### Tasks in Teams

The Tasks in Teams app delivers an integrated view of personal and team tasks. The app is called *Tasks by Planner and To Do* for now, but Microsoft plans to rename it to simply *Tasks* at some unstated future date.

As shown in Figure 9-15, the Tasks app divides tasks into:

- **My tasks:** these items are personal (not shared with anyone else) and include tasks created in the app, To Do, and Outlook. These tasks are managed by To Do.
- **Shared tasks:** these tasks come from Planner plans associated with the teams or groups the user belongs to. Each group or team that you're a member of that has an associated plan appears here. All the plans owned by the team or Group appear under the team, together with the name of the tab and the channel name. For example, Book planning is the name of the tab in the General channel of the Office 365 for IT Pros team. The name of the channel is often the same as the name of the plan, but it

doesn't have to be. If a channel is hidden, then any plan linked to a tab in that channel is also hidden. Tasks shows these plans as hidden lists. You can see the tasks in a hidden list by clicking the list to expose details of the plan and then selecting the plan.

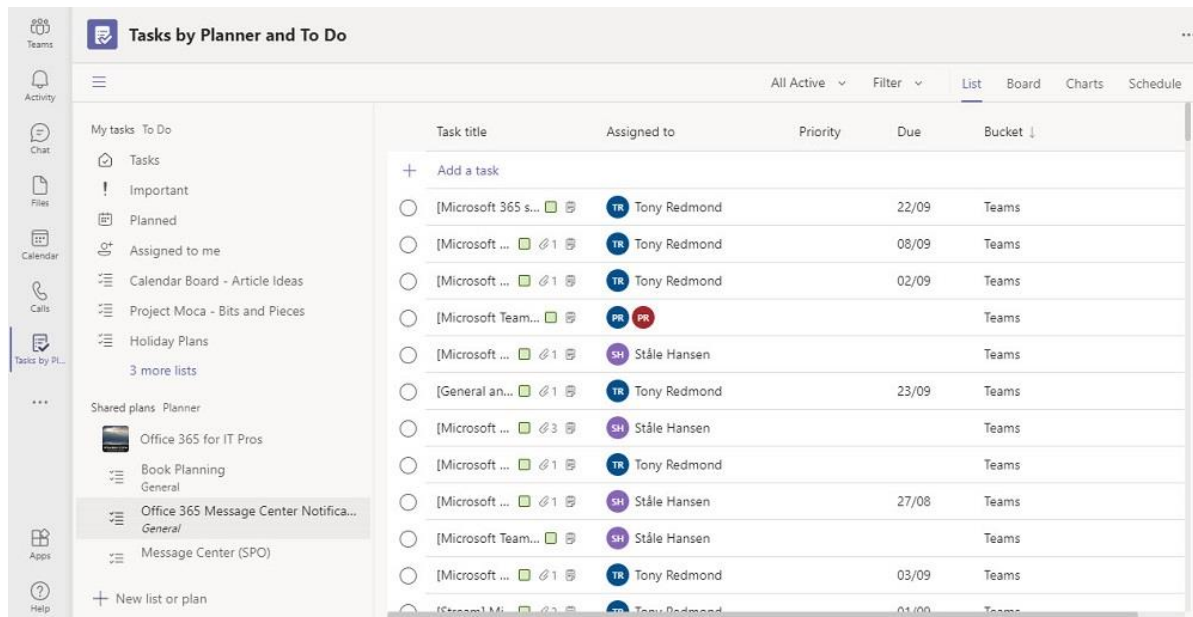


Figure 9-15: Planner and To Do Tasks are available in the Tasks app for Teams

When you open a list to work with tasks, the app presents the appropriate user interface to interact with the tasks. For personal tasks, you see the To Do interface; if the Planner manages the task, you see the Planner interface (Figure 9-16). You can't move or copy a personal task to Planner or make a Planner task a personal item.

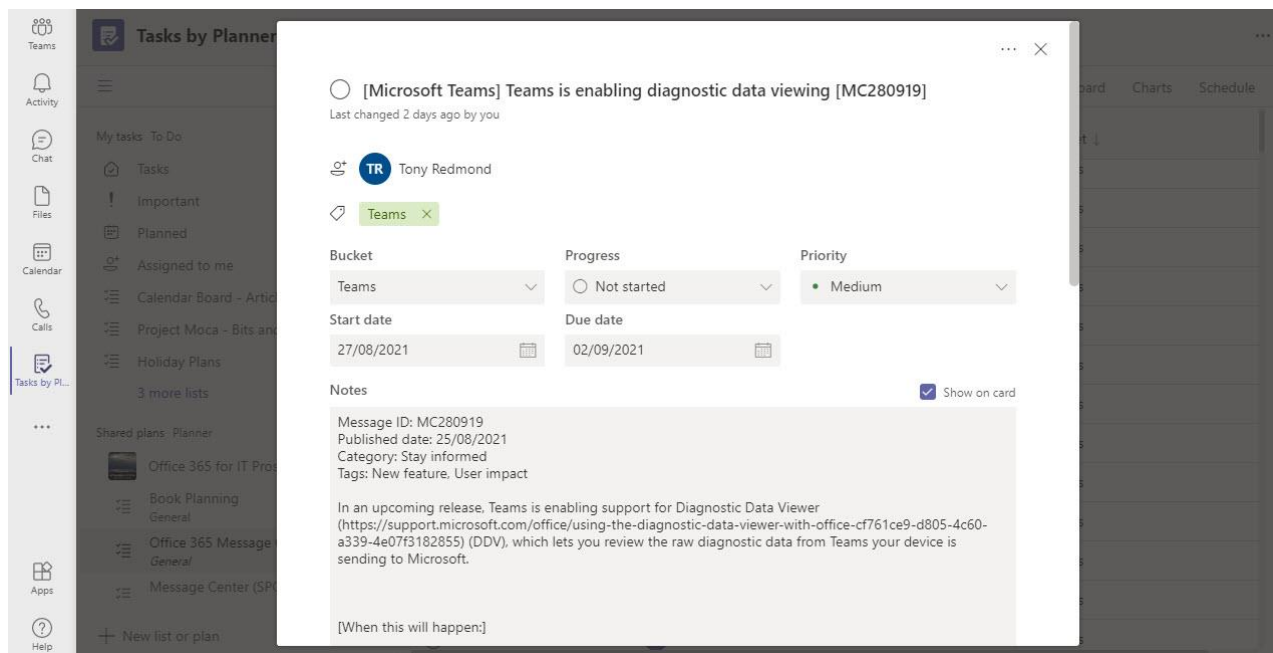


Figure 9-16: Working with a Planner task through the Teams app

Working with a task happens in the same way as in the native app and Teams writes any changes made through the app back to the native repository using Graph APIs. Due to the need to refresh client-side caches, a slight delay occurs – before the change appears in the native app, but this shouldn't be a problem because people don't usually rush to check To Do or Planner after updating a task in Teams. One minor difference

between the Planner browser app and the app in Teams is that Teams doesn't send email notifications about task updates or closures.

Apart from delivering a single integrated view of personal and work tasks, the app can do some things that aren't available in the standalone apps. For example, in the list view, you can select multiple tasks and edit them at one time. This is a great way to assign new people to tasks (perhaps after someone has left a team) or to update the completion status or set a new due date for a set of tasks.

## Accessing Plans Through Teams Channel Tabs

Sometimes people need to concentrate on work items belonging to a single plan. This is best accomplished by adding a Planner tab to a team channel. As it turns out, according to Microsoft's telemetry, Planner is the most popular channel tab in Teams.

If a plan hasn't already been created for the underlying Group, Teams creates it and attaches the plan to the tab. If one or more plans already exist for the Group, you can select one and attach it to the tab. You can then work with the tasks in the plan through Teams or click the **Go to website** (globe) icon in the set of options at the top right-hand corner of the pane. This launches Planner and loads the tasks for the plan into a browser tab.

## Removing a Plan from a Channel

To remove a plan accessed through a channel tab, open the plan, select the arrow next to the tab, and then **Remove**. You now have a choice:

- **Remove the plan from Teams:** This is the default and leaves the plan accessible through the Planner browser interface. If needed, you can add the plan back to Teams.
- **Permanently remove the plan:** To do this, select the checkbox *Permanently delete this plan and all its tasks*. This removes the plan data from the Planner Azure service. The plan is then irrecoverable through any interface.

Unlike removing a plan through the Planner browser interface, removing a plan through Teams does not affect the Group or any other resource attached to the Group or team.

# Linking Planner and the Microsoft 365 Message Center

As an example of what's possible with Planner, Microsoft released a service to synchronize notifications about changes in Microsoft 365 services from the Message Center in the Microsoft 365 admin center (described in chapter 4) to create tasks in a target plan. Each time Microsoft announces a change in the service, a new task is created that can be assigned to the person or team responsible for communicating and managing that type of change within the organization. A Power Automate flow runs on a scheduled basis to synchronize information from the Message Center to Planner. The integration is enabled through an option above the list of messages in the Message Center (Figure 9-17). When you select the option, you can choose which types of Message Center notifications you want to be synchronized, and which plan you want them copied into. Synchronization occurs once a day.

## Message center

Each message gives you a high-level overview of a planned change and how it may affect your users, and links out to more detailed information to help you prepare. [Learn more about managing changes](#)

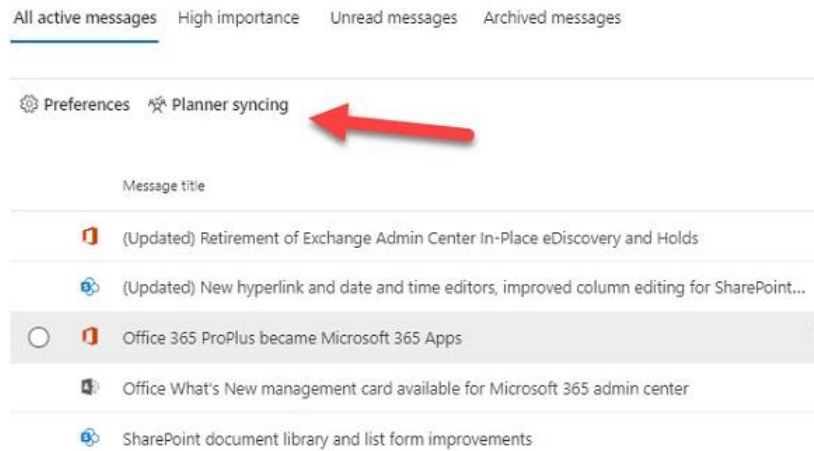


Figure 9-17: Planner synchronization of Microsoft 365 updates

Tasks synchronized from the Microsoft 365 Message Center have start dates assigned by Microsoft but no due dates. The start dates are associated with the work being done by Microsoft to introduce a new feature instead of when a feature might be available in a tenant. If you don't adjust task dates after synchronization, it's easy to end up with a schedule full of red (overdue) tasks. For this reason, it's important to assign tenant-specific dates for tasks when assigning tasks to users.

The practical value of this feature is that it lets you capture change notifications in a single plan so that you can review the changes and possibly act, such as assigning tasks to other users, sending out notification emails, making infrastructure changes, and so on. It can be difficult to remember to keep up with the volume of change notifications in the Message Center, which has poor search functionality and no way to deeplink to an individual notification. Copying these changes into Planner gives you a much more flexible way to track and manage these changes, and it's a nice touch that highlights how well different parts of the Microsoft 365 platform work together. To learn more, see [this blog post](#).

# Chapter 10: Managing Video

**Tony Redmond**

## The Cloud Video Platform

Stream is the Microsoft 365 enterprise video service that allows users to upload, view, and share videos securely. Given the proliferation of video-sharing apps like YouTube available to consumers and the preference of some users to consume information through videos, the role of video communications within businesses is growing. It is undeniable that many people are more accustomed to and prefer to learn from short-form videos than from reading memos and reports. The role of the new Stream is less about the management of Microsoft 365 video content; its focus has moved to enable the best possible playback of video and audio content stored in SharePoint Online and OneDrive for Business.

Information about the architecture and management of Stream Classic is available in Chapter 14 of the Companion Volume.

## Stream on SharePoint

Stream is in the middle of a transition from its original Azure-based storage to using SharePoint Online and OneDrive for Business (ODSP). The new Stream, also called Stream on SharePoint, is the version that uses ODSP storage. The move to ODSP solves multiple problems afflicting classic Stream including:

- Storage of video content in country-level data centers to satisfy local data residency rules.
- Inclusion of video content in information governance (retention) and information protection policies.
- Ability to search and find the spoken words captured in transcripts of Teams meeting recordings in eDiscovery cases.
- Access for external users through standard SharePoint Online external access and sharing controls.
- Video content can be included in backup processing for SharePoint Online and OneDrive for Business.
- Increased storage quota for video content, particularly in OneDrive for Business (important as Teams meeting recordings often consume substantial amounts of storage).
- Programmable access to video content through the Graph API.

Because the new Stream is built on top of ODSP, it follows that many of the services created for Stream classic are no longer necessary. For example, Stream doesn't use Azure blob storage, so its quota management system is defunct. ODSP permissions replace the Stream classic permissions model, and so on. The transition allows the Stream developers to concentrate on media management and playback instead of recreating wheels that are already available. In effect, Stream on SharePoint is really more of a set of services to manage video storage and playback instead of an app.

Stream hosted on SharePoint doesn't have a Stream portal in the same way that Stream classic or the older Office 365 Video application do. Instead, if an organization wishes to have a video portal (for example, to allow access to videos about the business), they need to use a different approach. For example:

- A SharePoint site tailored to highlight and feature selected videos.
- A channel in a team dedicated to the same purpose.
- Organization videos published through [Viva Learning](#).
- Videos published through a Yammer community.

As described later, the transition to the new Stream is complete for Teams meeting recordings. The latest information about the migration of content from classic Stream to ODSP is that the [migration tool is in private preview with some customers](#). Microsoft has not yet revealed the date for when the migration tool will become generally available. A preview of the new Stream client is available, and tenants can switch the Stream tile in the Microsoft 365 app launcher to point to the new Stream instead of Stream classic.

## Stream Licensing

Stream is included in all enterprise plans (including DoD, GCC, and GCC High), the education plans, and the front-line worker plans. It is also included in the Business and Business Premium plans. See [this page](#) for more licensing information. The Stream functionality available to Office 365 E3 and E5 tenants includes advanced features like speech-to-text and closed captions, transcript and caption generation, and searches.

## Stream Browser Client

Stream supports video upload and viewing through a wide range of browsers, including Microsoft Edge and the current versions of Chrome, Brave, and Safari. Currently, a preview version of the new Stream client is available. A setting in the SharePoint admin center controls the behavior of the Stream app tile in the Microsoft 365 app launcher. Three values are available:

- The default option is to **Automatically switch to Stream (on SharePoint)**. Microsoft controls this option and will set it after the migration of existing Stream content is complete.
- **Stream (on SharePoint)** directs users to the preview GUI for the new Stream. The user can switch to the classic Stream GUI if they want.
- **Stream (Classic)** forces people to use the classic Stream GUI.

The new Stream client is a composite of a browser interface to manage videos and the standard OneDrive web audio and video player to play videos and control their settings. If you play a video from the OneDrive for Business or SharePoint Online clients, you use the same player. The client uses Microsoft Search to find and display videos stored in Microsoft 365 apps, including those stored as attachments in user mailboxes. Because the new app is based on ODSP, it works differently from the classic Stream browser app. Until the migration of older Stream content is complete, the new app only has access to this content:

- Teams meeting recordings (both your recordings and ones shared with you).
- Videos that you upload to the new client (or OneDrive for Business). This includes video files uploaded through applications like Yammer and Teams.
- Videos shared with you by other users.
- Video file attachments for emails in your Exchange Online mailbox.

The preview of the new Stream client is missing some functionality that's available in the older client, including:

- Video recording (screen capture) direct from the client to a video stored in Stream.
- Trim a video. Trimming allows an owner to remove some extraneous content from the start and end of a video.
- Replace a video. Sometimes a trim isn't enough, and videos need a little more post-production work to iron out imperfections. Video owners can download a video, process it with a video editor like [TechSmith Camtasia](#) or Microsoft Clipchamp, and upload the updated version. The ability to replace a video is important because this retains the link used for the original video and avoids the need to replace it in places like Teams channel tabs or SharePoint pages.

Microsoft publishes [an online spreadsheet](#) to compare the functionality available in the new Stream client against the classic client. It describes what features are available, those under development, and functionality that will not be available in Stream 2.0.

The remainder of this chapter focuses on the new Stream client.

## Management Client

Unlike Stream classic, the new Stream doesn't have a switch to reveal administrative settings. Much of the need for those settings disappeared with the transition to ODSP because standard ODSP and Microsoft 365 settings control aspects like sharing, storage quotas, recovery of videos for deleted users, and retention. The ability to organize videos through Microsoft 365 Groups by storing videos in the SharePoint team site belonging to the groups is still there, even if a dedicated GUI to perform group-based video management is unavailable.

## Network Demand

Like any video application, Stream benefits from good network performance between servers and clients. Microsoft publishes articles to help administrators understand the characteristics of network performance for Stream, including [Video delivery and network overview](#) and [how to scale video delivery for Stream](#). The current advice is for the classic client, but it should not be very different from what's necessary for the new client.

## Stream and Other Microsoft 365 Apps

Like many other Microsoft 365 apps, Stream uses components from across Microsoft 365 and contributes functionality to other apps. Examples of how Stream integrates with other applications and services include:

- Stream classic can use Microsoft 365 Groups to manage access to videos. This functionality is not available in the new Stream app.
- You can link a Stream channel to a Teams channel tab to allow users to play videos from the channel. Stream channels are currently available only in the classic client. You can also paste a link to a Stream video in a channel conversation or personal chat for readers to access.
- Yammer can playback Stream videos. You can paste the link for a Stream video into a Yammer message and readers can then play the content back inline.
- The SharePoint web part for Stream can highlight a video, channel, or list of videos (from all of Stream) on a page.
- [Forms can collect feedback](#) through quizzes, polls, and surveys for videos through the *Interactivity* tab for a video. This functionality is only available in the classic client.

Other apps can use [Stream's implementation of oEmbed](#) to display videos or channels.

## Stream User Functionality

Stream's user functionality divides into these major sections:

- Uploading video content.
- Managing videos.
- Video playback.

The Stream client (Figure 10-1) uses the same design "language" as other Microsoft 365 apps. At the top of the client, the recommended list of videos contains files that you've worked with, recently watched, or those that Stream believes are of interest. One of the videos listed is in the recommended list because it is linked to a recent meeting (see the sharing section below). Filters control the videos shown beneath the recommended set, including one to display videos marked by the user as favorites. If the selected filter is set to All, Stream displays the full set of videos available to the user in Microsoft 365 locations, including any found in email



attachments. The videos include those that the user owns and those that they can access because of access gained through a sharing link, Teams chat, or direct access.

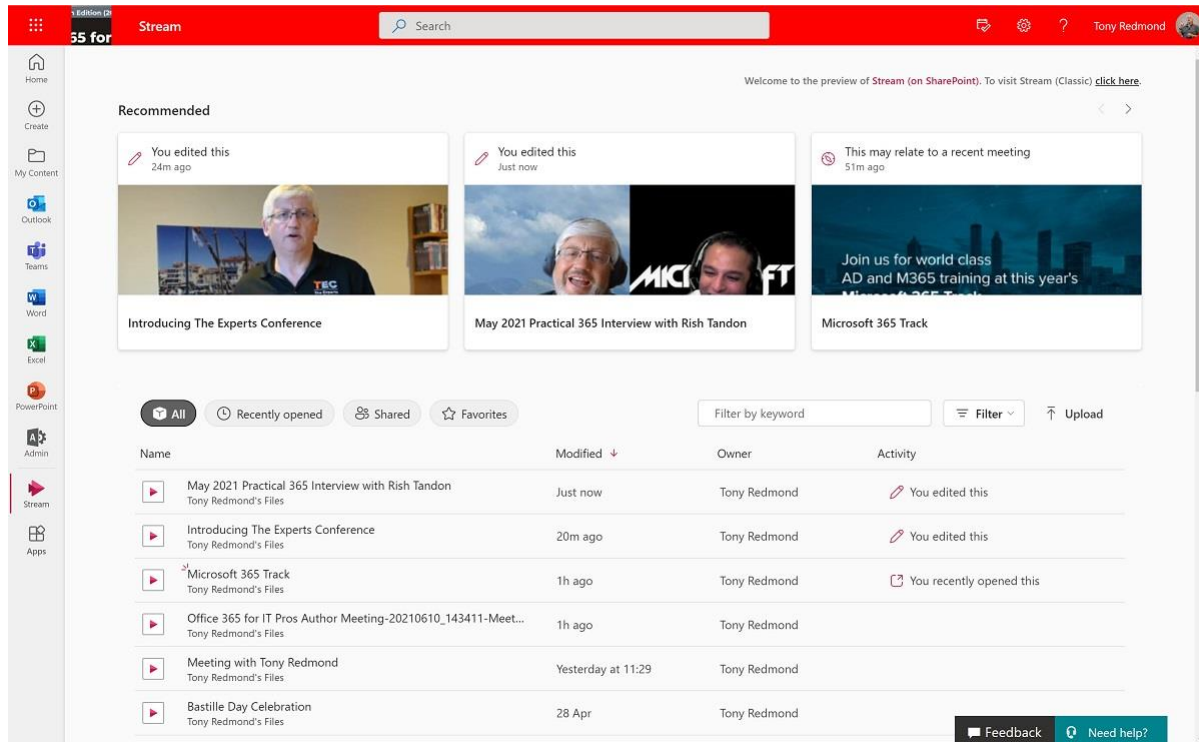


Figure 10-1: The new Stream client (preview)

No matter how someone accesses Stream content, Stream won't show them a video unless their account has the right to view it.

## User Settings

The user settings available for the new Stream browser client are minimal. Currently, they are limited to selecting a theme and deciding whether to use dark mode.

## Uploading and Accessing Videos

To upload a video, click the upload link and select the source files (in the [supported file formats](#) – audio-only files are not supported). You can also drag and drop a file to the client. A newly-uploaded video receives the default permission in the target site or account. Recordings of Teams channel meetings are available to team members while recordings of Teams personal calls are available to the participants.

After uploading a video, the owner can:

- Play the video.
- Open the location of the video. This option opens a new browser tab positioned in the OneDrive or SharePoint folder where the file is located.
- Share the video.
- Add a link to the video to a calendar meeting. Stream uses the OWA meeting request screen to compose the meeting invitation. Make sure that the recipients of the invitation have access to the video. The client uses meeting insights to highlight videos sent in meeting invitations before the event occurs to give the user the chance to review the content beforehand.
- Add a link to the video to To-Do. This action creates a bare-bones personal task in To Do.
- Mark the video as a favorite. This adds the video to the set listed through the Favorites link.
- Hide the video. The Stream client doesn't show the video in its list.

- Download a copy of the video.

Opening a video with the Stream video player allows the owner to manage some settings for the video (Figure 10-2). People who watch the video use the same player, but they can't change the settings.

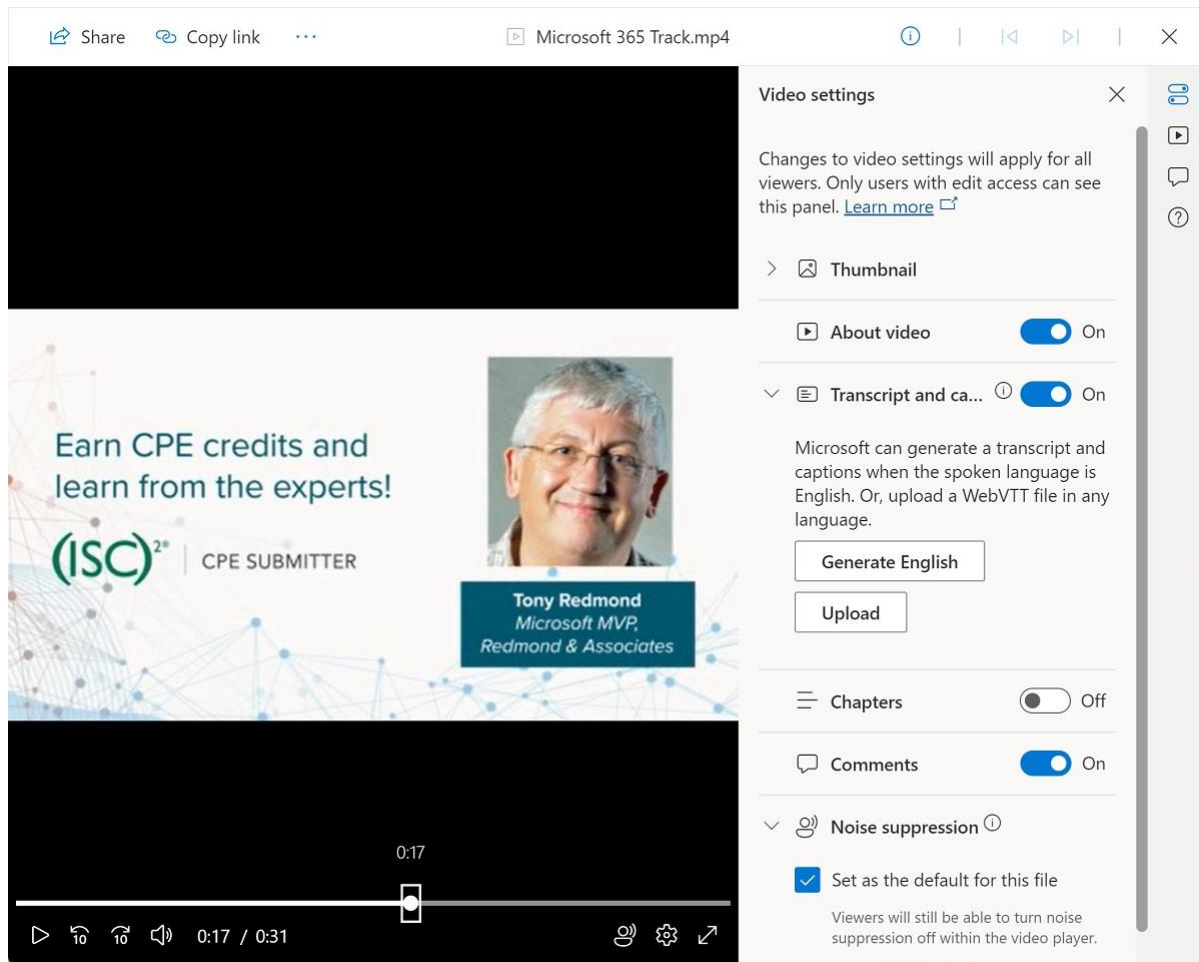


Figure 10-2: Options available in the Stream video player

The video settings are available in an item in an overlaid menu on the video player. The full set of menu items are :

- **Video settings:** Opens a pane to allow the video owner to turn on different options.
- **About the video:** If enabled by the owner, displays some details about the video. The owner can add some text to describe the video or rename the video if necessary. A basic text editor is available to enter content, including bulleted lists and hyperlinks. Pasting text from another source is supported.
- **Transcript:** Opens a pane to show the transcript (and chapter markers, if used).
- **Comments:** Opens a pane to show any comments posted for the video. Stream uses the same commenting facilities available for Office documents stored in ODSP.

The video settings are:

- **Thumbnail:** Select a frame from the video to use as the thumbnail displayed for the video by the Stream client. To select a frame, move the pointer in the seekbar (the time bar under the video) to select an appropriate frame.
- **About video:** Controls if any information is available for display to users when they watch the video. The text entered about a video is indexed and discoverable using SharePoint search. It's a good idea to include information about a video such as the time and place of the recording, important parts of the video, and perhaps some hyperlinks for users to find more information about the topic.

- **Transcript and captions:** Stream doesn't generate captions or a transcript for uploaded videos unless the owner requests this to happen. For videos where the spoken language is English, Stream can generate a transcript. Teams meeting recordings use the captions generated during the recording.
- **Chapters:** To help organize videos, creators can insert chapter markers in files to help users navigate within a video to find the information they're interested in. If the owner enables chapter markers, Stream displays the markers in the timeline shown by the player. See [this article](#) for more information.
- **Noise suppression:** Users can enable noise suppression if they want when playing a video. This control affects if noise suppression is the default for a video. It should be turned on if significant background noise is present.

## Storage Quotas

The storage consumed by video files is charged against the SharePoint Online (for videos owned by Microsoft 365 Groups) and OneDrive for Business (for personal videos) storage quotas. OneDrive for Business [offers 1 TB of storage for Microsoft 365 business accounts and "unlimited storage" for enterprise accounts](#). The impact of storing Teams meeting recordings is moderated by the application of an automatic expiration policy. Videos expire after 120 days unless the video owner overrides the expiration period by changing it (or removing the expiration) or applying a retention label to the file. Upon expiration, video files follow the normal OneDrive for Business recycle process.

The exact size of a video file depends on its format, quality, and length. As a guide, expect to use approximately 7.5 MB per minute of 1080p MP4 video with smaller amounts consumed for lower-quality video (the storage required per minute doesn't vary for Teams meeting recordings saved to One Drive for Business). Stream counts the original file size of the uploaded video against the quota and doesn't take other factors such as the size of transcoded videos and caption files into account.

## Other Ways to Get Video Content into Stream

Other ways are available to get video content in supported formats into Stream, including:

- With the Stream mobile app. This uploads the file to Stream classic.
- Making a screen capture video with Stream. This capability is not yet available with the new Stream client.
- Recording a Teams meeting.
- Uploading a video file to a Teams chat or Yammer conversation.
- Running a Live Event with Teams or Yammer.

In all cases, once the video is uploaded, Stream processes it to create the playback files, captions, and transcript.

## Sharing Stream Videos

Because the new Stream stores its files in ODSP, sharing a video is very much like sharing any other SharePoint or OneDrive file. To share a video, select a video and select Share from the [...] menu to expose the following sharing mechanisms:

- **Email.** ODSP generates a sharing link and inserts it into an email addressed to the link recipients. The sharing link settings control what the recipients can do with the video. For instance, the settings can block downloads. Recipients of the sharing link can be inside or outside the organization.
- **Copy link.** The user can copy the sharing link to the clipboard and share the link elsewhere as needed. For example, you can create a Teams channel tab to point to a specific video. Make sure that the permissions set on the link allow the intended audience to access the file.

- **Teams.** ODSP uses the Share to Teams feature to post the link to the video to a personal chat, group chat, or channel. This is the only method supported to share a video file stored in a user mailbox. The process is the same as that used to share an item from Outlook to Teams (see Chapter 12).

If you want to start a video playback at a specific point, you can add an instruction to the link. For example, this instruction tells Stream to start the video at the 180 second mark.

```
&nav={"playbackOptions":{"startTimeInSeconds":180}}
```

For instance, a link containing the instruction for a video stored in a user's OneDrive for Business account might look like this:

```
https://office365itpros-my.sharepoint.com/:v:/g/personal/james_ryan_office365itpros.com/ETbXnV4vB-pMsfv9LK81d9MBtTS9KQIBnZaHNq9LxcrfMQ?e=qoCLkw&nav={"playbackOptions":{"startTimeInSeconds":180}}
```

## Transcripts

A Stream transcript is made up of individual timecoded captions in Web Video Text Tracks, or WebVTT (VTT) format. For English language videos, Stream can generate a transcript by interpreting the voice track to create captions along with the corresponding timecode. For example, here are three segments from a transcript.

```
00:00:26.645 --> 00:00:30.965
We'll talk about it. I
have my viewpoint on it as well.
00:00:30.965 --> 00:00:35.015
OK, OK, so let's get things
going and what we might want to
00:00:35.015 --> 00:00:39.335
do is then if we can just take
some of the snippets out of this.
```

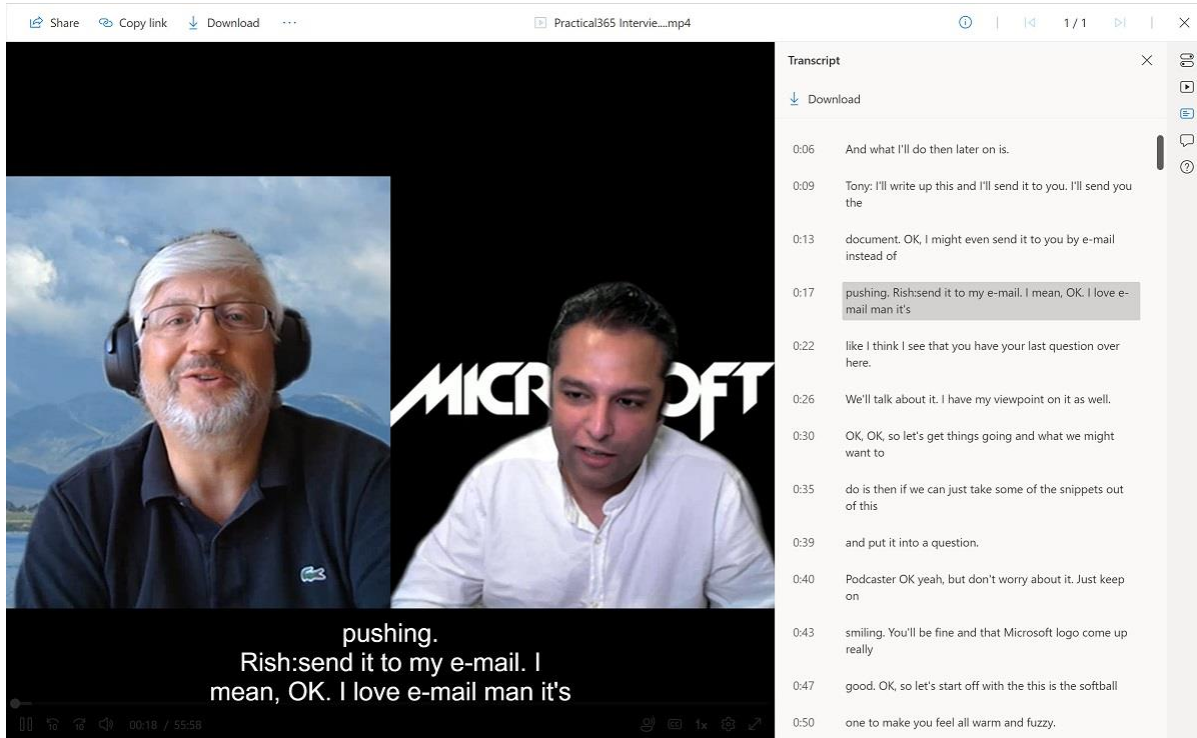


Figure 10-3: Playing a Stream video with its transcript

Figure 10-3 shows the transcript as displayed in the Stream viewer. You can see how the individual captions appear in the transcript, including the three shown above. Selecting a caption repositions the player at that

point in the video. You can also see how Stream displays the same text under the video when the user enables closed captions.

Automatic transcript generation is very helpful and, generally, it does a reasonable job of interpreting the spoken words in a video. However, the nature of automatic text generation is that it can always do with a little help. If you want to improve the quality (accuracy) of the captions:

- Download the transcript file and amend it with any text editor. Unless you use a VTT editor (here's a [free online example](#)), don't edit the time codes – just the text.
- Delete the existing transcript.
- Upload the amended file.

You can create misunderstood words, add speaker attribution (an example is shown in Figure 10-3), or include other details that might be useful to the watcher.

## Noise Suppression

Noise suppression isolates speech from background noise in the audio feed to make it clearer and more distinct. Video owners can update the details of videos to enable noise suppression. When noise suppression is enabled, viewers have the option to keep noise suppression on or disable it during video playback.

Most videos qualify to be processed for noise suppression. The criteria include:

- The video is two hours or shorter, and no larger than 3 GB.
- An audio track is available, but not when multiple audio tracks in different languages exist in a video.
- The video is not a recording of a Teams meeting. This is because noise suppression is automatically done when Teams meetings are recorded.

Noise suppression is not available for recordings of Live Events. You can't disable noise suppression on a tenant-wide level.

## Viewership Statistics

When people view videos using the web player, Microsoft 365 captures statistics to help video owners understand how engaged viewers are with the content. Not everyone watches a video from end to end, and it can be important to know what parts of a video attract the most engagement or where viewers begin to drop off. To see the viewership statistics for a video, use the Open File Location option to go to the video's storage location in OneDrive for Business or SharePoint Online and hover over the file to display its details card. Then click the Views link. You'll then see the number of views per week, who viewed the video, and its ability to retain viewers (Figure 10-4) over time.

Viewership retention is different from the number of viewers. If someone opens a video at any point, they count as part of the total viewership. Retention depends on how many of the total viewers saw a specific point in the video. For instance, Figure 10-4 reports that 80% of the viewership saw the video at the 14-second mark with 100% viewership attained before and after this point. That kind of graph indicates viewers dipping in and out of the video. The number of views includes "rewatches," or occasions when someone goes back and reviews a part of a video again. The statistics include these events when someone goes back to a part of a video during a viewing session. Usually (as seen in this case), viewership drops off toward the end of a video but it's more unusual to see a dip at the start. This is probably caused by people skipping ahead to avoid the opening credits. The viewership chart aggregates data for the entire lifetime of a video: you cannot choose to view data for specific periods.

Viewership data is tied to a specific version of a file. If you edit a video and upload the new content to Stream, statistics accrue for that version and don't inherit the data for the older video.

Stream adjusts the data points in the time chart depending on a video's duration. If a video is short (under 200 seconds), a data point appears for each second. For longer videos, Stream adjusts the graph by aggregating data over longer periods. For example, a 10-minute video uses 3-second data points.




Figure 10-4: Statistics for a video

## Stream Mobile Apps

Stream mobile apps are also available for [iOS](#) and [Android](#). The current versions of the mobile clients only work with the videos stored in classic Stream. Versions based on SharePoint storage are under development.

With the mobile clients, users can browse channels and groups to find and play videos (Figure 10-5), add videos to watchlists, like a video or make comments about its content, edit, delete, and share their videos, and download videos for offline access. Users can choose to minimize data consumption by streaming videos in standard definition format or use more data for higher fidelity. The automatic transcript (captions) is supported on mobile clients. [This page](#) has the latest details about the Stream mobile apps.

Being able to record and upload videos on mobile devices is an especially popular feature. When recording on a mobile device, the Stream mobile app supports these features:

- Swap between the available cameras on the device. The default camera selected for a new video is rear-facing, but you'll probably want to use the front-facing camera for in-person shots.
- Record multiple clips before uploading the video to Stream. When all the clips are recorded, you can drag and drop them into the order you want the clips to appear in the video.

- Include photos from the device in a clip. For example, you could take a picture of a new product and tape a commentary for the picture.
- Annotate (draw), add emojis, or apply filters to a clip.
- Trim clips by removing content from the front or end of a video.

When the creator uploads the video, Stream processes it and prepares the content for publication. Before the upload, the creator should make sure to assign the language to a video to allow Stream to generate an automatic transcript. Users can also upload videos from the device's camera roll or a cloud storage app like OneDrive.

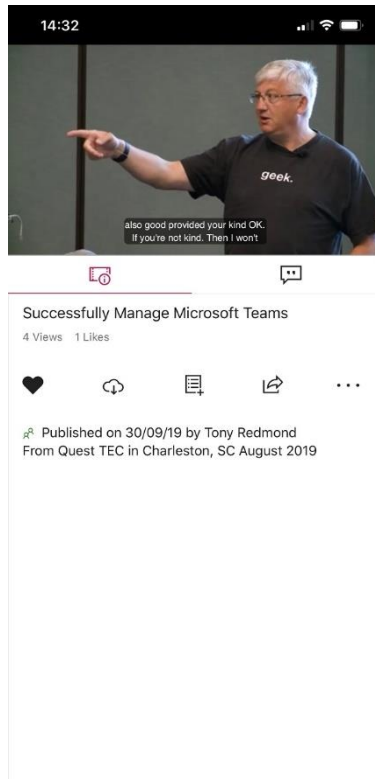


Figure 10-5: Watching a video on Stream for iOS

Other mobile applications can be used to create videos, many of which offer functionality over and above what's available in the Stream mobile app. The advantage of the Stream app is its integration with Microsoft 365, which makes it easy to upload and share videos. See [this page](#) for more information about recording and editing Stream videos on Android and iOS devices.

## Stream and Teams

Teams makes use of functionality from many Microsoft 365 and Azure apps. In the case of Stream, Teams uses the following features:

- Recording of Teams meetings.
- Publishing Stream videos via a channel tab.
- Recording of Live Events. After the end of a live event, the video stream automatically transitions from a live feed to an on-demand recording. The recording is available online and accessible to attendees and other users for 180 days. If the organizer of a live event wants to keep the recording for longer, they must download the video and upload it to Stream. The storage of a live event recording counts against the Stream quota for the tenant.

Teams is the first workload to transition to the new Stream architecture. Stream stores all new Teams meeting recordings in OneDrive for Business (personal meetings) or SharePoint Online (channel meetings). Older recordings and other videos remain in Azure until Microsoft makes migration tools available to move this content to SharePoint Online and OneDrive for Business.

## Recording Teams Meetings

Recordings can be made for Teams private meetings (including group chats), channel meetings, ad-hoc (“meet now”) meetings, personal (1:1) meetings (including VOIP and PSTN calls), and live events (see above). The organizer and participants of Teams meetings can record meetings if the Teams meeting policy assigned to their accounts enables this option. The person who starts a recording is deemed to be the owner. Meeting organizers can update the options for individual meetings so that a recording starts automatically when a meeting commences (the first person with the presenter or owner role joins the call).

A separate setting in the meeting policy enables meeting organizers to generate automatic transcripts using the Teams desktop client. Transcript generation was originally available only for meetings held in U.S. English (or rather, meetings detected by Teams as being in U.S. English). Today, Teams supports an expanded range of languages for live caption and transcript generation including other variants of English, German, French, Italian, Spanish, Russian, Hindi, Arabic, Finnish, Danish, Swedish, Japanese, Korean, Chinese, and Dutch.

When a Teams meeting is recorded, Teams adds a bot to the call to capture the voice and video stream from participants (even if the meeting is voice-only). Recordings use a 3x3 view (the last nine speakers). Once the meeting ends or the recording is halted, Teams stores the recording as an .MP4 file as follows:

- For **personal meetings**: the Recordings folder of the OneDrive for Business account of the meeting organizer. The exception is when a participant with co-organizer or presenter rights records the meeting. In this case, that person is the owner, and the recording goes into their OneDrive for Business account. However, the meeting organizer has edit rights to the recording. The same occurs if meeting options dictate automatic recording. In this case, the person holding the presenter role who first joins the meeting becomes the owner of the recording, and Teams creates the MP4 file in their OneDrive for Business account and grants the meeting organizer edit access.
- For **channel meetings**: the Recordings folder for the channel in the SharePoint document library belonging to the team. For example, General\Recordings. The exception is for meetings in private channels. The sites used for private channels don’t need individual folders as exist for normal channels in the team site, so recordings go into the root of the document library of the private channel’s site. The person who starts a channel recording has edit rights to the file.
- For **1:1 calls and group chats**: the Recording folder of the OneDrive for Business account of the person who records the meeting.

If the meeting is left active, the recording ceases after four hours. Access rights for recordings stored in OneDrive for Business do not include external meeting participants, including guest accounts. If external access is needed for a recording, the owner must grant access.

Teams also posts a link to the recording in the meeting chat or channel. The MP4 file can be edited or processed by whatever tools are available to trim or otherwise adjust the recording. If available, meeting transcripts are available in the same place.

Different processing applies for recordings of 1:1 calls. These meetings can only be recorded if both participants use Teams clients and the meeting policy assigned to one of the participants allows them to record the meeting. Recordings are unsupported if one of the participants uses a dial-in number or connections through federation (external access). It’s also important to understand that only participants from the host tenant have (read) access to the recordings by default. External participants, including guest users,



cannot access the recording unless the meeting organizer shares the recording with them. Table 10-1 summarizes how meeting participants can access Teams meeting recordings.

<b>Type of meeting</b>	<b>Who can view recording</b>	<b>Access to recording</b>
Group chat	All chat participants from the host tenant.	Recording posted in meeting chat and in the OneDrive for Business account of the person who starts the recording.
Private or ad-hoc meeting	All participants from the host tenant.	Recording posted in the meeting chat and in the OneDrive for Business account of the person who starts the recording. The recording is also available in meeting recap. Meeting participants from the same tenant have read-only access rights to the recording. External participants, including guest accounts, must be assigned permission to access the recording.
Channel meeting	Members of the team owning the channel.	Recording posted in the meeting chat in the channel and the Recording folder of the channel folder in the SharePoint site belonging to the team.
1:1 meeting	Both participants (if in the same tenant).	Recording posted in the chat and in the OneDrive for Business account of the person who starts the recording. The other party has read-only access if they have a tenant account.
VOIP and PSTN calls	Call owner.	Recording and transcript are available in the call history and call details.

Table 10-1: Access to Teams recordings

No matter how a recording is made, only the owner can download or remove the video. The owner can upload the recording to Stream if they wish to share the content to a wider audience (except externally, as Stream doesn't support this type of sharing).

**Sharing Permissions in OneDrive:** Permissions for meeting recordings stored in OneDrive are limited to internal users, even if guests participate in the call. If it is necessary to share a recording with an external user, the owner (the person who started the recording) must update the sharing list to include that user. Channel recordings stored in SharePoint can be accessed by any team member. Another thing to remember is that only the owner can download the recording from OneDrive. Other participants in personal meetings receive view-only permission and are blocked from downloading.

## Recordings for VOIP and PSTN Calls

Teams users who place calls to VOIP or PSTN recipients can record the call and generate a transcript if the calling policy assigned to their accounts have the following settings enabled:

- *AllowCloudRecordingForCalls*: Controls if the user can record PSTN and VOIP calls. By default, the setting is True.
- *AllowTranscriptionForCalling*: Controls if the user can generate a transcript of a call. By default, the setting is False.

For example, to update the default Teams calling policy to allow users to record calls and generate transcripts, the command is:

```
[PS] C:\> Set-CsTeamsCallingPolicy -Identity Global -AllowCloudRecordingForCalls $True -AllowTranscriptionForCalling $True
```

Meeting recordings are stored in the OneDrive for Business account of the user who starts the recording and is available with the transcript in the call history and call details.

## Retention of Teams Recordings

Most recordings of Teams meetings age rapidly. Time and activities overtake the matters discussed and the value of the recording diminishes over time. For this reason, many organizations want to remove Teams meeting recordings sooner than other content. Microsoft has committed to delivering a special retention policy to process Teams meeting recordings which can be used by all tenants with Teams licenses. This feature isn't yet available, so for now only tenants with the necessary Office 365 E5 or Microsoft 365 E5 compliance licenses can create an auto-label policy to apply retention labels to the recording files stored in SharePoint Online and OneDrive for Business. Alternatively, if you don't have the licenses for auto-label policies, you can train users to apply retention labels manually.

Background labeling processes find content identified by policy criteria (a keyword query, sensitive information type, or trainable classifier). In the case of Teams meeting recordings, a keyword query of *(ProgID:Media AND ProgID:Meeting)* locates the MP4 files. These are programmable identifiers set on recordings when Teams creates the files. You can test the effectiveness of the query by using it with SharePoint search to find Teams meeting recordings available to you.

When specifying the locations covered by the policy, make sure that you include all the locations where recordings might be found:

- OneDrive for Business (personal meetings).
- Microsoft 365 Groups (channel meetings). This includes the SharePoint sites connected to teams.
- SharePoint Online. This covers any copies of recordings moved from the SharePoint sites connected to teams.

When the auto-label policy matches the query against a file, it applies the retention label specified in the policy settings. The retention period in the label dictates how long the recording remains in place. Once the retention period (say, six months) lapses for recordings, background processes remove the files. See Chapter 17 for more information about retention policies and labels.

Compliance administrators can track auto-label activity through the Activity Explorer in the Microsoft 365 compliance center. Another way to check the progress of auto-labeling is to run the PowerShell script [downloadable from GitHub](#) to report the recording files which receive a retention label together with information about how long it takes auto-labeling to happen after the creation of recordings.

## Auto-Expiration of Teams Meeting Recordings

Microsoft says that 96% of Teams meeting recordings are not rewatched 60 days after the original meeting and 99% are not watched after 110 days. To help preserve storage quota, Teams marks an auto-expiration date on each recording. Unlike retention labels, which need Office 365 E3 or above, auto-expiration works for all SKUs which include Teams.

Originally, Teams used a 60-day retention period. Following customer feedback, Microsoft increased this to 120 days, which is the expiration period for recordings created by enterprise users. Recordings for users with Office 365 A1 licenses have a 30-day retention period. It doesn't matter if the recording is in SharePoint Online or OneDrive for Business. Users can change the expiration period for individual recordings by updating file properties through the file details pane (selecting preset values of 14, 30, or 60 days, a custom date, or *Never Expire*). Organizations can set a default expiration period for newly created recordings using the Teams meeting policy assigned to user accounts. The auto-expiration settings can be managed by updating meeting policies in the Teams admin center. Alternatively, you can use PowerShell. For example, to set the default

expiration period for recordings of meetings made by people assigned the *VIP User Meeting Policy*, run the command:

```
[PS] C:\> Set-CsTeamsMeetingPolicy -Identity "VIP User Meeting Policy"
-NewMeetingRecordingExpirationDays 180
```

The minimum retention period is one day, and the maximum is 9,999 days (which should be enough for anyone). You can also set retention to -1, meaning that recordings of team meetings never expire. The expiration period for A1 users can only be reduced from the default 30 days.

Background processes run to evaluate recordings and check their expiration date. If the process detects an expired file, the process moves the file into the recycle bin and clears the expiration date field. Users receive email notifications when expired recordings move into the recycle bin. If necessary, they can rescue important recordings from the recycle bin for up to 90 days after deletion. Once moved back from the recycle bin, the recording has no expiration date set and will therefore not be evaluated for deletion again.

To help users understand when a recording approaches expiration will see visual indications in:

- Beside the link to the meeting recording in the meeting chat. Anyone with view access to the recording sees the expiration notice.
- Two weeks before expiration, a red icon appears beside the MP4 files for files in the Recordings folder of OneDrive for Business accounts (personal meetings) or SharePoint Online sites (channel meetings).

Auto-expiration applies only to new recordings. Existing recordings do not have an expiration period. Auto-expiration is only available for meeting recordings and cannot be used with other file types held in OneDrive for Business and SharePoint Online.

If you move a file with an expiration date to another site, the expiration date stays with the file and remains active. Copying a file creates a new file without an expiration date. The same applies if you download a meeting recording and then upload it to a different site to create a new file. Retention labels take precedence over auto-expiration. In other words, if a recording has a retention label, that's what governs how it is kept (or removed). For this reason, you should assign suitable retention labels to important recordings if you want to make sure that Microsoft 365 will not remove the files.

## Searching Spoken Words in Teams Meeting Transcripts

When Teams generates a transcript for a meeting, it stores the spoken words in the Exchange Online mailbox of the person who starts the transcript (for personal meetings) or group mailbox (for channel meetings). After the meeting finishes, a background process copies the information from Exchange Online to match the words with segments in the video recording. This process accomplishes two goals. First, the playback of the meeting recording can display the transcript in a separate pane. Second, Microsoft Search indexes the spoken words to make it possible for users to search meeting transcripts based on words spoken during the meeting and captured in the transcript. Searching for spoken words in meetings can be done using SharePoint Online search or in OneDrive for Business. In both cases, the search is scoped to only reveal words from meetings the user is entitled to access.

Including spoken words from meetings in Microsoft Search indexes will eventually mean that it will be possible for eDiscovery investigators to include this information in their searches. This isn't possible today, probably because of the need for work to link the transcript and video recording intelligently in search results (for both preview and export).

## Commenting on Videos

Stream videos stored in OneDrive or SharePoint support the ability for users to comment on their content. Although users typically receive view-only permission for video and audio files, Stream makes an exception to

allow users to post comments via a button included in the set of overlaid options available during video playback. People with edit permission for a video can remove comments or prevent the posting of comments by updating video settings.

Apart from the per-video control over comments, organizations can remove the ability for people with view-only permission to post comments through a tenant setting. To do this, run the *Set-SPOTenant* cmdlet from the SharePoint Online module and change the value of *ViewersCanCommentOnMediaDisabled* from the default (False) to True:

```
[PS] C:\> Set-SPOTenant -ViewersCanCommentOnMediaDisabled $True
```

## Including Stream in a Teams Channel Tab

To make it easy for people to find videos, team owners and administrators can create channel tabs to connect to a video. To create the connection, enter the URL (retrieved with the Share option in Stream or from the browser address bar) for the video into the Teams dialog screen. The tab name is set to be the name of the video, but you can rename the tab if necessary.

## Stream Audit Events

Stream classic generates specific audit events as users interact with videos through the classic clients. In Stream 2.0, SharePoint Online generates the audit events because videos are processed just like any other file stored in a SharePoint site or OneDrive account.

You can search audit data with the audit log search feature in the Microsoft Purview Compliance portal or using the PowerShell *Search-UnifiedAuditLog* cmdlet. Chapter 21 explains how to search using either method, but for illustration purposes, the code below shows how to search the audit log and return the count of videos upload and modification actions by individual users.

- When a client uploads a video, it initially creates a file with a ~tmp prefix. For example, ~tmp7C\_Microsoft 365 Track.mp4. A *FileUploaded* event captures this action.
- Stream processes the video, updates its settings, and renames the file (in the example above, the renamed file is Microsoft 365 Track.mp4). These actions generate two *FileModified* audit events. One logged for the video owner (rename the file), the other for app@sharepoint (update the video settings).

The code below removes the events generated by the system app@sharepoint account from the reported data.

```
[PS] C:\> $EndDate = (Get-Date).AddDays(1); $StartDate = (Get-Date).AddDays(-30)
[array]$Records = (Search-UnifiedAuditLog -Operations FileUploaded, FileModified -StartDate
$StartDate -EndDate $EndDate -Formatted -ResultSize 5000)
If (!$Records) {Write-Host "No audit records found - exiting!"; break}

$StreamRecordings = [System.Collections.Generic.List[Object]]::new()
ForEach ($Rec in $Records) {
    $AuditData = $Rec.AuditData | ConvertFrom-Json
    If (($AuditData.SourceFileExtension -eq "mp4")) {
        $RecordingFileName = $AuditData.SourceFileName
        $DateLoc = $RecordingFileName.IndexOf("-202")
        If ($DateLoc -eq -1) {$Topic = $RecordingFileName} Else
            {$Topic = $RecordingFileName.SubString(0,$DateLoc)}
        $DataLine = [PSCustomObject] @{
            Workload      = $AuditData.Workload
            Date          = $Rec.CreationDate
            User          = $Rec.UserIds
            Recording     = $RecordingFileName
            Topic         = $Topic
        }
    }
}
```

```

        Site           = $AuditData.SiteURL
        FullURL        = $AuditData.ObjectId
        Folder         = $AuditData.SourceRelativeURL
        Operation      = $Rec.Operations }
    $StreamRecordings.Add($DataLine)
} #End If
} #End For

$UploadedFiles = $StreamRecordings | ? { $_.User -ne "app@sharepoint" }
$UploadedFiles | Sort {$_.Date -as [DateTime]} -Unique -Descending | Out-GridView

$UploadedFiles | Group-Object User | Sort Count -Descending | Format-Table Name, Count

```

Name	Count
Jane.Nix@office365itpros.com	74
James.Ryan@office365itpros.com	22
John.Hubbard@office365itpros.com	1
James.Joyce@office365itpros.com	1
Ben.Owens@office365itpros.com	1

To see a summary of activity for individual video files, change the *Group-Object* command slightly to:

```
[PS] C:\> $UploadedFiles | Group-Object Topic | Sort Count -Descending | Format-Table Name, Count
```

Some Stream-specific events are available such as actions logged for transcript creation, deletion, and access. This code will find those audit records.

```
[PS] C:\> $Operations = "FileTranscriptContentAccessed", "FileTranscriptCreated",
"FileTranscriptDeleted"
[array]$Records = (Search-UnifiedAuditLog -Operations $Operations -StartDate $StartDate -EndDate
$EndDate -Formatted -ResultSize 5000)
```

One issue in moving away from Stream-specific audit events to the more general-purpose events generated by SharePoint Online is that the Stream activities can be a small percentage of the overall set of audit events found by a search. SharePoint Online and OneDrive for Business generate many *FileModified* audit events as users work with Office documents, especially if the AutoSave feature is used as this can generate multiple events as AutoSave saves modified versions of files during a session.

# Chapter 11: Managing Microsoft 365 Groups

**Tony Redmond**

## Modern Groups and Distribution Lists

Microsoft has used different names for modern groups since their introduction in November 2014, including Office 365 Groups, Outlook Groups, Groups in Outlook, and now [Microsoft 365 Groups](#). To simplify matters, we use “Groups” unless we need to refer specifically to another type of group, such as a mail-enabled security group.

[In April 2017](#), Microsoft said that “*more than 10 million people rely on [Groups](#) in Outlook every month to work together and get things done.*” At that time, Office 365 had roughly 100 million active users, so Groups enjoyed 10% penetration. The number of people who rely on Groups is now much higher because the primary mission for Groups is the membership service for over 20 Microsoft applications, including Teams, SharePoint, Yammer, and Stream. In particular, the dramatic growth in Teams drove the creation of a huge number of Groups. While in some ways Groups continue to be a “better distribution list” for the email-centric community, the role of Groups in membership management across Microsoft 365 is its more important function.

## An Identity and Membership Service

The basic idea behind Groups is to bring people, information, and applications together to enable better communication and collaboration. We can summarize how Groups work in three steps:

1. Users create and manage groups using their client of choice. The client depends on the purpose for which they want to use the group. For instance, if they believe that group members prefer to communicate via email, they should create the group with Outlook. If they prefer using chat or threaded conversations, they might choose Teams or Yammer.
2. The new group object exists in both Azure AD and Exchange Online. Azure AD is the *system of record* and is responsible for synchronizing information about the new group to other workloads across Microsoft 365 (for example, Planner and SharePoint Online) to make those applications aware that a new group is available. In some respects, because both types of group control access to resources, you can think of a Microsoft 365 group as having some of the same characteristics as a security group. The big difference between the two types is that a Microsoft 365 group has both members and owners while a security group has just members. In addition, a workload must recognize Groups as a valid security mechanism before Groups can control access to resources.

The Exchange Online directory holds email-related information about Microsoft 365 Groups, including their proxy addresses. When you use Exchange clients or cmdlets from the Exchange Online PowerShell module to update Groups, a dual-write mechanism commits the updates simultaneously in Azure AD and Exchange Online. Following a successful update, Azure AD synchronizes the changes with other application directories.

3. Group members select the tools they need to get work done from the range of applications that support Groups. For example, a team might choose to connect a plan via a channel tab.

We will explore how these steps work in practice through the rest of this chapter.

## Applications that use Microsoft 365 Groups

A wide range of Microsoft 365 applications uses Groups as a membership service. Table 11-1 lists some of the major applications which use the Groups service.

<b>Group</b>	<b>Focus</b>	<b>Storage</b>
Outlook	Email-based conversations within public or private teams. Limited to "more than" 1,000 members. Dynamic membership supported.	Exchange Online mailbox
Microsoft Teams	Personal chats and channel-based conversations. The membership limit for an individual team is 25,000.	Azure Cosmos DB
Microsoft Planner	Task-based plans used by public or private teams. Shares the same limit as Teams.	Azure data service
Yammer Communities	Open, idea sharing, collaboration at scale. The membership limit for these groups is up to 50K if synchronized to an on-premises AD and higher if not.	Yammer
SharePoint sites	Structured document and information management.	SharePoint
Stream	Controls access to video content in the Microsoft Stream service.	Stream (moving to OneDrive for Business and SharePoint Online)
Power BI	Controls access to workspaces.	Power BI

Table 11-1: The various kinds of groups available within Microsoft 365

The only real limit for groups is imposed by [Azure AD](#), where "any number of objects can be members of a single group," or, if you synchronize with an on-premises directory, an Azure AD group object is limited to 50,000 members.

## Core Principles

Figure 11-1 illustrates how three core principles combine to form the architecture of Microsoft 365 Groups. These are:

**Azure AD is the directory of record.** When you create a new group, two writes occur in tandem. The first is to Azure AD to create the object that becomes the definitive identity for the new group. The second is to the EXODS directory to create the group mailbox. This process ensures that the new group can begin to function immediately. After Exchange Online creates the new group mailbox, it updates the group object in Azure AD with information such as its email address. Provisioning processes create other components like the SharePoint Online document library and the shared OneNote notebook.

Azure AD propagates information about group properties and members using a "forward synchronization" process to the directories used by individual applications to enable those applications to recognize and respect the membership of groups. For instance, if a user updates the membership of a group using OWA, writes occur to update the membership information held in Azure AD and EXODS. Notifications occur to make other group-enabled applications aware of the change. You can think of group membership as a form of "access all areas" pass that is valid for access to all the resources available to the group.

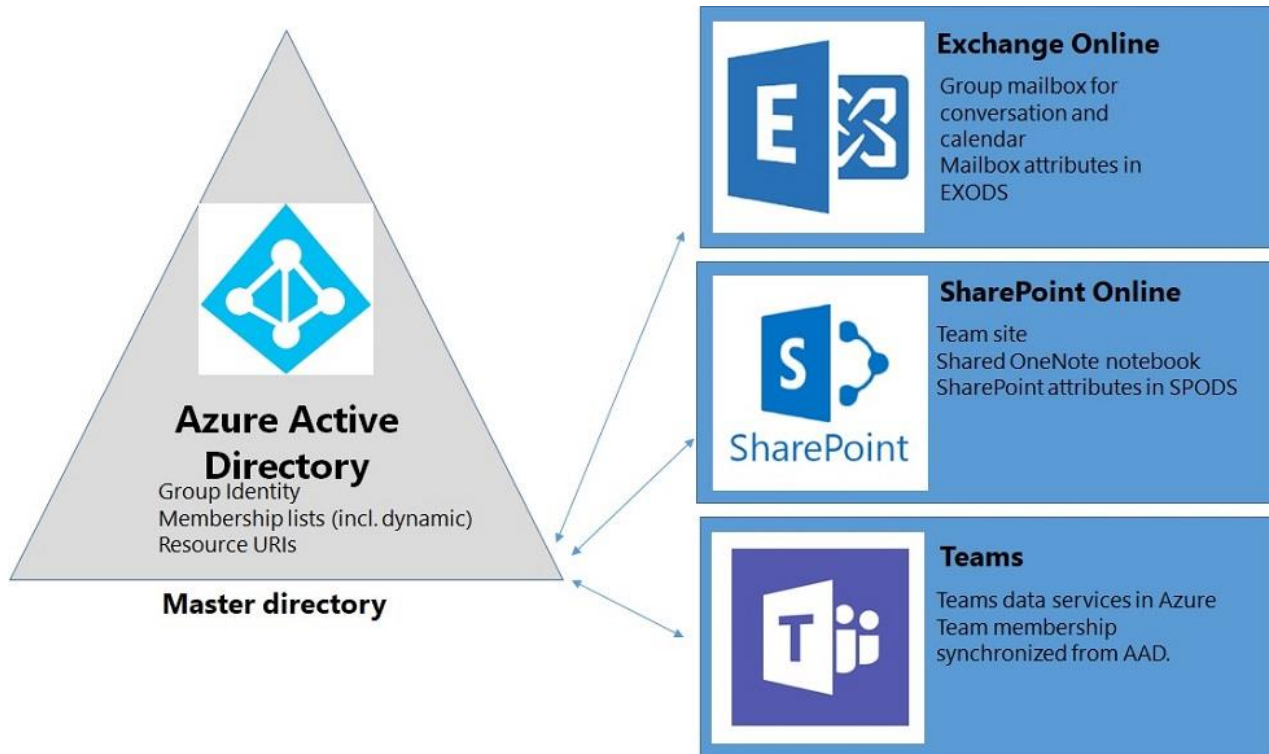


Figure 11-1: The loosely coupled Groups architecture

Although notifications are the primary mechanism used to inform workloads of changes to which they should respond, applications might not respond at once to the notifications as they might have other items to process. The forward synchronization process, which runs in the background on a timed basis, acts as a backstop by making sure that all workloads apply any outstanding updates consistently. The focus of the synchronization process is scale rather than immediacy. In other words, a change might take one or two minutes (or sometimes much longer, depending on the load on the service) before replication fully occurs across all directories. Although the change might not be immediate, it will happen. This is important because as more applications use Groups as a membership service, the interaction between those applications and Azure AD becomes more complex. Typically, if an application has a separate directory (like Teams or Yammer), the application updates both directories and relies on synchronization to spread the change to other workloads.

Viewed in the context of Azure AD, Groups are mail-enabled groups that have members and owners. The Groups service manages membership for apps like Teams and Planner and therefore controls access to the resources available through these apps, some compare Groups to security groups (in the loosest sense).

In hybrid environments, tools like AAD Connect can synchronize group membership to an on-premises Active Directory where the groups show up as distribution lists. You need to pay some care and attention to the synchronization process to make sure that everything works as expected. See the Identities chapter for more information.

**Federated resources are drawn from multiple applications.** Microsoft 365 spans a growing set of applications, each of which delivers specific functionality. Each application controls its data stores. Federation allows Microsoft to select from the functionality available in existing applications and expose that functionality to users through new applications and their clients. Thus, a group has the following components:

- A mailbox to store threaded conversations from Exchange Online. The Microsoft 365 substrate uses the group mailbox for other purposes, such as holding compliance records generated for Teams and Yammer. The group mailbox includes a shared calendar, which Teams uses for both the team and channel calendars.



- A SharePoint Online team site, including a document library to allow group users to share documents and other files.
- A shared OneNote Online notebook. The notebook is in the Site Assets document library of the SharePoint site.
- Other applications can add components to the mix. For example, Teams adds several Azure-based data services to hold its conversations, chats, and graphics.

A Yammer-based group uses similar components, with the difference being that the group mailbox holds only compliance and substrate data. When Teams creates a group, members communicate through chats and only use the group mailbox for the calendar. Stream uses groups to control what videos are available to group members. Likewise, Power BI can use a group to control the workspaces that are available to users.

The single identity for the group means that members automatically gain access to all the application resources available to the group to create an integrated end-to-end user experience across Microsoft 365. Federation allows Microsoft to introduce new applications to Microsoft 365 more easily and with less disruption to other applications than would be the case if each application depends on its resources. Members of the group can then use whatever functionality is available in an application. The functionality and its data remain under the control of each application.

**Loose coupling with groups and applications.** Apart from their dependency on membership and identity services, applications that use Groups operate independently of Groups and can exist independently of each other. Applications are aware of each other because they share the common linkage through Groups and respect the common identity represented by a group. When changes occur in an application, other applications become aware of the changes through a synchronization process. For example, if a user creates a new group with Outlook, the notification that arrives from Azure AD to Exchange Online forces the provisioning of a new group mailbox to hold group conversations and the calendar. The presence of a new group object in Azure AD makes SharePoint Online aware that some work is necessary to update itself when someone tries to use the new group.

Although it is easy to see how Groups can replace some usage of email distribution lists, they cannot replace all distribution lists. For example, you cannot nest a group inside other Microsoft 365 groups. In addition, Exchange distribution lists support all types of email-enabled objects in their memberships, like mail users and mail contacts. Because Microsoft 365 uses groups to control access to applications and data, only Azure AD users and guest accounts can be members.

The fact that a group has a single identity with a single set of permissions is a major part of the value of Groups. It means that groups can act as a means of access to different Microsoft 365 applications, which is how applications like [Forms](#) and Power BI use Groups. Forms, for instance, uses Groups as the way to share forms between users in what it calls "Group forms." Because a group is a single identity, when members join a group, they gain access to all the components available to that group with the same rights to the content as other members. If you need to implement granular access to information, you should use whatever mechanisms are supported by individual applications rather than trying to transform the team identity of a group into a set of individual custom permissions.

## Groups Basics

Before we examine the details of how Groups work, we should first consider some of the basic ideas behind how Microsoft views these objects and how this affects their evolution as a major Microsoft 365 component. Here are some of those ideas:

**A focus on self-service:** Microsoft wants Groups to be easy for users to set up on an on-demand basis. By default, anyone can create a new group (as we will discuss later, you can control this through settings in the Azure AD policy for Groups). Users creating objects that consume system resources without oversight fly in

the face of traditional IT practices where IT exerts control over what a user can and cannot do. The idea is that by making groups easy to create, users can collaborate as they see fit and work together in ad-hoc or formal teams. Remember that resources are more abundant inside Microsoft 365 than is the case in most on-premises deployments, so a reduced focus on the need to control resources is an understandable position to take.

**Public by default:** Originally, Microsoft made groups public by default to encourage people to join and take part in as many groups as possible. It seemed like a good idea in 2015. After several years, the experience of deployments and customer feedback moved Microsoft to believe that groups should be private by default. The change aligns Groups with Teams as both apps use private as the default access type. It has the side effect of making groups less discoverable, but it is easy to set a group to be public after creation if you want to follow the original philosophy.

**Simple permissions:** The history of IT is full of stories about people getting into trouble with permissions. Either they do not have the right permissions to do the job, or their account has permissions that are unnecessary, which means that the user ends up accessing data that they should not. Once a user is a member of a group, they have full access to the contents available to that group, no matter what application provides and manages that content. If someone leaves the group, they lose access to the information belonging to the group. It is as simple as that. The downside of this approach is that non-members have complete access to some information held in public groups.

**Sharing is easy:** Users can share personal documents stored in their OneDrive for Business account with other people, including external users (if allowed by the tenant sharing policy). Members of groups can share documents from the group document library with other users and other groups just as easily. Thus, you should never put a document into a group document library unless you want to grant full control over that information to every other member of the group.

**A shared memory:** As discussed before, administrators often bolted public folders to distribution lists to record contributions for later review. Groups incorporate a shared memory from the start in that once someone joins a group, they enjoy full access to all the information that has accumulated in the group since its initiation. Nothing is secret or hidden from group members. The shared memory makes it easy for new members to quickly familiarize themselves with the work of the group without the need for an existing member to send them messages or documents or update them with details of group meetings. Everything is present for the new member to explore.

**A growing experience:** Members might begin by taking part in group conversations, just as they have used distribution lists for decades. Over time they can become more fully involved in the work of the group and begin to use the other facilities to interact with other people. Since their introduction, Microsoft has steadily added new features to Groups. More importantly, as shown in their use by the Teams and Planner applications, Groups have become the keystone for Microsoft 365 collaboration.

# Group Components

## Members and Owners

- The functionality enabled through Groups is based on the single identity represented by the group objects stored within Azure AD. The single identity and simple permissions model create a low-touch need for administration. Membership information exists as sets of links that connect back to user accounts. Three types of members exist:
- **Owners:** These users can update group properties and add members and owners to the group. A user must be a group member before they can be an owner.

- **Members:** These users have full access to all the resources of the group. However, they must connect to the group to access conversations. Members can add new members to a public group or request that an owner adds someone to a private group.
- **Subscribers:** These are members who have opted to receive copies of conversations and group calendar events via email. A group does not have to have any subscribers, which is the normal case for the groups used by Microsoft Teams where conversations occur through chats rather than email. To ensure that they can participate in group conversations, Exchange adds guest accounts to the subscriber list when they join a group. This also applies to groups enabled for Teams.

Group membership can be dynamic or static. You can create dynamic Groups through the Azure AD admin center where the membership types are “dynamic user” (dynamic group) and “assigned” (static or normal group membership). When you create a dynamic group, you must also specify the query for Azure AD to find group members. Dynamic groups can also be created using PowerShell or the Groups Graph API. You cannot change the membership type for a group from dynamic to static afterward. Dynamic groups can be constructed using the membership of other Azure AD groups, including Microsoft 365 groups with assigned or dynamic membership and distribution lists.

Azure AD user accounts can be present in all categories of membership. Guests can only be present as Members and Subscribers and cannot be group owners. Members can elect to take part in conversations via email, in which case they join the set of subscribers. Except for dynamic groups, administrators and group owners manage the three types of members by manipulating the links belonging to the group. You can do this through group-enabled applications like OWA or Teams, the Outlook mobile apps, or by running the set of PowerShell *\*-UnifiedGroupLinks* cmdlets as discussed in the PowerShell chapter.

**No way to Print Group Membership:** No Microsoft application includes a method to print or report the details of the membership of a group. As [explained in this article](#), it's easy to create a report with PowerShell.

## Ownerless Groups Policy

Most of the Microsoft 365 apps prevent the removal of a group's last owner. However, the deletion of an account (for example, when someone leaves the company) might result in some groups falling into an ownerless state. Although this is not a good condition, the group continues to function. New members can continue to join if the group is public but membership requests to join private groups stay unapproved because no owner exists to process requests.

You can check for ownerless groups with PowerShell.

```
[PS] C:\> [array]$Groups = (Get-UnifiedGroup -ResultSize Unlimited | Select DisplayName, ManagedBy) $Groups | ? {$_.ManagedBy.Count -eq 0}
```

After finding an ownerless group, administrators can add a new owner by updating the group through the Microsoft 365 admin center or EAC or by running the *Add-UnifiedGroupLinks* cmdlet.

The Group Ownership Governance policy (aka the ownerless group policy) automates the process of finding and fixing ownerless groups. Configured in the Microsoft 365 Groups section in Org settings in the Microsoft 365 admin center, the policy uses background processes to periodically check for ownerless groups. For each ownerless group, Microsoft 365:

- Examines the group membership and makes a random selection from the active tenant accounts in the membership to become potential group owners. The administrator sets the number of accounts that Groups can invite to become the group owner. This number can be from 1 to 90. Active means that the account performs some activity in the group in the last 90 days.
- Sends an invitation by email to the potential group owners to ask them to take on the role.

- Actions the response of those invited to become group owners.
- Continues the process for the number of weeks (up to 7) set in the policy.

Microsoft 365 captures audit records for all the actions involved in sending out and processing responses to invitations to become group owners.

The governance policy can process all or selected groups. It can also constrain the selection of new group owners to the members of a security group. This feature allows an organization to restrict group ownership to specific accounts instead of having random people chosen (potentially to become owners of very important groups). No special license is necessary if the policy allows random selection, but if opt for restricted selection, all the members of the groups which come within the scope of the policy require Azure AD Premium P1 licenses. See [this article](#) for more information.

## Group Privacy (Access Type)

A group can be private or public (the default for a new group is private). A private group is one where the group owners control the membership. This is the most suitable group to use when group content is confidential. If a group's membership is open, anyone can view its membership, but not the content unless they join the group. Users can ask to join a private group, and Exchange then sends their request to join to the group owners. Any of the group owners can allow or deny a request to join.

You can change the default access type for groups created by Outlook clients by updating the organization configuration. This setting does not affect groups created from non-Outlook endpoints.

```
[PS] C:\> Set-OrganizationConfig -DefaultGroupAccessType Public
```

Public groups are open to any user in the tenant. Anyone can join a public group if they wish. A user can enroll other users into public groups without their permission (the newly enrolled users can then decide whether they want to take part in the group).

You can change the access type for a group using the following methods:

1. **OWA:** Select the group, edit its properties, and switch the type. For more details, see the section about creating and updating groups with OWA later.
2. **Outlook:** Select the group, then **Edit Group** and update the setting.
3. **Teams:** Select a team and use **Edit Team** to change the setting.
4. **Outlook mobile:** Expand the Groups section in resources and select the group, edit its properties, and change the privacy setting.
5. **Planner:** Select a plan and then **Planning settings**, then choose the Privacy setting in the **Group** tab.
6. **EAC** or **Microsoft 365 admin center:** Select the group and edit its settings.
7. **PowerShell:** Run the *Set-UnifiedGroup* cmdlet and set the *AccessType* property to be either Private or Public.

```
[PS] C:\> Set-UnifiedGroup -Identity TestGroup -AccessType Private
```

If your tenant uses sensitivity labels for container management, the sensitivity label assigned to the group controls its privacy setting and you can only update the privacy setting by changing the sensitivity label.

## Email Conversations

Aside from sharing documents, the basic way members communicate within a group is through threaded conversations. The individual items in message threads exist in the Inbox folder of the group mailbox for Outlook-based Groups or the Yammer data store for Yammer communities. In either case, clients create a view to sort the items sent to the group and present them as threaded conversations. Users can contribute to

conversations through the clients or by sending an email. In either case, items flow through the normal transport processing pipeline.

The calendar is the other major folder used in a group mailbox. It functions much like any other calendar in a personal or shared mailbox, with the notable exception that you can't use granular permissions with group calendars.

**Accessing folders in group mailboxes:** Although group mailboxes have a complete set of default folders that can hold items, clients usually only expose the Inbox and Calendar folders. The folders function in much the same way as in other mailboxes. For instance, junk mail sent to a group goes into the Junk Email folder (but no one knows about these messages). The Sent Items folder stores copies of messages generated by clients, such as notifications to users that someone has assigned them a task using Planner. If a group is team-enabled, you also find compliance records for standard channel conversations in the *TeamsMessagesData* folder. If the group is subject to a hold, the folders retain copies of items until the hold lapses. You can view the complete set of folders in a group mailbox by running this command:

```
[PS] C:\> Get-MailboxFolderStatistics -Identity GroupName | Format-Table Name, ItemsInFolder
```

Unless you are curious and want to explore the mailbox, no real reason exists to access the other folders and you can safely leave them alone. However, if you want to, you can access the folders in a group mailbox by adding the mailbox as a shared folder with OWA. OWA treats the group mailbox exactly like a shared mailbox and you can open any of the folders and examine the items it holds.

## Auto-Reply for Group Mailboxes

Group mailboxes support auto-reply messages. This is useful when the need exists to inform people about some special processing for messages that arrive in a group mailbox. For example, to tell:

- Customers about the procedure to process messages sent to a group.
- Internal people to inform them that a group is team-enabled, and all conversations occur in the team. Although most team-enabled groups are hidden from Exchange clients and are not included in Exchange address lists like the GAL, this doesn't stop people sending emails to the SMTP addresses of the groups.
- People when a group (or team) is not in active use or archived.

Setting an auto-reply for a group mailbox is just like setting one for a user mailbox. In this example, we use the *Set-MailboxAutoReplyConfiguration* cmdlet to create an auto-reply. The message sent to external people tells them that the mailbox is not in use. Because the group is team-enabled, the message used for internal people includes a mailto: link to the email address for a team channel. The recipient can click on the link to send their message to Teams.

```
[PS] C:\> Set-MailboxAutoReplyConfiguration -Identity "Office 365 for IT Pros" -ExternalMessage "Sorry, this mailbox doesn't accept email" -AutoReplyState Enabled -InternalMessage 'Please! We use Teams for communication, so send your message to <a href = "mailto:943a5091.office365itpros.com@emea.teams.ms">Teams</a> and it will be dealt with there.'
```

## Moving Personal Email to Groups

Users can move items from personal mailboxes to Groups using drag and drop in either OWA or Outlook. You can move items from any mailbox to which you have access. The moved items go into the inbox folder in the group mailbox. If you reply to the item, the recipient list includes all the recipients from the original message plus the group.

## Group Document Libraries

Apart from conversations and the shared calendar, the most obvious difference between Groups and traditional email distribution lists is the Files functionality available through the SharePoint team site owned by a group. When you create a new group, the provisioning process uses a site template called **GROUP#0** to create a SharePoint Online site for the team site. You cannot change the template to apply other settings, such as assigning a storage quota (which you can update afterward using the *Set-SPOSite* cmdlet). All the limits applied to SharePoint sites apply to the sites used by Groups.

Users can add content to document libraries using the following methods:

- Drag and drop files from any location accessible to your PC (local drives, OneDrive for Consumers, file servers, etc.). However, you cannot drag and drop files from another SharePoint Online site or group document library. We will discuss later how to move files from a SharePoint Online site to a group document library.
- Upload files or the contents of entire folders.
- Create files from scratch using Office Online applications.
- Save files into the document library from desktop Office applications.
- Save attachments received by email into the document library using Outlook or OWA.
- Use the OneDrive Sync client to upload and manage files.

When a file is in a Group's document library, the group controls its permissions. All documents in the library inherit the permissions to allow group members equal access to the content. SharePoint checks permissions when a user requests access to a file by checking the group membership. The storage used by group document libraries counts against the overall allocation for the tenant.

If the group is public, any user in a tenant can join the group and access the files in the document library. It is also important to remember that the group owns the content in the group document library rather than the original authors. If someone leaves a group, all the content they authored or amended remains behind.

**Document Library Regional Settings:** The regional settings for document libraries created for Groups inherit their time zone and locale (country) from the account which creates the group. Group members see date and timestamps for documents as if they were in that time zone. You can change the site's time zone and locale (country) by updating the Regional Settings (in Site Information). Sites created in the SharePoint admin center use the regional settings defined in the Site creation settings.

## Implementing Groups

Microsoft 365 Groups are a platform for user-driven collaboration. In an ideal world, or a small tenant, this approach works, and users create the right number of groups, each used for its assigned purpose. Allowing everyone to create new groups in larger organizations is more problematic because the potential then exists that underused or unwanted groups will linger in the directory and absorb resources. The latter point is debatable because Microsoft controls the resources and tenants do not incur any cost when a user creates a group. In general, it is better to set out a well-defined set of rules to govern the creation of groups and communicate them to users upfront than to attempt to apply rules retrospectively and control a sprawl of unwanted groups.

Here are some simple questions to help show the need for a new group:

- **What purpose will the group serve?** If the person who wants to create a group cannot say why the group is necessary and how it differs from existing groups, then there is a fair chance that the new group is unnecessary. Sometimes you cannot avoid creating a new group, as in the case when a Planner creates a new plan (from scratch or by copying an existing plan), which results in the

automatic creation of a group to support the membership and conversations for the plan. See the Tasks chapter for more information about Microsoft Planner.

- **Who is the intended audience for the group?** Is the group to be public (anyone can join it) or private (restricted to a defined set of users)? Does the group need to support guest access?
- **What app will people use to access the group?** The available options include Outlook, Teams, and Yammer. Each app brings its own set of advantages and disadvantages to the table. The best answer depends on factors like the organization's culture, strategic direction for collaboration, and above all, the desired functionality. For example, an Outlook-based group is a good choice for a team that needs to receive and process inbound communications from customers. A team is a good choice for small working groups who need to work on projects. A Yammer community works well for cross-organization communication, especially when the organization is very large. The thing to remember is that it is important to select the right technology to meet the need of those who will use it instead of trying to force one answer for all.

## Group Controls

The default situation is that anyone in a tenant can create a new group. In the following pages, we'll get into the details of how to apply different levels of controls to restrict the creation of groups. For now, we can summarize by saying that controls over group creation are available at these levels.

**Azure AD (directory):** the Groups section of the Azure AD admin center has a setting to control if users can create new Microsoft 365 groups. The *Users can create Microsoft 365 groups in Azure portals, API or PowerShell* setting is in the General settings section for Groups in the Azure AD admin center. By default, the setting is On. If it is turned off, users won't be able to create new Microsoft 365 groups even if allowed by the Azure AD directory policy for Groups. A similar setting in the Azure AD admin center controls the ability of users to create new security groups.

**Microsoft 365 (all workloads):** The Azure AD directory policy for Groups is a method to control different aspects of groups, one of which is to disable the ability of all users to create new groups and replace this with a restricted set of people permitted to create groups. A security group is usually the best method to define the set of users. Restricting group creation is an Azure AD Premium P1 feature.

**App-level:** Outlook clients use a setting in the OWA mailbox policy to define if a user can create a new group through Outlook (including OWA and Outlook mobile). This control does not require any additional licenses.

It's up to each organization to decide what level of control they wish to exert over group control, including the use of third-party software to impose group creation request and approval workflows. You can also implement a group creation workflow using a Power Apps flow (many templates exist for this task).

## Azure AD Policy for Groups

The Azure AD policy for Groups stores several settings to control different aspects of Groups. Applications retrieve settings from the policy to know how they should interact with Groups. Table 11-2 lists the policy settings.

<b>Policy setting</b>	<b>Use</b>
<i>AllowToAddGuests</i>	Controls whether users can invite external accounts to become guest members of Groups (and Teams). By default, this setting is <i>True</i> , meaning that group owners can add guests.
<i>AllowGuestsToAccessGroups</i>	Controls whether guests can access content held in the SharePoint site belonging to Groups. By default, the value is <i>True</i> .
<i>AllowGuestsToBeGroupOwner</i>	Controls whether guests can be group owners. This setting is not operational.

<i>ClassificationDescriptions</i>	Comma-separated descriptive pairs for the classifications available in the tenant.
<i>ClassificationList</i>	Stores a comma-separated list of valid classification values to give users a visual indication of what is stored in different Groups.
<i>CustomBlockedWordsList</i>	A list of words that the tenant does not wish group owners to use in the names of new groups.
<i>DefaultClassification</i>	Sets the default classification to apply to new groups.
<i>EnableGroupCreation</i>	Controls whether users can create groups. By default, this value is <i>True</i> and any user can create a new group.
<i>GroupCreationAllowedGroupId</i>	Points to the object identifier (GUID) of a security, distribution, or Microsoft 365 group whose members can create new Groups. This group is used if the <i>EnableGroupCreation</i> setting is <i>False</i> .
<i>GuestUsageGuidelinesUrl</i>	Defines a URL displayed in client user interfaces to remind users about guidelines for how to share company information with guests. For example, <a href="http://mydomain.com/GuestUserAccessO365Groups.html">http://mydomain.com/GuestUserAccessO365Groups.html</a>
<i>PrefixSuffixNamingRequirement</i>	Used by the group naming policy to define if it should apply a prefix and a suffix to the display names of new groups.
<i>UsageGuidelinesUrl</i>	Defines a URL displayed in client user interfaces to give guidelines about the creation and usage of Groups. For example: <a href="http://mydomain.com/HowtoUseO365Groups.html">http://mydomain.com/HowtoUseO365Groups.html</a>
<i>EnableMIPLabels</i>	Tells Groups to use sensitivity labels instead of text-based classifications.

Table 11-2: Settings in the Azure AD Groups policy

While you can configure the *AllowGuestsToBeGroupOwner* setting to *True* to allow guests to become owners of groups, this capability is not exposed in the Groups or Teams user interfaces.

**Get Azure AD Right Too:** It's amazing how often you need to look up Azure AD for something to do with a group. For instance, you might want to know where the owner of a group works or who their manager is. If the directory isn't populated accurately, you'll get bad results. It is always best to ensure that the information held in Azure AD is as accurate as possible (it will never be 100%) before building any script or other procedure that depends on a directory lookup.

## Licensing Groups Functionality

If you use Groups or Teams without changing any setting in the Azure AD Groups policy, the only license needed is the one included in plans like Exchange Online, Teams, or Office 365 E3. Things become more complex if you decide to utilize some of the more advanced features enabled and controlled through policy settings because Microsoft requires accounts to have Azure AD Premium P1 licenses to benefit from the added functionality. Microsoft's view is that these features reduce administrative effort or make groups more productive for organizations, so they price the features accordingly as extensions of the basic Azure AD capabilities bundled with every tenant.

The general rule is that any account that benefits from a feature must be licensed. Sometimes this rule is easy to understand, as in the case of dynamic Groups where the requirement is to license the accounts found by the queries used to populate group membership. For instance, if you create a dynamic group based on a query that finds all users with mailboxes in the tenant, it means that you need licenses for all those accounts. In other cases, the logic is harder to follow. For example, if you implement a policy to control the creation of new groups, you must license every account in the group allowed to create new groups plus the administrators who manage the policy. On the other hand, if you use the expiration policy to control how long groups exist before their owners must reconfirm that the groups are needed, you only need licenses for the members belonging to the groups to which you apply the expiration policy.



Table 11-3 summarizes the licensing requirement for the features you can activate or control through the Azure AD Groups policy. See this [support article](#) for more information.

Feature controlled by a group policy setting	Requires Azure AD Premium P1
Allow non-admin users to create groups	No
Restrict the ability to create new groups to a defined set of accounts	Yes
Use the expiration policy to force the renewal of groups after a set period	Yes
Create and use dynamic Groups	Yes
Apply group naming policy	Yes
Apply default classification to new groups	Yes
Provide URL with usage guidelines to internal users	Yes
Provide URL with usage guidelines to guests	Yes
Allow guests to access groups	No
Force new groups to have classification selected from a predefined list	Yes

Table 11-3: License requirements for features controlled by group policy settings

The requirement for Azure AD Premium licenses to use certain features makes them less attractive than if Microsoft bundled the functionality into the standard plans as it does for Microsoft 365 Business Premium plan. However, you might have other reasons to acquire Azure AD Premium licenses, such as buying subscriptions for the Enterprise Security and Mobility suite or some Microsoft 365 plans, so the problem might not be as big as it seems.

## Guest Access Licenses

[Guest access to Microsoft 365 applications](#) is included in all Microsoft 365 Business Standard and Premium plans and Enterprise subscriptions, which means that you do not need to assign licenses to guest users to access applications like Groups, Teams, and Planner. However, Microsoft's [guidance for B2B collaboration licensing](#) sets out the rules for guest access to "paid" (premium) Azure AD features like conditional access policies. Briefly, Microsoft allows tenants 50,000 free authentication activities for premium activities monthly. Access past this threshold is charged against an Azure subscription taken out by the tenant. See [this article](#) for more information.

## Creating and Updating the Groups Policy

By default, a tenant does not have an active Groups policy, and applications use default settings. To change any settings, we first create a directory setting object to hold the policy settings from the correct template. Azure AD loads default values into the new policy and you can then update those settings. Some of the examples described here use cmdlets in the Azure AD Preview module.

Several directory settings templates exist in the instance of Azure AD used by a tenant, including one to hold group policy settings. A different template is available to control guest access for specific groups (explained later). To see the available templates in a tenant, use the *Invoke-MgGraphRequest* to query the templates API:

```
[PS] C:\> Get-MgDirectorySettingTemplate | ? {$_.DisplayName -like "*group*"} | fl DisplayName, Values, Id
```

```
DisplayName : Group.Unified
Values      : {NewUnifiedGroupWritebackDefault, EnableMIPLabels, CustomBlockedWordsList,
              EnableMSStandardBlockedWords...}
Id          : 62375ab9-6b52-47ed-826b-58e47e0e304b
DisplayName : Group.Unified.Guest
```

```

Values      : {AllowToAddGuests}
Id         : 08d542b9-071f-4e16-94b0-74abb372e3d9

```

The directory template used to create the Groups policy for the tenant has the display name "Group.Unified" while the one used to amend settings for a specific group is called "Group.Unified.Guest." The templates have identifiers of *62375ab9-6b52-47ed-826b-58e47e0e304b* and *08d542b9-071f-4e16-94b0-74abb372e3d9* respectively. The identifiers are the same in all tenants.

To create a new Groups policy to hold customized settings for the tenant, we run the *New-MgDirectorySetting* cmdlet to create the new policy from the template.

```

[PS] C:\> $PolicyId = (Get-MgDirectorySettingTemplate | ? {$_.DisplayName -eq "Group.Unified"}).Id
New-MgDirectorySetting -TemplateId $PolicyId

```

When the *New-MgDirectorySetting* cmdlet runs, it creates a new directory settings object. Any application that can access Azure AD can then check the settings to know what it needs to do to respect the policy. For example, when a user tries to create a new group, the client can verify that the user can take this action. To view the policy settings, run the *Get-MgDirectorySetting* cmdlet. In this case, many of the settings have non-default values.

```

[PS] C:\> Get-MgDirectorySetting | ?{$_.DisplayName -eq "Group.Unified"} | ForEach Values

```

Name	Value
CustomBlockedWordsList	
ClassificationDescriptions	General Usage:Anyone can access,External Access:Available outside the company,Internal Only:Must not be shared with external people,Confidential: Can only be disclosed with management permission
DefaultClassification	General Usage
PrefixSuffixNamingRequirement	
AllowGuestsToBeGroupOwner	False
AllowGuestsToAccessGroups	True
GuestUsageGuidelinesUrl	
GroupCreationAllowedGroupId	A3c13e4d-7083-4448-9224-287f10f23e10
AllowToAddGuests	True
UsageGuidelinesUrl	http://office365itpros.com/GroupGuidelines.html
ClassificationList	General Usage,External Access,Internal Only,Confidential
EnableGroupCreation	False

We will discuss how to update individual settings in the Groups policy using the *Update-MgDirectorySetting* cmdlet over the next few pages. Remember that the names of policy settings are case-sensitive, so be sure to type the names exactly as shown here (for example, "AllowToAddGuests" rather than "Allowtoaddguests"). If you use the wrong name, Azure AD cannot update the policy. Another thing to consider is that Exchange Online distributes mailboxes for a tenant across multiple servers that can run inside different Microsoft data centers, so it can take an hour or so before new or updated settings are active across a tenant. Finally, if you make a mistake and want to start over with a new group policy, you can do so by removing the directory settings object:

```

[PS] C:\> $PolicyId = (Get-MgDirectorySetting | ? {$_.DisplayName -eq "Group.Unified"}).Id
Remove-MgDirectorySetting -DirectorySettingId $PolicyId

```

In the rest of this section, we cover how to update the settings in the Groups policy to control classifications, group creation, naming, and group guidelines.

## Container Management for Groups, Sites, and Teams

Groups, Teams, and SharePoint sites are *containers* for information associated with Groups. Some host internal discussions, some hold confidential information that the tenant wants to keep restricted within the organization, and some exist to share information with guests. To help users understand how sensitive or confidential the information stored within a container is, Groups introduced the notion of classifications. A

classification is a text-only marker (like “Secret” or “Confidential”) assigned to a group and displayed by applications. A text-only marker can deliver a visual reminder to users when they work with important data, but the marker is no more useful than a sticker applied to a paper document. If you’ve deployed sensitivity labels in the tenant, you can replace classifications with labels. When you assign a sensitivity label to a group, Azure AD stamps settings defined in the label on the group. Today, the label settings for groups cover:

- **Privacy:** The group is private or public.
- **Guest access:** Whether the group owner can add guests to the membership.
- **Unmanaged devices:** Controls how group members can access content in the group’s SharePoint site when they connect with an unmanaged device.

Microsoft is gradually extending the number of label settings that apply to groups, such as external sharing behavior for SharePoint. Sensitivity labels have other settings, such as encryption and marking, but these settings are not applied to containers or to the individual items of data held in the containers. If you want to protect individual documents or messages, you can assign sensitivity labels to those items.

A sensitivity label doesn’t need to have container management settings. Any label without these settings is excluded by client applications when they display the list of labels available for assignment to a group, team, or site. For debugging purposes, you can check the set of labels by entering the following into a browser (replacing *tenant* with your tenant’s name):

```
https://tenant.sharepoint.com/_api/GroupSiteManager/GetGroupCreationContext
```

A tenant can choose to use classifications or sensitivity labels to mark groups. They cannot use both. In either case, when you apply a classification or sensitivity label to a group, the marking is synchronized across workloads. Clients connecting to the workloads can display the marking prominently in their user interfaces. (Figure 11-2).

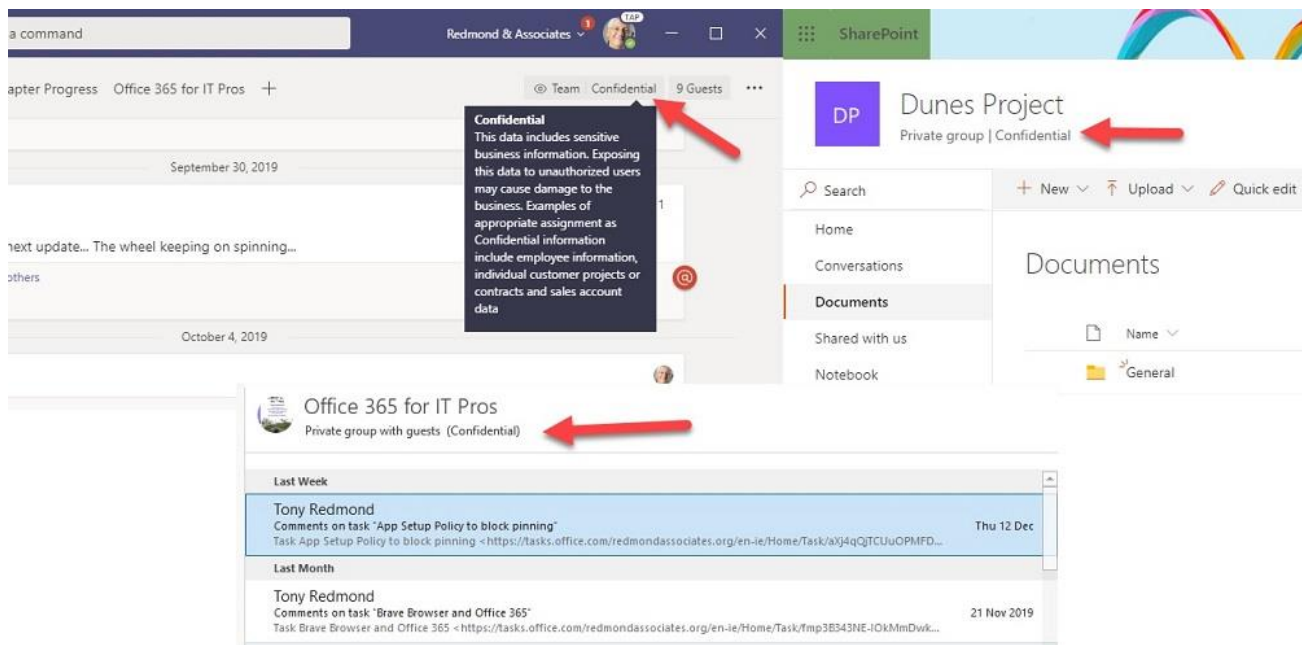


Figure 11-2: Sensitivity markings in Teams, SharePoint, and Outlook

To assign or update sensitivity labels for groups you can use:

- **Client applications:** OWA, Outlook, Outlook mobile, Teams, Teams mobile, SharePoint browser client.
- **Administrative portals:** SharePoint admin center (you can add a sensitivity column to see the labels applied to sites), Azure AD admin center, Teams admin center.
- **APIs:** PowerShell (*Set-UnifiedGroup* and *Set-SPOSite*). Support is also available to [assign sensitivity labels to Groups via the Graph API](#).

To use sensitivity labels instead of classifications, follow the directions in the Information Protection chapter to enable sensitivity label support for SharePoint Online. Then, update the groups policy to enable sensitivity labels.

```
[PS] C:\> $TenantSettingsId = (Get-MgDirectorySetting | ? {$_.DisplayName -eq "Group.Unified"}).Id
$templateId = (Get-MgDirectorySettingTemplate | ? {$_.DisplayName -eq "Group.Unified"}).Id
Update-MgDirectorySetting -TemplateId $TemplateId -DirectorySettingId $TenantSettingsId -Values
(@{'name'='EnableMIPLabels';'value'='true'} | ConvertTo-Json)
```

Like updates to any policy, changes can take some time to become active in all applications. Before making the switch, make sure that you have defined an appropriate set of sensitivity labels to use with groups.

## Creating Group Classifications

If you don't choose to use sensitivity labels, you can use text-only classifications as visual markers for groups. The *ClassificationList* setting in the Groups policy holds a list of classifications defined as a comma-separated list within quotes without spaces between the values. For example:

*"Confidential,External Access,Top Secret"*

The PowerShell commands shown below demonstrate how to update the classification settings for a tenant. The following steps occur:

- Fetch the existing settings.
- Update the *DefaultClassification* setting to be "Confidential." The default classification must exist in the classification list. Azure AD applies the default classification to a new group if the owner does not select another classification. Using a default classification is a feature requiring Azure AD Premium P1 licenses.
- Update the variable holding the *ClassificationList* setting with a comma-separated list of classifications. Some tenants use over 100 classifications, but it is often wiser to have a more restricted set so as not to confuse users.
- Update the variable holding the *ClassificationDescriptions* setting with more expanded explanations of what each classification means.
- Run the *Update-MgDirectorySetting* cmdlet to update the Groups policy with the new settings.

```
[PS] C:\> # Fetch current directory setting and extract values
$TenantSettings = Get-MgDirectorySetting | ? {$_.DisplayName -eq "Group.Unified"}
$Values = $TenantSettings.Values
($Values | ? Name -eq 'DefaultClassification').Value = "Confidential"
($Values | ? Name -eq 'ClassificationList').Value = "General Usage,External Access,Internal
Only,Confidential"
($Values | ? Name -eq 'ClassificationDescriptions').Value = "General Usage:Anyone can
access,External Access:Available outside the company,Internal Only:Must not be shared with external
people,Confidential:Can only be disclosed with management permission"
Update-MgDirectorySetting -DirectorySettingId $TenantSettings.Id -Values $Values
```

## Assigning and Changing Classifications

Usually, when creating a group, the owner selects the value that best describes the information held in the group from the list defined in the policy. Placing a classification on a group is optional and does not affect how the group works. Adding or updating a new classification or removing a classification from the list does not affect classifications placed on existing groups.

The *ClassificationDescriptions* setting is associated with the classification list. Although you do not need to provide descriptions, it is usually a good idea to give people a hint about which classification is the most appropriate for a new group or team. The classification might be enough, but if not, you can give some extra information for each classification in the form of a short description. Azure AD will not complain if you input long descriptions (say, more than 100 characters), but the space reserved to display descriptions in the user

interface of different applications might, so it is best to keep descriptions short and snappy. Make sure to include a description for each classification in the *ClassificationList* setting.

You can find the classifications assigned to groups with a simple PowerShell command:

```
[PS] C:\> Get-UnifiedGroup | ? {$_.Classification -ne $Null} | Format-Table DisplayName, Classification
```

To change a group's classification after creation, tenant administrators or group owners can edit the property through a client. For example, in the Teams desktop or browser clients, use **Edit team** and then change the setting. In the SharePoint Online browser app, the classification is available in **Site information**, while in OWA and Outlook you can change the classification through the **Edit group** option. Alternatively, use the *Set-UnifiedGroup* cmdlet to make the change. Make sure that you use a value from the list in the policy as otherwise, the cmdlet will fail:

```
[PS] C:\> Set-UnifiedGroup -Identity Hydra -Classification "General Usage"
```

Groups created with PowerShell do not have a classification unless one is specified when running the *New-UnifiedGroup* cmdlet. To fix the problem, you can run this command to assign a default classification to all unclassified groups:

```
[PS] C:\> Get-UnifiedGroup | ? {$_.Classification -eq $Null} | Set-UnifiedGroup -Classification "Internal Only"
```

## Switching from Classifications to Sensitivity Labels

No automated migration process is available to replace classifications assigned to groups with sensitivity labels. To start, you need to understand how classifications are used within the organization and figure out which sensitivity label is the most appropriate replacement. After you know how to replace the classifications, you can update groups manually or write some PowerShell to read classifications from groups and assign the most appropriate label. The code below uses a simple *Switch* statement to select the label to assign based on the classifications assigned to a group. After selecting the label, the script calls the *Set-UnifiedGroup* cmdlet to update the group. The classification for each group remains unchanged.

```
[PS] C:\> # Switching starts...
$PublicLabel = "2fe7f66d-096a-469e-835f-595532b63560"
$InternalLabel = "27451a5b-5823-4853-bcd4-2204d03ab477"
$SecretLabel = "81955691-b8e8-4a81-b7b4-ab32b130bff5"
$ConfidentialLabel = "1b070e6f-4b3c-4534-95c4-08335a5ca610"
# Find groups in the tenant that haven't already been assigned a sensitivity label
$Groups = Get-UnifiedGroup -ResultSize Unlimited | ? {$_.SensitivityLabel -eq $Null}
If ($Groups.Count -eq 0) { Write-Host "Congratulations - you've switched over to sensitivity labels"
}
Else {
  ForEach ($Group in $Groups) {
    Switch ($Group.Classification)
    {
      "General Use"      {$LabelToApply = $InternalLabel}
      "External Access" {$LabelToApply = $PublicLabel}
      "Internal Only"   {$LabelToApply = $SecretLabel}
      "Confidential"    {$LabelToApply = $ConfidentialLabel}
      Default           {$LabelToApply = $InternalLabel }
    }
    Write-Host "Processing" $Group.DisplayName
    Set-UnifiedGroup -Identity $Group.DistinguishedName -SensitivityLabelId $LabelToApply
  }
}
```

As you can see, we define variables to hold the GUIDs for several sensitivity labels. You can discover the GUIDs for labels by running the *Get-Label* cmdlet (after connecting to the compliance endpoint).

It takes a little while for the new label settings to synchronize from Exchange Online to SharePoint Online and Teams. To check that the right label is assigned to a site, you can run the *Get-SPOSite* cmdlet and examine the *SensitivityLabel* property. For example:

```
[PS] C:\> Get-SPOSite -Identity https://office365itpros.sharepoint.com/sites/BankingTeam | Format-Table Title, SensitivityLabel
```

```
Title          SensitivityLabel
-----          -
Banking Team   f5b1ba01-59f5-4ba0-b73b-f60e348cdc6e
```

## Controlling Group Creation

If a tenant allows unrestricted creation of Groups, each licensed user in the tenant can exploit twenty-two separate ways available across a spectrum of clients and applications to create up to 250 groups. A thousand-user tenant therefore might end up with 25,000 user-created groups to add to the groups created by administrators. Allowing people to create the groups they think they need through a self-service model is in line with Microsoft's original view that control over groups should be user-led. Although the model gives users the power to decide what groups they need to collaborate with their colleagues, it's easy to see how chaos can result and why, in most cases, it is better when administrators put some thought and planning into the process of group creation and maintenance. This is especially true in large enterprises.

Previous experience with user-controlled creation of shared objects, such as the mayhem that often occurred in public folder hierarchies when anyone could create top-level folders in the early versions of Exchange, proves that if you allow users free rein to create new objects, you can expect a rapid expansion of those objects, many of which duplicate the purpose of other objects. The usual upshot of creating many groups for no good reason is user confusion. People do not know which groups to use for what reason, especially if you do not implement a naming policy and groups with duplicate names arise from Teams, Yammer, and Outlook. This can lead to chaos in the GAL with groups scattered amongst other mail-enabled recipients. Creating groups often leads to groups used for a period and then left to decay. It is to impose order on potential madness that tenants decide to exert control over who can create new groups. The group expiration policy helps to clean up unused groups and should be used whenever possible (if you have the necessary Azure AD Premium licenses). We discuss how the expiration policy works later.

Apart from the setting in the Azure AD admin center described above, administrative interfaces are not subject to the same controls over group creation as imposed on clients. Users holding administrative roles for the tenant, including Exchange Administrator, can create new groups using PowerShell or an administrator portal even when blocked by application-level controls. On the other hand, those holding administrative roles cannot create groups using clients like OWA or Teams unless they are allowed by policy. The administrative roles that can create new groups are:

- Global Administrator.
- Team Service Administrator (in the Teams admin center or Azure AD admin center).
- User Administrator (in the Microsoft 365 admin center or Azure AD admin center).
- Exchange Administrator (in the Exchange admin center or Azure AD admin center).
- SharePoint Administrator (in the SharePoint admin center or Azure AD admin center).
- Directory Writers (in the Azure AD admin center).
- Groups admins (in the Microsoft 365 admin center and Azure AD admin center).

In addition, groups created through administrator interfaces do not come within the scope of the group naming policy.

When a group creation policy is in place, Teams users cannot create new teams unless they're allowed by policy. However, if they are owners of Microsoft 365 Groups that are not team-enabled, they can create new teams from those groups.

## Using the Azure AD Policy to Control Group Creation

Follow these steps to create the settings to control the creation of new groups through the Groups policy.

**Create a group for the set of authorized users:** You need to know the set of users who can create new Groups. To define the set of authorized users, you create a group to hold their names. This can be a distribution list, a security group, or a group. A security group is the best choice for two reasons. First, if you create new groups through SharePoint, you need a security group to control this process. Second, unless you mail-enable the group, a security group does not appear in the GAL and is therefore less likely to be known to users. If you restrict group creation and do not specify a group for authorized users, only administrators can create new groups.

**Add the set of authorized users to the group:** Because Groups only exist in the cloud, the users who create new groups must have accounts homed in the cloud. Make sure that you add all the people you want to be able to create new groups as members of the security group. Being an owner of the group is insufficient. Include those who hold administrative roles so that they can create new groups from applications. Make sure that the members of the group have Azure AD Premium P1 licenses.

**Prepare to edit the Azure AD policy:** Because no user interface exists for this purpose, you must create the policy and populate its settings using PowerShell. The first step is to retrieve the object identifier for the group that you just created to hold the set of authorized users. This is easily done by running the *Get-MgGroup* cmdlet:

```
[PS] C:\> Get-MgGroup -Filter "displayName eq 'GroupCreationControl'"
```

Note the value of the *Id* property for the group you want to use.

**Use PowerShell to update the Azure AD policy:** Two controls are used for group creation:

- Set the **EnableGroupCreation** setting to *False* to show that only authorized users can create groups. To reset and allow anyone to create groups, update the setting to *True*.
- Populate the **GroupCreationAllowedGroupId** setting with the object identifier of the group holding the list of authorized users. This is the *Id* property retrieved using the *Get-MgGroup* cmdlet as described above. No check is done to establish what the identifier points to when you use it in a policy setting, so it is easy to make a mistake and end up in a situation where only an administrator can create a group. For this reason, you should avoid typing the identifier in manually and instead copy the value from the group properties shown in the Azure AD admin center or use PowerShell to retrieve the value.

**Members Only!** It's critical to understand that only the members of the nominated control group can create new Microsoft 365 groups. A common mistake is for an administrator to create the control group (and so become the owner) and then assume that they can create new groups. Administrators can create new groups, but only through administrative interfaces. They will be unable to create new groups using applications like Teams because they are not a member of the control group.

Together, the two policy settings tell applications that only the members of the authorized group can create new groups. In this example, we fetch the object identifier for the new group used to control group creation and the current settings of the Groups policy. The values are placed into variables and used to update the two settings with the necessary values. You can retrieve the group identifier of the security group with the *Get-MgGroup* cmdlet as described previously. We then write the new values for the settings back into the policy.

```
[PS] C:\> $GroupId = (Get-MgGroup -Filter "displayName eq 'GroupCreationControl').Id
$TenantSettings = Get-MgDirectorySetting | ? {$_.DisplayName -eq "Group.Unified"}
$Values = $TenantSettings.Values
($Values | ? Name -eq 'EnableGroupCreation').Value = "false"
($Values | ? Name -eq 'GroupCreationAllowedGroupId').Value = $GroupId
Update-MgDirectorySetting -DirectorySettingId $TenantSettings.Id -Values $Values
```

After making the changes, you can check that the settings are as you expect. This code retrieves the name, identifier, owners, and current membership of the group used to control group creation.

```
[PS] C:\> $Values = (Get-MgDirectorySetting | ?{$_.DisplayName -eq "Group.Unified"}).Values
$GroupId = $Values |?{$_.Name -eq "GroupCreationAllowedGroupId"} | Select -ExpandProperty Value
[array]$Owners = (Get-MgGroupOwner -GroupId $GroupId)
[array]$OwnerNames = $Null
ForEach ($Owner in $Owners) {
    $OwnerNames += (Get-MgUser -UserId $Owner.Id).DisplayName }
[string]$OwnerNames = $OwnerNames -join ", "
Write-Host ("The name of the group defined to control group creation is {0} and its identifier is {1}. Its owners are {2}." -f (Get-MgGroup -GroupId $GroupId).DisplayName, $GroupId, $OwnerNames)
Write-Host ""
Write-Host "The accounts allowed to create new Microsoft 365 groups are:"
$Members = Get-MgGroupMember -GroupId $GroupId
ForEach ($Member in $Members) {(Get-MgUser -UserId $Member.Id).DisplayName }
```

Exchange Online stores details of some group policy settings in its organization configuration, so the following command also reveals who can create new groups:

```
[PS] C:\> $Members = Get-MgGroupMember -GroupId (Get-OrganizationConfig).GroupsCreationWhiteListedId
ForEach ($Member in $Members) {(Get-MgUser -UserId $Member.Id).DisplayName }
```

**Test that the new policy works:** A user included in the membership of the authorized user group should be able to create new groups from any application that supports the policy, including Planner, Teams, Dynamics, Stream, SharePoint Online, and OneDrive for Business as well as the Outlook and Teams mobile apps. Applications signal errors when a user who is not allowed to create a group tries to create a new group (something like *"The group couldn't be created. Your admin hasn't given you permission to create a new group"*).

**Lack of Granularity in the Creation Policy:** The use of a single policy to control group creation makes things simple. The downside is that once you restrict creation, the decision applies to all applications that use Groups for membership services. For instance, you might want to control the creation of new Plans but not Teams, but the same policy applies to both applications, as it does to Stream, SharePoint, and the other group-enabled applications. You should consider this point when you decide to apply restrictions to group creation.

A script demonstrating how to control the group creation settings in the Azure AD policy [is downloadable from GitHub](#).

## Changing the Group for Authorized Group Creators

Over time, you might want to swap the group specifying the set of users allowed to create groups for another group, perhaps for testing purposes. You can change the group specified in the policy by updating the *GroupCreationAllowedGroupId* setting in the policy with the identifier for the group. Like the code used to establish a group to control creation, we retrieve the identifier for a group called "Authorized Users" and use it to update the policy.

```
[PS] C:\> $GroupId = (Get-MgGroup -Filter "displayName eq 'Authorized Users').Id
$TenantSettings = Get-MgDirectorySetting | ? {$_.DisplayName -eq "Group.Unified"}
$Values = $TenantSettings.Values
($Values | ? Name -eq 'GroupCreationAllowedGroupId').Value = $GroupId
Update-MgDirectorySetting -DirectorySettingId $TenantSettings.Id -Values $Values
```



## OWA Mailbox Policy and Group Creation

Originally, OWA was the first client to support Groups and the *GroupCreationEnabled* setting in the OWA mailbox policy was the original mechanism to control group creation. This was acceptable when only Exchange Online clients could access Groups. Although many other applications now use Groups, the OWA mailbox policy is still in force today. The rule is as follows:

- Users allowed by the Azure AD policy for groups to create groups can always do so in non-Outlook apps (like Teams and Planner).
- Users allowed by the Azure AD policy for Groups to create groups can only do so in Outlook or OWA if the *GroupCreationEnabled* setting in the OWA mailbox policy assigned to their mailbox is *True*.

In other words, if you want to create groups through OWA or Outlook, the OWA mailbox policy assigned to your mailbox must have *GroupCreationEnabled* set to *True*. To understand the situation with policies and mailboxes in a tenant, this code looks for OWA mailbox policies that allow group creation and then reports the set of mailboxes assigned any of those policies:

```
[PS] C:\> [array]$OWAPolicies = Get-OWAMailboxPolicy | ? {$_.GroupCreationEnabled -eq $True} |
Select -ExpandProperty Identity
Write-Host ""
Write-Host "OWA Mailbox policies allowing group creation:"
Write-Host ""
$OWAPolicies
[array]$Mailboxes = Get-CasMailbox | ? {$_.OWAMailboxPolicy -in $OWAPolicies } | Select DisplayName,
OWAMailboxPolicy
Write-Host ""
Write-Host "The OWA Mailbox policy assigned to these mailboxes allows them to create Microsoft 365
Groups:"
$Mailboxes
```

The use of the OWA mailbox policy has persisted for longer than expected. It is now an anachronism in the context of modern group management. Nevertheless, while the policy setting persists and is respected by clients, it's important to make sure that the correct values exist in the policies assigned to users who are allowed to create groups by the Azure AD Groups policy. One way to do this is to assign an OWA mailbox policy with *GroupCreationEnabled* set to *True* to the members of the group allowed to create groups. Here's how to do this with PowerShell. The steps are:

- Find the group defining the accounts that can create groups set in the Azure AD Groups policy.
- Find the membership of the group.
- Loop through the membership and use the *Set-CASMailbox* cmdlet to update their mailbox with an appropriate OWA mailbox policy. In the example below, the OWA policy called "OWAFullAccess" has *GroupCreationEnabled* set to *True*.

The code is:

```
[PS] C:\> $TenantSettings = Get-MgDirectorySetting | ? {$_.DisplayName -eq "Group.Unified"}
$GroupId = $TenantSettings.Values | ? {$_.Name -eq "GroupCreationAllowedGroupId"} | Select -
ExpandProperty Value
$AuthorizedUsers = (Get-MgGroupMember -GroupId $GroupId | Select -ExpandProperty Id)
ForEach ($Mbx in $AuthorizedUsers) {
    Set-CASMailbox -Identity $Mbx -OwaMailboxPolicy OwaFullAccess }
```

Remember that this is a one-time operation. If you update the membership of the Azure AD group defined in the policy, you'll need to run the script again to make sure that all members have the appropriate OWA mailbox policy.

## The Groups Admin Role and Group Creation

The *GroupCreationAllowedGroupId* setting in the Azure AD policy for Groups allows tenants to dictate which users can create new groups. However, being allowed to create new groups by policy does nothing to allow

people to manage groups thereafter, so this work had to be done by tenant administrators. The situation is manageable in small tenants but becomes more problematic as the number of groups grows.

The Groups admin role is designed to solve the problem. This is a standard Microsoft 365 administrative role that can be assigned to user accounts through the Microsoft 365 admin center, Azure AD admin center (where the role is called Groups Administrator), or PowerShell. When assigned, the Groups Admin role allows the holder to manage the following actions for Groups:

- Create, edit, delete, and restore Microsoft 365 groups and Azure AD security groups.
- Create, edit, and delete group creation, expiration, and naming policies.

Groups admins can manage groups and group policies through administrative interfaces such as the Microsoft 365 admin center or PowerShell. Holding the role does not allow users to create new groups through client interfaces like OWA. If you want administrators to be able to create groups using all applications, you must add them to the group defined to control group creation in the Azure AD policy for Groups.

An [Azure AD custom role](#) can also be used for fine-grained control over group creation.

## Azure AD Group Naming Policy

Exchange Online has a naming policy that applies to distribution lists. Settings in the Azure AD for Groups policy can similarly control the display name generated for new groups created by applications like Teams and Yammer. You can update the policy with the Microsoft Graph or with PowerShell (see below) or you can use the Azure AD admin center (Figure 11-3).

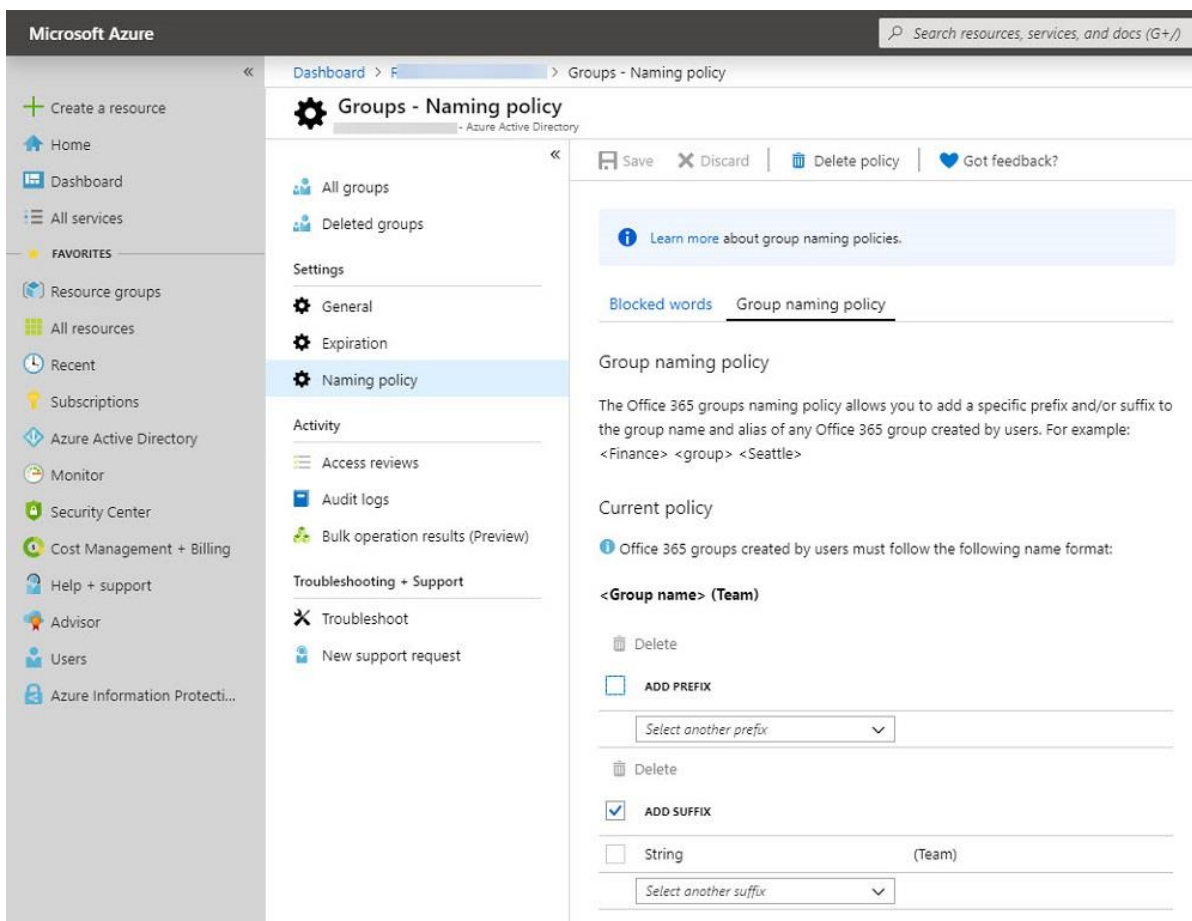


Figure 11-3: Working with the Azure AD policy for Groups in the Azure AD admin center

## Settings for the Groups Naming Policy

Two settings in the Groups policy control the form of display names users can give to groups when they create new groups or update the display name of existing groups:

- The ***PrefixSuffixNamingRequirement*** setting is a pattern controlling whether the group name has a prefix and suffix. The pattern can include fixed strings or reference user attributes such as the department, company, or office of the user who creates the group. Attribute values come from the information held in Azure AD from the account of the user creating the group. You can combine multiple attributes in a display name. Although combining multiple attributes to create a display name might seem like a good idea, it is best to favor simplicity over complexity when selecting what Azure AD attributes to use to construct a display name. The policy ignores values that are missing in the directory. If the tenant uses a naming policy for Exchange distribution lists, it is wise to use different prefixes for the two types of groups. Finally, the display name created for a group can be up to 256 characters. If you tend to use very long names for groups, you must ensure that the application of prefixes and/or suffixes to display names does not exceed the 256-character limit.
- The ***CustomBlockedWordsList*** holds a set of words that users cannot include in group names. These might be words that you consider offensive or have other reasons to exclude. Companies often include business terms that they do not want people to use in group names such as “CEO” to avoid the possibility that users communicate with the wrong person or group. The check against blocked words is case-insensitive and you do not have to enter the blocked words alphabetically. For practical purposes, you can enter as many words as you like in the list as Azure AD sets no upper limit. The Azure Active Portal allows you to upload a CSV containing blocked words to add to the policy.

The naming policy does not apply when administrators create groups. The assumption is that these users know what kind of display name to give to a group.

### Deciding on a Prefix or Suffix for Group Names

The original thinking about group names followed that of Exchange distribution lists and largely preferred using a prefix instead of a suffix to mark groups. The logic behind the choice is that this allowed users to find all the groups together in the Global Address List (GAL) and other address lists.

Useful as this approach is for applications that use the GAL, the reasons for using a prefix are less valid for applications like Teams, Yammer, and SharePoint which also create groups but don't use the GAL. A suffix is often a better approach for these applications because it is less visually intrusive as the marking moves to the end of the screen space reserved by the application to list group names (like the list of teams shown to the left of the Teams client). The choice to use a prefix or suffix is heavily influenced by what applications are in use within the tenant.

### Naming Policy in Action

Here is an example of how to update the settings used for the group naming policy. In this case, we want to add the prefix “O365Grp-” to the names of newly-created groups (and existing groups when a group owner updates their properties). We also specify a set of custom blocked words.

```
[PS] C:\> $TenantSettings = Get-MgDirectorySetting | ? { $_.DisplayName -eq "Group.Unified" }
$Values = $TenantSettings.Values
($Values | ? Name -eq 'PrefixSuffixNamingRequirement').Value = "O365Grp-[GroupName]"
($Values | ? Name -eq 'CustomBlockedWordsList').Value =
"Sexy,Stupid,Giggles,Funny,CFO,CEO,Shit,Payroll"
Update-MgDirectorySetting -DirectorySettingId $TenantSettings.Id -Values $Values
```

A small change is made to use a suffix instead of a prefix. For example, let's assume that we want to add the suffix “ (Team)” to the display names of new groups. The *PrefixSuffixNamingRequirement* setting for the policy is then:

```
$Settings["PrefixSuffixNamingRequirement"] = "[GroupName] (Team)"
```

After a naming policy is in effect, clients apply the policy when users create new groups or edit the display name for an existing group. [Many of the applications](#) which create Groups support the naming policy. Clients check the name entered by the user rather than the name as it will be after the application of the prefix and/or suffix dictated by the policy (what Microsoft refers to as the “decorated” name). Another point to remember is that the check for blocked words is for complete words rather than substrings. This is to avoid problems where a blocked word is part of another legitimate word (think “ass” as part of “class”).

The screenshot shows the 'New group' form in Office Web App (OWA). On the left, there is a 'New group' header and a description: 'Working together on a project or a shared goal? Create a group to give your team a space for conversations, shared files, scheduling events, and more. Read group usage guidelines'. Below this is an illustration of four people's heads. On the right, the form fields are:
 

- Group name:** 'Stupid silly payroll'. Below it, a preview shows 'O365Grp-Stupid silly payroll' with a red error message: 'The name can't contain 'Stupid''.
- Email address:** 'stupidssillypayroll'. Below it, a preview shows 'O365Grp-stupidssillypayroll@Office365ITPros.com' with a red error message: 'The address can't contain 'Stupid''.
- Description:** A text box with the placeholder 'Tell people the purpose of your group'.
- Settings:**
  - Privacy:** 'Private - Only approved members can see what's inside'.
  - Classification:** 'Confidential'.
  - 'Send all group conversations and events to members' inboxes. They can stop following this group later if they want to.'

 At the bottom are 'Create' and 'Discard' buttons.

Figure 11-4: OWA detects a blocked word in the name of a new group

Some differences exist in the detail of how a client applies the naming policy. In general, web-based clients like OWA (Figure 11-4) and Teams generate a preview of the group name after the application of the policy together with warnings when a user types in a blocked word for a group name. Other clients, like Outlook desktop and Outlook mobile clients, enforce the policy without giving so many visual clues as users type in names. Instead, these clients flag errors when they check the policy before trying to create or edit a group.

### Updating Group Display Names with PowerShell

When you implement a naming policy for groups, its effect is not retrospective, and the naming policy will not apply to existing groups. If you want to bring those groups into line with the policy, you must update the display name for older groups. One solution is to use PowerShell to make a one-time pass to find and update non-compliant group names. The example [script available on GitHub](#) extracts the group naming policy for the organization and uses it to figure out the display names for groups that don't currently follow the policy. The example works when either a prefix or suffix is used to create the display names for groups. Because it's PowerShell code, you can include other conditional processing. For instance, team-enabled groups might use one form of display names while those not used with Teams might have another.

If you implement a group naming policy, you might end up with some groups having display names set by policy and others named outside the policy guidelines. For example, if administrators create new groups or teams, the new objects take the display name as entered by the administrator. All other groups or teams in

the organization follow a different convention, meaning that you do not achieve the goal of having a common naming scheme for groups. If this happens, you can use code like that shown above to adjust the names of the non-compliant groups.

## Group Policy Settings in Exchange Organization Configuration

For convenience (mostly for OWA), copies of some group policy settings are held in the Exchange Online organization configuration. You can view these settings by running the *Get-OrganizationConfig* cmdlet:

```
[PS] C:\> Get-OrganizationConfig | Format-List Group*,Data*
GroupsCreationEnabled      : False
GroupsCreationWhitelistedId : 12cb915b-2365-4bed-baf6-6257b3543273
GroupsUsageGuidelinesLink  : Http://office365exchange.com/GroupGuidelines.html
GroupsNamingPolicy         : O365Grp-[GroupName]
DataClassifications        : <DataClassifications
Default="Confidential"><Classifications><Classification Name="General Use"
Description="Anyone can access" /><Classification Name="External Access" Description="Available
outside the company" /><Classification Name="Internal Only" Description="Must not be shared with
external people" /><Classification Name="Confidential" Description="Can only be disclosed with
management permission" /></Classifications></DataClassifications>
```

The group settings in the organization configuration are read-only and cannot be updated with the *Set-OrganizationConfig* cmdlet.

## Creating New Microsoft 365 Groups

If permitted by policy, a user can create new Microsoft 365 Groups through the following interfaces:

- **Admin portals:** The Microsoft 365 admin center, Exchange admin center (which also includes the option to upgrade a distribution list), Teams admin center (create a new team), and Azure AD admin center.
- **Email clients:** OWA, Outlook for Windows or Mac (Microsoft 365 Apps), Outlook for iOS and Android.
- **Applications integrated with Groups:** These applications include Power BI, Dynamics 365, Teams, Stream, Yammer, Flow, and Microsoft Planner.
- **SharePoint Online:** Create a new team site or link an existing team site to a Group; you also have the option to create a group from a hub site.
- **OneDrive for Business:** Create a new group-enabled site.
- **Programmatically:** With PowerShell using the *New-UnifiedGroup* (Exchange Online module), *New-MgGroup* (Microsoft Graph PowerShell SDK), or *New-Team* (Microsoft Teams module) cmdlets; or through the [Microsoft Graph Groups API](#).

The list of possible ways to create a group grows all the time. Depending on the apps and configuration used in your tenant, you might have other methods available for group creation than those listed above.

Administrators most likely create new groups through the Microsoft 365 admin center, EAC, or PowerShell while users are likely to use OWA or Outlook. To create a new Group from the Microsoft 365 admin center, navigate to the **Active teams and groups** page in the **Teams and Groups** section and select **Add a Group**. Make sure that you select *Microsoft 365* as the group type to create. The group creation wizard then gathers basic properties for the new group. These properties are:

- **Group name.** This is the display name visible to users through the directory. Make sure that the name conveys the intention and purpose of the group because it is highly likely that this is the sole information people will use when they decide whether to join the group. It is sensible to check the directory beforehand to ensure that the name you want to use does not clash with an existing object such as another group, a team, a distribution list, or a shared mailbox. You can rename a group after

creation by editing its properties or by running the *Set-UnifiedGroup* cmdlet to update the *DisplayName* property. The display name for a group can be up to 256 characters and can contain periods.

- **Description.** A free-text description to inform users about the intended use of the group. You can add whatever you like here but the text should tell users and administrators why the group exists and its intended use. In a multi-lingual organization, you might like to include translations of a text explaining the group's purpose in multiple languages. If a group holds confidential information, it is wise to make that fact known to group members here.
- **Owners and members.** You must add at least one owner to manage the group and its membership (it's generally better to have at least two owners in case one leaves the organization). You can add some group owners and members when you create the group or build out the group membership afterward by either adding new members or, for dynamic groups, by changing the query used by Azure AD to calculate its membership. You cannot nest a group inside another Microsoft 365 group, nor can you add a traditional distribution list to become a group.
- **Group settings.** Administrative interfaces such as the Microsoft 365 admin center allow administrators to assign an email address to a new group. User clients, such as OWA or Outlook, collect details of the group name for Exchange Online to create a unique SMTP email address and **Alias** (or mail nickname) for the new group. Exchange Online derives the alias by removing any spaces from the group name before checking against the directory to make sure that the value is unique. If not, the user must change something in the alias to make it unique (for instance, appending a number to its end). You can add periods to the alias if the period has other text on both sides. For example, "Sales.Group.1" is OK, while "Sales.Group.1." is not (see the documentation for the [New-UnifiedGroup](#) cmdlet for more information). Exchange Online then combines the default email domain for the tenant with the alias to form the SMTP address, which also must be unique. For example, if the alias is *SalesProfessionals* and the tenant's default email domain is *contoso.onmicrosoft.com*, the SMTP address will be *SalesProfessionals@contoso.onmicrosoft.com*. After creation, an administrator can add more email addresses to the group or change the primary email address. For instance, groups can have an address from a vanity domain. Azure AD prohibits the use of certain well-known or "highly-privileged" reserved aliases to avoid any possibility that users create groups with these words. Prohibiting the assignment of reserved aliases to groups is not unique to Azure AD (Google Workplace also includes this restriction). The reserved aliases include "abuse", "admin", "administrator", "hostmaster", "majordomo", "postmaster", "root", "secure", "security", "ssl-admin", and "webmaster." Only administrators can create groups that include these terms in identifiers (the Microsoft 365 admin center and EAC won't allow administrators to create distribution lists or mail-enabled security groups with a reserved alias).
- **Sensitivity.** If your organization uses sensitivity labels for container management, you can assign a label to the group. As explained in the Information protection chapter, the new group inherits several settings from the label
- **Privacy.** A group can either be Public, which means that anyone can join a group and access the resources available to the group (notebook, mailbox, calendar, document library, plan, and so on), or Private (the default), in which case only nominated members can access the group resources. Azure AD does not mark private groups in any special way within the directory, so they appear alongside public groups. For this reason, it is a good idea to consider giving a confidential group a name that does not create interest (or gossip) on the part of users. For instance, a group called "Planning 2020 Layoffs" is likely to attract unwanted attention! If users want to join a private group, they can apply to do so. A request then goes via email to the group owners, who then decide whether to add the user to the group membership. After creating a private group that needs to remain confidential, you can hide its existence from address lists to prevent the group from appearing in the GAL. Remember that you can change a group access type later.

- **Team-Enabled.** The Settings section includes the option to create a team for the new group. Use this option when a group's function is to manage the membership of a team rather than hosting Outlook conversations, so you should choose it to make the group available in Teams. It's entirely possible to access a group through Outlook and Teams, but it's more usual to choose one or the other. If you choose to team-enable a new group, make sure that the chosen group owners have Teams licenses as otherwise, they won't be able to access Teams. To help, when you come to assign group owners, the owner picker shows you which accounts have Teams licenses.
- **Role assignment:** You can [assign an Azure AD administrative role](#) to the group. Only private groups should hold Azure AD administrative roles. Assignment is permanent and cannot be reversed, so if a mistake is made, you must remove the group and recreate it.

Click **Next** to review the settings for the new group and then **Create group** to complete the process. After a short pause, the new group is ready for use.

## Managing Groups in the Microsoft 365 Admin Center

Some group settings are updatable only after the creation of a new group. These include:

- **Language.** The default language is the one for the user who creates the group. You can change the language at any time. Exchange Online uses the selected language in messages sent to group members, such as in the footer of copies of conversations. It has no impact on the content of group discussions or anything else belonging to the group.
- **Send copies of group conversations to members' inboxes.** The default is "On" for groups created using the Microsoft 365 admin center or an Outlook client. This means that the group adds new members automatically to its subscriber list. Subscribers receive copies of conversations and meetings via email and do not need to access the group mailbox unless they want to. Members can contribute to conversations by replying to the messages they receive. Individual users can disable or enable this setting as they like. Groups do not generate daily or weekly digests of conversations. Apart from guest members, groups used with Teams do not add new members to the subscriber list because their conversations are in Teams.

You can update settings through the Microsoft 365 admin center, Exchange admin center, Outlook clients, or PowerShell. In the Groups section of the Microsoft 365 admin center, you can do the following:

- Filter to only show all groups, team-enabled groups, or dynamic groups.
- Add group owners and members.
- Edit group settings, such as a group's display name, description, and whether to hide it from Exchange address lists.
- Allow external addresses to send emails to a group.
- Assign a new primary SMTP address for a group.
- Enable a group for Teams.
- Recover deleted groups.
- Export details of the groups to a CSV file. If a filter is applied, only the filtered groups are included in the export file.

## Updating Group Membership

After making sure that the new group is set up correctly, the next step is to build the group membership. To do this in the Microsoft 365 admin center, go to the Groups section, and select the group. Under the *Members* tab, you can add group owners and members, including guests.

Using a GUI works well when you need to add a small number of users. Scrolling up and down within a large list to find the right person to add to a group can lead to mistakes, as can searching the GAL to find the right person when they have a common surname. For instance, many organizations have multiple users named

“John Smith” or “Tom Jones.” The Microsoft 365 admin center shows the user principal names for group members, but these can also be similar and hard to decipher. Large tenants often solve the problem by including organizational prefixes like departments in display names to help people find the right user.

OWA automates the process somewhat by allowing group owners to import members from distribution lists or other groups. When this happens, Exchange Online expands the membership of the input group, and adds eligible accounts to the group membership, so it is an effective way to add many members quickly. It is possible to add members to groups programmatically, including the conversion of some types of traditional email distribution lists to become Groups. Updating group membership through PowerShell is not to everyone’s taste, but it can be the most efficient way to approach the task.

**Limits:** All software has limits and Groups are no different. Apart from the limits imposed by the underlying applications such as SharePoint Online storage quotas or the number of items that a folder can store in an Exchange Online mailbox (well over 100,000), Microsoft documents some specific limits for Groups. The most important are:

- A group can have up to 100 owners. The client user interfaces (including PowerShell) will prevent an attempt to add more. If you hit the limit, you must demote an existing owner before you can add a new one.
- An individual user cannot create more than 250 groups. This limit exists to speed up retrieval of the groups owned by an individual from Azure AD. However, the limit could become a factor in large tenants where central control governs the creation of new groups. The simple workaround is to make sure that you reassign ownership for newly-created groups to a different user and remove the account that created the group from its owner list. The limit does not apply to global tenant administrators.
- Outlook Groups with more than 1,000 users are supported with the caveat that access to the group mailbox (for conversations and calendar) will be slow. By contrast, the reduced number of connections to the group mailbox allows team-enabled groups to support 10,000 members.

A tenant can support up to 500,000 Groups. These limits apply to all Groups, including those used with Teams, Yammer, Planner, Dynamics 365, and Power BI. The number can be increased by Microsoft if necessary.

## Welcome Notifications

When you add someone as a member of a group using the Outlook apps or the admin tools (including PowerShell), the system generates a welcome notification for the user and sends it to them via email. The notification tells the person that they are now a member of the group, whether the group is public or private, and how many others are members. A link in the message (View group in Outlook) opens the group conversations using OWA. Exchange Online creates the notification in the language of the user (based on mailbox settings), but the footer information in copies of conversations received by group subscribers is in the language assigned to the group.

The exception to this rule is when a Teams client creates a group to manage the identity and membership of a team. The conversation activity for these groups takes place within Teams, so as users join these teams, Exchange Online generates the welcome message for the team instead of the group. The content of the two types of messages is similar, with the big difference being that the link in the Teams welcome message loads the target team in a browser.

## Creating Microsoft 365 Groups from the Azure AD Admin Center

Although you can create a new Microsoft 365 group from the [Azure AD admin center](#), this is not recommended. To create a group, select Active Directory, then Groups, and then create a new group. Set the membership type to Assigned and the group type to Microsoft 365. After you add members, Azure AD signals



Exchange Online to create the new group mailbox and SharePoint Online to create the new team site. The process of fully-provisioning group resources takes longer than it does when creating groups elsewhere too.

The synchronization process eventually completes, and the new group is available. However, because the Azure AD admin center focuses exclusively on managing Azure AD, it does not update the email-related properties managed in the Exchange Online directory. For example, you cannot assign the primary SMTP address for a new group to anything but the service domain for the tenant. By comparison, when you create a new group with OWA, Outlook, or EAC, you can populate those attributes (like if members receive updates via email) to create a fully-functional group.

## Updating Group Properties

Once the group exists, you can update its properties. Group maintenance operations like updating properties or membership are usually performed through the Microsoft 365 admin center or EAC. However, the admin centers don't support access to some group properties. For instance, you cannot add a MailTip or change the property which hides team-enabled groups from Exchange clients. You can manage these properties through PowerShell, which is one reason why some administrators script the creation of groups to ensure that their properties comply with the standards set for the tenant. For example, this command updates several properties for a new group, including a new primary SMTP address based on a vanity domain, a new alias, modification of the access type for the group, notes about the purpose of the group and allowing external users to contribute to group conversations by sending emails to the group.

```
[PS] C:\> Set-UnifiedGroup -Identity ProductMarketingGurus -PrimarySmtpAddress ProductMarketigGurus@Office365ITPros.com -Alias ProductGurus -AccessType Private -Notes "Product Marketing Discussions - Private Group" -RequireSenderAuthenticationEnabled $False
```

## Non-Latin Group Names

Exchange Online generates the primary email address for new groups based on the group identifier (alias), which is in turn based on the group name. When clients run in non-English languages, group owners can input non-Latin characters in the group name. Behind the scenes, clients generate the identifier and primary email address from the group name. Some clients use the [internationalized domain name](#) (IDN) standard to deal with non-Latin characters. Although this is generally a good thing and creates valid identifiers and email addresses, it can result in the creation of email addresses that are not user-friendly. To avoid the problem, use a group identifier (alias) that only includes basic Latin characters. Alternatively, you can assign a new primary SMTP address to the group post-creation.

## Email Address Policies and Groups

By default, Exchange Online assigns a newly created Group a primary SMTP address belonging to the tenant's default domain plus a secondary address for the service domain (the one ending in onmicrosoft.com). You can assign extra SMTP addresses to groups after creation, but you might want to assign SMTP addresses from separate domains depending on who creates the group.

Email address policies control the primary SMTP addresses Exchange Online assigns to new groups. Any of the known domains for the tenant are usable for this purpose. For example, this email address policy dictates that all new groups receive their primary SMTP address from the Office365itpros.com domain. The groups also receive a secondary proxy address from the service domain:

```
[PS] C:\> New-EmailAddressPolicy -Name Groups -IncludeUnifiedGroupRecipients -EnabledEmailAddressTemplates "SMTP:@office365itpros.com" -Priority 1
```

See [this page](#) for more information. Chapter 3 in the companion volume covers the use of address policies within a tenant.

## Assigning Send As and Send On Behalf Of Permissions to Groups

How to assign the *Send As* and the *Send On Behalf Of* permissions to mailboxes is explained in Chapter 6. You can also grant these permissions to Groups, or rather to the group mailboxes, to let group members (and non-group members) send messages as if they were the group (*Send As*) or on behalf of the group (*Send On Behalf Of*).

You can assign permissions by editing the group properties in EAC or through PowerShell. With EAC, edit the group and go to the group delegation settings. Enter the names of the users to whom you wish to assign the *Send As* and *Send On Behalf Of* permissions and then **Save**. Two cmdlets manage the permissions through PowerShell. The *Add-RecipientPermission* cmdlet assigns the *SendAs* permission and the *Set-UnifiedGroup* cmdlet assigns the *Send On Behalf Of* permission. Here is an example of how to set the two permissions:

```
[PS] C:\> Add-RecipientPermission Microsoft365Gurus@Office365ITPros.com -AccessRights SendAs -Trustee Kim.Akers@Office365ITPros.com
```

```
[PS] C:\> Set-UnifiedGroup -Identity "Microsoft 365 Gurus" -GrantSendOnBehalfTo @{"Add="John.Smith@Office365ITPros.com"}
```

After users have been assigned permissions, they can use [OWA or Outlook to send messages](#) as if they were the group.

## Controlling Group Email Traffic

For new groups created with OWA, the default setting is to subscribe members to the group. This means that Exchange Online distributes copies of any message sent to the group to all members, including tenant users and guests. Members can control the communications they wish to receive through **Manage Group Email** in the cogwheel menu, selecting whether to receive all messages, replies to their posts, calendar events, or no messages. They can also subscribe using the link in the group card displayed by Outlook.

The group's *AutoSubscribeNewMembers* property is *True* to subscribe members or *False* if not. You can change the *AutoSubscribeNewMembers* settings by:

- Editing the group properties with OWA or Outlook
- Updating the **Subscribe members** setting for the group in the Microsoft 365 admin center
- Running the *Set-UnifiedGroup* cmdlet. For example:

```
[PS] C:\> Set-UnifiedGroup -Identity Office365ITPros -AutoSubscribNewMembers:$False
```

Groups created by Teams set the *AutoSubscribeNewMembers* property to *\$False* because members of these groups don't communicate via emails. To see a list of groups that do not distribute copies of conversations and events to members, use the command:

```
[PS] C:\> Get-UnifiedGroup | ? {$_.AutoSubscribeNewMembers -eq $False } | Select Displayname
```

To receive copies of group conversations and events, members join the group's subscribers list. The *Get-UnifiedGroupLinks* cmdlet can report the subscriber list for a group:

```
[PS] C:\> Get-UnifiedGroupLinks -Identity Office365ITPros -LinkType Subscribers
```

No way exists to view the subscribers of a group except via PowerShell. The membership view presented by Outlook clients lists individual users as either an owner or a member and doesn't display any information about subscribers.

When a group is "chatty" with many conversations, some members might decide that they do not wish to receive updates via email and are quite happy to access the conversations in the group mailbox instead. Tenant users can control the updates they receive for a group through email through settings available in OWA, Outlook for Windows, Outlook for Mac, and Outlook mobile. Guests can ask a group owner to remove

them from the group's subscriber list if they do not want to receive updates. Even after removing a group member from the subscriber list, they continue to receive any conversations or replies posted to the group where replies include the special @ mention in the text (for instance, @Tony). The same is true when you explicitly add someone to the TO: or CC: list of a message sent to a group. Remember that the settings to control updates sent by email are specific to a group. If you want to stop updates arriving in your Inbox for every group you belong to, you must remove your account from the subscriber list for each group.

In OWA, the group card allows users quick access to different group resources (like conversations and files) and settings including the ability to control subscriptions. To access a group's card, find a message sent to the group and click on the group name in the message header. You can then move the **Follow in inbox** slider in the group card from **Off** to **On** (Figure 11-5) to control whether you receive updates via email. Behind the scenes, moving the slider to **Off** removes your account from the group's subscriber list while moving it to **On** adds your account to the subscriber list. Outlook for Mac and Outlook for iOS also support group cards with the **Follow in Inbox** setting for a group.

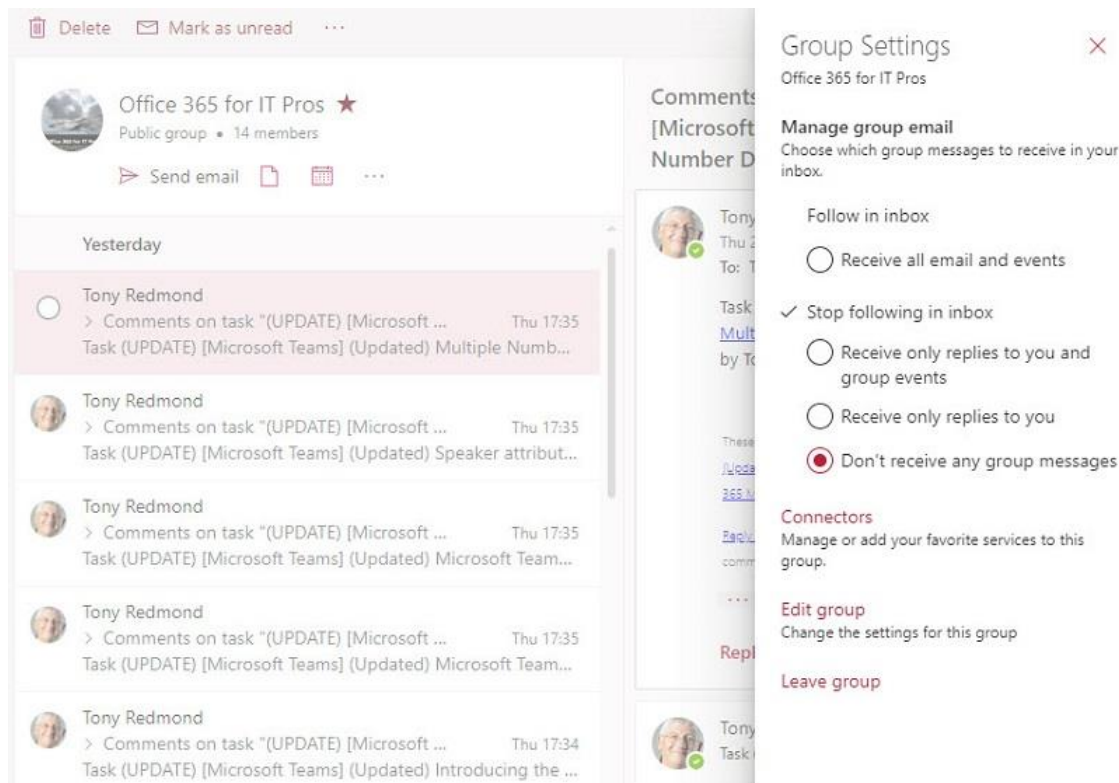


Figure 11-5: The Follow in Inbox setting controls if a user receives updates for a group by email

## Groups and Transport

Every group has an SMTP email address. Members and others (if allowed by the group settings) can use that address to route email to groups just like any other recipient. The messages eventually arrive at the Exchange Online mailbox server that hosts the active copy of the group mailbox. At this point, the Exchange transport service processes the messages before delivery to the group mailbox. All messages are subject to whatever processing the tenant implements through transport rules.

In addition to sending emails to groups, clients can post new conversations and reply to existing conversations. Depending on the client, the mechanism to post to a group is slightly different. Outlook handles a message posted to a group conversation much like any other message. Upon submission, the normal Outlook synchronization mechanism picks up the change in the local folder and dispatches the message to the Exchange Online transport service. Eventually, Exchange delivers the message to the group

mailbox where it shows up in a conversation. The item posted to the group stays in the user's Sent Items folder.

By default, the Exchange Online transport service does not deliver copies of messages posted by a user to their Inbox. If users want to receive copies of messages, they can either:

- Add their email address as a CC or BCC recipient.
- Use OWA settings to enable the *Send me a copy of email I send to a group* option (in the Groups section). An administrator can do this for a user by running the *Set-MailboxMessageConfiguration* cmdlet to set *EchoGroupMessageBackToSubscribedSender* to True. For instance:

```
[PS] C:\> Set-MailboxMessageConfiguration -EchoGroupMessageBackToSubscribedSender $True -Identity Jane.Sixsmith
```

The first approach works on an on-demand basis. In other words, the user receives copies of messages they post when they add their email address as a recipient. The second works for all groups where the user is a member of the group's subscriber list. Selecting *Receive all email and events* in the Follow in Inbox option described above adds a user to the subscriber list for a group.

OWA and mobile clients do not support offline working as the Outlook desktop clients do. Because all posts must go through the transport service, the possibility exists that a short delay might be obvious between the time when a user sends a post to the group and the post showing up in a conversation. The delay comes about through the need for the transport system to process the message through its pipeline to allow rules and routing to occur. To give the user confidence that their post succeeded, OWA and mobile clients create a "fake" post in the conversation that is only visible to the submitter. Behind the scenes, email processing continues and if a transport rule decides for some reason to reject the post, it does not deliver it to the group mailbox. The submitter will receive a non-delivery notification to tell them that an administrative rule blocked their contribution. At this point, OWA removes the fake post so that it no longer appears in the conversation. The transport service usually processes messages very quickly and the time elapsed between submission and removal is very short, but if delays occur, users might report a "disappearing item" if they ignore the non-delivery notification. It is all by design.

## Managing Groups with Outlook Clients

Users allowed to create new Groups can do this with Outlook (desktop, OWA, and mobile). Group owners can also update the settings of their groups through these clients. In this section, we review how to create a new group using Outlook desktop and the management options available in OWA.

The dependency on the group mailbox for conversation storage means that Outlook Groups have a lower membership limit than other applications which use the Microsoft 365 Groups membership model. However, the limit is soft and based on practice, and Outlook groups don't break if they have larger memberships. The limit exists because we know that concurrent access to a shared mailbox becomes problematic as the number of connections grows. Not every member will access the mailbox (to read conversations or look at the calendar) at the same time, so 2,500 members is a "safe" limit, assuming normal usage patterns. If you have groups where members seldom access the mailbox, you can have more members. There are groups in production today that surpass this number of members by a considerable margin because their owners know only a small percentage of the membership ever access the group concurrently.

## Creating and Editing Microsoft 365 Groups with Outlook (for Windows)

Three methods exist to create a new group with Outlook:

- Right-click on **Groups** in Outlook's resource navigation tree and then **New Group**.
- Click the **New** drop-down list in the **Home** section of Outlook's ribbon and then select **Group**.
- Click **New Group** in the group ribbon.

Figure 11-6 shows the information about the new group. You can select public or private for the privacy setting and select a value from the set of classifications defined in the Azure AD policy for Groups (if set). If the tenant uses sensitivity labels for container management, the selected label controls group settings like privacy and guest access. You can update the group to change the settings afterward. When all details are input, click **Create** to create the new group. After Outlook creates the group, you can add members by inputting the names of individual users, guest accounts, or groups. If you add a group as a member, Outlook expands its membership and adds any cloud mailboxes found in the membership to the membership of the new group. The account used to create an Outlook group joins the group automatically as an owner. Later, you can add or remove group members by selecting **Group Settings**.

Figure 11-6 is a screenshot of the "Create Group" dialog box in Outlook. The dialog has a title bar with a question mark and a close button. The main content area is titled "Create Group" and contains the following fields and options:

- Name:** A text box containing "Company Group Initiative". Below it, the text "Group name: Company Group Initiative" is displayed.
- Email address:** A text box containing "CompanyGroupInitiative" with a green checkmark on the right. Below it, the text "Group email address: CompanyGroupInitiative@office365itpros.com" is displayed.
- Description:** A text box containing "A group to coordinate all aspects of the company growth initiative project." To the right of the text box is a link for "Group Usage Guidelines".
- Sensitivity:** A dropdown menu currently showing "Limited Access".
- Privacy:** A dropdown menu currently showing "Public - Anyone in your organization can see group content." Below this dropdown, the text "The privacy setting is determined by the sensitivity you select." is displayed.
- Checkboxes:** A checked checkbox with the text "Send all group email and events to members' inboxes. They can change this setting later."
- More Settings:** A link labeled "More Settings".
- Create Button:** A button labeled "Create" with the text "You'll be able to add members after you select Create." next to it.

Figure 11-6: Creating a new group with Outlook for Windows

Group owners can update the properties of a group. You can change the name of the group, its description, its membership, privacy setting, the language used in messages generated by the group, and whether people from outside the group can send emails to the group to take part in group conversations. Other properties, like the primary SMTP address, are either invisible or you cannot change through Outlook. The component used to edit group settings is shared with OWA.

## OWA Group Management

The Manage Groups section at the bottom of the OWA's folder list and the Groups section in OWA's People section allow group members and owners to see details of the groups they belong to (including group membership), perform actions like leaving the group, and access the different resources associated with a group. Owners can edit the settings of a group. For example, in Figure 11-7, we see that the group is under the control of the Groups expiration policy and that it is due to expire in 18 days. The owner can go ahead

and renew the group or, if the group has served its function and is obsolete, allow it to expire and enter the removal cycle. During the removal cycle, owners have 30 days during which they can recover a deleted group. To do this, open the *Deleted* link under Groups, select the group you want to restore, and click the **Restore** button.

Notice that OWA puts groups with outstanding actions at the top of the list. One of the flagged groups has a pending join request while the others are awaiting renewal. To change group settings, use the **Edit** option.

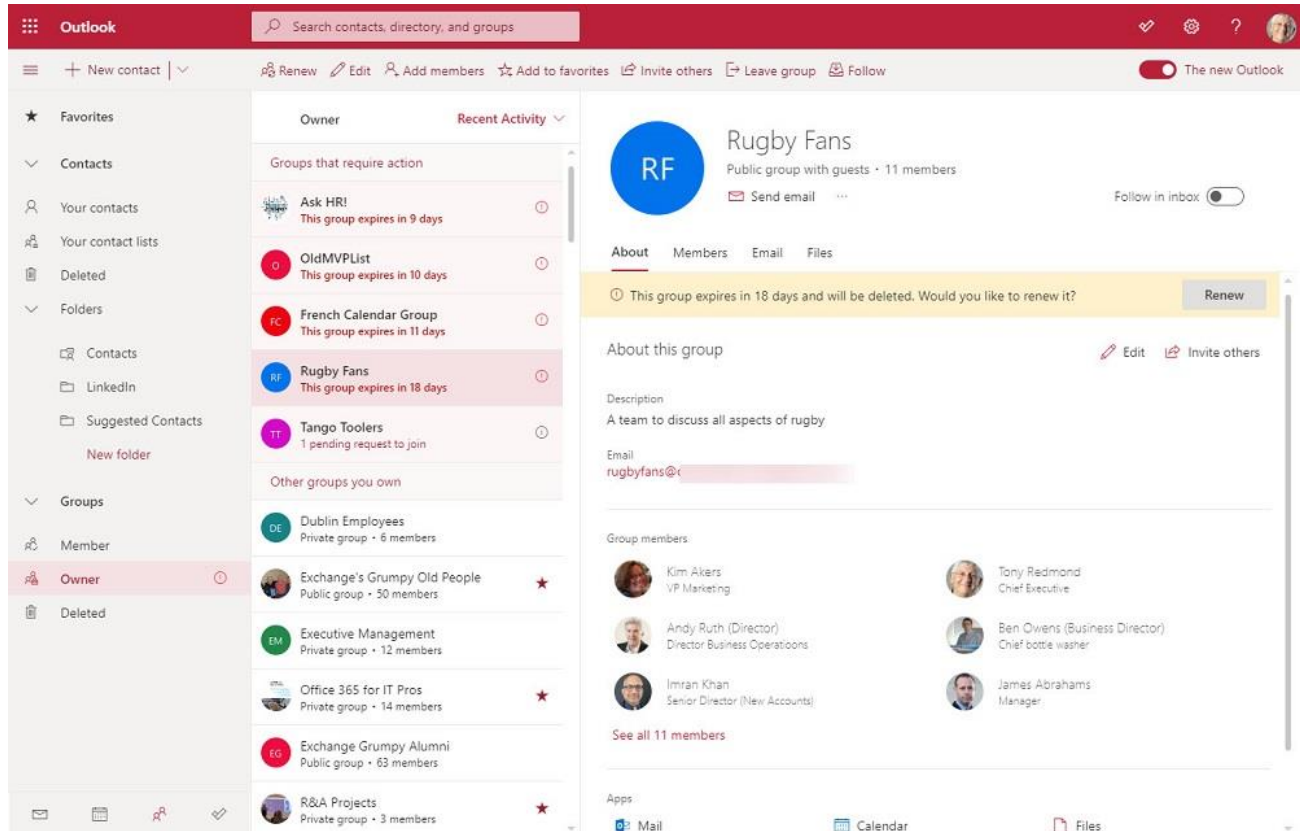


Figure 11-7: Managing the settings for a Group with OWA

## Guest Access to Microsoft 365 Groups

Microsoft 365 applications that support guest user access leverage the [B2B collaboration framework](#) for Azure AD to create and manage guest accounts. Microsoft 365 Groups use [Azure AD guest accounts](#) to enable access for external people to resources controlled by the tenant. Groups, Teams, SharePoint, and Planner support guest access. If an organization's Yammer network runs in Microsoft 365 native mode, [guests can be added to Yammer communities](#) using Azure B2B collaboration, but only if the guests have Microsoft accounts. A single guest account can be a member of multiple groups and teams. Any application that supports Azure B2B Collaboration can use existing guest accounts to grant access to their resources.

Apart from the membership limits that exist for different types of groups, there is no specific limit on the number of guests who can join a group (or team). For example, you can have a group with 1 local member (the owner) and 2,499 guests. See the Identities chapter for more information about the management of guest accounts in Azure AD.

### Guests Blocked by Risky Sign-in Policy

Tenants can control the type of sign-ins accepted by Azure AD by configuring an [Azure AD sign-in risk policy](#) that blocks suspicious sign-ins once the circumstances surrounding a sign-in reaches a set threshold. If their

accounts are deemed to be at risk, guest users attempting to sign-into the tenant can be blocked by policy. When this happens, the user's account cannot be unblocked in the tenant hosting their guest account. Instead, an administrator in their home tenant must [unblock the account](#) by taking an action such as resetting the account password or dismissing the risk detections.

## Adding Guests to a Group

Apart from administrative interfaces, you can use OWA, Outlook, or the Outlook and Teams mobile apps to add guest members to groups. A group owner or a tenant administrator can add a guest while other group members can request an owner to add a guest. For public groups, members can add guests to the group unless barred by policy (for instance, the sensitivity label assigned to the group might not allow guest members). In either case, when you add a guest to a group, an invitation is generated and sent to the guest to join the group. To start, open the group and click the *members* link (which shows the current number of members) to display the complete group membership, including any guests. Now click the **Add Members** icon and enter the email address of the guest you want to add. If the group settings allow guest access, the email address is checked against the tenant directory and your email address autocomplete list (Figure 11-8).

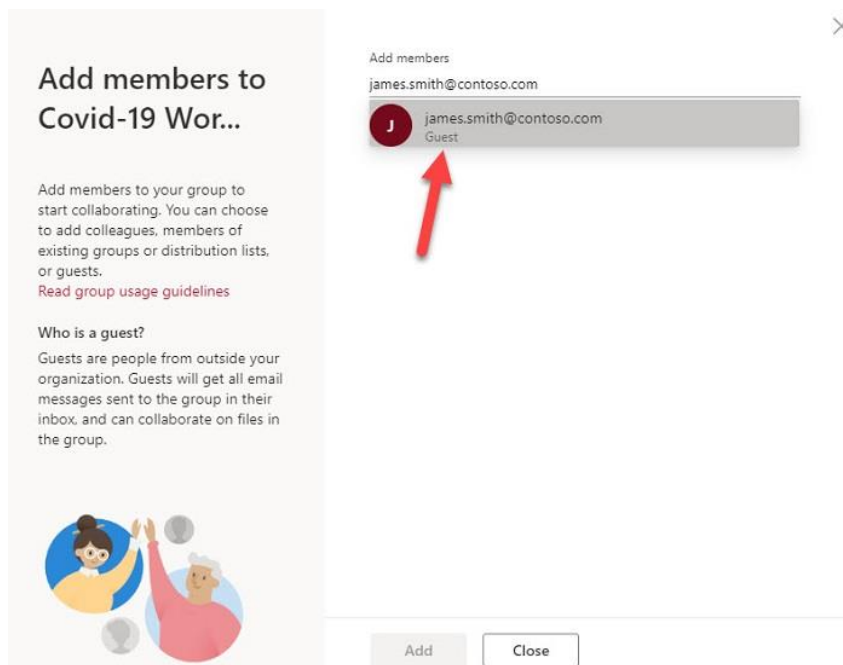


Figure 11-8: Adding a guest to a group with OWA

You are not restricted to using addresses already known to the directory as any valid email address is acceptable. If you input the email address of an account belonging to the tenant, they join the group as a regular member instead of a guest. Azure AD adds the guest to the group's membership and generates the email invitation after you click **Add**.

If group settings (for the tenant or a specific group) block guest access, OWA displays an error to say that you must contact an administrator to add a guest. OWA also flags an error if you try to add an email address for a guest that already exists for an unsupported object (like an email-enabled public folder).

**Guests and the Microsoft 365 admin center:** You can edit the membership of a group via the Microsoft 365 admin center and can add a guest to a group's membership there, providing that the guest account already exists in the directory. You can create new guest accounts in the Azure AD admin center or with an application that supports Azure B2B Collaboration, such as Groups or Microsoft Teams.

When OWA lists group members, it adds a "Guest" suffix to the display name of guest members as a visual reminder that this member is external. OWA also displays MailTips to warn users whenever they post new

conversations. These visual indications serve as a warning that members should not share or discuss confidential material within the group.

## How Guests Join a Group

When a guest is added to the membership of a group, Groups checks Azure AD to confirm whether a guest account for the email address exists in the tenant. If not, Azure AD creates a new guest account. An invitation then goes to the guest to let them know that they are now a member of the group. Should the guest wish to decline the invitation, the message has a link to allow the guest to leave the group. If not, they can click the link to access the group shared document library and confirm their membership.

If a guest does not receive (or loses) the email notification to tell them about their new membership status, the group owner can send them the URL for the document library. The guest can input the URL into a private browser session and go through the sign-in process to connect to the group. The private session makes sure that existing cached credentials don't interfere with the process. The URL for a group document library is in this form:

*<https://mytenant.sharepoint.com/sites/groupalias/>*

Another approach is to remove the user from the group and then add them again to force the regeneration of the invitation. This is perhaps the easiest way to solve the problem of lost invitations.

Some guests who receive invitations to join a group find that they cannot log into Microsoft 365 to access the group using the email address specified in the invitation. This might be because they have the same address registered elsewhere in Microsoft 365 (as a user in another tenant) and for a Microsoft consumer service, like OneDrive personal. In this case, the solution is to either use a different address to invite the guest or ask them to [rename their personal Microsoft consumer account](#).

A guest can also receive an invitation to a file held in a group document library through the action of sharing a document belonging to the group as a "cloudy attachment" (one where the attachment is a link to the content held in a OneDrive for Business or SharePoint Online library rather than a complete copy of the document). Clicking the attachments invokes the invitation to access the document.

If the guest chooses to accept the invitation, two conditions can occur. First, they might already have a Microsoft (MSA) or Microsoft 365 account linked to the email address used for the invitation (including a consumer account for a service like Outlook.com). If so, they can use this account to authenticate with Microsoft 365 and access the group files.

If they do not have a Microsoft account, a redirect occurs to [invitations.microsoft.com](https://invitations.microsoft.com) and the user can sign-up for a Microsoft account, including creating a password, which in turn becomes the guest account in the tenant. Creating a fully functional guest account needs the guest to prove that they own the email address associated with the account. Azure AD sends a verification code to the email address specified by the person extending the invitation. When prompted, the user gives the code to authenticate their address and complete the activation process, the guest account can then access the group.

**Finding group files:** Guests who have accounts in other tenants cannot add external groups to their set of favorite groups. Instead, to ensure fast access, they should bookmark the link contained in the invitation message as a favorite in their browser. Clicking the bookmark will then bring the guest to the Files section of the group.

## What Guests Can Access in Groups

Applications supporting the Azure B2B collaboration framework dictate what happens after someone accepts an invitation to access content managed by the application. In most cases, guests access group resources via browsers or mobile clients. Traditional desktop clients don't usually support guest access. You can't, for



instance, configure Outlook to open a group in another tenant to see the conversations in the inbox or the group calendar. Instead, Exchange Online uses an indirect access model to allow guests to participate in group interactions via email. They send emails to the group to contribute to conversations and receive responses via email, including copies of calendar events posted to the group calendar. Guests can then decide if they want to add these events to their calendars.

External people who are not group members can participate in conversations via email if the group's *RequireSenderAuthenticationEnabled* property is *False*. Generally, it is a bad idea to allow open access to a group because it creates the possibility that an attacker who obtains the group's email address will use it to send malware to the group, which can end up being copied to group members.

The *RequireSenderAuthenticationEnabled* setting does not affect guest members because they authenticate through their guest accounts. The exception to this rule is where the primary SMTP address of a guest is different from the email address used for their invitation. When this happens, attempts to respond to a conversation fail because the email address on the inbound item does not match the address of a group member. An easy workaround is to set *RequireSenderAuthenticationEnabled* to *False*, which allows any external user to email the group. Alternatively, you can create a guest to allow the person to log in using their normal UPN and send messages to the group using their primary SMTP address.

Guests can access to the SharePoint Online sites belonging to the groups to which they belong. They can use the SharePoint browser interface to open, add, edit, and remove items from the sites. In addition, people can share individual documents in document libraries from other SharePoint sites with guest accounts.

Table 11-4 summarizes the access available to guests to resources managed by Groups.

<b>Workload</b>	<b>Feature</b>	<b>Guests can access?</b>
Exchange Online	Contribute to conversations in the group mailbox	Yes (via email)
Exchange Online	View items in the group shared calendar	No direct access, but can receive updates via email
Exchange Online	Search group conversations	No
Exchange Online	Browse Global Address List for tenant (note: guests can use the Azure AD admin center to view directory information if enabled by the tenant)	No. (Note: Guests don't appear in any default Exchange address list)
SharePoint Online	Access a single document	Yes – via specific sharing invitation
SharePoint Online	Access the complete team site owned by the group/team	Yes
SharePoint Online	Search group documents	Yes
OneNote	Access group shared notebook	Yes
Information Protection	Members can post encrypted conversations to group conversations but cannot access messages protected with sensitivity labels. Guests cannot read files protected with sensitivity labels in the document library unless they are explicitly granted access by the labels.	No (apart from encryption)
Planner	Access plan associated with group	Yes
Teams	Access channel conversations in the team owned by the group and files stored in the group's SharePoint site. Conduct personal chats with other tenant members. Access private channels.	Yes
Power BI	Access Power BI workspace associated with group	No
Dynamics 365	Access customer data associated with group	No
Groups	Browse Azure AD for groups to join	No

Groups	Perform group management	No
--------	--------------------------	----

Table 11-4: Group functionality available to guests

## Restricting Guest Access to Confidential Material

A basic principle of the Microsoft 365 Groups membership model is that every member enjoys the same access to resources belonging to the group such as its document library, plans, and regular team channels. Owners maintain group membership and settings and have the same access to group resources as other members. The model is simple and easy to understand and manage, but some potential downsides exist that tenant administrators need to understand.

Consider files stored in a group's SharePoint document library. When a guest member joins a group, they receive the same access rights to documents and other files held in the document library. They can view, print, add, edit, check out, and perform all the file management actions available to other group members, including the ability to remove items. Two issues might occur:

- Guests will remove information from libraries.
- Guests can access confidential material.

The first issue is easily handled by using retention labels to prevent permanent deletion of important information. The second can be dealt with by using specific locations to store confidential information that external people should not access.

For instance, let's assume a project team needs some advice about a complex contract document from an external legal advisor. Three approaches are available:

- Allow the advisor full access to the group where the contract is stored along with other project files.
- Create a specially-defined group or team.
- Create a private or shared channel in a team and invite the external advisor to access the information through the channel.

The answer depends on the relationship with the external advisor and the need of the external advisor to access other supporting information that might not exist in a designated location.

If you store confidential material alongside other files in the document library, can you make sure that guests cannot access this content? The answer is yes, but you must apply sensitivity labels with encryption to protect the files from external access. As explained in the Information Protection chapter, a sensitivity label with encryption uses rights management to define the rights that users have over items. Users must be able to authenticate their access before SharePoint Online will decrypt and expose the content. Authors always hold full rights over files while other users might only be able to read the material. If a label does not grant rights to external users, they cannot open files protected by that label.

Sensitivity labels do not encrypt document metadata, so guest users can see the title and author information for protected files in a document library, even if they cannot open the content. Protected files circulated as email attachments to external people will also be inaccessible if the external users do not have the rights to open the files.

In all cases, it's important to understand the reasons for granting access to guests to a group or team and what resources they have access to if they become members before issuing any invitations.

## Mail User Objects Created for Guest Accounts

The creation of a mail user object in the Exchange Online directory (EXODS) happens the first time the guest joins a group or team in a tenant. The mail user object is a convenient holder of the guest account's email address to allow Exchange Online to route email to the guest. Exchange Online creates a mail user object even if a mail contact object with the same email address already exists. SharePoint Online and OneDrive for

Business also create guest accounts with associated mail user objects for sharing links. However, if you update an existing sharing link to add new guests, it doesn't result in the creation of mail user objects for the added accounts (perhaps because there's no need to send emails to them).

To mark their special purpose, mail user objects for guest accounts have a *RecipientTypeDetails* value of *GuestMailUser*. Do not remove one of these special mail user objects. If you do, you'll also remove the associated guest user account and the guest will lose any access they have to resources within the tenant, including sharing links and memberships of teams and groups.

By default, Exchange Online hides the mail user object for guest accounts from address lists. However, you can update the *HiddenFromAddressLists* property for a mail user object to force its inclusion in the GAL. Remember that it will take a day or so before newly unhidden objects appear in the OAB, but once unhidden, the objects show up in the online GAL and users can then address messages to the guest like any other recipient. Make sure that a mail contact for the guest isn't already present as un hiding the mail user object will then create a duplicate in the GAL.

```
[PS] C:\> Set-MailUser -Identity Julia.Foran_Contoso.com#EXT# -HiddenFromAddressListsEnabled $False
```

The presence of the mail user objects allows us to add guest accounts to groups using the *Add-UnifiedGroupLinks* cmdlet. In this example, we pass the alias created for the mail user object to tell Exchange Online which guest we want to add to the membership. As you can see, the alias is based on the UPN for the guest account.

```
[PS] C:\> Add-UnifiedGroupLinks -Identity MyGroup -LinkType Member -Links  
Julia.Foran_Contoso.com#EXT#
```

Alternatively, you can use the *New-MgGroupMember* cmdlet to add the guest to the group's membership.

```
[PS] C:\> $NewUser = (Get-MgUser -UserId  
John.Doe_outlook.com#EXT#@office365itpros.onmicrosoft.com).Id  
$GroupId = (Get-MgGroup -Filter "startsWith(displayname, 'Office 365 Adoption')" ).Id  
New-MgGroupMember -GroupId $GroupId -DirectoryObjectId $NewUser
```

Guests participate in group conversations via email, so we must make sure that they receive copies of group conversations via email. If the *AutoSubscribe* property for the group is *\$True*, new members automatically join the subscriber list when they become a member. If not, we need to add them as a subscriber.

```
[PS] C:\> Add-UnifiedGroupLinks -Identity MyGroup -LinkType Subscriber -Links  
John.Doe_outlook.com#EXT#
```

## Giving Guests a Face

Given the collaborative nature of Groups and Teams, it makes sense to update guest accounts with suitable photos. It's much better when people get a visual reminder of whom they work with. Applications like Planner, OWA, SharePoint, Teams, and OneDrive display photos if they are available for guest accounts. For example, the left-hand screen in Figure 11-9 shows a conversation within Teams where a guest account has a photo. This is a much nicer visual reminder about the person than the anonymous circle with two letters used for guests without photos.

Administrators can't upload a photo to a user or guest account through the Microsoft 365 admin center, but you can through the Azure AD admin center. Go to the Users section, select the guest account, and upload a JPEG or PNG file. The file must be less than 100 KB. If users can't give you a suitable photo, you might be able to use their profile picture from their LinkedIn.com account, as these photos are correctly sized and well under the 100 KB limit. If you must downsize a large picture, a 100 x 80-pixel size usually works. To have a better-populated directory, it is a good idea to update other information about the guest into Azure AD at the same

time, such as their first name, last name, display name, job title, company name, and contact details. Applications can then include this information in contact cards.

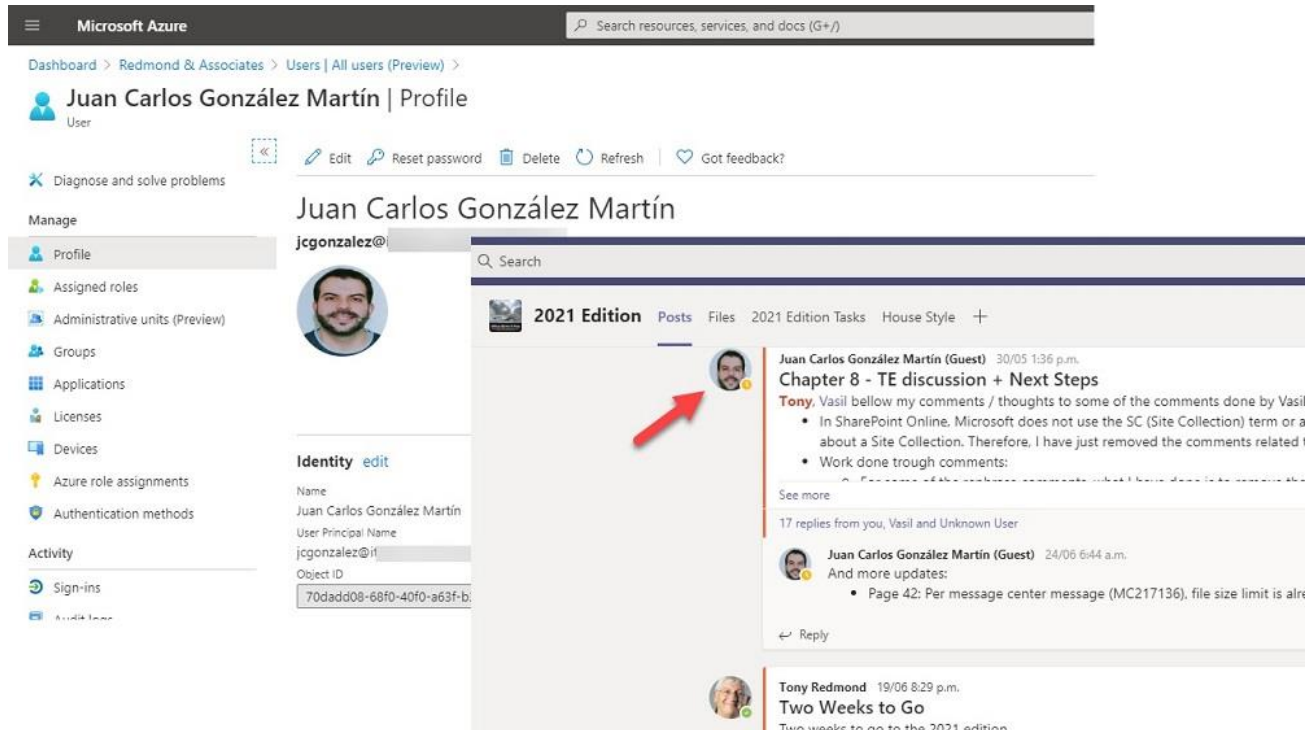


Figure 11-9: Using photos for Guest accounts

After updating guest accounts with photos, the background jobs responsible for synchronizing Azure AD with Microsoft 365 applications will make the images available in apps. This process might take up to a day to complete. To force Teams to synchronize, always update another attribute of the account, like the display name. This “tickles” Teams to let the app know that it should synchronize the account. Browser clients pick up newly available photos first. Mobile and desktop clients must synchronize their local caches with the workload directory. This can add another day or so before new images appear.

Another advantage of adding photos to guest accounts is that Outlook Mobile displays the photos when displaying messages from guests. Outlook and OWA can retrieve thumbnails from LinkedIn based on a person’s email address, but Outlook Mobile depends on thumbnails stored in user and guest accounts.

## Adding Guest Photos with PowerShell

You can also update the picture for Azure AD accounts using the `Set-MgUserPhotoContent` cmdlet. For example:

```
[PS] C:\> Set-MgUserPhotoContent -UserId a7eae252-3220-4254-91b2-c73918989a75 -InFile
c:\temp\GuestPhoto.jpg
```

You can then take this technique further to update all guest accounts in the tenant. For instance, let’s say that you want to flag guest accounts with a special image unless a photo is already present for an account. The idea is that the image will give people a visual clue of the guest’s status. This code does the job, even if it’s not very fast.

```
[PS] C:\> [array]$Guests = Get-MgUser -Filter "UserType eq 'Guest'" -All | Sort DisplayName
ForEach ($Guest in $Guests) {
    $PhotoExists = $Null
    $PhotoExists = Get-MgUserPhoto -UserId $Guest.Id -ErrorAction SilentlyContinue
    If (!$PhotoExists) {
        Write-Host "Photo does not exist for" $Guest.DisplayName "- updating with default guest logo"
        Set-MgUserPhotoContent -UserId $Guest.Id -Infile C:\Temp\DefaultGuestPicture.jpg
        Sleep -Seconds 3 }
}
```

```
Else { Write-Host "Photo found for" $Guest.DisplayName }
}
```

## Adding a Photo to Your Guest Account in Other Tenants

Applications do not allow guest users to update their photos, but this can be done with PowerShell. First, create a new PowerShell session and run the *Connect-AzureAD* cmdlet to connect to the target domain that hosts the guest account (unfortunately, the cmdlets in the Microsoft Graph PowerShell SDK don't allow guests to update their photos). Use the tenant service domain (the one with onmicrosoft.com) to connect.

```
[PS] C:\> Connect-AzureAD -TenantDomain remotedomain.onmicrosoft.com
```

If you don't know the service domain for a tenant, you can connect using the tenant GUID (a service like [What's My Tenant ID](#) tells you the GUID to use). For example, this command connects to Microsoft's tenant:

```
[PS] C:\> Connect-AzureAD -Tenant 72f988bf-86f1-41af-91ab-2d7cd011db47
```

In either case, you will be asked to authenticate. Use your normal account (it is linked to the guest account). Then, if you need to find the service domain, run the *Get-AzureADTenantDetail* cmdlet as follows:

```
[PS] C:\> $ServiceDomain = (Get-AzureADTenantDetail | Select -ExpandProperty VerifiedDomains |
?{$_.Name -Like "*onmicrosoft*"}).Name
```

The user principal name for your guest account in the tenant is usually composed of:

- The user principal name for your account in your home tenant, replacing the @ sign separating the username from the domain with an underscore.
- The string #EXT#@.
- The service domain of the target tenant.

Taking *Tony.Redmond@office365itpros.com* as an example, the user principal name for the guest account is *Tony.Redmond\_office365itpros.com#EXT#@remotedomain.onmicrosoft.com*.

Run the *Set-AzureADUserThumbnailPhoto* cmdlet to update the guest account. These commands create a variable to store the user principal name for the guest account, update the guest account with a photo, and check that photo data is present for the account.

```
[PS] C:\> $UserId = Tony.Redmond_office365itpros.com#EXT#@remotedomain.onmicrosoft.com
Set-AzureADUserThumbnailPhoto -ObjectId $UserId -FilePath c:\temp\tr.jpg
Get-AzureADUserThumbnailPhoto -ObjectId $UserId
```

It can take several days before the photo shows up in all applications.

## Managing Guest Account Properties with PowerShell

Details of guest accounts are available by running the *Get-MgUser* cmdlet. Here are three examples:

1. You can pass the user principal name of the guest account. The format of the user principal name assigned to a guest account is *username\_domainname#EXT#@servicedomain*.
2. You can search Azure AD for a guest account using a property of the account like the display name.
3. You can look for all guest accounts known in the tenant by looking for accounts with the *UserType* property set to "Guest".

```
[PS] C:\> Get-MgUser -UserId Michael_yandex.com#EXT#@office365itpros.onmicrosoft.com | Format-Table
UserPrincipalName, DisplayName
```

UserPrincipalName	DisplayName
-----	-----
michael_yandex.com#EXT#@office365itpros.onmicrosoft.com	Flayosc Worker

```
[PS] C:\> Get-MgUser -Search 'DisplayName:Michael' -ConsistencyLevel Eventual | Format-Table
UserPrincipalName, DisplayName
```

```
UserPrincipalName          DisplayName
-----
michael_yandex.com#EXT#@office365itpros.onmicrosoft.com  Flower Guru
```

```
Get-MgUser -Filter "Usertype eq 'Guest'" -All | Format-Table DisplayName, ProxyAddresses
```

Azure AD processes guest accounts like any other account. The *Update-MgUser* cmdlet updates the properties of an account, including guests. For example, here is how to change an account's display name:

```
[PS] C:\> $User = Get-MgUser -Search 'DisplayName:Michael Van Horenbeeck' -ConsistencyLevel Eventual
Update-MgUser -UserId $User.Id -DisplayName "Michael Van Hybrid"
```

Other attributes of the Azure AD object for a guest account include:

- **ObjectType:** User.
- **AccountEnabled:** True.
- **CreationType:** Invitation. You cannot change this.
- **Mail:** The SMTP email address for the guest.
- **MailNickName:** A modified version of the email address to show that this is a guest. For example, *Flowers\_outlook\_com#EXT#*. This attribute is a unique alias for the account.
- **OtherMails:** A multi-valued attribute used for different purposes for different sorts of accounts (for instance, accounts enabled for MFA store their "rescue" email address in this attribute). Guest accounts hold their SMTP email address in this attribute. SharePoint Online uses the attribute to control external file sharing, so if you remove this attribute for an account, you remove the ability for the user to access any files shared with them.
- **UserPrincipalName:** Used in the same manner as the UPN for other Azure AD accounts.
- **UserType:** Guest.

If you use the *Get-UnifiedGroupLinks* cmdlet to examine the membership list for a group that has guests, you'll see that the name assigned to guests is based on the first part of their user principal name.

```
[PS] C:\> Get-UnifiedGroupLinks -Identity BRK3001 -LinkType Members
```

```
Name          RecipientType
-----
TRedmond      UserMailbox
amitgutta.microsoft.com#EXT#  MailUser
chris.fiessinger_microsoft.com#EXT#  MailUser
e.zenz_microsoft.com#EXT#  MailUser
benny.niaulin_share-gate.com#EXT#  MailUser
t.redmond_live.com#EXT#  MailUser
flowers_outlook.com#EXT#  MailUser
```

In a hybrid environment where group membership synchronizes with an on-premises Active Directory via AADConnect, the synchronization process excludes guests because guest accounts do not exist in on-premises directories.

## Resetting Membership for Guests

Occasionally the creation process for a guest account does not complete as planned. The guest account exists in Azure AD and the guest shows up in a group's membership list, but whenever the guest attempts to access the group files, they receive an error telling them that their account is not in the tenant directory.

The quickest and simplest solution is to recreate the guest account. To do this, you must remove the account through the Microsoft 365 admin center or with PowerShell. Here's how to do the job with PowerShell:

```
[PS] C:\> Remove-MgUser -UserId t.redmond_live.com#EXT#@Office365ITPros.onmicrosoft.com
```

Now go back to a client and re-add the guest to force the creation of a new guest account. Everything should now work! The big downside with this method is that permanently removing a guest account also removes the membership for that user to all Groups and Teams and removes sharing permissions that the user might have in SharePoint and OneDrive for Business sites. For this reason, you should not remove a guest account without first thinking about the potential consequences.

## Mail Contacts and Guest Accounts

A basic rule of email transport is that an addressable object must have a unique email address. In other words, you cannot have two objects in a directory that share the same email address. Guest accounts are an exception because they can share an email address with a mail contact. If you add a new guest account and give the same email address as a mail contact, Groups creates the guest account based on the properties of the mail contact. You cannot do the reverse and create a mail contact using the email address of an existing guest account.

The need to support two objects with the same email address exists because the two objects serve different purposes. A mail contact exists to allow users to send emails to external contacts, individually or through a distribution list. Guest accounts also have email addresses. However, guest accounts don't appear in address lists (see below) but can gain authenticated access to application resources in a tenant.

### Including Guests in Exchange Address Lists

By default, guest accounts (or rather, the mail user objects created for guest accounts) do not appear in Exchange address lists. You can see the list of guest accounts known to Exchange Online with the command:

```
[PS] C:\> Get-Recipient -RecipientTypeDetails GuestMailUser | Format-Table DisplayName,
HiddenFromAddressListsEnabled, PrimarySmtpAddress
```

DisplayName	HiddenFromAddressListsEnabled	PrimarySmtpAddress
Chris Burger	True	Chris@burger.com
Jon Vickers	False	flayosc32@outlook.com
Benjamin N Smith	True	benjamin.n.smith@contoso.com
James Tracey	True	

We can see the guests that appear in Exchange address lists (*HiddenFromAddressListsEnabled* is *\$True*) and guests who have not yet accepted an invitation to join the tenant (its *PrimarySmtpAddress* is blank). If you want an account to show up in Exchange address lists (to allow users to email the guest), update the mail user object for the guest:

```
[PS] C:\> Set-MailUser -Identity vasil_michev.com#EXT# -HiddenFromAddressListsEnabled $False
```

The guest account appears in online Exchange address lists soon afterward. It will be included in the offline GAL the next time Exchange Online generates OAB updates and Outlook clients download and process updates.

Even if you do not update guest accounts to appear in address lists, you can still include them in distribution lists by updating the group membership with PowerShell. Here are two examples of how to add a guest to a distribution list:

- The first uses a lookup against Azure AD to fetch the registered email address for a guest user.
- The second uses the name of a mail user object created for a guest account. You could also pass the full user principal name for the account.

```
[PS] C:\> Add-DistributionGroupMember -Identity DL1 -Member (Get-MgUser -UserId
stale.hansen_cloudway.no#EXT#@office365itpros.onmicrosoft.com).Mail
Add-DistributionGroupMember -Identity "External Suppliers List" -Member
ExternalPerson_outlook.com#EXT#
```

## Guests in SharePoint and OneDrive for Business Pickers

If you want guests to show up in the people pickers used by SharePoint Online and OneDrive for Business, you can run the *Set-SPOTenant* cmdlet to update behavior for the tenant or *Set-SPOSite* for a specific site. For example:

```
[PS] C:\> Set-SPOTenant -ShowPeoplePickerSuggestionsForGuestUsers $True
```

# Controlling Guest Access to Groups

Four distinct pieces come together to control guest access to Groups:

- Azure AD must allow invitations to external users to join groups.
- The settings for Microsoft 365 Groups in the Microsoft 365 admin center must allow people outside the organization to access group content (see below).
- SharePoint Online must allow sharing of content stored in SharePoint and OneDrive for Business sites with external users.
- The Azure AD policy for Groups must allow guests to become members of groups.

These settings must be in place before Teams supports guest user access.

Because guest access for Outlook Groups primarily focuses on group document libraries, it follows that a prerequisite for sharing to occur is that SharePoint external sharing must be enabled for the tenant. This setting is available in the Sharing section of the SharePoint admin center. If a tenant does not allow sharing, guest access to Groups cannot work. SharePoint Online allows you to restrict sharing with users in specific domains (a whitelist). However, applications ignore this whitelist when adding guests to groups. If you want to prevent guests from domains outside the whitelist from becoming members of groups, you should implement the block policy available for groups and conduct a periodic check of group memberships. We will get to these topics shortly.

The Azure AD settings for the tenant must also allow invitations to go to guests. For instance, if you disable the “Members can Invite” setting in the **User Settings** section for Azure AD in the Azure portal, group owners cannot send invitations to people outside the organization.

## Group Settings for Guests

By default, Microsoft 365 tenants support guest accounts as members of Microsoft 365 Groups. If you want to disable this capability, edit the Microsoft 365 Groups settings in the org-wide settings section of the Microsoft 365 admin center (Figure 11-10) and uncheck both:

- **Let group owners add people outside your organization to Microsoft 365 Groups as guests.** Unchecking this setting removes the ability of group owners to add new guests. Administrators can still add guests to groups.
- **Let guest group members access group content.** Unchecking this setting removes the ability of guest members to access any group resources which are not explicitly shared with them. Existing guest members of groups will lose access to group resources.

A separate Microsoft Teams setting (*Allow guest access in Teams*) controls if team owners can add guests. You can have a situation where the tenant supports guest access outside Teams. For instance, you might want to allow guests to share SharePoint Online documents but not permit guests to be members of Teams.





### Microsoft 365 Groups

Choose how guests from outside your organization can collaborate with your users in Microsoft 365 Groups. [Learn more about guest access to Microsoft 365 Groups](#)

- Let group owners add people outside your organization to Microsoft 365 Groups as guests
- Let guest group members access group content  
If you don't select this, guests will still be listed as members of the group, but they won't receive group emails or be able to access any group content. They'll only be able to access files that were directly shared with them.

Figure 11-10: Microsoft 365 admin center settings for guest access to Groups

Optionally, you can use PowerShell or sensitivity labels to update the settings for a group to allow or deny guest access to that group. We discuss how to do this later.

## Using the Groups Policy to Control Guest Access

The Azure AD Policy for Groups applies to all applications that support guest access to groups, including Teams and Planner. The policy settings are independent of the sharing controls for SharePoint Online, which means that you can disable guest access as a member of groups while keeping the ability for users to issue invitations to external users to share specific items in SharePoint libraries and lists.

The Azure AD policy for Groups includes two settings to control guest access to groups. The *AllowToAddGuests* setting controls whether group owners can add guest user accounts to the membership of groups and teams. By default, the value of the setting is *True*, meaning that owners can add guests to any group in the tenant where the group settings do not block guests (we'll discuss how to block guest access to individual groups shortly). The *Get-MgDirectorySetting* cmdlet exposes the relevant policy settings:

```
[PS] C:\> $Settings = (Get-MgDirectorySetting | ? {$_.DisplayName -eq "Group.Unified"} | Select -ExpandProperty Values )

$Settings | ? {$_.Name -Like "*Allow*"}

Name                               Value
----                               -
AllowGuestsToAccessGroups          True
AllowToAddGuests                   True
```

### Stopping Guests Joining Group Membership

PowerShell cmdlets in the Azure AD module manage the settings in the Azure AD Policy for Groups. As an example, these commands update the *AllowToAddGuests* setting in the policy to *False* to stop owners from adding guest members. The organization-wide block on adding guest users set in the policy applies even when individual groups have a group-specific *AllowToAddGuests* setting of *True*. In other words, the organization-wide setting trumps an individual group setting.

```
[PS] C:\> $TenantSettings = Get-MgDirectorySetting | ? {$_.DisplayName -eq "Group.Unified"}
$Values = $TenantSettings.Values
($Values | ? Name -eq 'AllowToAddGuests').Value = "false"
Update-MgDirectorySetting -DirectorySettingId $TenantSettings.Id -Values $Values
```

Administrators can add guests to groups even if the *AllowToAddGuests* policy setting is *False*. The control applies only to owners. Administrators can add a guest using:

- **PowerShell:** By running *Add-UnifiedGroupLinks*, *Add-TeamUser*, or *New-MgGroupMember*.
- **Admin centers:** Through the Microsoft 365 admin center, the Azure AD admin center, the SharePoint admin center, or the Teams admin center.

Guest accounts that already exist in the tenant directory can join group membership immediately. However, if you need to add a new guest account to a group, you must first create the guest account through the Azure AD admin center. Once the guest account exists in the tenant directory, it can be added to a group's membership. The guest won't be able to access any resources until they redeem their invitation to join the tenant.

## Turning Off Guest Access

If you update the *AllowGueststoAccessGroups* setting in the Azure AD Policy for Groups to *False*, guests lose access to the groups to which they belong after a brief period to allow cached directory settings to refresh. Think of this as "flipping a switch" to turn access on or off if the need suddenly arises to limit sharing for a tenant. You can restore access later by updating the Groups policy. Guests will continue to have access to individual files for which they have received sharing invitations.

## Blocking Guest Access for Selected Groups

Individual groups inherit the settings of the organization policy and use these values unless overridden by settings for individual groups. When a group owner tries to add a new guest to the membership of a group or team, Groups checks the settings of the target group to check if it can proceed with the action. If allowed by the overall organization policy to control guest access to groups, the next check is against the setting for the target group to see if it supports guests. If both checks pass, Azure AD can add a guest to the group's membership.

Two methods exist to control guest access for individual groups:

- The method described below uses PowerShell to manipulate template settings for the group to block guest access.
- The settings for sensitivity labels include the ability to control guest access. When an administrator or owner applies a sensitivity label to a group, the group inherits the container management settings, including that for guest access.

In both cases, the group's *AllowToAddGuests* setting is updated to block or allow guest access.

### Blocking Guest Access to a Group with PowerShell

The *Get-MgGroupSetting* and *Update-MgGroupSetting* cmdlets retrieve and update group settings. The process to update the settings for an individual group is very like the one used to manipulate a policy setting in the organization policy. The first step is to retrieve the object identifier for a group and store it in a variable. We can do this by running the *Get-UnifiedGroup* cmdlet to retrieve the *ExternalDirectoryObjectId* property for the group. The object identifier is a GUID used to find and update the group object within Azure AD. For example:

```
[PS] C:\> $ObjectId = (Get-UnifiedGroup -Identity "Sales Professionals").ExternalDirectoryObjectId
```

Now, check that a settings object exists for the group. A group will have a settings object if it has a sensitivity label or if someone created the object using PowerShell:

```
[PS] C:\> $GroupSettings = Get-MgGroupSetting -GroupId $Group
```

If the *\$Settings* variable does not hold a value after running the command, you know that the group does not yet have a settings object. This is fine because it simply means that the group obeys the organization-level settings. To block guest access for the group, we create a settings object for the group as follows:

- Create a new directory setting template object using the values of the "Group.Unified.Guest" template object from Azure AD.

- Create a new group setting and populate the setting with a hash table containing the name *AllowToAddGuests* and the value *\$False*.
- The *ObjectId* pointing to the group comes from the command above.

Here's the code:

```
[PS] C:\> $GroupTemplateId = (Get-MgDirectorySettingTemplate | ? {$_.DisplayName -eq
'Group.Unified.Guest'} | Select -ExpandProperty Id)
New-MgGroupSetting -GroupId $ObjectId -TemplateId $GroupTemplateId -Values
(@{'name'='AllowToAddGuests';'value'='false'} | ConvertTo-Json
```

Note that you can't override the setting for guest access imposed by the sensitivity label assigned to the group with PowerShell.

After the block is in place, any attempt to add a guest generates an error. Existing guests remain in place and are unaffected by the block. To make sure that no guests have access to the group, you must remove existing guests from the group membership after you block guest access.

To reset a group to allow guests, you reverse the process and either set the value for the *AllowToAddGuests* setting to *\$True* and then update the settings object again or remove the policy from the group. This code uses the *Get-MgGroupSetting* cmdlet to read the values from the group and the *Update-MgGroupSetting* cmdlet to update the settings.

```
[PS] C:\> $GroupId = (Get-UnifiedGroup -Identity "Sales Professionals").ExternalDirectoryObjectId
$GroupSettings = Get-MgGroupSetting -GroupId $GroupId
Update-MgGroupSetting -GroupId $GroupId -TemplateId $GroupTemplateId -Values
(@{'name'='AllowToAddGuests';'value'='false'} | ConvertTo-Json) -DirectorySettingId
$GroupSettings.Id
```

Administrators can always add a guest to a group if necessary. Here is an example of a command to add a guest to a group with PowerShell. For this command to work, the guest account must already exist in the tenant directory.

```
[PS] C:\> Add-UnifiedGroupLinks -Identity "Confidential Group" -LinkType Member -Links
SomeUser_outlook.com#EXT#
```

The exception is for dynamic groups, where Azure AD calculates the group membership by running a query against the directory. You must adjust the query to make changes to the membership of a dynamic group as you cannot add an individual member.

See the PowerShell chapter for information about how to use a script to find groups matching certain criteria and block guest access to those groups.

## Restricting Guest Accounts to Certain Domains

Tenants can deploy an Azure AD policy called the *B2BManagementPolicy* to manage the domains that guest users can come from. To access the policy, go to the **External Identities** section in the Azure AD admin center, select **External collaboration settings**, and scroll down to **Collaboration restrictions**. Three values are available:

- Allow invitations to be sent to any domain (most inclusive): This is the default setting, and it means that invitations to join Groups and Teams can go to users in any other domain.
- Deny invitations to the specified domains: You can create a policy to block invitations going to specific domains. Microsoft believes that using a deny list is the most common scenario as most organizations know the domains with whom they do not wish to share information. For example, you might decide to block guest users with consumer email addresses, and block domains like Gmail.com, Yandex.com, Outlook.com, Yahoo.com, and so on. Including the domains of competitors in a deny list

is also sensible. When a deny list is in place, Azure AD blocks any attempt to invite someone to join a group or team if the invitee has an email address in one of the blocked domains.

- Allow invitations only to the specified domains (most restrictive): You can create a policy with an allow list, meaning that invitations can only go to domains in the list. Azure AD blocks any attempt to generate invitations to any other domain that is not on the list.

Applications do not implement blocks. This is done by Azure AD when an application tries to create a new guest account from a prohibited domain.

To create a policy with a deny list, click the deny list button and enter the domains to include and then Save (Figure 11-11). You can add up to 60 domains to a block or allow list. After you populate a list, the policy applies to all applications that use Azure B2B collaboration for external sharing, including Groups and Teams.

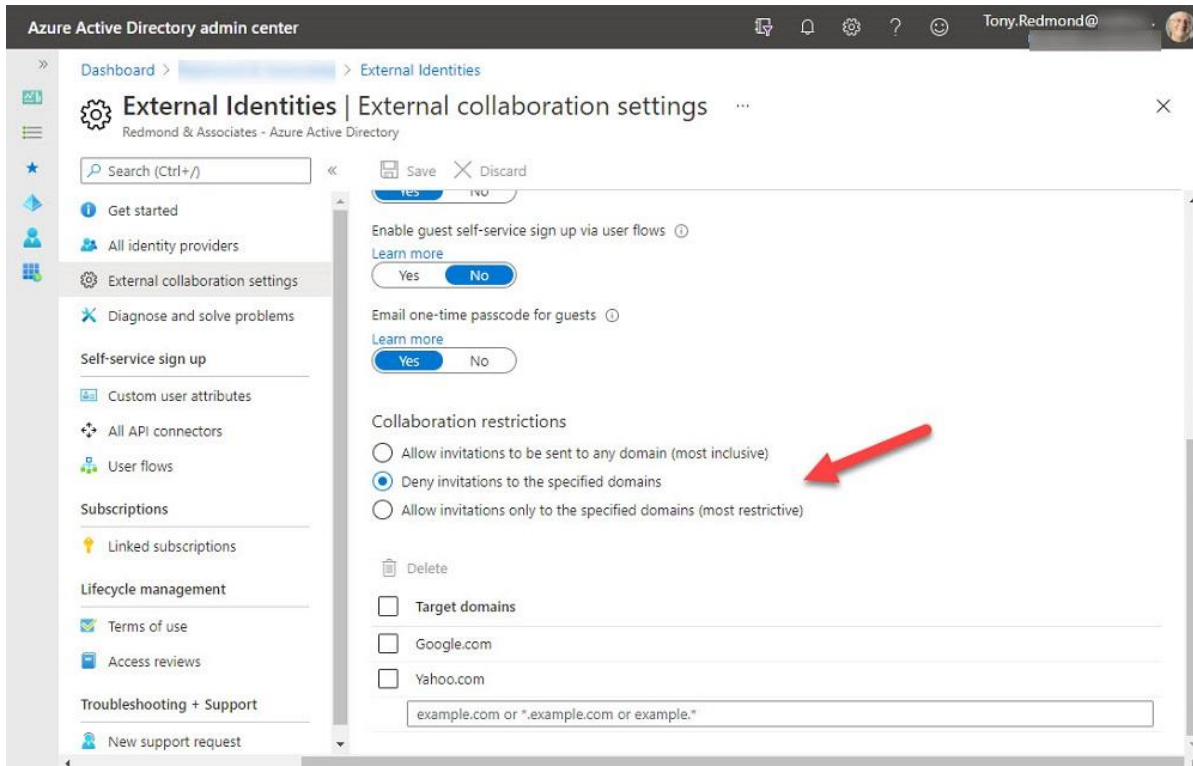


Figure 11-11: Adding domains to the deny list in the Azure B2B collaboration policy

Remember that the policy stops only the creation of new guest accounts. If you want to remove guest accounts from certain domains, you need to take separate action. We'll consider how best to do this shortly.

A tenant can only have a single deny or allow list. If you think that an allow list will work better than a deny list that is already in place, you must remove the current deny list and then create the allow list. The external collaboration list is separate from [a similar list to restrict sharing](#) for OneDrive for Business and SharePoint Online. One list blocks the ability to issue invitations to join groups or teams, and the other blocks sharing invitations for documents or folders.

If you create a deny or allow list to control access to certain domains, you should synchronize the domains with the [SharePoint Online restricted domains setting](#). Failure to do this can lead to inconsistencies. For example, a guest user might discover that they can view conversations in a team but cannot access the SharePoint document library.

### Finding the Source Domains for Guest Accounts

Before you create your collaboration policy, it's a good idea to check the domains where guests currently in the directory come from. This is easily done with a few lines of PowerShell. This code:

- Finds guests in the tenant.
- Extracts the domain from the user principal name for each guest and stores it in a list.
- Sorts the list and creates a hash table to count the occurrences for each domain, and then sorts the hash table in descending order.
- Outputs the result.

```
[PS] C:\> $Domains = [System.Collections.Generic.List[Object]]::new()
[array]$Guests = (Get-MgUser -Filter "UserType eq 'Guest'" -All | Select DisplayName,
UserPrincipalName, Mail, Id | Sort DisplayName)
ForEach ($Guest in $Guests) {
    $Domain = ($Guest.Mail.Split("@")[1])
    $Domains.Add($Domain)
}
$DomainsCount = @{}
$Domains = $Domains | Sort
$Domains | ForEach {$DomainsCount[$_]++}
$DomainsCount = $DomainsCount.GetEnumerator() | Sort -Property Value -Descending
$DomainsCount
```

Name	Value
----	-----
microsoft.com	59
outlook.com	11
quest.com	6
hotmail.com	5
gmail.com	4
emea.teams.ms	4

## Using PowerShell to Manage Domain Deny or Allow Lists

We can manage the Azure AD policy for external collaboration with PowerShell. For example, to see details of the current B2B Collaboration policy, use this code to run a Graph request against the policy endpoint:

```
[PS] C:\> $Uri = "https://graph.microsoft.com/beta/legacy/policies"
$Policy = Invoke-MgGraphRequest -Uri $Uri -Method Get
$B2BPolicy = $Policy.Value | ? {$_.displayName -eq 'B2BManagementPolicy'}
```

Name	Value
----	-----
deletedDateTime	
isManagementRestricted	
alternativeIdentifier	
definition	{{"B2BManagementPolicy":{"InvitationsAllowedAndBlockedDomainsPolicy":{"BlockedDomains...
keyCredentials	{}
id	14c3fc40-25fd-4f32-a0f0-c3fdd22f7de8
displayName	B2BManagementPolicy
type	B2BManagementPolicy
isOrganizationDefault	True
createdDateTime	08/12/2021 13:59:32

You can update the definition part of the policy with PowerShell, but it's easier and likely to be more accurate to update the policy through the Azure AD admin center.

## Finding Guests from Blocked Domains

Even with an Azure B2B Collaboration policy in place, some unwanted guests might already exist in group memberships. These unwanted guests might have joined groups before you implemented the policy or slipped in because a domain is not in the policy's blocked list. To fix the problem, you can scan the memberships of groups with guests to identify guests from blocked domains. The script works by first reading the set of blocked domains from the definition property in the Azure AD B2B Collaboration policy.

```
[PS] C:\> $Uri = "https://graph.microsoft.com/beta/legacy/policies/14c3fc40-25fd-4f32-a0f0-
c3fdd22f7de8/definition"
$Data = Invoke-MgGraphRequest -Uri $Uri -Method Get
```

```
$Policy = $Data.Value | ConvertFrom-Json
$BlockedDomains =
$Policy.B2BManagementPolicy.InvitationsAllowedAndBlockedDomainsPolicy.BlockedDomains
```

The script then finds Groups (including Teams) with guest members.

```
[PS] C:\> $Groups = Get-UnifiedGroup -ResultSize Unlimited -Filter {GroupExternalMemberCount -gt 0}
```

Finally, the code checks the guest members in each group to find if any come from blocked domains.

```
[PS] C:\> $Members = Get-UnifiedGroupLinks -Identity $Group.ExternalDirectoryObjectId -LinkType
Member | ? {$_.RecipientTypeDetails -eq "GuestMailUser" } | Select ExternalEmailAddress,
DisplayName, ExternalDirectoryObjectId
ForEach ($Guest in $Members) {
    $Domain = $Guest.ExternalEmailAddress.Split("@")[1]
    If ($BlockedDomains -contains $Domain) {
        Write-Host ("Found guest user {0} ({1}) in group {2}" -f $Guest.DisplayName, $GuestEmail,
$Group.DisplayName) }
}
```

Finally, the script generates a report of the guest accounts from blocked domains and outputs it as a CSV file. You can [download the full script from our GitHub repository](#). A blunter way to solve the problem is to look for guest accounts from a specific domain and remove them from Azure AD. This action removes the account from the membership of any group they belong to and stops access to any document shared through a sharing link.

```
[PS] C:\> $BadAccounts = [System.Collections.Generic.List[Object]]::new()
[Array]$Guests = (Get-MgUser -Filter "UserType eq 'Guest'" -All | Select Displayname,
UserPrincipalName, Mail, Id | Sort DisplayName)
ForEach ($Guest in $Guests) {
    If (($Guest.Mail.Split("@")[1]) -eq "Gmail.com") {
        $BadAccounts.Add($Guest.UserPrincipalName) }
} #End Foreach
If ($BadAccounts) { Write-Host ("Removing {0} bad accounts" -f $BadAccounts.count) }
ForEach ($BadAccount in $BadAccounts) { Remove-MgUser -UserId $BadAccount }
```

## Removing and Recovering Groups

The nature of all IT systems is that some groups become unwanted soon after creation or that the use of a group declines over time to a point where it falls into disuse. You might then decide to remove the group from the tenant. As noted previously, you can remove a group from the Microsoft 365 admin center, EAC, OWA, Outlook, a mobile app, or from a group-enabled application like Teams. When you remove a group, the action removes all the connected resources belonging to the group from SharePoint, Planner, Teams, Yammer, and so on. In addition, group owners or tenant administrators can remove a group by running the *Remove-UnifiedGroup* or *Remove-Team* cmdlets. In this example, we remove a group and suppress the prompt to continue that the cmdlet usually needs before it proceeds:

```
[PS] C:\> Remove-UnifiedGroup -Identity "Offshore Traders" -Confirm:$False
```

In addition to explicit removals by owners or administrators, if you operate an expiration policy, Azure AD removes groups automatically when they expire after a set period (explained later). In all cases, a removed group and any associated resources first enter a soft-deleted state and stay there for 30 days. During this time, an administrator or group owner can restore the group to make it accessible again to users. Alternatively, they can force the permanent removal of the group, which means that the resources of the group become irrecoverable. Table 11-5 lists the data recovered for a soft-deleted group.

<b>Group Resource</b>	<b>Data restored</b>
Azure AD Object	Group object, properties, and membership.
Exchange Online mailbox	Group mailbox holding conversations and group calendar. Also, the SMTP email address. Note – if the SMTP address has been assigned to another group, you cannot restore the original group until you update the SMTP address for that group.
SharePoint Online Team site	Group document library, including the shared OneNote notebook and the folders used by channels within Teams. The site is retained by SharePoint Online for up to 93 days. During this time, the site URL cannot be reused.
Planner	Group shared plan, including the plans created for channels within Teams.
Teams	Team channels and associated messages and metadata such as team and channel properties.
Yammer	Conversations in Yammer data store (for Yammer communities).

Table 11-5: Data restored when recovering a soft-deleted group

## Auditing Group Deletions

When you (soft) delete a group, Azure AD notes the details in a “Delete Group” audit record in the audit log. Later, when it removes a group permanently, Azure AD captures a “Hard Delete Group” audit record. Permanent deletion happens when the *Microsoft Online Service Garbage Collector process* runs, so the exact period when a group is in a soft-deleted state will vary from 30 days to perhaps a few days afterward.

The report and auditing chapter contains many examples of how to interrogate the audit log to understand when events occur. For instance, to look back and see when the garbage collector permanently removes groups, we look for “Delete Group” operations (audit records for both soft-delete and hard-delete actions are returned). We can then parse the audit data from the events found using the *Search-UnifiedAuditLog* cmdlet to generate a report using code like:

```
[PS] C:\> $StartDate = (Get-Date).AddDays(-90); $EndDate = (Get-Date)
[array]$Records = (Search-UnifiedAuditLog -StartDate $StartDate -EndDate $EndDate -Operations
"Delete Group")
$Report = [System.Collections.Generic.List[Object]]::new()
  ForEach ($Rec in $Records) {
    $AuditData = ConvertFrom-Json $Rec.Auditdata
    If ($AuditData.ResultStatus -eq "Success") {
      $ReportLine = [PSCustomObject]@{
        Timestamp = Get-Date $AuditData.CreationTime -Format g
        User      = $AuditData.UserId
        Action    = $AuditData.Operation
        Status    = $AuditData.ResultStatus
        GroupId   = $AuditData.Target.id[1]
        Group     = $AuditData.Target.id[3]}
      $Report.Add($ReportLine)}
$Report = $Report | Sort GroupId -Unique
$Report = $Report | Sort {$_.Timestamp -as [DateTime]} -Descending
$Report | Format-Table Timestamp, User, Group -AutoSize
```

The script sorts the report data twice. The first sorts by the group identifier to remove duplicate audit records which sometimes turn up in the audit log. The second sorts by date. The output is something like:

TimeStamp	User	Group
18/05/2022 16:37	Tony.Redmond@office365itpros.com	2021 Edition Book
15/04/2022 10:01	ServicePrincipal_1342cefb-7a89-4ee2-af90-c8443053e1e8	Plastic Production (Team)
30/03/2022 17:46	Administrator@office365itpros.com	March 2023 Sales Operations
25/03/2022 17:49	Administrator@office365itpros.com	Sun Seekers
15/03/2022 11:10	James.Ryan@office365itpros.com	Analytics Anonymous

These actions are all soft deletes. The hard deletes to remove the groups permanently happened a month later. The system removals are listed with "Certificate" as the User.

## Recovering a Deleted Group

As explained earlier, during the 30-day soft-deleted period, group owners can restore groups with OWA. Administrators can recover groups through the Microsoft 365 admin center, EAC, or with PowerShell. To recover a group with the Microsoft 365 admin center, select **Deleted groups** under the Teams & Groups section.

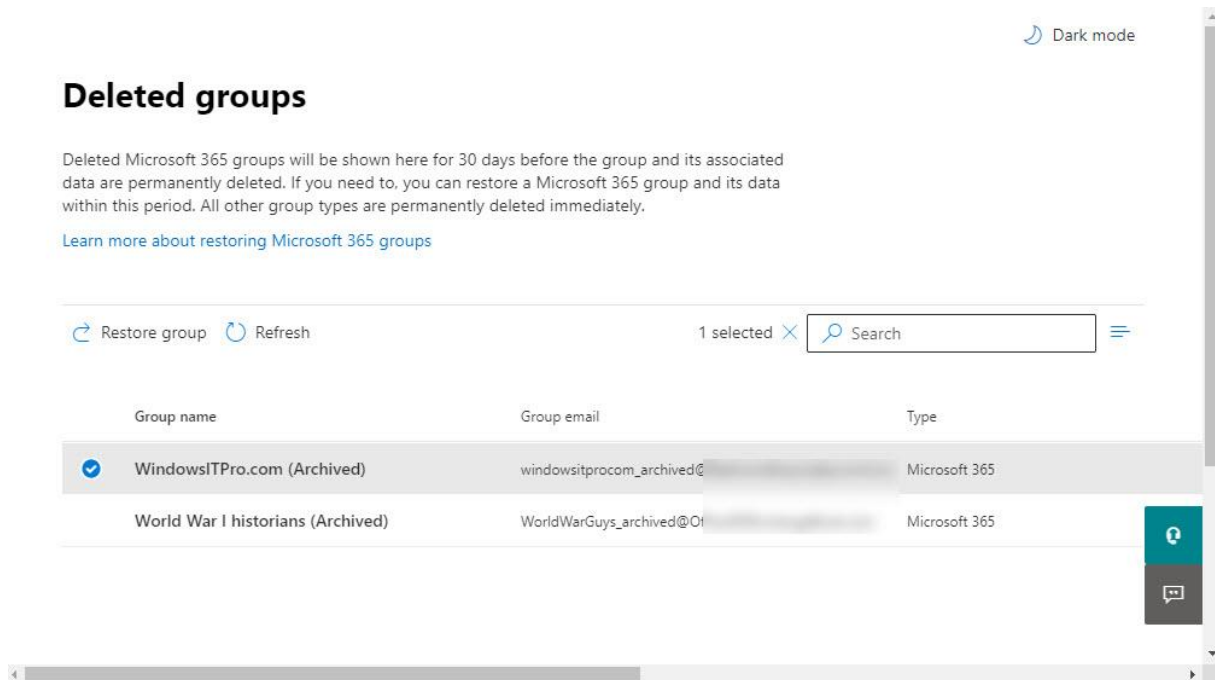


Figure 11-12: The Microsoft 365 admin center lists a set of deleted groups

In Figure 11-12 we see a set of deleted groups that have not exceeded their 30-day soft-deleted retention period. To restore a group, select it and either:

- Use the **Restore group** option in the group header.
- Click the group name to bring up the details pane, which has some basic information about the (display name, description, and email addresses), and click the **Restore group** button.

Both options perform the same processing to restore the deleted group. After a short delay, the group object is restored in Azure AD and begins the process of notifying the associated workloads to reconnect. As discussed later, it takes a little while for resources like the SharePoint site, a plan, or a team to be reconnected.

### Recovering a Deleted Group with PowerShell

The PowerShell command to list the set of deleted Microsoft 365 Groups is:

```
[PS] C:\> Get-UnifiedGroup -ResultSize Unlimited -IncludeSoftDeletedGroups:$True | ?
{$_ .WhenSoftDeleted -ne $Null} | Sort WhenSoftDeleted | Format-Table DisplayName,
PrimarySmtAddress, WhenSoftDeleted
```

You can mimic the steps taken by the Microsoft 365 admin center to recover a deleted group as follows. First, retrieve the object identifier for the group you want to recover:

```
[PS] C:\> $Object = (Get-UnifiedGroup -Identity "Group to Recover"
-IncludeSoftDeletedGroups).ExternalDirectoryObjectId
```

Now, use the identifier to recover the group.

```
[PS] C:\> Undo-SoftDeletedUnifiedGroup -SoftDeletedObject $Object
```



## Reconnecting Group Resources

Because many different applications use Groups, the recovery of a deleted group involves several complex operations. Recovery of a simple group (one that uses Exchange and SharePoint resources only) usually happens within a few minutes and should be complete within an hour. Depending on the current load within the service, recovery of a group that is enabled for Teams or Planner might take up to 24 hours before all the data is reconnected. The added delay is due to the need to synchronize multiple directories and often to ensure that data for Teams and Planner stored in Azure data services are restored correctly. It is also the case that when you restore a dynamic group, it can take up to 24 hours before Azure AD runs the query to calculate the group membership.

The first stage in recovery is to restore the group using the option available in the Microsoft 365 admin center or the Azure AD admin center. Once the group object is restored, Azure AD notifies each of the attached workloads to tell them to reconnect. As the workloads respond to the instruction, the full set of group resources come online to construct a fully-functional group.

Restoring the directory objects is only the first stage in the recovery process and does not mean that users can use the membership represented in the directory to access group contents. To perform further validation and before telling users that they can work with the group, you should check each of the workloads used by the group using clients to confirm that the expected conversations, files, plans, and chats are present. If you use a browser, make sure that you clear the browser cache or create a private session to avoid any problems caused by stale data. Teams is usually the last application to have its data restored and available to users.

If a workload is not ready an hour or so after the restore, wait for another hour, and check again. Because the restoration process involves many different connections, it could be that one of the steps in the process is waiting for another step to complete.

### Time Limit for Recovery

You have 30 days to recover a group following its deletion. After this period elapses, you might still be able to recover some information for a deleted group. For example, administrators might be able to recover files from the SharePoint recycle bin. If a hold is active for the group, you can run a content search to find information in the group mailbox and document library and export those items. However, you will have individual items and cannot reassemble them into a functioning group because the group object no longer exists in Azure AD.

## Recovering Individual Documents and Items

The possibility always exists that someone will remove a document in error, which then brings the question of how best to restore the now-deleted document. The first place to look is in SharePoint's recycle bin where items stay for up to 93 days after deletion (across the first and second stages). If the item isn't in the recycle bin and you have a backup, you might be able to recover the document from the backup. To prevent SharePoint from deleting items after 93 days, you can assign a retention label to important documents or apply a retention policy to entire sites. Users can delete documents with retention labels or those that come within the scope of a retention policy, but SharePoint keeps the retained items in the preservation hold library, and administrators can restore documents from there.

When someone removes an email conversation from a group, the item goes into the Deletions folder in the group mailbox and stays there for 14 days, after which the Managed Folder Assistant permanently removes the item from the mailbox database. You cannot use the standard Recover Deleted Items feature in Outlook or OWA because it is not available for group mailboxes. During the 14-day retention period, an administrator can recover deleted items by running the *Get-RecoverableItems* and *Restore-RecoverableItems* cmdlets to view and restore deleted items as these cmdlets work with group mailboxes.

Backup solutions from ISVs offer the ability to backup and restore documents selectively, which is a better outcome than having to restore a complete site. AvePoint has a backup solution for Groups that can restore conversations.

## Group Expiration Policy

Although it is possible to restrict the ability to create new groups to a select set of users, even in the most tightly managed tenant, some groups eventually reach their best-by date and become disused. If this happens, it is good to track down those groups, recover anything valuable stored in the group resources, and then remove the groups to reduce GAL debris. Groups are not unusual in this respect. Experience shows that the same falloff in usage over time happens for shared mailboxes, distribution lists, public folders, and other objects shared by teams of people.

To help administrators manage potentially obsolete groups, Azure AD supports the group expiration policy to control how long groups can exist within a tenant before the groups need to be renewed. If groups expire, Azure AD can automatically remove them from the tenant. The expiration policy can apply to some or all groups, no matter which application creates a group, or how people use the group and its resources.

A disabled group expiration policy exists for every tenant. An enabled expiration policy is an Azure AD Premium feature requiring licenses for every member of a group coming within the scope of the policy. Administrators also need an Azure AD Premium P1 license to configure the policy. If you access the Azure Portal without this license, you will not see the UI controls for the policy. However, you can manipulate the policy settings with PowerShell even if your account does not have a license. To work with the expiration policy, go to the [Azure AD admin center](#) and navigate to **Expiration** under Group Settings to see the current policy settings. You can now enable the policy, define settings and the groups that come under the scope of the policy, and then click **Save** to make the policy effective.

The settings in the group expiration policy are:

- The **group lifetime**: The default is 365 days, but you can select a higher or lower value. For example, you could use 730 days to make groups expire after two years. Consider setting a high number of days in the policy initially so that older groups do not expire when you enable the policy and to allow group owners to become used to the idea of expiring groups. For example, Microsoft started with a group lifetime of 360 days and subsequently reduced the period to 180 days.
- The **default notification address**. This handles the situation where a group has no owner. You can specify a single address (usually an administrator) to receive these notifications, the address of a distribution list or group, or the SMTP addresses for multiple recipients (separated by semi-colons).
- The **groups that come under the scope of the policy**. Initially, the value is **None**, meaning that the policy is in the default disabled state and groups do not expire. Selecting **All** means that every group in the tenant comes within its scope. The **Selected** button allows you to apply the policy to one or more groups, up to a maximum of 500 groups. You pick the target groups for the policy from a list of all groups in the tenant (for example, you might decide to exclude all the groups used by Teams from the policy and only include the groups that Teams doesn't use). This is easy to do in a small tenant but can become very tiresome for larger tenants which might span thousands of groups. In this case, you can use PowerShell to script to include the right groups within the scope of the policy. We explore how to manage the expiration policy through PowerShell in the PowerShell chapter. Figure 11-13 illustrates the settings of a group expiration policy that applies to the selected groups.

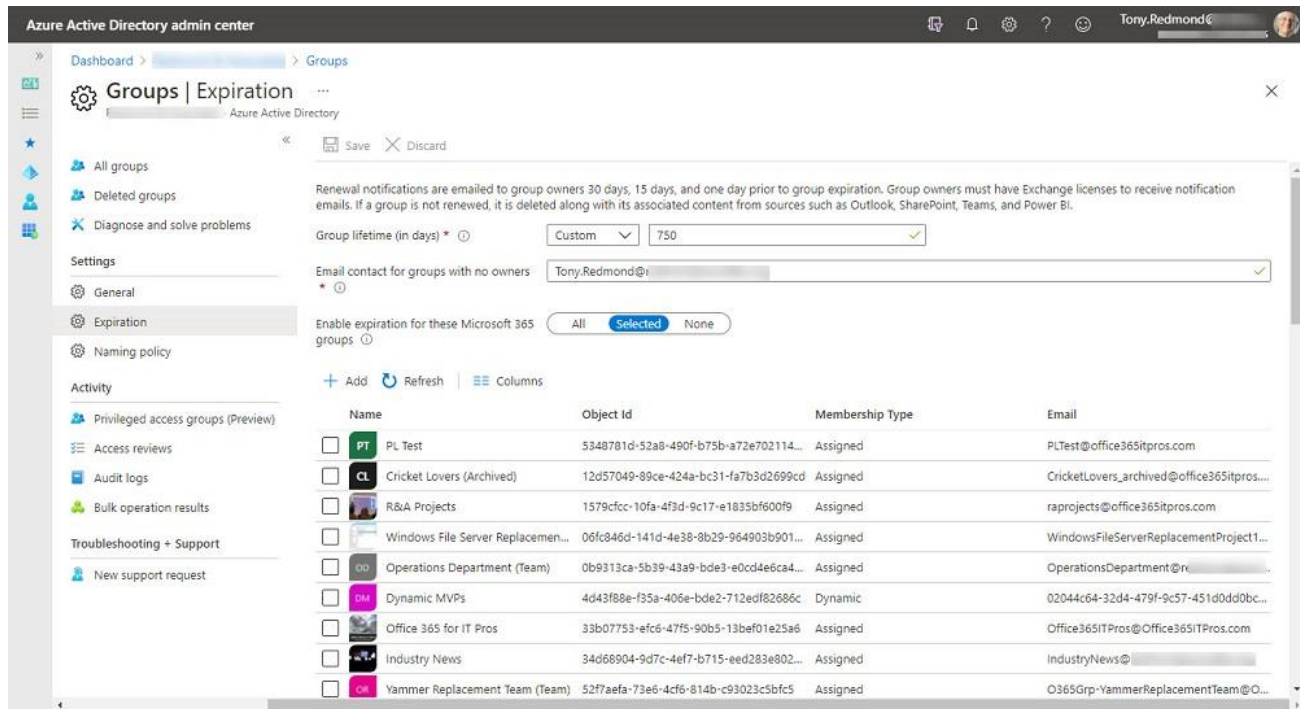


Figure 11-13: Updating the group expiration policy settings

If you change the scope of the policy from **All** to **Selected**, the policy ceases to apply to the groups outside the selected set.

## Automatic Renewal

Originally, Groups used an age-based renewal mechanism with the sole criterion for renewal being the time since the original creation date of the group and then, following a renewal, the date of the group's last renewal. Group owners had to renew groups manually, even if the groups are very active. An example of where this might cause a problem is when you use a policy covering all groups in the tenant, including a group used to limit the users who can create new groups. Eventually, that group will expire like every other group. If you ignore the notification, Azure AD removes the group when its expiration period elapses, which means that no one except an administrator will then be able to create a group.

The expiry policy uses a check for group activity. When the time comes to renew a group, Groups looks for evidence that the group is active. A background process called the Activity Tracking Service monitors group activities logged in the Microsoft Graph when users perform activities such as:

- SharePoint Online (group document library): View, edit, download, share, or upload files.
- OWA: Join a group, read, or send messages to group conversations, or like a message. Note: sending emails to contribute to a group conversation is not considered a renewal activity. The interaction must happen through OWA.
- Teams: Interact with a channel in the team.
- Yammer (network configured in Microsoft 365 mode): View a post in a Yammer community or use Outlook to respond to a Yammer interactive message.

It's obvious from this list that group renewal never considers some common activities when calculating group activity. For example, the process ignores adding tasks to Planner plans. Even in the workloads which are covered, many activities go unremarked, like participation in Teams meetings, using apps in channels, or assigning retention or sensitivity labels to SharePoint documents.

Many signals for the selected set usually exist for an active group, and when the time comes for renewal, Azure AD renews the group for a further period without the need for owner intervention. If the renewal

process finds no activity signals for a group, the group owners receive group expiry notification messages and must go through the normal renewal process. At the Ignite 2020 conference, Microsoft said that 79% of all groups coming within the scope of an expiration policy automatically renewed because of their activity. That's an impressive number, but the fact that 21% of groups didn't auto-renew because of a lack of activity gives pause for thought as it means that over a fifth of all groups created did not have enough activity (as measured by the limited set of signals) to allow Exchange Online to auto-renew them.

## Logging the Renewed Date

The *renewedDateTime* property of a group holds the last renewal time for a group (or the time of creation). Azure AD maintains the property to know when a group expires. After you renew or restore a group, Microsoft 365 updates its *renewedDateTime* property with the current date and time to give a new starting point for the group's expiry countdown. Records for the update to group properties are logged in the Audit logs section of the Azure AD admin center. In Figure 11-14, we can see the record captured when a group's renewal date was updated to 20 July 2019. Records noting group updates are also captured in the audit log.

Audit logs

Columns Refresh Download Export Data Settings

Service: Core Directory Category: GroupManagement Activity: All Status: All Target: Enter target name or upn Initiated By (Actor): Enter actor name or upn

Date: Last 7 days Show dates as: Local UTC

Apply Reset

DATE	SERVICE	CATEGORY	ACTIVITY	STATUS	TARGET(S)	INITIATED BY (ACTOR)
7/20/2019, 11:41:24 AM	Core Directory	GroupManagement	Update group	Success	Ignite 2017	Microsoft Approval Managem...
7/20/2019, 11:41:23 AM	Core Directory	GroupManagement	Update group	Success	Ignite 2017	Microsoft Approval Managem...
7/20/2019, 11:41:23 AM	Core Directory	GroupManagement	Update group	Success	Ignite 2017	Microsoft Approval Managem...

Details

Activity Target(s) Modified Properties

TARGET	PROPERTY NAME	OLD VALUE	NEW VALUE
Ignite 2017	RenewedDateTime	["2017-07-30T10:43:34Z"]	["2019-07-20T09:41:23Z"]
Ignite 2017	Included Updated Properties		"RenewedDateTime"
Ignite 2017	TargetId.GroupType		"Unified"

Figure 11-14: Azure AD logs a group renewal

## Expiry Notifications

No one wants a system to remove data without warning. When an expiration policy is active, Azure AD checks the last renewed date for every group covered by the policy. The auto-renewal process takes care of groups that are active (see the discussion below). For groups that have no activity signals in the Graph, warning notifications are sent to group owners to tell them when their groups need renewal.

Exchange Online sends expiry notifications using email irrespective of how groups are created or used. Group owners must check their mailbox regularly for expiry notifications. If not, they might overlook a notification to say a group is about to expire and Azure AD will remove the group unexpectedly. Teams exposes the expiration date for a team in its settings (*Manage team > Settings > Team expiry*) and allows team owners to renew the underlying group for the lifetime defined in the policy. As explained earlier, group owners can also renew a group through the **Manage groups** section of OWA.

Three warning notifications are sent before Azure AD removes a group:

- 30 days before the group expires.
- 15 days before the group expires.
- One day before the group expires.

For example, if the expiry interval is 365 days (one year), the timeline in Table 11-6 applies. (You cannot change these intervals as they are hardcoded).

<b>Days</b>	<b>Action</b>
1	Group created (or renewed).
335	First expiry notification sent to group owners.
350	Second expiry notification sent to group owners.
364	Final warning sent to group owners.
365	Expiry period reached. Azure AD soft-deletes the group.
395	The 30-day soft-delete retention period expires. Azure AD removes the group permanently.

Table 11-6: Timeline for Group Expiry

When they receive notifications, group owners decide whether to renew the group or let it expire. To help owners decide whether they wish to renew the group, notifications include links to:

- **Outlook:** Open the group and view the conversations that are taking place to see the last time that anyone was active. If the group is used for Teams, none of these events are captured.
- **SharePoint:** Open the document library in a browser to reveal the documents stored there. Because Teams uses the SharePoint library to hold its files, including a folder for each channel in a team, this is a valuable indicator of activity in a team.
- **Teams:** If the group is team-enabled, open the team, and expose the channels, conversations, and apps inside the team.
- **Group details:** Open Azure AD in a browser to reveal information about group members and owners.

## Group Renewal

To renew a group, an owner clicks the **Renew group** button in the notification message. This brings them to the Azure portal to renew the group. Three things can happen:

1. If the group still exists and has not expired, Azure AD renews the group and signals success. An "Updated Group" audit record is captured because the renewal updates the group's *RenewedDateTime* property. The same action occurs when a group is renewed from Teams.
2. If the group is soft-deleted, Azure AD restores it and sets a new expiration date.
3. If the group is hard deleted (permanently removed from the tenant), the owner sees an error message. The group is no longer recoverable.

If you decide not to renew a group and let it proceed to deletion, consider preserving any valuable information that exists in these resources before the deletion happens. This is not an automatic process, and it will take time and effort to retrieve information.

After the expiration period expires, Exchange Online soft-deletes the group and sends a final email to the group owners to inform them about the group deletion. All group resources are removed at this point – the group mailbox, Yammer community, team, plan, and SharePoint site, but kept in a restorable state. A group owner can restore the group for up to 30 days following the deletion. The notification message shows the drop-dead date, and once this period passes, the group becomes irrecoverable.

Normally, Azure AD soft deletes expired groups. However, special processing occurs if groups expire at the point when a tenant enables the policy. In this case, a special form of notification is sent to the owners and treats these notifications as second reminders. In effect, even though their groups are already technically expired, the owners have 15 days to renew these expired groups.

## Restoring Expired Groups

When a group expires, the expiration process moves the group into a soft-deleted state and its owners get a confirmation that Azure AD removed the group. The notification includes a **Restore group** button that the owner can click to restore the group and bring it back from the dead. Group owners can also restore a group through OWA.

Like any other soft-deleted group, Azure AD permanently removes the group after 30 days, so you have limited time to restore a group. Alternatively, you can use the steps described earlier to recover a group.

## Dynamic Microsoft 365 Groups

Microsoft 365 Groups support both static (fixed) and dynamic membership. Static membership is where you add and remove people from group membership via a client or programmatically (for example, with PowerShell). Dynamic groups compute their membership by running a query against Azure AD. Because of the processing load required to maintain dynamic group membership, Azure AD restricts the number of dynamic groups permitted per tenant to 5,000.

Dynamic membership is useful when you have groups that change often and whose membership can be determined by reference to one or more attributes of user accounts. For example, the people who work in the Madrid office, or everyone in the Engineering department. Teams, Outlook, and Yammer support dynamic groups.

The available methods to create a dynamic group are:

- Create the group from the Azure portal, setting its type to be "Microsoft 365" and then setting its membership type to be "Dynamic User."
- Create the group using PowerShell by running the *New-MgGroup* cmdlet. Due to the difficulties of making sure to build correctly-formed complex queries, we recommend using this method only for dynamic groups that use simple queries.

In all cases, when a group has dynamic membership, you cannot update group membership through OWA, Outlook, or the PowerShell *\*-UnifiedGroupLinks* cmdlet, and the only methods available to change membership are:

- Alter the Azure AD query to find another set of users.
- Update the properties of individual accounts in Azure AD so that they come under the scope of the query used to calculate group membership.

As noted above, although you can create a dynamic group with PowerShell, the usual approach is to manage these groups through the Azure AD admin center. Go to the [Groups section](#) and click **New Group**. Make sure that you select *Microsoft 365* for the group type and *Dynamic User* for the membership type and enter a name and description for the new group. The final step is to configure rules for the dynamic query used to find group members.

You have a choice of using a simple rule or an advanced rule to form the query. A simple rule checks against the value of an attribute like Department, City, or Country. An advanced rule combines checks against multiple properties to find the accounts to include in the group and includes some of the more complex checks using operators like *-in*, *-notin*, and *-any*. The *-in* and *-notin* operators are especially valuable in comparing the value of user attributes against a list of values (such as countries or department codes). The maximum size of an advanced rule is about 3,000 characters.

In Figure 11-15, we see an advanced rule based on a modified version of one of the [examples suggested by Microsoft](#). The intention is to find accounts assigned Office 365 E5 licenses, and it's done using a query of:

```
user.assignedPlans -any (assignedPlan.servicePlanId -eq "2f442157-a11c-46b9-ae5b-6e39ff4e5849" -and assignedPlan.capabilityStatus -eq "Enabled")
```

The query checks the service plans assigned to accounts and looks for an enabled license for the plan identified by the GUID. In this case, the GUID is for Microsoft 365 advanced auditing, included in Office 365 E5. It's highly unlikely that Office 365 accounts have disabled advanced auditing, so it's a good check. See the section in Chapter 23 covering license management with PowerShell to understand how to find plan GUIDs.

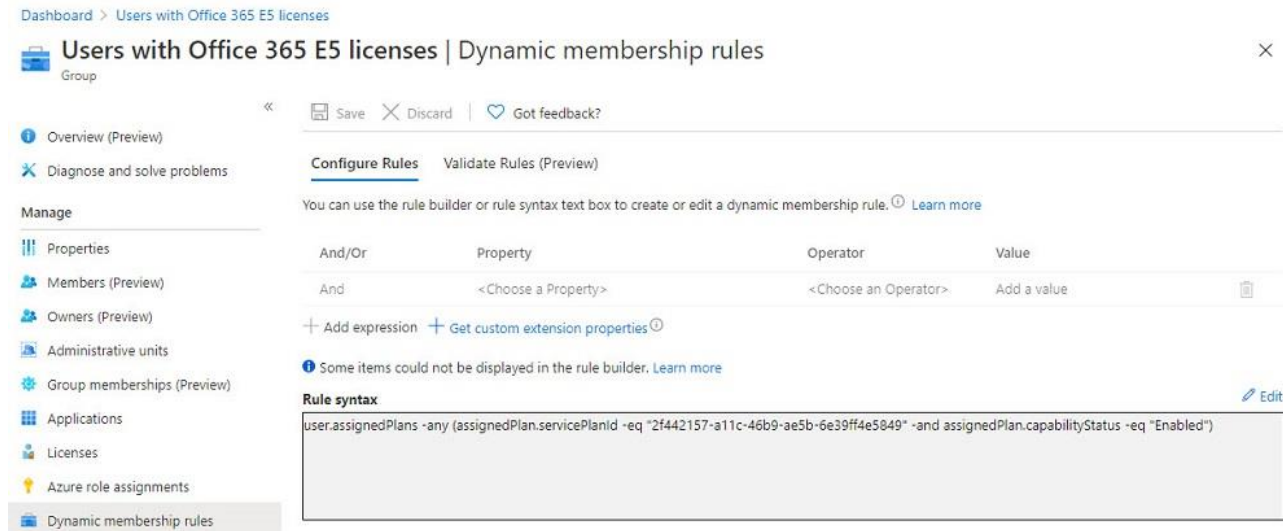


Figure 11-15: Creating rules for dynamic membership of a Microsoft 365 group

The **Validate Rules** option allows you to input accounts to see if they match the criteria set for the group, which is a useful way of checking the effectiveness of the query.

When the query is complete, save it, and create the group. After creation, it can take a little time before the new group shows up everywhere within Microsoft 365. This is because of the need to synchronize information between Azure AD and other workloads. It will also take Azure AD some time to calculate the membership of the group. The two biggest factors are the overall size of the tenant directory and the number of matching entries. Azure AD can easily support dynamic groups of more than 10,000 members and Microsoft has reported (Ignite 2020) that they've created and used a dynamic group containing more than 100,000 members. When Azure AD completes its scan for matching objects, you will see "update complete" as the membership process status in the overview properties for the group together with a timestamp when the Azure AD last calculated group membership calculation. A list of members is available through the **Members** link in group properties. Azure AD does not include disabled accounts when calculating dynamic group membership.

If you update the query for a dynamic group, Azure AD will perform another scan to recalculate group membership. Although the exact time needed for this process to complete depends on the current load on the service, the size of the tenant directory, and the number of matching objects, you can expect that a change to the membership rules should be active within a few hours.

## Checking Dynamic Membership

If you are uncertain about the query or want to confirm that it will generate the expected results, you can use the *Get-MgUser* cmdlet to execute a similar query against Azure AD and check the results that it reports. If the results from both queries match, you know that everything is working as expected. Here is an example of how to query Azure AD with an advanced query:

```
[PS] C:\> Get-MgUser -All -Filter "Department eq 'IT' and Country eq 'Ireland'" | Format-Table DisplayName, Department, Country
```

DisplayName	Department	Country
Andy Ruth (Director)	IT	Ireland
David Pelton (Sales)	IT	Ireland
Luka Abrus (Sales)	IT	Ireland

You can then compare the membership calculated with the filter with the data reported by running the *Get-UnifiedGroupLinks* cmdlet against the group.

```
[PS] C:\> Get-UnifiedGroupLinks -Identity "IT Ireland" -LinkType Member
```

Or, if you want to go back to Azure AD, you can read the membership from the group:

```
[PS] C:\> $GroupId = (Get-UnifiedGroup -Identity "Marketing Department").ExternalDirectoryObjectId
$Members = (Get-MgGroupMember -Group $GroupId)
ForEach ($Member in $Members) { Get-MgUser -UserId $Member.Id | Select DisplayName }
```

## Dynamic Group Owners

Like any other group, a dynamic group must have at least one owner. The ownership of a dynamic group is static in that these members are not affected by the query used to compute the normal group members. Instead, group owners are added and removed individually to the group using the Azure AD admin center or with PowerShell.

## User Access to Dynamic Groups

User access to dynamic Groups works identically to groups that have static membership. Every time someone tries to access a group resource, Azure AD checks whether they have the necessary permission and if so, the user gains access. Like regular Groups, messages sent to a dynamic group exist as conversations in the group mailbox. Groups automatically adds members as subscribers so that they receive copies of all contributions to group conversations via email. One difference between a static and dynamic group is that Groups does not add the user who creates the group to the member list. This is because the query for the group defines the membership. Unless the account that creates the group comes within the scope of the query, the group owner is not a member and therefore cannot access group resources. Likewise, the group owner is not part of the subscriber list.

Groups does not notify members when they lose access to dynamic groups due to a change in the query that populates the group membership or because their account details have changed. The first time someone is likely to discover that they no longer are a member is when they try to access some group resources.

**Hybrid Dynamic Groups:** Groups are unique to Microsoft 365. Only Azure AD can calculate the membership of dynamic groups. For this reason, hybrid deployments always route messages addressed to these groups to Exchange Online. If you enable group writeback with the AADConnect tool, the directory synchronization process ensures that the value of the *targetAddress* attribute for dynamic groups uses the service domain (i.e., onmicrosoft.com) rather than any vanity domain belonging to the tenant. By routing messages sent to these groups via onmicrosoft.com, Exchange Online processes the message through its transport service and queries Azure AD to discover the group membership, and then delivers copies of the message to individual group members. More information about using the AADConnect writeback feature in hybrid deployments is available in the companion volume and [online](#).

## The Cost of Dynamic Groups

Using dynamic groups means that you need to buy Azure AD Premium P1 licenses for administrators who create and update dynamic groups plus any user whose account falls under the scope of a filter used by a dynamic group. In other words, if you create a dynamic group that has a filter that resolves to every person in the organization, you need to buy an Azure AD Premium P1 license for all those users. In most cases, it is better to use a standard dynamic distribution list to communicate with large sets of employees because you



avoid any licensing issues and will not exceed the supported limit for group membership. If you want to use something like an “All Employees” group to foster a sense of collaboration rather than simply a means of sending out information, a public Yammer-based group might be a better choice.

Microsoft does not enforce the licensing requirement for dynamic groups. Those who create or update the queries for dynamic groups need a license before the Azure AD admin center allows them to change queries, but the query used in groups returns all matching accounts found in the directory even if some do not have the necessary license. Microsoft might enforce the licensing restriction in the future. When that happens, queries will return licensed accounts only.

Because dynamic groups come at a cost, you should be careful to decide which groups should be dynamic, and which can stay with static membership. Before creating a new dynamic group, ask whether the membership is likely to change over time. If you expect a low turnover in group membership, then static membership is a better choice as it will not incur an extra cost and the membership is not dependent on the directory. On the other hand, organizations that manage many groups whose membership is volatile (such as the students who sign up for a class) and can calculate group membership by querying one or more of the supported directory attributes for dynamic membership might be able to justify the cost.

**Using Dynamic Groups with Teams and Planner:** You can use a dynamic group with Teams but not with Planner, which depends on fixed team membership. Dynamic groups are intended for groups whose membership is highly volatile and whose communication is email centric.

## Comparing Email Dynamic Distribution Lists and Dynamic Microsoft 365 Groups

Experience in using Exchange dynamic distribution lists suggests that dynamic membership is a popular feature. Some obvious differences exist in the implementations of the two types of dynamic groups. Table 11-7 compares the two ways to create objects with dynamic membership inside Office 365.

<i>Attribute</i>	<i>Dynamic distribution list</i>	<i>Microsoft 365 dynamic list</i>
Resolved against	Exchange Online Directory (EXODS)	Azure Active Directory
Used for	Email	Teams, Outlook Groups, Yammer
Purpose	Send email	Send email and manage access to Microsoft 365 resources like SharePoint sites
Supported objects	Any mail-enabled recipient type (including hybrid objects)	Azure AD user and guest accounts
Licensing	Included in Exchange Online	Azure AD Premium P1
Syntax for query rules	OPATH	ODATA
Filters based on	<a href="#">Exchange object attributes</a>	<a href="#">Azure AD object attributes</a>

Table 11-7: Comparing Exchange Online dynamic distribution lists and dynamic Microsoft 365 groups

## Groups and Compliance

As explored elsewhere in this book, compliance technology exists to help companies follow the regulatory and legal frameworks that apply to their business activities. Groups use Microsoft 365 compliance features rather than workload-specific features. For example, you can include group mailboxes and the group document libraries in simple content searches, or the searches associated with eDiscovery cases. Retention policies can apply to content held in group mailboxes and document libraries and group owners can apply

retention labels to conversations to keep those items for compliance purposes. Table 11-8 outlines the functionality supported by Groups in different search and hold scenarios.

	<b><i>Include group mailbox</i></b>	<b><i>Include group document library</i></b>	<b><i>Create in-place hold on data</i></b>
Retention policies	Can keep or remove selected group conversations or Team channel messages	Can keep or remove selected files in group document libraries	Can impose in-place hold to keep items for defined periods
Content searches	Yes	Yes	Yes
Core and Premium eDiscovery cases	Yes	Yes	Yes

Table 11-8: Compliance features applied to Groups

Yammer-based groups hold their discussions in the Yammer data store. Compliance records for eDiscovery are generated if the Yammer network is configured in Microsoft 365 native mode.

**Controlling Group Creation and Compliance:** One aspect of compliance sometimes overlooked is the desirability of controlling group creation. If you allow everyone to create groups, you might end up with an unmanageable mess. Not only will obsolete and unwanted groups clutter up the tenant, but their presence does make it harder to figure out what groups hold information needed for compliance purposes. If compliance is of concern, consider applying a policy to control group creation so that you know what groups exist, their purpose, and whether they should come under the control of compliance policies.

## Yammer and Groups

Microsoft bought Yammer for \$1.2 billion in June 2012 to gain a presence in the enterprise social networking market. Facebook defines the concept of a social network to many people, a feeling emphasized by the film of the same name that charts the development of Facebook from an application used to connect college students in an electronic yearbook to a platform that supports billions of users today. Facebook is a public social network that is open to anyone who cares to join. An enterprise social network, like Yammer, offers much the same communication and collaboration facilities as you find in Facebook, but inside the confines of a single enterprise. Of course, you can stretch the concept of a single enterprise to a public space, as in the case of Microsoft's original Office 365 Network where Yammer hosted over 88,000 members discussing technical issues and ideas in the many groups that reflect the different interests of customers. Microsoft transitioned the Office 365 Network from Yammer to a new platform in September 2016.

Microsoft enables Yammer for all enterprise tenants, and the app supports browser, Apple iOS, Android, and Mac clients. Once activated, you can use Yammer in diverse ways. Some find that Yammer is an excellent way to introduce users to cloud applications. Yammer behaves very much like other social networking networks and users who are familiar with those applications can quickly become productive with Yammer. Given a well-curated set of groups for people to go to find and share information, Yammer can quickly become a success factor in a roll-out, especially when the number of potential users who might contribute to conversations is more than can be accommodated by Groups or Teams.

Yammer supports sharing across groups (in other words, you can cross-post a conversation to have it appear in multiple groups). It is also possible to search for content across the groups to which you belong. And of course, Yammer comes with the compulsory marks of approval seen in other social networking products - the infamous "Like" and emoticons. Yammer also boasts integrations with various other applications such as Dynamics 365 and Salesforce.com.

Yammer brings its distinct blend of collaborative capabilities to the mix. It is collaboration writ large, created with the enterprise in mind. Yammer is best when used by hundreds or thousands of contributors meeting in groups dedicated to different topics, especially when contributors are geographically remote from each other. Some examples of Yammer communities are cross-company knowledge sharing or distributing information on behalf of specific parts of the business, such as an HR discussion group or an IT suggestions group.

Any consideration of collaboration should take Yammer into account. You should lay out the needs of the organizations and map those needs against the different Microsoft 365 applications available. You can then decide which type of group is best suited to specific scenarios. Although the question seems a tough one to answer, the reality is that most organizations will find the choice relatively straightforward. Most tenants will standardize on a certain type of collaboration and use it wherever possible. Those who have used email in the past will continue to do so with Outlook-based groups while organizations that deployed Yammer will continue to use Yammer communities. Teams has replaced email for many internal communication scenarios and is often easier to deploy than Yammer, so it's another factor to take into account.

## The Evolution of Yammer

Anyone looking at the collaboration options available will probably conclude that some overlap exists. Why use Groups instead of Yammer or vice versa? How do these applications compare to Teams, which some commentators speculate will eventually replace Yammer? These applications allow users to contribute to discussions and preserve contributions so that they are accessible to other users who join the debate after it starts. The applications support the sharing of files and are accessible through a browser interface – and that is pretty much where the similarities end. Yammer's focus is on social networking capabilities for large enterprises (the largest customer Yammer networks support well over 100,000 members) while Groups and Teams support much smaller memberships.

Because Outlook groups can act as distribution lists, they are easy to deploy. Users who interact with Outlook Groups through email and never access the shared resources will probably think of them as another form of distribution lists. You send messages to groups and copies arrive in the mailboxes of any member who subscribes to the group. Given the large population of Exchange Online mailboxes, Outlook groups have a lot of room to grow.

Yammer is different. You can certainly interact with Yammer communities through email, but the experience is less seamless. Although Microsoft enables Yammer for enterprise plans and does the heavy lifting of implementation, the decision to deploy Yammer inside a company needs more thought and preparation than simply turning it on. Such an approach will invariably result in a couple of active groups (for some period) and a lot of stagnation. Any project to introduce a new form of collaboration to an enterprise, even though people are aware of the power of social networking through their exposure to consumer versions, needs planning, evangelism, leadership, endorsement from executive level, and a whole lot of energy applied to get the network going, including the identification of suitable ways to use the network to solve real-life business problems. An electronic chat room is nice to have but worthless in the long term; an active social network that brings people together on an ongoing basis to debate, refine, and drive solutions to identifiable and quantifiable business issues is invaluable.

## Yammer Network Modes

Yammer organizes its communities into a network. Older Yammer deployments could support multiple networks within a single tenant, but now it's more common to find a single network per tenant. The network mode describes the functionality available to the network and can be one of three modes:

- [Native mode for Microsoft 365](#). As of January 2020, all new Yammer enterprise networks start in native mode (and can't change to any other mode).

- Non-native mode. All Yammer external networks are in this mode. External networks are those which support people outside the organization.
- Hybrid mode.

Best results with Microsoft 365 are possible when a Yammer network runs in native mode:

- Microsoft 365 administrative tools can manage Yammer users and communities. Groups manage the membership of Yammer communities and Yammer respects the group creation policy. A special group is created for the “*All Company*” community with no membership. All tenant administrators are added as owners of this group. In the case of other communities, Yammer network admins must have the appropriate Microsoft 365 permissions to manage the membership of communities where they are not listed as group owners. The Groups used by Yammer are inaccessible to clients of other group-enabled applications (like Outlook).
- Each community can access resources like a SharePoint Online team site. The files uploaded to a Yammer community are stored in a document library on its team site.
- All communities can use Microsoft 365 Live Events.
- Yammer generates compliance records for conversations. The compliance records for a discussion are in the group mailbox belonging to the Yammer community and can be discovered by content searches.

Moving a Yammer network to native mode is not something that you do on a whim. It is a one-time irreversible process performed using the Microsoft 365 alignment tool. The tool checks all the settings in a network and makes whatever changes are needed to transition to native mode. For instance, all Yammer users must have Azure AD accounts, all communities must be connected to Groups, and all files must be stored in SharePoint Online. If necessary, the tool makes changes to objects or removes items (like external users or users who can't be found in Azure AD) to make the network compliant. When the process is complete, you cannot change to another mode. See [this page](#) for more information.

## Yammer Communities

Yammer traditionally referred to its collection of users as “groups.” In late 2019, Microsoft announced that Yammer would now use “communities” instead as this word more accurately reflects the kind of knowledge sharing and community building that Yammer is good at. Communities still use Groups to manage their identities and membership and connect to group-provisioned resources like SharePoint sites.

When a Yammer community uses a Microsoft 365 group, you can work with the group using the Groups cmdlets in the Microsoft Graph PowerShell SDK and Exchange Online modules, or by using Graph API queries. For example, you can set the classification or access type for a Yammer community by running the *Set-UnifiedGroup* cmdlet. You can also recover a deleted Yammer community using the process to recover a deleted Outlook group. Although you can add members to a Yammer community using the *Add-UnifiedGroupLinks* cmdlet, it is best to manage membership through Yammer and this minimizes any directory synchronization lag.

Yammer communities in networks configured in Microsoft 365 native mode support Azure B2B collaboration guest user accounts. The functionality only works for external accounts in other Office 365 tenants. Microsoft hopes to support external access for Microsoft Service Accounts soon.

## Yammer Discussions

Leaving Microsoft Teams aside, tenants have a choice of technology to host group discussions: Yammer or Exchange (Outlook Groups). The major difference between Yammer communities and Outlook Groups is that Yammer keeps its conversations in the Yammer data store while an Outlook group stores its conversations as items in the Inbox folder in a group mailbox. A Yammer community that uses the Groups service is easily identifiable because it lists links to the group-connected resources provisioned for the group (Figure 11-16).

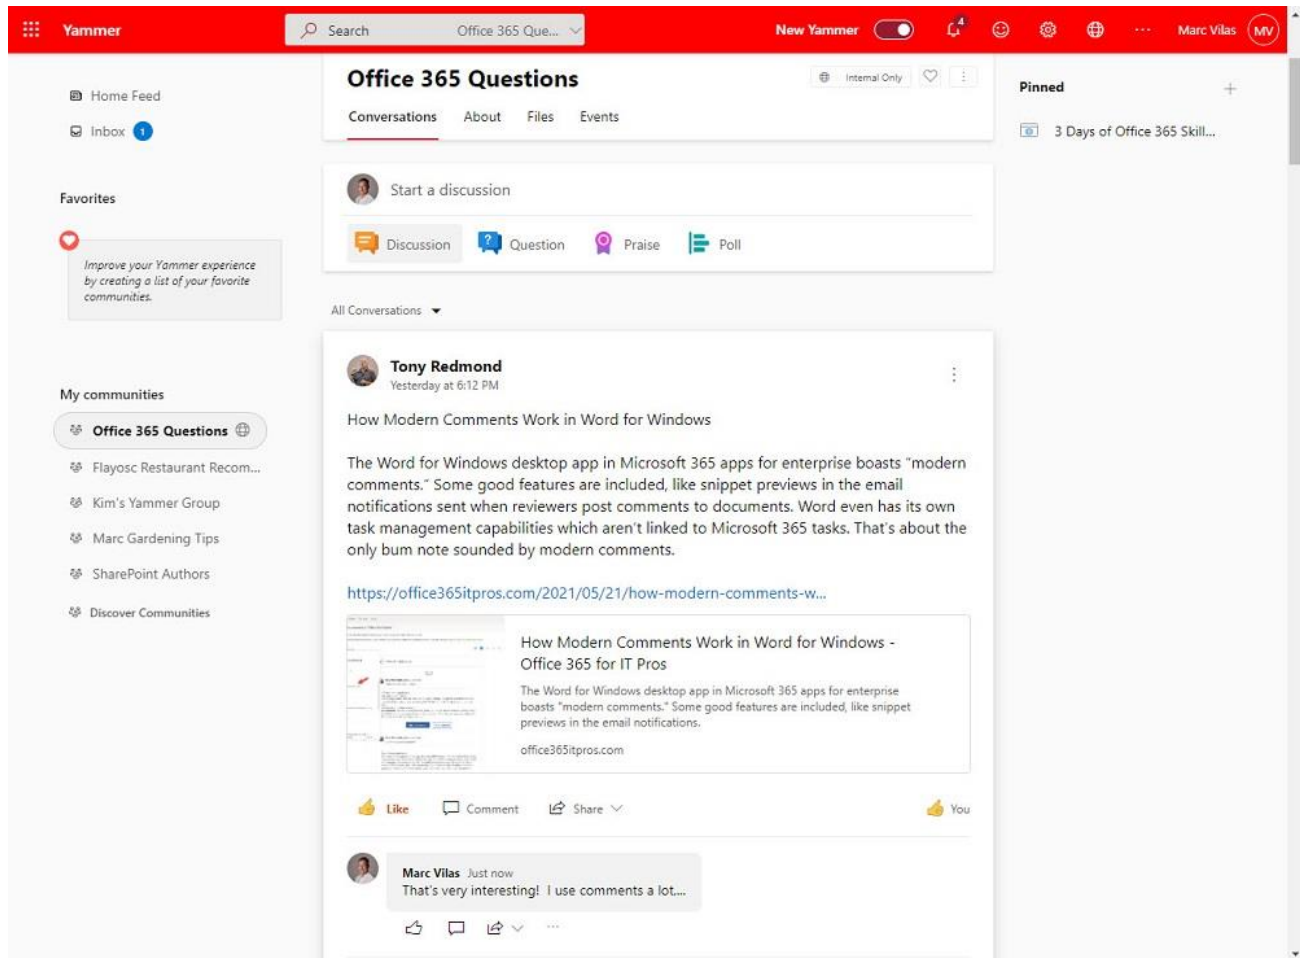


Figure 11-16: A discussion in a Yammer community

Yammer supports marking communities as “official,” as a way of increasing the status of a community and the discussions contained there. For instance, an organization might mark an HR Community as official to encourage people to ask questions relating to HR policies and regulations there. Among the other features Yammer supports are:

- Mark conversations as announcements to give them extra importance. Community members receive email notifications for announcements even if they disable email notifications for Yammer.
- Threaded conversations.
- Delegate access (users can allow others to post on their behalf in communities both people belong to).
- Feature a conversation (pin it to the top of the feed).
- Ask and answer questions. You can also change a regular conversation to become a question. Yammer uses different formatting to highlight questions in user feeds.
- View insights (statistics) about a conversation.
- Restrict posting of new topics in communities to administrators.

Although you cannot transform a Yammer community to become an Outlook group or vice versa, it is possible to run the two types of groups alongside each other. After all, both use a common group framework, and the only major differences are the storage of conversations and the clients that people can use to access a group. Outlook and OWA cannot access conversations in a Yammer community. Access to conversations in a Yammer community is only possible through the Yammer browser and mobile clients. Although they do not use Exchange Online for messaging, members of Yammer communities can still send messages and users can also send messages to Yammer communities.

## Actionable Notifications

As conversations happen in Yammer communities, members can choose to receive notifications via email. Traditionally, the notifications were simple messages with a clickable link to open a browser and go to the conversation. Notifications include a *“fully-interactive Yammer thread”* with the ability to:

- See the complete Yammer conversation.
- Like a comment in the conversation.
- Share the conversation with another Yammer community.
- Know how many people have already seen the conversation.
- Add people to a conversation.
- Post a comment to the conversation from OWA, including attaching files and images or @mentioning people in a comment.
- Vote in polls.

Actionable notifications only work in OWA. Other clients see the older-style notifications.

## Yammer Compliance Records

When a network is configured in native mode, Yammer stores compliance records in the *MessageIngestion\Yammer* folder in the non-IPM part of user and group mailboxes. The records are created by the Microsoft 365 substrate when people post new messages to Yammer communities or private conversations. Records are available for different forms of Yammer posts, including polls, questions, discussions, and so on. Graphics and GIFs included in posts are captured. Reactions to posts are the only element that the substrate does not capture.

The example below uses the *Get-ExoMailboxFolderStatistics* cmdlet to examine and report what it finds in the Yammer folder:

```
Get-ExoMailboxFolderStatistics -Identity "Company-Wide Debate" -Folderscope NonIPMRoot -
IncludeOldestAndNewestItems | ?{$_.FolderType -eq "Yammer"} | Format-Table Name, ItemsInFolder,
NewestItemReceivedDate
```

```
Name      ItemsInFolder NewestItemReceivedDate
-----
Yammer    6213 1 Aug 2021 12:12:17
```

When run against a user mailbox, the code reports the number of Yammer compliance records created for posts to communities and private messages sent by that user. Using compliance records is often used to track activity in groups. This script in the [Office 365 for IT Pros GitHub](#) repository generates an activity report for Yammer communities using the compliance records, including calculating how long it has been since anyone posted to a community.

## Data Residency for European Union Tenants

In May 2019, Microsoft announced that new European Union tenants could choose to have their data stored in Microsoft data centers in Europe. Before making this decision, organizations should be aware of the [current restrictions](#), some of which have a significant impact on Yammer functionality.

## Evaluating Yammer, Groups, and Teams

Occasionally, Microsoft 365 offers too many ways to get work done. Selecting the right tool for team collaboration is one of those times. Groups offer a solid choice for teams to work together based on email while Microsoft Teams delivers a chat-based workspace that works well for small to medium teams. Groups, Teams, and Yammer support guest user access (the Yammer network must run in Microsoft 365 mode).

Because many factors drive the choice of the best collaboration tools for an organization, it is hard to make a choice and say that any one tool will handle all circumstances and contexts. No collaboration tool ever invented in the history of technology has been able to make that claim in any credible sense. To help you figure out the best choice of tools for use inside your tenant, Table 11-9 lists some of the points that you can consider when comparing the different group/team-centric collaboration methods supported by Microsoft 365.

<i>Characteristic</i>	<b>Groups</b>	<b>Yammer Communities</b>	<b>Microsoft Teams</b>
Built on Microsoft 365 technologies	<i>Exchange Online</i> (conversations and calendar). <i>SharePoint Online</i> (team site) <i>OneNote</i> (shared notebook). <i>Planner</i> (plans) Connectors (if enabled).	<i>Yammer</i> (conversations). <i>Exchange Online</i> (group mailbox). <i>SharePoint Online</i> (team site). <i>OneNote</i> (shared notebook). <i>Planner</i> (plans).	Azure data services (messages and media items). <i>Exchange Online</i> (calendar). <i>SharePoint Online</i> (document library). <i>OneNote</i> (shared notebook). <i>Planner</i> (plans). Connectors, Bots, Apps, and Tabs (all if enabled).
Clients	Outlook for Windows or Mac (Microsoft 365 Apps), OWA, Outlook mobile.	Browser, Windows and Mac desktops, and mobile clients.	Browser, Windows, Mac, and Linux desktops, and mobile clients.
Scalability	Maximum of 2,500 members (dynamic groups are also supported).	The membership limit is > 50,000 (unless synchronized with an on-premises directory).	Membership limit of 25,000 per team. Dynamic groups are supported.
Access	Default private but can be set to public.	Default public (any user can join), can be set to private.	Default private but can be set to public.
Communication style	Choice of direct access to conversations or by sending email to the group. Contributions can be short or long. A group can host thousands of conversations.	Threaded discussions that occur over wide-ranging topics. Contributions are usually short and entered directly through the browser. Interaction through email is possible via actionable messages.	Support for personal and shared (channel) messages organized in threaded conversations, which are usually very short. Interaction through email is possible via actionable notifications. Email support for mailing to channels is also available.
Search	Outlook constrains search to conversations within a single group. Office.com supports searching across all groups and teams to find messages and documents.	Search across Yammer communities to which a user has access.	Search within all teams, a single team, and a channel (including associated files). Office.com search finds messages and documents.
Group purpose	Team-based collaboration centered on email and shared documents. Good choice for teams that work together on an ongoing basis for a prolonged period. Replacement for	Wider (up to the entire company) collaboration and sharing of ideas. A good choice is to host open discussion forums, groups dedicated to sharing knowledge and	Teams working on projects. A good choice for teams that need to work on a specific problem or issue. A team can have up to 200 channels to use to segment discussions.

	traditional email distribution lists.	information with a broader community, and when a “community of interest” form of collaboration is required.	Private channels are available for a subset of team membership. Shared channels are available for collaboration with external users.
Management	Microsoft 365 admin center and EAC. A wide range of PowerShell cmdlets and Graph APIs are available to interact with mailbox and group.	Yammer administration console. No PowerShell support, apart from being able to use some of the Groups or PowerShell SDK cmdlets to manage group properties. No Graph API support for access to Yammer content.	Teams admin center. A Teams PowerShell module is available, and management of team-enabled groups can be done using the cmdlets in the Exchange Online module. Good support in Graph APIs.
External access	Support for tenant and guest users with some limitations for users with on-premises mailboxes.	Support for tenant and guest users (for networks configured in Microsoft 365 mode).	Support for tenant and guest users (and external users with shared channels) with some limitations for users with on-premises mailboxes.
Support for Microsoft 365 compliance and data governance features	Support for content searches, eDiscovery cases, retention labels, and retention policies. Actions are captured in the audit log.	If the network runs in Microsoft 365 mode, Yammer creates compliance records for discussions. Files stored in SharePoint are indexed and discoverable. Actions are captured in the audit log.	Compliance records are captured in Exchange mailboxes. These are indexed and discoverable and can be included in content searches and eDiscovery cases. Teams retention policies are supported. Teams supports DLP. Actions are captured in the audit log.
Data sovereignty	Supported in all Microsoft 365 data center regions. Multi-geo deployments are supported.	Hosted in U.S. and EMEA data centers.	Hosted in most Microsoft 365 data center regions. Multi-geo deployments are supported.

Table 11-9: Comparing different Microsoft 365 collaborative applications

Some guidelines are necessary to help administrators and users select the best form of collaboration technology for a specific task. The following might serve as a basis for discussion:

- An Outlook group is an excellent choice when a team requires a collective repository where they can share documents, a OneNote notebook, and a calendar, and discuss topics through email. It is likely that the group has a well-defined purpose and might be part of the organizational structure, like a department or location. Because these groups often serve as a direct replacement for traditional distribution lists, the email traffic can be relatively high compared to the number of files generated in the document library. Together, the files and conversations held in the group represent its collective history and are valuable in terms of bringing new members up to speed with the group’s activities. Groups and the Outlook apps offer the most comprehensive support for the Microsoft 365 data governance framework.
- Microsoft Teams is a good choice for small to medium projects where sets of people come together to execute specific tasks. Deadlines are likely to be in place for the team to achieve and this



encourages a certain kind of focused and rapid conversation between team members that Microsoft refers to as “high-velocity chats”. The chats tend to be short and informal. The intention behind a team is to get work done rather than to accumulate knowledge. When the project or work item completes, the team members might disband and come together in other teams.

- Yammer communities are an effective way to share knowledge broadly within large communities. They work well as public forums within companies, where Yammer can be the basis for communities who share a common interest. Unlike Teams and Groups, which configure their groups to be private or public, Yammer communities often host unstructured audiences that flex over time. The information held in Yammer communities is often not as specific as found in Groups or Teams but is no less valuable. Indeed, the knowledge developed through the contributions of the higher number of users supported by Yammer often has long-lasting value for the entire company.

Considerable overlap exists between the various methods. An Outlook group might be as good a choice as a team in some circumstances. Some companies use nothing but Yammer and are quite happy with their choice. Some will find that Teams revolutionize the way that small groups work. All of which proves that collaboration is uniquely individual to a company or organization. No one size fits all, which is why choice exists within Microsoft 365.

## Distribution Lists

Distribution lists or distribution groups have existed in email systems since the early days of the technology. A distribution list is a collection of known email recipients which gives users the ability to send emails to everyone on the list through a single recipient. The membership of a distribution list is composed of mail-enabled recipients registered in the Exchange Online directory and accessible through the GAL:

- User and shared mailboxes.
- Mail-enabled public folders.
- Other distribution lists, including dynamic distribution lists.
- Mail contacts. You can include a Teams channel in a distribution list by creating a mail contact using the channel’s email address.
- Mail users (including those created for guest accounts).
- Microsoft 365 Groups (these objects can only be added using PowerShell).

A single distribution list can span up to 100,000 members. The Exchange Online transport system calculates the total recipient count after expanding the distribution list, including any nested distribution lists and dynamic distribution lists in the membership. If you use Azure AD Connect to synchronize with an on-premises directory, a lower maximum of 50,000 members applies. After it expands the set of individual recipients, Exchange figures out how best to route a copy of the message to each recipient. It is at this point that Exchange imposes other limits, such as the 25 MB maximum message size that can be sent to distribution lists of up to 99,999 members.

If your tenant uses large distribution lists (more than 5,000 members), Exchange forces you to configure delivery management restrictions for those lists to lessen the possibility of people sending messages to the lists in error. Sending a 5 MB message to a distribution list containing 100,000 recipients imposes a huge load on the service. Limitations include specifying a set of mailboxes allowed to send messages to large lists or imposing moderation on those lists so that messages must be approved before Exchange distributes copies to all members.

Normally, the sole purpose of the distribution list is to act as a convenient way to send messages to multiple people. Alternatively, a distribution list can be a mail-enabled security group (MESG), which means that you can use it to send mail and control access to resources, such as access to a shared mailbox. You cannot change a distribution list to make it a mail-enabled security group or vice versa.

The perceived wisdom expressed by Microsoft and many others is that Microsoft 365 Groups are more functional and productive than distribution lists and that users will benefit if organizations upgrade distribution lists to become Groups. For this reason, Microsoft offers multiple methods to convert distribution lists. Notwithstanding the features available to Groups, valid reasons still exist to continue using distribution lists, including:

- The ability to include different mail-enabled recipient types in group membership. For instance, mail contacts, mailboxes, and other groups.
- No licensing requirement to use dynamic distribution lists.
- Easier interoperability with other third-party email systems.
- Ability to nest distribution lists.

Microsoft is aware of where the functionality of Groups lags behind the capabilities of distribution lists and is working to close the gap, so this is an area to monitor over time.

## Creating a New Distribution List

The Groups section of the EAC offers the opportunity to create a new Microsoft 365 group, a distribution list, mail-enabled security group, or a dynamic distribution list. The group type tells EAC which UI to display to gather information about the properties of the new group. When you choose a distribution list, the GUI displays some text to convince you to create a new Microsoft 365 group instead, which is fine if you need the extra features like a SharePoint Online site.

To simplify the creation of distribution lists, just a few settings are needed, divided into Basics and Settings. The two basic settings are:

- **Name:** The *Name* property of the distribution list must be unique. If you don't supply a value for the *DisplayName* property, it takes the value assigned to the *Name* property. The *DisplayName* is the one listed in the GA and it's important to give some thought to its composition. Ideally, the name should convey the reason why the group exists, how to use it, and what its membership might be. Groups created by administrators are not subject to the organizational group naming policy as this only applies to user-created groups. We will discuss how to create a group naming policy and how to restrict users from being able to create groups later.
- **Description:** An optional free-text description to explain the purpose of the distribution list. In some cases, administrators note who requested the creation of the list and capture some background for its expected use.

The settings dictating the basic operation of the distribution list are:

- **Group mail address** (Primary SMTP address): To ensure uniqueness, I typically create the username part of the mail address from the name of the distribution list (separating words with full stops). The domain can be chosen from any of the tenant's registered domains. The list also receives an alias from the tenant's service domain.
- **Communication:** A checkbox to note if people outside the tenant can send emails to the list.
- **Joining the group:** This can be open (anyone can join), closed (members must be added by a list owner), or owner approval (people can apply to join but must be approved).

Once the EAC creates the distribution list, you can go ahead and add its membership and amend other details of the list (Figure 11-17). The person who creates the distribution list is automatically an owner but is not a member. The ideal situation is to have two or more owners. This is more important with a security group because, apart from administrators, only owners can add or remove members from a security group. You can choose to add owners to the membership so that they receive messages sent to the group, but this is not necessary. The ownership of a group can include other groups as well as individual users.

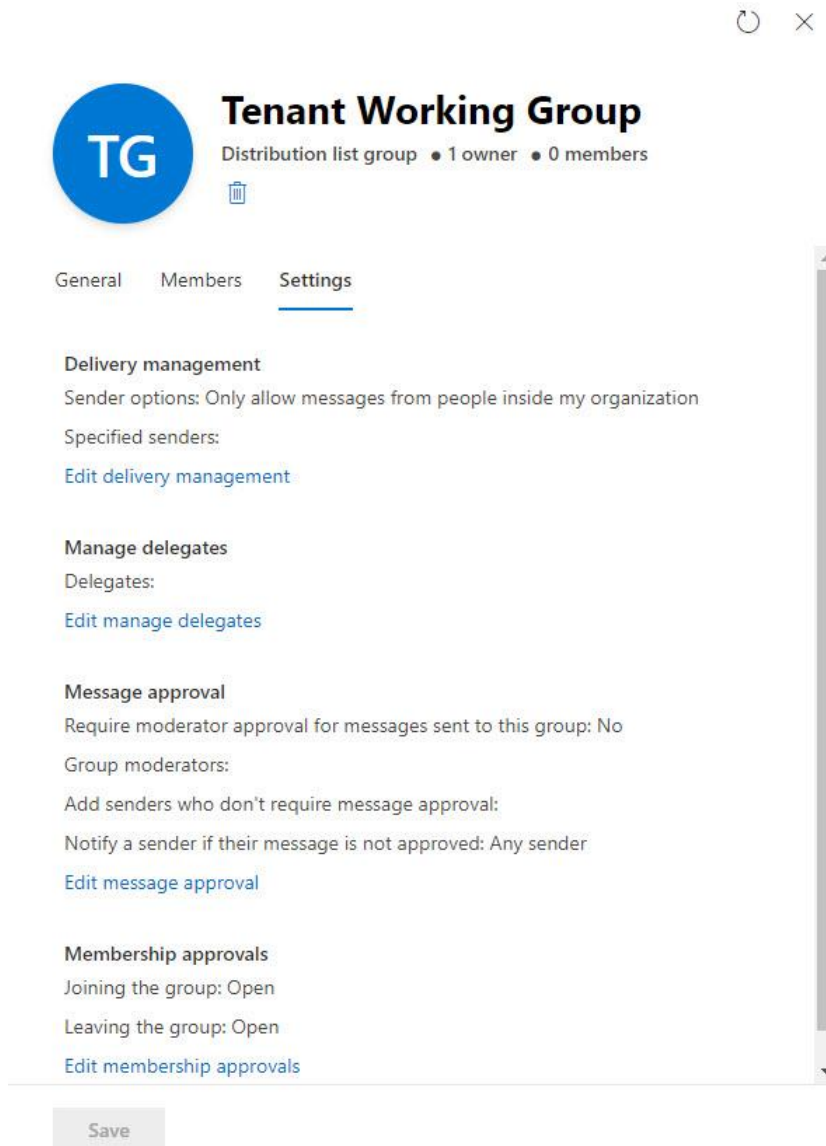


Figure 11-17: Amending the settings for a distribution list

Users can browse available groups that they belong to or which they own through the **Distribution groups** section of OWA options. Options are also available to join a group and leave a group. Joining a group involves a user first searching the directory to find the group and then creating a request to join it. Exchange automatically approves requests to join open groups and rejects them for closed groups. Otherwise, Exchange sends requests to join these groups to the group owner(s) for approval.

The *MemberJoinRestriction* and *MemberDepartRestriction* properties control how users can join or leave distribution lists. You can update a distribution list with the *Set-DistributionGroup* cmdlet. This example sets the group properties so that an owner must approve requests to join but a user can leave without approval.

```
[PS] C:\> Set-DistributionGroup -Identity "All Company" -MemberJoinRestriction ApprovalRequired -MemberDepartRestriction Open
```

The options to control user-initiated joining and leaving of a group do not exist for either mail-enabled security groups or dynamic distribution lists. Mail-enabled security groups are security principals that can grant access to other resources. It makes sense that only the group owners should be able to control who enjoys that access. Exchange Online calculates the membership of a dynamic distribution list by running a query against the Exchange directory service. An administrator can decide whether the results of a query will

find a specific object by adjusting the object's properties (for instance, they change the value of an attribute used by the query), but unless the owner is also an administrator, they cannot influence the query results.

Other properties of distribution lists often updated with PowerShell include:

- **RequireSenderAuthenticationEnabled:** The default is *\$True*, meaning that Exchange Online will deliver messages to the group only when the sender can authenticate (the sender belongs to the tenant). Exchange rejects messages generated by external senders unless this property is *\$False*. Normally, you do not want external people sending emails to internal distribution lists, but some good exceptions exist. For example, when a distribution list route messages from external parties (like customers or partners) to a public folder or Teams channel.
- **RejectMessagesFrom, RejectMessagesFromDLMembers, RejectMessagesFromSendersOrMembers and AcceptMessagesOnlyFrom, AcceptMessagesOnlyFromDLMembers, AcceptMessagesOnlyFromSendersOrMembers:** Exchange Online allows for reasonably granular control over who can send to a distribution list. You can block specific users or members of other distribution lists with the Reject\* properties or block everyone except those specified in the Accept\* properties. The EAC populates the Accept\* properties when you edit a distribution list to update its delivery management parameters. All the objects specified in these lists must be known to Exchange Online and can include mail contacts and mail users. If you try to add a user via their SMTP address, Exchange Online checks the address against the Exchange directory service and rejects it if the address does not match a mail-enabled object.

When updating a reject or accept list for a group, you must update the relevant *\*SendersOrMembers* property and Exchange Online will sort out the individual object types and update the properties accordingly. For instance, here is how to update a group so that it will only accept messages from two mailboxes and the members of another group:

```
[PS] C:\> Set-DistributionGroup -Identity "Executive Committee"
-AcceptMessagesOnlyFromSendersOrMembers "Ben Owens", "CEO Mailbox", "Vice Presidents"
```

- **GrantSendOnBehalfTo:** Controls who has the right to send messages on behalf of the distribution list.
- You can also assign the **SendAs** right to a user for a distribution list. For example, here is how to assign the *SendAs* permission for the Executive Committee distribution list to the CEO Mailbox user.

```
[PS] C:\> Add-RecipientPermission -Identity ExecCommittee -Trustee "CEO Mailbox" -AccessRights
SendAs
```

**Note:** To protect resources, Exchange Online [enforces limits](#) on the number of recipients that can exist for an individual message (500) and the total number of recipients that a mailbox can send messages to in a single day (10,000). These specific limits exist to prevent tenants using Exchange Online from as a bulk email service. When Exchange Online checks messages against these limits, distribution lists, and dynamic distribution lists count as a single recipient, even if the use of these groups means that a message addresses many more recipients than allowed by the limit after the transport service expands groups into individual addressees.

## Blocking BCC Delivery to Distribution Lists

Distribution lists can often become very chatty, which leads members to create inbox rules to filter messages sent to the list to stop emails from cluttering up their inbox. People can get around inbox rules by addressing the distribution list as a BCC recipient. Inbox rules can process TO and CC recipients, but not BCC recipients because these recipients are not present in the message header.

To solve the problem, you can set the *BccBlocked* property of a distribution list to `$True`. By default, the value is `$False`. When set to `$True`, the transport service blocks any message sent to the distribution list as a BCC recipient and the sender receives a non-delivery notification with code 5.7.138. To block BCC deliveries to a distribution list, run the *Set-DistributionGroup* cmdlet:

```
[PS] C:\> Set-DistributionGroup -Identity "Message Board Posting" -BccBlocked $True
```

## Reporting Distribution List Membership

You can use the EAC, the People app, or Outlook to view the membership of a distribution list, but neither interface does a good job with distribution lists of more than 20 or so recipients. The best experience is through the People interface in OWA, which makes it easy to scroll through large memberships. However, none of the interfaces provided by Microsoft allow you to print off the membership of a distribution list, so unless you invest in a third-party reporting product (see the discussion about standard usage reports in Chapter 21), it's necessary to create custom reports using PowerShell. For example, this one-liner extracts all the members of the Executive Committee distribution list and outputs some of their details into a CSV file.

```
[PS] C:\> Get-DistributionGroupMember -Identity "Executive Committee" | Select-Object Name,
DisplayName, PrimarySmtpAddress | Export-CSV "c:\temp\ExecutiveCommittee.csv" -NoTypeInformation
```

A script to create a report of distribution list membership [is available on GitHub](#).

## Updating Distribution List Membership with PowerShell

The membership of a distribution list is composed of backward links to the EXODS objects for the mail-enabled recipients that make up the membership of groups. EXODS is used instead of Azure AD because some mail-enabled objects which can be members of distribution lists, like mail-enabled public folders, do not exist as Azure AD objects. To view the membership of a distribution list, use the EAC or the Microsoft 365 admin center. If you view it through the Azure AD admin center, objects not in Azure AD are missing.

So much for reporting. Updating group membership is more interesting. We can use the *Update-DistributionGroupMember* cmdlet to update the complete group membership at one time, overwriting anything that might already exist. As shown below, the input to the cmdlet is a comma-separated list of identities (aliases, display names, or email addresses) for the group members. Make sure that Exchange can resolve all the identities. If even one is erroneous, the update will fail.

```
[PS] C:\> Update-DistributionGroupMember -Identity "Blog Writers" -Members @("TRedmond", "Bowens")
```

Alternatively, you can use the *Add-DistributionGroupMember* cmdlet to add a single object to a group or *Remove-DistributionGroupMember* to remove a member. For example:

```
[PS] C:\> Add-DistributionGroupMember -Identity "Blog Writers" -Member VanHybrid
Remove-DistributionGroupMember -Identity "Blog Writers" -Member VanHybrid -Confirm:$False
```

Here's another example of a script to add a mailbox to a set of distribution lists. A single-column CSV file holds the aliases of the target distribution lists. After checking that the input mailbox is valid, the *Import-CSV* cmdlet imports the set of distribution lists from the file and stores it in an array. A loop then moves through the array to add the mailbox to each of the target lists.

```
[PS] C:\> $CheckName = Read-Host "Enter Name of mailbox to add"
Try {
    $Mbx = Get-ExoMailbox -Identity $CheckName -ErrorAction Stop | Select -ExpandProperty
    PrimarySmtpAddress}
Catch {
    Write-Host "No mailbox can be found called" $CheckName; break }

$DLs = Import-CSV "C:\Temp\InputDLs.csv"
```

```
ForEach ($DL in $DLs) {  
    Try {  
        Add-DistributionGroupMember -Identity $DL."DL Alias" -Member $Mbx -ErrorAction Continue }  
    Catch {  
        Write-Host "Couldn't add" $Mbx "to DL" (Get-DistributionGroup -Identity $DL."DL  
Alias").DisplayName }  
}
```

Sometimes people change jobs and need to transfer membership of several distribution lists to another person. This is easily scripted with PowerShell, as evident [in this example](#).

## User Updates of Distribution Lists

Users with cloud mailboxes can edit the membership of the distribution lists that they own using Outlook or OWA. However, distribution list owners should select the online Global Address List rather than the Offline Address Book (OAB) as this will ensure that up-to-date group membership is available and that updates occur quickly. You can still update group membership through the OAB, but everything seems to happen much slower.

Things are much trickier in a hybrid environment where those with Exchange Online mailboxes cannot manage on-premises distribution lists, even if they are a registered owner. The same is true for those who have on-premises mailboxes as they cannot edit distribution lists created in the cloud. Permissions and the ability of Outlook to support cross-platform editing of the membership of distribution lists are the main reasons. In hybrid environments, you are better off using OWA to edit distribution list membership unless you are positively sure that the mailbox and group are on the same platform.

## Including a Microsoft 365 Group in a Distribution List

You can include a Microsoft 365 group in the membership of a distribution list, but only through PowerShell. The Microsoft 365 group is treated like a nested distribution list composed of mailboxes and guest accounts (mail contacts). To add a group, run the *Add-DistributionGroupMember* cmdlet. For example:

```
[PS] C:\> Add-DistributionGroupMember -Identity P365.Authors -Member ExchangeMVPs
```

## Removing a Distribution List

The *Remove-DistributionGroup* cmdlet removes a distribution list. Once you remove a list, it's irrecoverable and will need to be recreated if removed in error. This command is an example of removing a distribution list without using a confirmation prompt:

```
[PS] C:\> Remove-DistributionGroup -Identity P365.Authors -Confirm:$False
```

## Including a Teams Channel in a Distribution List

As discussed in Chapter 12, the email integration for Teams creates email addresses used to send emails to a channel in a team. Messages sent to a channel appear as new conversations in the channel. Teams also captures a copy of messages delivered to the channel in the SharePoint site belonging to the team. To add a team channel email address to a distribution list, create a new mail contact with the address and add the mail contact to the distribution list.

## Including a Guest Account in a Distribution List

Groups, Teams, Planner, and SharePoint are among the apps which use Azure B2B Collaboration to share resources with external people using guest accounts. Exchange Online represents guest accounts as mail users, but the mail user objects for guest accounts don't show up in address lists, so you can't add them to a distribution list through EAC or a client UI. Instead, you can add them to a distribution list via PowerShell. This code adds the mail user object for the guest account JohnSmith@outlook.com to a distribution list:

```
[PS] C:\> Add-DistributionGroupMember -Identity MyDL -Member JohnSmith_outlook.com#EXT#
```

## Discovering Distribution Lists and Groups Someone Belongs to

You can discover the distribution lists to which a user belongs by examining their mailbox properties through EAC or the Microsoft 365 admin center. Here's how to do it with PowerShell:

```
[PS] C:\> $Dn = (Get-ExoMailbox -Identity Kim.Akers).DistinguishedName
Get-DistributionGroup -Filter "Members -eq '$Dn'"
```

You can replace *Get-DistributionGroup* with the *Get-ExoRecipient* cmdlet to report the different types of groups to which someone belongs. We sort the output to report membership of the different types of groups based on the *RecipientTypeDetails* (group type) information.

```
[PS] C:\> Get-ExoRecipient -Filter "Members -eq '$Dn'" -Properties RecipientTypeDetails | Sort
RecipientTypeDetails | Format-Table DisplayName, RecipientTypeDetails
```

You can use a variation of this command to focus on membership of a specific group type. For instance, this command creates a list of Microsoft 365 Groups to which the user belongs because we specify that we only want to see data with *RecipientTypeDetails* set to *GroupMailbox*.

```
[PS] C:\> Get-Recipient -Filter "Members -eq '$Dn'" -RecipientTypeDetails GroupMailbox | Format-
Table DisplayName
```

## Dynamic Distribution Lists

Traditional email distribution lists work very well. However, some well-known deficiencies exist. Maintenance of group membership to ensure accurate delivery of messages remains an issue, especially for groups with frequent membership changes. Dynamic distribution lists do not have fixed membership. Instead, the Exchange transport service calculates the membership by resolving a query (recipient filter) against the EXODS directory to return the current set of members. The recipients can be any type of mail-enabled object supported by Exchange, including user and shared mailboxes, public folders, mail contacts, and other distribution lists.

In the past, the transport service performed resolved the directory query on an on-demand basis as messages arrived addressed to the list. Because the query fetches the latest recipient information from the directory, this approach guarantees that the recipient set used to address messages is always up to date. However, resolving on-demand queries can create a heavy demand on system resources, especially when the recipient filter is complex.

Taken together with the fact that most list memberships do not change very often, in late 2021, Microsoft moved Exchange Online to a timed calculation model like that used for Azure AD dynamic groups. This means that the transport service calculates the memberships of dynamic distribution lists:

- At least once daily.
- After the recipient filter changes.
- Following the creation of a new dynamic distribution group.

In the case of the last two scenarios, it can take up to two hours before group membership is available for use. You can see when the transport service last updated the membership for dynamic distribution lists by running the *Get-DynamicDistributionGroup* cmdlet and examining the *CalculatedMembershipUpdateTime* property. For example:

```
[PS] C:\> Get-DynamicDistributionGroup | Format-Table DisplayName, CalculatedMembershipUpdateTime
```

DisplayName	CalculatedMembershipUpdateTime
Company-DDG	12/12/2021 20:00:36
Dublin users	12/12/2021 20:00:36
Office 365 Gurus	12/12/2021 20:00:36
AuthorizedGroupCreators	12/12/2021 20:00:36

Dynamic Distribution Lists are simple, robust, and work well when you need to communicate with a changeable set of mail-enabled recipients. The key to a successful deployment is to make sure that EXODS (and by extension, Azure AD) is populated with accurate data to allow recipient filters find the correct recipients. With that thought in mind, let's discuss how recipient filters work.

## Recipient Filters

Recipient filters are OPATH queries executed against the directory to return a set of objects. A query can vary from very simple lookups such as "all the users in the London office" or "everyone who works for the Accounting department" or "all full-time employees" to more complex queries involving lookups against multiple recipient attributes. The weak point for a recipient filter is the accuracy of the data in the queried repository. If the properties of EXODS objects are not populated with accurate data, then incomplete or erroneous sets of recipients will be found when the transport service uses the queries to find the set of recipients to route messages. On the other hand, if EXODS is well maintained (or Azure AD, as many attributes synchronize from Azure AD to EXODS) with the information required to support effective queries, dynamic distribution lists work very well.

Once you know the query to find the desired recipient set, creating a new dynamic distribution list is not difficult. The EAC interface helps administrators to compose the query used to determine distribution list membership. Each condition that you add to the query is a rule that checks an attribute of the objects that fall under the scope of the query (all recipient types, user mailboxes, and so on) against one or more text values. In Figure 11-18, we check the Company attribute for user mailboxes to look for a specific value. If multiple acceptable values exist for an attribute, input each value separated by a comma.

Groups > Add a group

**Assign users**

Group owners have unique permissions. They can add or remove members, delete conversations from the shared inbox, and change group settings. Group owners can also rename the group, update the description, and more.

Owner

Members  
 Specify the type of recipients that will be members of this group.

All recipient types

Only the following recipient types

- Users with Exchange mailboxes
- Mail users with external email addresses
- Resource mailboxes
- Mail contacts with external email addresses
- Mail-enabled groups

Membership in this group will be determined by the rules you set below

Add another rule

Back Next Cancel

Figure 11-18: Defining settings for a dynamic distribution list

**Dynamic Distribution Lists and Azure AD:** Dynamic Distribution Lists are an Exchange object and are only registered in EXODS, the Exchange Online directory. They do not appear in Azure AD because they



are not Azure AD objects like accounts, distribution lists, security groups, Microsoft 365 groups, and dynamic Microsoft 365 groups.

## How Recipient Filters Work

Dynamic distribution lists use one of two types of recipient filters:

- **Precanned filters** are used when a DDL is created through the EAC. Precanned filters are restricted to queries against a small number of object properties such as the department, city, and 15 customizable attributes.
- **Custom filters** are created when you use PowerShell to define a recipient filter. Custom filters are more powerful and flexible than precanned filters because a wider set of properties can be included in a query. Once you apply a custom filter to a DDL, you won't be able to edit the filter through the EAC.

The EAC UI restricts queries to a limited set of mailbox properties selected because they are the most popular:

- StateOrProvince.
- Company.
- Department.
- Any of the 15 single-value custom attributes (*CustomAttribute1* through *CustomAttribute15*).

When you save a dynamic distribution list with either EAC or PowerShell, Exchange writes its recipient filter into the list properties after applying some extra processing to the filter to ensure that it finds the right objects. This means that a relatively simple recipient filter as viewed in the EAC GUI looks very different when viewed through PowerShell using the *Get-DynamicDistributionGroup* cmdlet. For example:

```
[PS] C:\> Get-DynamicDistributionGroup -Identity "Microsoft 365 Gurus" | Select RecipientFilter,
RecipientFilterType

RecipientFilter      : (((((RecipientType -eq 'UserMailbox') -and (CustomAttribute1 -like 'Exchange'))
--and (Alias -ne $null))) -and (-not(Name -like 'SystemMailbox{*}')))) -and (-not(Name -like
'CAS_{*}')))) -and (-not(RecipientTypeDetailsValue -eq 'MailboxPlan')))) -and (-
not(RecipientTypeDetailsValue -eq 'DiscoveryMailbox')))) -and (-not(RecipientTypeDetailsValue -eq
'PublicFolderMailbox')))) -and (-not(RecipientTypeDetailsValue -eq 'ArbitrationMailbox')))) -and (-
not(RecipientTypeDetailsValue -eq 'AuditLogMailbox')))) -and (-not(RecipientTypeDetailsValue -eq
'AuxAuditLogMailbox')))) -and (-not(RecipientTypeDetailsValue -eq
'SupervisoryReviewPolicyMailbox')))) -and (-not(RecipientTypeDetailsValue -eq 'GuestMailUser')))) -
and (-not(Name -like 'SystemMailbox{*}') -and (-not(Name -like 'CAS_{*}')) -and (-
not(RecipientTypeDetailsValue -eq 'MailboxPlan')) -and (-not(RecipientTypeDetailsValue -eq
'DiscoveryMailbox')) -and (-not(RecipientTypeDetailsValue -eq 'PublicFolderMailbox')) -and (-
not(RecipientTypeDetailsValue -eq 'ArbitrationMailbox')) -and (-not(RecipientTypeDetailsValue -eq
'AuditLogMailbox')) -and (-not(RecipientTypeDetailsValue -eq 'AuxAuditLogMailbox')) -and (-
not(RecipientTypeDetailsValue -eq 'SupervisoryReviewPolicyMailbox'))))
RecipientFilterType: Custom
```

The recipient filter shown above looks for user mailboxes that have “Exchange” in *CustomAttribute1* and excludes all the various types of system mailboxes. The filter seems very complex and difficult to construct, but you don't have to insert these exclusions as Exchange creates them automatically.

When you see *Custom* as the *RecipientFilterType*, you know that PowerShell was used to create a custom recipient filter (see example below). If you see *Precanned*, you know that the recipient filter is based on the pre-prepared queries (like “all mail-enabled recipients”) available through EAC. Only dynamic distribution lists with precanned filters can have their recipient filters updated through the EAC. Although the other properties of dynamic distribution lists which have custom filters can be maintained through EAC, PowerShell must be used to edit their custom recipient filters.

## Checking the Effectiveness of a Recipient Filter

The EAC doesn't include a way to check the effectiveness of a recipient filter. In other words, how to be sure that the filter will find the right set of recipients when the Exchange transport service resolves it against the directory to compute the list membership. However, you can check with PowerShell by inputting the filter to

the *Get-Recipient* cmdlet. For example, here's how to check the recipient filter for the Microsoft 365 Gurus dynamic distribution list:

```
[PS] C:\> Get-Recipient -RecipientPreviewFilter (Get-DynamicDistributionGroup -Identity "Microsoft 365 Gurus").RecipientFilter
```

Name	RecipientType
Jeff.Guillet	UserMailbox
Tony.Redmond	UserMailbox
Paul.Robichaux	UserMailbox

In this instance, if you expect the query should return three mailboxes, it works. If not, the fault is either in the query (it searches for the wrong data) or the underlying data (does not contain the right values). Testing the results of a query is important not only because the results will reveal any flaw in your logic and identify whether the filter returns the correct recipient set when the query executes against the directory. It will also tell you how many recipients are in the selected set. Equipped with that knowledge, we can construct some guidelines to help administrators maintain dynamic distribution lists. For example, your organization might use the following guidelines for both standard and dynamic distribution lists:

- Any distribution list with more than 200 recipients should have a MailTip to highlight the fact to users. Exchange Online automatically displays a MailTip to tell the sender how many people will receive their email, but you can add another MailTip to urge additional caution when sending to very large distribution lists.
- Any distribution list with more than 500 recipients should be moderated to ensure that people do not inadvertently broadcast to distribution lists such as "The Entire Company". Moderated distribution lists are also a good way to stop email storms that result when users reply-all to very large distributions.

After you create a new dynamic distribution list or update the recipient filter for an existing list, the Exchange transport service resolves the recipient filter against the directory to compute the set of members. This process can take up to two hours. When it is complete, you can run the *Get-DynamicDistributionGroupMember* cmdlet to check what recipients Exchange will use when someone addresses a message to the list.

```
[PS] C:\> Get-DynamicDistributionGroupMember -Identity "Microsoft 365 Gurus"
```

Name	RecipientType
Tony.Redmond	UserMailbox
Jeff.Guillet	UserMailbox
Paul.Robichaux	UserMailbox

Even if the members do not appear in the same order, the memberships reported by *Get-Recipient* and *Get-DynamicDistributionGroupMember* for the same list should be the same.

Finding the set of dynamic distribution lists a user belongs to is a two-step process. After fetching the distinguished name of their mailbox, we use the distinguished name to check the membership of each dynamic distribution list as it passes through the pipeline:

```
[PS] C:\> $Dn = (Get-ExoMailbox -Identity Chris.Bishop).DistinguishedName
Get-DynamicDistributionGroup | ? {(Get-DynamicDistributionGroupMember -Identity
$_PrimarySMTPAddress | ? {$_DistinguishedName -eq $Dn})}
```

## Creating and Managing Dynamic Distribution Groups with PowerShell

PowerShell is a good way to create lists with both simple recipient and complex custom recipient filters. Our example requires a relatively simple recipient filter, as shown below where two precanned attributes are specified in parameters passed to the *New-DynamicDistributionGroup* cmdlet. The first checks for a value in

*CustomAttribute9* using the *ConditionalCustomAttribute9* parameter; the second limits the set of recipients returned to mailbox users by specifying that value in the *IncludedRecipients* parameter. After the new list is created, the *Set-DynamicDistributionGroup* cmdlet updates the list to add a manager (owner) and a MailTip. The person specified as the list manager must exist in the GAL.

```
[PS] C:\> New-DynamicDistributionGroup -Name DynamicPowerApps -DisplayName "Power Platform Experts DL" -ConditionalCustomAttribute9 PowerPlatform -IncludedRecipients MailboxUsers -PrimarySmtpAddress PowerPlatform@office365itpros.com -Alias Power.Platform
Set-DynamicDistributionGroup -Identity Power.Platform -ManagedBy James.Joyce@office365itpros.com -MailTip "People with real expertise in Power Platform"
```

## Complex Recipient Filters

Recipient filters support [a much wider set of filterable properties](#) than the set exposed by the EAC when composing query rules, but you can only use these properties by creating a custom recipient filter and writing the filter into the dynamic distribution list using PowerShell. Take the example where we want to create a dynamic distribution list to communicate with users whose mailboxes have been placed on litigation hold. To accomplish the goal, we follow these steps:

- Build the filter and test its effectiveness by using it with the *Get-Recipient* cmdlet.
- Update an existing dynamic distribution list or create a new dynamic distribution list with the tested filter.

For example, here's how to define and test a filter to find mailboxes on litigation hold based in Dublin.

```
[PS] C:\> $Filter = "((LitigationHoldEnabled -eq '$True') -and (StateOrProvince -eq 'Dublin') -and (RecipientType -eq 'UserMailbox'))"
```

```
Get-Recipient -RecipientPreviewFilter $Filter
```

Name	RecipientType
TRedmond	UserMailbox
James.Ryan	UserMailbox

Another example is a filter that looks for accounts with a certain job title that the organization has not disabled (the *ExchangeUserAccountControl* setting for a mailbox is set to *AccountDisabled* when the owning Azure AD account is blocked). This requires a filter that's slightly more complex to cover the range of job titles that might be in use. The need to specify multiple variants of the property is explained by the lack of wildcard support preceding a value in recipient filters (like *"\*architect"*). You can use wildcards after the value (like *"architect\*"*). The net result is that a filter for multiple variants of a value can end up with many *-or* clauses:

```
[PS] C:\> $Filter = "((Title -eq 'Architect') -or (Title -eq 'Senior Architect') -or (Title -eq 'Principal Architect') -and (ExchangeUserAccountControl -ne 'AccountDisabled'))"
```

Note that system values like *\$True* and string literals are enclosed in single quotes. If you don't do this, OPATH will signal an error. After testing the filter and verifying the accuracy of the chosen filter, we can use it when creating a dynamic distribution list. In this case, we use the filter to find accounts on litigation hold in Dublin:

```
[PS] C:\> New-DynamicDistributionGroup -Name "Dublin Mailboxes on Litigation Hold" -DisplayName "Dublin User Mailboxes on Litigation Hold" -Alias DublinLitigation -PrimarySmtpAddress Dublin.Litigation@Office365itpros.com -RecipientFilter $Filter
```

```
Set-DynamicDistributionGroup -Identity DublinLitigation -ManagedBy Legal.Assistant@Office365itpros.com -MailTip "People under litigation hold in Dublin"
```

Dynamic distribution lists which address sets of people like those under litigation hold might be deemed confidential. In this situation, you can hide them from address lists by setting their *HiddenFromAddressLists* property to *\$True*:

```
[PS] C:\> Set-DynamicDistributionGroup DublinLitigation -HiddenFromAddressListsEnabled $True
```

Anyone who needs to send messages to a hidden dynamic distribution list can do so by using its SMTP address. Messages sent to dynamic distribution lists with queries that don't find any recipients go into a void. Exchange doesn't generate a non-delivery notification for the sender because the address for the dynamic distribution list is valid: the problem is that the list query doesn't return any recipients, which is a valid condition. This is a good reason for testing the recipient filter for a dynamic distribution list to make sure that it finds the correct set of recipients.

## Finding Inactive Distribution Lists

Another common request is to know if any distribution lists are inactive and are therefore candidates for removal. Exchange Online does not include a way to find and report inactive distribution lists, so we must create one with PowerShell. The key points to remember are:

- A distribution list is active when people use it to address messages.
- Evidence of distribution list activity can be found in the message tracking logs by running a message trace to find events noting the expansion of distribution list memberships.
- Exchange Online keeps message tracking logs online for up to 10 days, after which the information is moved into data repositories and kept there for an extra 80 days. Thus, online searches can only look back 10 days to find expansion events. See Chapter 7 for more information about running message traces from the Exchange admin center.

With these points in mind, we can write a script to collect expansion events from the message tracking logs for the last 10 days and store the results in a table. We can then check the distribution lists in the tenant against the table to discover if we find a match. If we do, we know that the distribution list was used in the last ten days. If not, it was inactive at that time. Apart from reporting each list as it is checked, the script also outputs the results to a CSV file.

```
[PS] C:\> $EndDate = Get-Date; $StartDate = $EndDate.AddDays(-10); $Messages = $Null; $Page = 1
Write-Host "Collecting message trace data for the last 10 days"
Do
{
    $PageOfMessages = (Get-MessageTrace -Status Expanded -PageSize 5000 -Page $Page -StartDate
$StartDate -EndDate $EndDate | Select Received, RecipientAddress)
    $Page++
    $Messages += $PageOfMessages
}
Until ($PageOfMessages -eq $Null)

$MessageTable = @{}
$MessageTable = ($Messages | Sort RecipientAddress -Unique | Select RecipientAddress, Received)
$DLs = Get-DistributionGroup -ResultSize Unlimited
Write-Host "Processing" $DLs.Count "distribution lists..."
$Results = ForEach ($DL in $DLs) {
    If ($MessageTable -Match $DL.PrimarySMTPAddress) {
        [pscustomobject]@{Name = $DL.DisplayName ; Active = "Yes"}
        Write-Host $DL.DisplayName "is active" -ForegroundColor Yellow }
    Else {
        [pscustomobject]@{Name = $DL.DisplayName ; Active = "No"}
        Write-Host $DL.DisplayName "inactive" -ForegroundColor Red }
}
$Results | Export-CSV c:\Temp>ListofDLs.csv -NoTypeInformation
```

Given that message traces give us a limited ten-day window to detect inactive distribution lists, this is not a practical technique for a production-quality solution. Nevertheless, the method gives us the basis to develop the technique further into something that might work. For instance, you could run a script every ten days and merge the results over a few months to give a more precise view of inactive and active lists.

## Controlling User-Created Distribution Lists

The set of OWA options controllable on a per-mailbox basis includes a “distribution groups” link to allow users to manage distribution lists they belong to and those that they own. To see the set of distribution lists owned by a user, select **OWA Options**, then **General**, and then **Distribution groups**. OWA decides what options are available for a user to work with distribution lists by evaluating if the “*MyDistributionGroups*” and “*MyDistributionGroupMembership*” roles are in the user role assignment policy assigned to their mailbox.

In Figure 11-19 we can see that both roles are present in the default role assignment policy used by Exchange Online. Note that users can only create standard distribution lists – they cannot create security groups or dynamic distribution lists. As explained earlier, the ability for a user to create a new Microsoft 365 Group is controlled by a setting in the Azure AD policy for Groups.

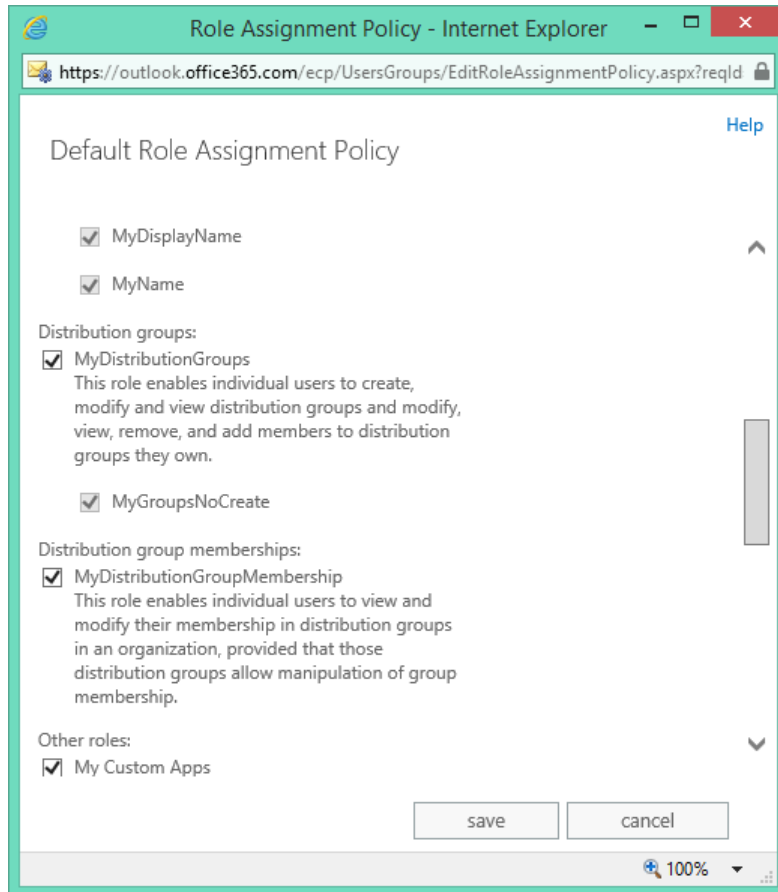


Figure 11-19: Distribution list entries in the default user role assignment policy

Even if your tenant deploys a distribution list naming policy, to prevent “directory anarchy”, many organizations also like to control who can create new distribution lists, and do not want users to create new distribution lists without oversight. You can remove the ability of users to create new distribution lists by unchecking the *MyDistributionGroups* option in the default role assignment policy as OWA will not then display the UI. However, a subtler approach is also possible that allows users to see the distribution lists to which they belong and to edit the membership of distribution lists that they own. Given that some organizations are migrating distribution lists to Microsoft 365 Groups, it makes sense to leave users with the ability to work with distribution lists that they already have while blocking their ability to create any more.

To achieve the goal, you must create a new user role assignment policy and amend it to remove the setting which allows users to create new distribution lists. Here are the steps:

- Go to the Permissions section of classic EAC, click **user roles**, and then **[+]** (the plus sign) to create a new user role assignment policy. Give the policy a suitable name and description and check all the

boxes to allow users access to the various items of functionality. In effect, you start by cloning the default user role assignment policy. In this example, the new policy is "Restricted Group Management".

- Using PowerShell, define a new management role that is based on the existing *MyDistributionGroups* role. We will call the new management role *MyGroupsNoCreate*.

```
[PS] C:\> New-ManagementRole -Name "MyGroupsNoCreate" -Parent MyDistributionGroups
```

- The new management role still allows users to create new distribution lists. To remove that ability, we have to remove the reference (or role entry) to the *New-DistributionGroup* cmdlet from the *MyGroupsNoCreate* management role. This command breaks the link.

```
[PS] C:\> Remove-ManagementRoleEntry MyGroupsNoCreate\New-DistributionGroup -Confirm:$False
```

- A user role assignment policy consists of a set of connections between management roles and the policy. The new user role assignment policy that we created in the first step still contains a reference to the standard *MyDistributionGroups* role. We need to remove it before we can add the altered management role that we have just created.

```
[PS] C:\> Remove-ManagementRoleAssignment "MyDistributionGroups-Restricted Group Management" -Confirm:$False
```

- Now we add a management role assignment to link the *MyGroupsNoCreate* role to the "Restricted Group Management" policy.

```
[PS] C:\> New-ManagementRoleAssignment -Name "MyGroupsNoCreate-Restricted Group Management" -Role MyGroupsNoCreate -Policy "Restricted Group Management"
```

- We can then check that the correct roles exist in the "Restricted Group Management" policy. You should find a link to the *MyGroupsNoCreate* in the list and no reference to *MyDistributionGroups*. Remember that *MyGroupsNoCreate* has all the functionality of *MyDistributionGroups* except the link to the *New-DistributionGroup* cmdlet. Because they cannot run the cmdlet, a user to whom we assign this policy can do everything except create a new distribution list.

```
[PS] C:\> Get-ManagementRoleAssignment -RoleAssignee "Restricted Group Management" | Format-Table Name, Role -AutoSize
```

To prove that everything works, we apply the Restricted Group Management policy to a mailbox:

```
[PS] C:\> Set-Mailbox -Identity 'Nancy Anderson' -RoleAssignmentPolicy "Restricted Group Management"
```

Then, after waiting ten minutes or so to ensure that any cached permissions are clear, we log in to validate that when the user accesses the **Distribution groups** option in the **General** section of OWA options, they still see the set of distribution lists that they belong to in addition to the distribution lists that they own. What we want to see is that the plus sign (+) icon that allows them to create a new distribution list is missing, as it is in Figure 11-20.

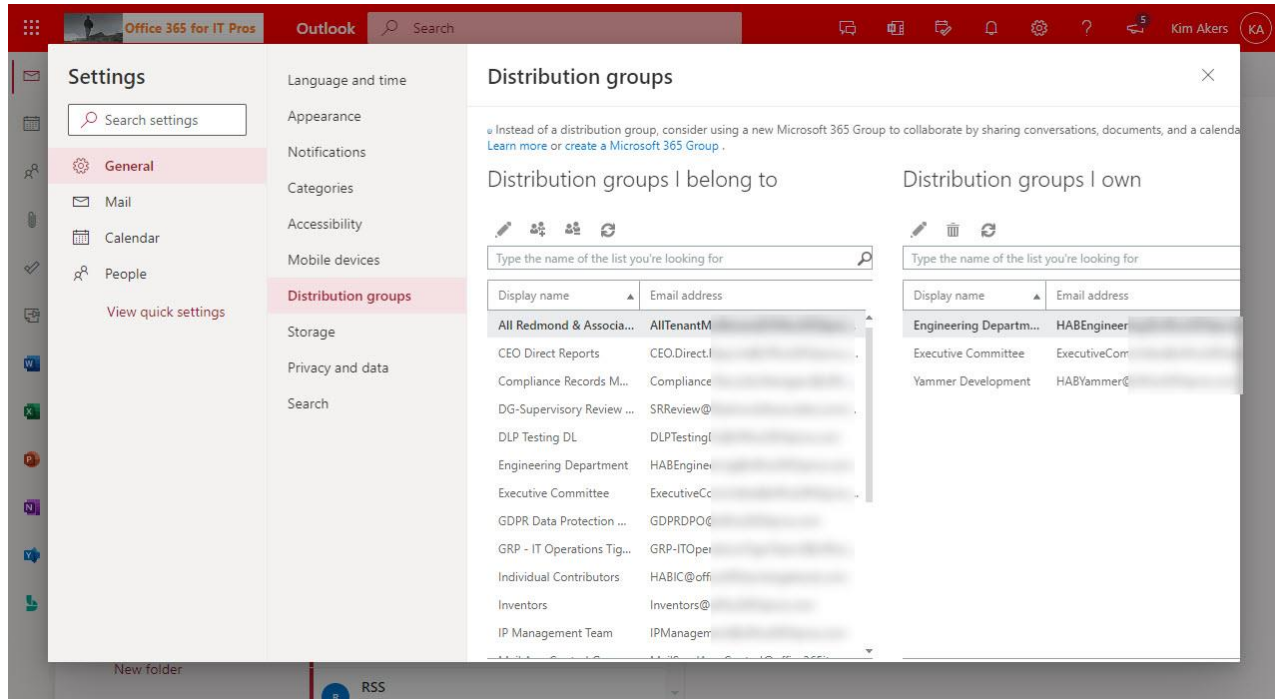


Figure 11-20: The effect of the restricted user role assignment policy

## Distribution List Naming Policy

If you allow users to create new distribution lists, you can expect that some will not be as disciplined as you might like when naming their distribution lists. No one likes to see entries such as “Billy-Bob’s Fishing Aficionados” show up in the GAL. The name does not convey a sense of the true business purpose behind its existence. To help guide users, organizations can create a distribution list naming policy. The policy can include:

- A prefix to apply to the user-supplied name for a new distribution list. The prefix can be multi-part and consist of text strings and the values of attributes for the user account stored in Azure AD. The policy forces the insertion of a blank character if the user account does not store a value for the chosen attribute.
- A suffix to add to the user-supplied name for a new distribution list. Again, you can use both text and attribute values.
- Barred or forbidden words that users cannot include in the name for a new distribution list. For instance, you can prohibit potentially offensive terms by including them in the policy.

To configure the distribution list naming policy, go to the **Groups** section of the EAC, select the distribution list tab and then **Add naming policy**. Figure 11-21 shows the screen to define the components of the naming policy. The parameters selected indicate that all names should be prefixed with “DL-” (the text string DL plus a hyphen) followed by the value of the Department attribute for the user’s account and then another string value containing a space. The policy should generate names for distribution lists like *DL-Sales Members*, *DL-Accounts Team 1*, *DL-CEO Office Executives*, *DL-Finance Volleyball team*, and so on.

One reason why organizations use a naming policy is to ensure that distribution lists appear in a single section within the GAL. Adding the department name ensures that the distribution lists belonging to a department are found together. Behind the scenes, Exchange Online stores details of the distribution list naming policy as settings in the tenant organization configuration. You can view the settings with the *Get-OrganizationConfig* cmdlet:

```
[PS] C:\> Get-OrganizationConfig | Format-Table Distr* -AutoSize
```

```
DistributionGroupDefaultOU DistributionGroupNameBlockedWordsList DistributionGroupNamingPolicy
```

{XXX, Cheese}
DL-<Department> <GroupName>

×

## Edit group naming policy

Policy
Blocked words

Add prefixes and suffixes to new group names.

**Current policy**  
DL-<StateOrProvince> "<GroupName>

**Create a policy**  
Choose a prefix to add to the beginning of the group names

Text	DL-		×
Attribute	Department		×
Text			×

[Add prefix](#)

**AND**  
Select a suffix to add to the end of group names

[Add suffix](#)

Save

Figure 11-21: Creating a distribution list naming policy

These properties have the following meanings:

- The *DistributionGroupBlockedWordsList* property holds words that cannot be used in a group name.
- The *DistributionGroupNamingPolicy* property defines the pattern for names. In the example shown below, the policy consists of four elements:
  - A text string "DL-."
  - <Department> means that the Department value stored in the user's Azure AD account is inserted.
  - Another space is then inserted.
  - <GroupName> means the name of the distribution list as supplied by the user.
- The *DistributionGroupDefaultOU* property is only used in on-premises deployments. However, it still appears in the Exchange Online organization configuration.

It is possible to create or amend the naming policy through PowerShell. For instance:

```
[PS] C:\> Set-OrganizationConfig -DistributionGroupNamingPolicy "DL-<Department> <GroupName>"
```

To remove the distribution list naming policy, set the value of *DistributionGroupNamingPolicy* to *\$Null*:

```
[PS] C:\> Set-OrganizationConfig -DistributionGroupNamingPolicy $Null
```

The distribution list naming policy applies to user-created distribution lists and has no retrospective effect. You must update the names of existing distribution lists if you want them to follow a new policy. To do this, use PowerShell to find distribution lists with non-compliant names and then update their names according to



the naming policy. The old EAC allows administrators to create distribution lists without applying the policy: the new EAC does not.

if an administrator creates a distribution list with PowerShell, they can pass the *IgnoreNamingPolicy* parameter to instruct Exchange Online to not apply the policy. For example:

```
[PS] C:\> New-DistributionGroup -Name "Tiger Team" -Members Jill.Smith@Office365itpros.com, Ben.Owens@office365itpros.com, Kim.Akers@office365itpros.com -IgnoreNamingPolicy
```

The distribution list naming policy does not apply to dynamic distribution lists or mail-enabled security groups because users cannot create these types of groups. The policy does not apply to Microsoft 365 Groups, which use a separate naming policy.

Defining a distribution list naming policy for Exchange Online can also affect the way that on-premises distribution lists synchronize in hybrid deployments. To make everything consistent, Exchange applies the naming policy to distribution lists belonging to the on-premises organization when they synchronize with Azure AD objects with AADConnect. Thus, an on-premises distribution list might have a different name in the on-premises directory than that shown in Azure AD unless you make sure to apply the same naming policy in both environments.

## Migrating On-Premises Distribution Lists to Exchange Online

Exchange Online doesn't include any out-of-the-box method to migrate a distribution list from an on-premises Exchange organization to Exchange Online. When an on-premises organization is synchronized with Exchange Online in a hybrid deployment, the suggested method is to remove the distribution list from on-premises and recreate it as a brand-new object in the cloud. For anything but simple lists with just a few members, this is a tiresome process, but it reflects the fact that transferring a distribution list to the cloud can be quite complex because:

- The owner(s) of the distribution list might not have their account(s) in the cloud. An on-premises user cannot manage a cloud-based distribution list.
- Objects for mail-enabled members of the distribution list might not exist in the cloud. For example, a mail contact in the on-premises environment might not be synchronized to the cloud.
- The distribution list might hold other distribution lists.
- The proxy addresses for the on-premises distribution list must be transferred to the new list.
- Some properties of distribution lists refer to other directory objects that must exist in the target environment before they can be used. For example, the property that controls the ability of users to send emails on behalf of the list.

It is possible to write a PowerShell script to concurrently connect to Exchange on-premises and Exchange Online and perform the processing to transfer a distribution list. The script must:

- Check that all the prerequisites are satisfied for the transfer to go ahead. For example, are all the members of the list known in the cloud.
- Create the target distribution list in Exchange Online.
- Read information about the source distribution list from Exchange on-premises and update the properties of the target distribution list.
- Assign a new proxy address to the source distribution list and transfer it to the target distribution list.
- Update the membership of the target distribution list with the membership of the source list.
- Hide the source distribution list from address lists so that the only list that is visible to users is the target. Eventually, if the transfer worked and no problems are found, the old list is removed.

An example of a migration script for distribution lists is [available on GitHub](#) (the documentation [is here](#)). As with any script, you should test it carefully and adapt the code where necessary to meet the needs of your deployment.

# Recipient Moderation

Moderation (or “message approval”) means that a message must first be reviewed and approved by a nominated moderator before Exchange Online can deliver it to the addressee. Up to ten moderators can be assigned responsibility to control messages sent to the following recipient types:

- User Mailboxes.
- Shared Mailboxes.
- Distribution Lists.
- Dynamic Distribution Lists.
- Mail Contacts.
- Mail Users.
- Mail-enabled Public Folders.
- Group mailboxes.

Teams supports moderation for messages posted to channels by restricting the ability to post to certain members.

Moderation for individual recipients (mailboxes, mail contacts, or mail users) is intended to “protect” the recipients against inappropriate email, often because the recipients are sensitive in some way. For example, it’s reasonably common to impose moderation on the mailboxes of senior executives so that an executive assistant or another moderator can review inbound messages before passing approved items to the executive. Moderation for distribution lists is often imposed to ensure that posts to very large distributions only appear when it is appropriate to share the content with so many people. In other words, you might not want to have everyone in the company be allowed to send messages to the 50,000-recipient “All Company” distribution list. In all cases, you can arrange for select users to bypass the requirement for moderation so that their messages are delivered to the recipient without going through the approval process.

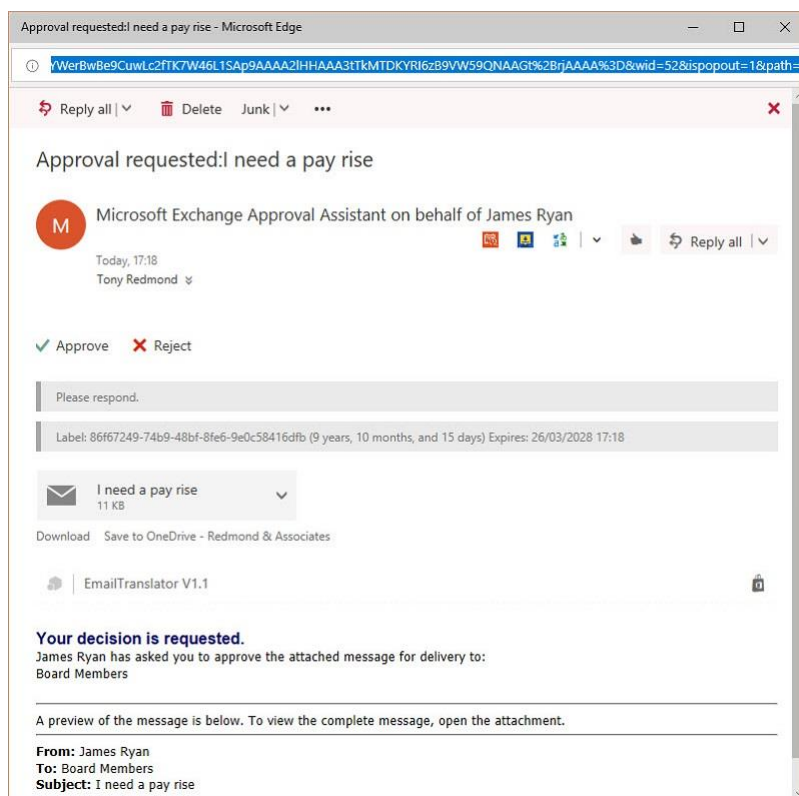


Figure 11-22: A message arrives for moderation

A moderated recipient can have multiple moderators, in which case all the nominated moderators receive copies of messages for approval (Figure 11-22). However, no quorum must be available before a message can be delivered as a simple approval from a single moderator suffices. If a message is declined by a moderator, Exchange Online returns it to the original sender. Logically, a moderator can send a message to the recipient that they are moderating without gaining further approval as can the owner of a distribution list.

A message awaiting moderation is held in an arbitration mailbox (you have no control over which arbitration mailbox is used) and should be processed within two days. However, the process that expires messages waiting in the arbitration mailbox runs on a weekly workcycle, which means that senders will receive notification that moderation has not occurred any time between two and nine days after the original message is sent – a tenant administrator can do nothing to accelerate the approval process. It is also interesting to note that Exchange Online throttles the number of expired moderation notifications to [300 per hour](#), which means that in periods when heavy usage is made of message moderation, some senders might not receive notifications that their messages have expired while awaiting moderation.

To set up moderation for a distribution list through EAC, edit the object's properties and select Message approval in the settings (Figure 11-23). Here you can set the flag to enable moderation and input the names of the users to act as moderators. Notice that you also can input the names of distribution lists or users who will bypass moderation and select who receives notification if their messages are not approved.

← ×

## Edit message approval

Specify if message sent to this group need to be approved, and choose moderators to approve or reject messages.

Require moderator approval for messages sent to this group

**Group moderators**

Search for users to add

TRedmond ×

**Add senders who don't require message approval:**

Search for users to add

B BoardMembers-convert ×

**Notify a sender if their message isn't approved:**

Only sender

Only senders in your organization

No notifications

Save changes

Figure 11-23: Editing the moderation settings for a distribution list

The EAC does not display the message approval UI for mailboxes, groups, contacts, mail users, or public folders, so you must set up moderation for these objects using PowerShell. For example, this command enables moderation for a mailbox and sets up three moderators. We also select a distribution list whose members bypass moderation.

```
[PS] C:\> Set-Mailbox -Identity "CEO Mailbox" -ModerationEnabled $True  
-ModeratedBy "Steve Smith", "Bill Jones", "CEO Assistant"  
-BypassModerationFromSendersOrMembers "Executive Committee"
```

Similar properties to enable moderation can be set through the *Set-DistributionGroup*, *Set-DynamicDistributionGroup*, *Set-MailContact*, and *Set-MailUser* cmdlets. When an object is enabled for moderation, Exchange Online creates an automatic MailTip that is displayed to users to inform them that a message will be moderated and might be delayed when they address a message to the object.

**Moderation for nested distribution lists:** Distribution lists (but not dynamic distribution lists) support the *BypassNestedModerationEnabled* parameter. If set to *\$True*, any nested distribution lists that also need moderation are governed by the decision of the moderator of the distribution list to which a message is addressed. If a moderator approves a message, it will be delivered to the members of all the nested distribution lists too. The default value of the flag is *\$False*, meaning that delivery to each moderated distribution list must be approved, but in most cases, it's reasonable to set the flag to *\$True* and allow the first moderator to control approval.

## Migration from Email Distribution Lists

Microsoft is keen that tenants should replace distribution lists with Groups. To make the move easy for tenants, five methods are available:

- Administrators can migrate individual distribution lists using the function available in the EAC.
- Group owners can migrate individual distribution lists using the function available in OWA.
- Administrators can run the PowerShell scripts created by Microsoft to migrate batches of distribution lists at one time.
- Administrators can run the *New-UnifiedGroup* cmdlet to convert an individual distribution list to a group.
- Tenants can create their versions of PowerShell scripts to perform custom migrations.

None of these methods work with distribution lists belonging to on-premises organizations. To convert these groups, you must first move them to Exchange Online, and then, when the groups are "cloud-owned", you can go ahead and convert them. Of course, it might be simpler to recreate the on-premises distribution lists as Groups and then remove them from the on-premises organization.

Email distribution lists have been around for a long time. Different forms of distribution lists exist, some of which are easy to convert while others are very difficult to move to Groups today. It is safe to say that it is relatively easy to code the steps to migrate a simple distribution list managed by Exchange Online that only has cloud mailboxes in its membership. Things become more challenging when the need arises to process distribution lists with the following characteristics:

- **Dynamic distribution lists:** Groups supports dynamic Groups, but no conversion facilities are available to migrate dynamic email distribution lists to become dynamic Groups. This could be due to some of the difficulties created by the fact that the two types of groups support different sets of member objects. For example, you can create a dynamic distribution list that includes several types of mail-enabled objects that are not all supported by Groups.
- **Nested distribution list:** It is common to meet distribution lists where the membership contains other (nested) distribution lists. For example, an "All Company" list might include several other groups, each of which might include people who work in a business unit or other division of the company. In turn, each divisional list might include other distribution lists for its operating units that hold the actual mail-enabled recipients. Nested distribution lists are not supported by Microsoft's migration toolset.

- **Mail-enabled universal security groups:** These groups are not migrated because Groups cannot act as a security principal.
- **Distribution lists that include more than mailboxes:** Groups only support user and guest accounts as members. They do not support some of the other kinds of mail-enabled recipients often found in email distribution lists, such as public folders.
- **Distribution lists with advanced settings:** Lists with send on behalf of settings or are hidden from address lists, or that have restrictions on who can join a group.

Microsoft summarizes the situation by saying that its toolset can deal with “*cloud-managed, simple, non-nested distribution lists.*” Email distribution lists that Microsoft’s migration toolset cannot process must stay in place until they are either not needed and you can remove them safely or Groups evolve to a point where they can serve as a replacement. In some cases, as in the example of nested or dynamic distribution lists used by organizations to send messages to a very large number of users, that point might never come. Remember too that some clients still in use do not support Groups, so do not plan to move away from distribution lists until all clients are prepared.

In this context, migrating an email distribution list to a group is only a question of moving members from one group to another. Any other assets used with an email distribution list, such as an associated public folder, are unaffected by the migration.

## Migration from the Exchange Administration Center

The classic version of the EAC supports the bulk migration of one or more distribution lists to Groups. The process is simple. If EAC detects that the tenant has some distribution lists that it considers eligible for migration, you can invoke a wizard to do the work. Eligibility means that the distribution lists are cloud-based rather than on-premises objects, do not include nested lists, the membership is composed of mailboxes only, and [meets other criteria](#) (see above).

### Bulk upgrade DLs to groups in Outlook

Upgrading DLs to groups in Outlook gives both existing and new members access to prior conversations and a simpler way to manage membership. [Learn more about bulk upgrades](#)

Showing Available for upgrade Refresh

<input type="checkbox"/>	DISPLAY NAME	MEMBE...	CREATED	LAST MODIFIED	EMAIL ADDRESS
<input type="checkbox"/>	Tiger Team	3	27/01/2022 15:30:02	27/01/2022 15:30:45	TigerTeam@F
<input type="checkbox"/>	Board Members	1	21/01/2015 12:46:22	21/04/2021 21:24:31	BoardMembe
<input type="checkbox"/>	Users Blocked for Weekend Access	2	09/04/2021 14:52:18	17/11/2021 18:55:06	Block.Weeker
<input type="checkbox"/>	Teams Policy Assignment - Ireland Workers	7	12/03/2021 13:12:21	17/11/2021 18:55:06	TeamsGroupl
<input type="checkbox"/>	Teams DLP Users	7	12/11/2020 15:16:21	17/11/2021 18:55:06	TeamsDLPUs
<input type="checkbox"/>	Teams Frontline Manager Policy Assignments	4	19/03/2021 13:13:07	17/11/2021 18:55:06	TeamsFrontLi
<input type="checkbox"/>	Trading Controlled Group	5	20/03/2020 15:17:04	17/11/2021 18:55:06	TradingContr
<input type="checkbox"/>	Users Who Don't Use MyAnalytics	3	23/03/2020 12:15:53	17/11/2021 18:55:06	Office365-Nc
<input type="checkbox"/>	Tenant Working Group	0	23/05/2021 21:29:25	17/11/2021 18:55:06	Tenant.Worki
<input type="checkbox"/>	Operations Alpha Team	3	05/07/2016 17:48:11	17/11/2021 18:55:00	OperationsAl
<input type="checkbox"/>	Planner Gurus	6	08/05/2017 14:06:09	17/11/2021 18:55:00	PlanGurus@c
<input type="checkbox"/>	Office 365 eBook Author Team	2	17/02/2016 19:02:37	17/11/2021 18:55:00	O365BookTe

0 DLs selected

Start Upgrade

Cancel

Figure 11-24: Bulk conversion of distribution lists in the classic EAC

The conversion process is painless. Select the groups that you want to migrate (Figure 11-24) and click **Start Upgrade**. The conversion wizard processes each group in turn in the background, running a command like that shown below to convert each distribution list in turn. To be sure that the wizard converts the right

distribution list, Groups references the group with its unique identifier instead of the alias, display name, or distinguished name.

```
[PS] C:\> New-UnifiedGroup -DlIdentity 'bf6aba7a-6dbc-4f2f-8e94-ae97289945d7'  
-ConvertClosedDlToPrivateGroup:$true -DeleteDlAfterMigration:$true
```

If successful, the wizard removes the old distribution list and a brand-new group exists, complete with the membership and other properties of the old group including its email addresses. The fact that the new group has the email addresses previously assigned to the old group means that users can reply to messages sent to the old group with a guarantee that Exchange can route those replies to the correct destination.

To check the progress of the conversion process, select “upgraded DLs” from the drop-down list and then **Refresh**. You might need to do this several times as EAC does not signal the successful conversion of a distribution list. The appearance of the new group on the list is the only way to know when the new group is ready to go.

When you convert a distribution list, you might consider changing the display name of the new group to reflect that it now offers added resources to its membership. This step is manual and is not performed by the conversion wizard.

The modern version of the EAC also supports upgrading distribution lists to Microsoft 365 groups, but only for individual distribution lists. The same process takes place to perform the upgrade.

**Upgrade Failure with Email Address Policies:** If you use email address policies (see the companion volume) to assign SMTP email addresses from a specific domain to newly created groups, you cannot use EAC or PowerShell to convert distribution lists to Groups. EAC will accept distribution lists for conversion and never do anything thereafter and any attempt through PowerShell throws the error that the distribution list cannot be migrated “*due to an email policy set by your IT administrator.*” The only solution is to convert the distribution list through a multi-step manual process as described later.

## User-Controlled Conversion

Many users manage distribution lists, which they might control on behalf of a business unit or just to help do their job better. If they are owners of groups, they can transfer the members of a distribution list to a new group by using OWA to edit the membership of the group and include the distribution list as a new member. When OWA saves the group, it expands the membership of the distribution list and adds individual members to the group. Although this is an effective way of moving members from email distribution lists to Groups, the obvious downside is that the old distribution list stays behind, leading to “GAL clutter”.

OWA allows owners of distribution lists to convert their groups. Behind the scenes, a background job periodically scans for distribution lists that are cloud-managed and only hold objects that can be members of Groups. The job uses other factors to find suitable distribution lists for conversion, such as those that have fewer than a hundred owners. When the processing completes, a list of convertible distribution lists owned (or managed by) the user is available for display by OWA. However, the choice to migrate distribution lists only appears when the user can create new groups as it obviously would not make sense to allow people who cannot create new groups to attempt a migration.

The set of eligible distribution lists appears at the top of the Groups section in OWA. The “DL” prefix appears before the name of the distribution lists to show that they are not groups. To migrate a distribution list, click on it. OWA then displays a screen with some details of the selected group and a big “Upgrade” button. Click to go ahead with the migration. A new group replaces the old distribution list and takes on all the properties of the source list.

## Microsoft's PowerShell Scripts

PowerShell is a terrific way to automate the process of converting distribution lists. Microsoft has [scripts available in the download center](#) to migrate all eligible distribution lists found in a tenant:

- **Get-DIEligibilityList.ps1**: Scans the tenant to find distribution lists to create a report detailing the distribution lists that can be migrated to Groups along with those that cannot (for one of the reasons given above). The output of the script is a text file (DIEligibility.txt) that you can review.
- **Convert-DistributionGroupToUnifiedGroup.ps1**: Converts all eligible distribution lists in a tenant to Groups.

The scripts are effective but can only deal with distribution lists that have the same characteristics discussed above.

## The DIY Approach to Converting Distribution Lists

Another example of how to approach the task of converting an email distribution list to a Microsoft 365 group is posted on [GitHub](#). Like all PowerShell scripts, the code is easily altered to meet the needs of a certain scenario. The basic set of tasks that a conversion script needs to perform include:

- Check whether a group with the same alias exists. The script could exit if this is the case or continue by migrating the distribution list to a new group with a different alias.
- Check whether the distribution list being migrated is a simple distribution list (type equals *MailUniversalDistributionGroup*). You could migrate the membership of a universal mail-enabled security group to a group but not the security principal. There is probably some good reason why the security principal exists, so it is probably best not to migrate these groups unless you are sure that the security element is resolved.
- Gather information about the owners (managers) and members of the input distribution list.
- Update the alias of the input distribution list with a new value to allow the reuse of the existing alias for the new group.
- Check whether the new group should be public or private. One way to do this is to look at the *MemberJoinRestriction* property of the input list. If this is "Closed" or "ApprovalRequired", then it is reasonable to assume that a private group should be the result.
- Create the new group by running the *New-UnifiedGroup* cmdlet.
- Migrate the properties of the input distribution list to the new group. A direct match of all properties does not exist, but a reasonable number of properties are common to both distribution lists and Microsoft 365 groups.
- Add the members of the new group by reading the membership information from the input distribution list and using the *Add-UnifiedGroupLinks* cmdlet to update the membership of the group. Some checking is necessary to ensure that only mailboxes are added to the group. Mail-enabled public folders and shared mailboxes are supported for distribution lists but not for Groups. In addition, you must expand the membership of nested distribution lists to figure out the set of valid members and then add them to the group.
- Because distribution lists act as a single address to distribute emails to members, it is reasonable to expect that the substitute group would behave comparably. Accordingly, you can also add the members of the group as subscribers so that they receive new copies of conversations via email.
- Add the owners of the new group by running the *Add-UnifiedGroupLinks* cmdlet. Users must be members of the group before they can be added as owners.
- Hide the old distribution list from the GAL so that users pick up and start using the replacement group.

The solution gathers the necessary PowerShell commands together into a script, like the version [available on GitHub](#). Since that script was written, Microsoft has updated the *New-UnifiedGroup* cmdlet so that it can convert a distribution list to a group if it is an “eligible” group.

A script built to migrate groups can be as complex as you think necessary to handle the kind of distribution lists found in your tenant. The caveat is that you should not migrate every distribution list without thinking. In some cases, a simple distribution list is the right tool for the job, and you do not need to create the added overhead and complexity that goes along with a group (document library, group mailbox, and so on).

## Comparing Groups, Distribution Lists, and Shared Mailboxes

Although Microsoft dedicates an impressive amount of development effort to support and enhance the capabilities of Groups, situations do exist when you should use a distribution list, a dynamic distribution list, or a shared mailbox rather than a group. You can use the following questions to help guide your choice:

- Is the intended use of the group to distribute email to a set of people? If so, then a standard distribution list is probably the best choice. Remember, if you need to, you can convert a standard distribution list to a group.
- Does the group span more than Exchange Online mailboxes? Groups support a limited set of recipient types (Exchange Online user mailboxes and guest users) while standard distribution lists support all mail-enabled recipient types – mailboxes, mail users, mail contacts, public folders, resource mailboxes, and other groups.
- Do you want to have access to more than the Inbox and Calendar folders and allow group members to move items between mailbox folders? If so, a shared mailbox is a better solution. The same is true if you want to use other mailbox features, such as assigning categories to messages.
- Do you want to preserve earlier contributions and conversations shared between members of the group? You can do this with the old-fashioned combination of a public folder (to hold the content) and distribution list or with a shared mailbox, but the added functionality available to a group means that it is a better long-term solution.
- Do you want dynamic membership of groups based on Azure AD attributes? You can have dynamic membership of Groups, but only with extra cost. In addition, if you want to have groups whose membership has more than Microsoft 365 user and guest accounts, then dynamic distribution lists are the right choice.
- Is the communication flow within the workgroup purely internal and the ability to send outbound emails is unnecessary? If so, Teams might be a better choice.

The arguments for Groups are often based on the use of a common identity to access resources drawn across from the service. They are much less email-centric than standard distribution lists are, so the question often comes down to the fundamental choice of whether you just need email or to exploit that common identity across Exchange Online, SharePoint Online, Teams, Planner, Power BI, and so on.



# Chapter 12: Teams Basics

**Tony Redmond**

## Workgroup Collaboration

Teams is Microsoft 365's workgroup application for consumer, enterprise, government, and education customers. The central concept behind Teams is that it delivers workspaces that bring people together to collaborate through chats, files, meetings, and associated applications. Chats, or conversations, can be private (often 1:1 but can involve more users) or public, in which case they occur in channels (ways to organize discussions) within a team. A chat is usually a more dynamic exchange of views than is the norm with email. Microsoft calls this "high-velocity" communication because many interactions might occur between the participants over a brief period. In addition, the chats tend to be shorter than email, with many contributions consisting of just a few words. Applications like Facebook, Twitter, and WhatsApp have popularized interactive chat in the consumer space, and companies like Bloomberg and Reuters have long offered real-time chat applications for the financial world, and unsurprisingly, this modality is popular in the world of personal communications.

Apart from conversations, Teams is a strong meeting and calling platform. People connect with audio or video streams to share information in a very natural fashion. The movement from in-person to virtual meetings was accelerated by the Covid-19 pandemic, generating billions of "collaboration minutes" monthly along with associated recordings, transcripts, captions, and files. Teams meetings bring together hundreds of participants in a highly interactive format, complete with screen and application sharing, chat, and recordings. Teams Live Events cater for much larger events such as company-wide "all hands" meetings where the number of presenters is typically limited. The calling side of Teams allows organizations to replace traditional phone systems. More on this topic in the chapter about Teams Devices and Calling.

Microsoft's [Teams Adoption Guide](#) sets out their vision and intent for Teams. We'll explore the concepts and ideas described in the guide throughout this chapter. We also describe the Teams architecture and how Teams integrates different Microsoft 365 components to deliver its functionality along with lots of practical information about how to use Teams. We develop the information presented here to explore how to manage Teams in a Microsoft 365 tenant in the next chapter.

## Teams User Base

At the Ignite 2018 conference, Microsoft announced that Teams was the fastest-growing business application in the company's history. Teams achieved that status just two years after its launch. Accenture is the largest Teams deployment and has a user community of more than 527,000. In December 2020, Accenture said that their use of Teams included [900 million minutes of audio conferencing and 90 million minutes of video conferencing](#) monthly. Microsoft put the number of audio minutes consumed by Accenture monthly at over one billion just a month later when reporting their [FY21 Q2 results](#). Microsoft said that the U.S. Department of Veterans Affairs also had more than 500,000 users.

In October 2021, Microsoft said that 138 organizations have more than 100,000 Teams users and more than 3,000 organizations have more than 10,000 Teams users (in January 2021, the numbers were 117 and 2,700 organizations respectively). 93 of the Fortune 100 use Teams. While these numbers are big and the number of large Teams deployments is growing fast, Teams still has plenty of room to grow into the Office 365 user base (Azure AD serves well over 200,000 organizations, most of which use Office 365).

For the overall user base, Microsoft changed their method of reporting [Teams usage in July 2021](#) from daily active users to monthly active users. They claimed that Teams has “nearly 250 million” monthly active users, but did not specify the breakdown across commercial users, education users, and personal users.

An active user is someone who does something with an application. For Teams, it means that they send a message, participate in a meeting, and so on. A big difference exists between active users and licensed users. Organizations often acquire licenses well before they deploy a product, and because Teams is bundled into many Microsoft 365 products, the number of people who can use Teams is potentially much higher.

**Teams and U.S. Government Cloud:** Due to the need to achieve compliance with several government security standards, the introduction of Teams features in the GCC, GCC High, and DoD clouds often lags several months after commercial availability. It’s not clear if the numbers for Teams users cited by Microsoft include the massive [U.S. Department of Defense contracts](#).

## Teams as the Modern Outlook

When Microsoft launched Outlook 97, they delivered an application that brought together the tools an office professional needed to get their work done. The notion of velocity was somewhat different then as it might take several hours to send and receive emails, especially over dial-up connections. While products like IRC and VAX Notes existed, the idea of unstructured, fluid, rapid, dynamic chats between team members chats was rare in an era when network bandwidth was scarce. Real-time video calling to the desktop or a smartphone was many years in the future. Even with its now ancient roots, since its introduction, Outlook has been the standard yardstick for measuring the effectiveness of office productivity tools. Microsoft sometimes refers to Teams as “Modern Outlook.” This doesn’t mean that Teams replaces Outlook. Rather, Teams is a new take on an integrated application that reflects the way that a growing part of the workforce work today. Teams lacks the email and personal calendar capabilities that Outlook has and is an application that focuses on internal rather than external communications. Instead, Teams includes other features to attract people to consider making this application the place where they spend a lot of their working time:

- Threaded conversations arranged in channels and organized in teams. Channels form the basis of Teams communications for groups of people within and outside the tenant. Conversations are persistent, meaning that they endure over time unless messages are removed by users, administrators, or by policy. Members who join a team can reference earlier conversations because they persist in the channels. Private and shared channels are exceptions as users must join the channel membership before they can access the messages in these channels.
- Chats (aka instant messaging), including federated interoperability with Teams users in other tenants and Skype consumer users. You can even chat with yourself!
- Video and audio meetings.
- Access to SharePoint Online document libraries and OneDrive for Business to store files.
- Extensibility through connectors, bots, apps, actionable messages, and tabs.

Table 12-1 compares some aspects of Outlook and Teams. This is not an attempt to say that one application is better than the other because both applications are firmly rooted in their time. In some ways, the difference between the two reflects the change in how people work now compared to when Microsoft created the first Outlook client. Another difference between the two is that Outlook is a personal tool that extends into some areas of sharing while Teams builds sharing into its core.

Another difference is that Outlook supports connections to mailboxes on-premises and cloud platforms (including non-Microsoft email servers). Teams focuses exclusively on Microsoft 365. The way Microsoft builds Teams by weaving together Microsoft 365 components makes it impossible to contemplate it working in most on-premises organizations.

Extensibility and adaptability are core to the Teams paradigm, which results in a far broader range of options to make Teams work in the way that individual teams need the platform to behave (in this respect, Teams follows the same approach taken by SharePoint). Through extensions built by partners and with automatic updating of clients, Teams adapts to changing circumstances in a way that Outlook has never could. Perhaps this is the major difference between the two applications.

	<b>Outlook</b>	<b>Teams</b>
<i>Communications</i>	Email for both internal and external communications.	Conversations are organized in personal chats and channels. Support for inbound email only. No outbound email.
<i>Collaboration/sharing</i>	Public Folders.	Files (SharePoint)/OneDrive for Business.
<i>Calendar</i>	Personal and shared.	Team calendar app (synchronized with Outlook). A <a href="#">channel calendar app</a> is also available.
<i>Sharing</i>	Shared mailboxes.	Private, shared, and regular channels and group chats.
<i>Integrated applications</i>	None.	SharePoint, OneDrive, Planner, OneNote, Power BI, Shifts, Whiteboard, Forms, and other integrated applications.
<i>Video/Instant Messaging</i>	None.	Video and audio meetings with chat, transcription, recording, and integrated meeting notes (including large meetings and webinars).
<i>Extensibility</i>	Outlook add-ins, connectors, and actionable messages.	Bots, connectors, tabs, apps, and actionable messages.
<i>Clients</i>	Outlook desktop (Windows and Mac), OWA, and Outlook mobile apps.	Desktop (Windows and Mac), browser, and mobile platforms.

Table 12-1: Comparing Outlook and Teams

Even with all its quirks and flaws, after over twenty years of polishing, Outlook still fits well into the work habits of hundreds of millions of people (perhaps because they have used Outlook for years). The growing number of Teams users is evidence of its acceptance in the modern world of work. Microsoft sometimes refers to Outlook and Teams as two hubs for teamwork, and that's about the best way to look at the two.

Intelligently used, Teams focuses on rapid, brief, informal conversations while Outlook continues as the go-to client when people need to communicate outside the tenant or communicate in a more formal and structured manner inside the tenant. Some refer to this mix-and-match approach as "right-sizing" email, meaning that email continues for all manner of communications while Teams handles some of the less formal interaction that typically happens within tight groups focused on a common purpose. Companies can choose to continue focusing on email as their primary method for collaboration, move to Teams, or mix and match the two styles to meet the needs of different sets of employees within the organization.

## Teams Versions

Teams is available in five mainstream versions:

- **Enterprise:** This is the version that's part of enterprise (and equivalent academic and government) plans and is the focus of the discussion here. The Microsoft 365 Business plans also include Teams. The enterprise version of Teams is the only one that supports the Microsoft Phone system.
- **Teams for Your Personal Life (aka Teams Personal):** Originally [launched in preview in June 2020](#), the version of Teams aimed at helping people to organize aspects of their personal lives is available for the Teams mobile clients and Windows desktop and browser clients. The clients support different

features tailored to the operating system, but all clients support personal chat (up to 250 participants), video and audio calls, and calendar scheduling of meetings. The application uses a Microsoft personal account for access to the Outlook.com mailbox and calendar and OneDrive shared file storage. Users can switch between work and personal accounts as easily as switching between different tenants. The Teams personal product requires the use of Microsoft Service accounts (MSAs). Windows 11 includes a version of Teams personal which includes chat and calling capabilities. This client is based on a new Teams architecture which Microsoft hopes to deliver in the enterprise desktop client sometime during 2022.

- **Exploratory/Trial:** If you have a Microsoft 365 plan that doesn't include Teams, Microsoft offers the [Teams Exploratory license](#). Users with a valid tenant account can sign up for the exploratory experience and receive a license for the software components needed for Teams, some of which, like Exchange Online, Forms, Planner, and Stream, their account might already hold. Tenant administrators can stop users from signing up for the Teams Exploratory license by disabling the option to let users install trial apps and services in the **User owned Apps and Services** section under Settings in the Microsoft 365 admin center.
- **Free:** Available to [organizations with 300 or fewer users](#) (some free organizations created since 2020 can technically support up to 500,000 users) who do not have an existing Office 365 or Azure AD tenant, Teams Free uses the same architecture running on the same infrastructure as its enterprise counterpart. The user interface is mostly the same and supports the full range of Teams clients. Users of the free version use accounts from vanity domains that are not associated with Azure AD domains or consumer domains such as Outlook.com or Gmail. However, behind the scenes, every Teams Free deployment has an Azure AD tenant to store its accounts and groups. Teams Free supports chats, audio calling, guest user access, online meetings, calling, and the Office Online apps, and includes 10 GB of storage for files within channels (SharePoint Online) and 5 GB per user (OneDrive for Business) for personal storage. Teams Free supports the installation of many apps. The target markets for the free version include SMBs, non-profit organizations, and schools.
- **Essentials:** Introduced in December 2021 for \$48/user per year, Teams Essentials is very like Teams Free. The major differences are higher participant numbers for meetings and group chats (300 versus 100), longer meetings, more storage, and Microsoft support. For an extra \$12/user per year, the Microsoft 365 Business Basics product includes Teams along with Microsoft 365 functionality (Exchange Online, SharePoint Online, Stream, Planner, etc.).

The big differences between the Enterprise and Free versions are:

- *Scalability.* The membership limit for teams in the enterprise version is 25,000 (tenant and guest accounts).
- *Storage:* The storage available for the free version tops out at 610 GB (for 300 users), after which customers must buy extra storage from Microsoft.
- *Access to Phone System features* such as PSTN calling and scheduled meetings (because the free version lacks an Exchange Online mailbox to use to schedule meetings).
- *Compliance.* The enterprise version captures records for compliance and audit purposes and supports retention and data loss prevention policies.
- *Office Collaboration:* The desktop versions of the Office applications include added collaboration and security features (like support for protected documents).
- *IT Controls:* The enterprise version includes tools for the administration of Teams, including automation through PowerShell and Microsoft Graph.
- *Extensions:* The enterprise version of Teams supports extra first- and third-party apps, like Stream, Planner, and Dynamics 365.

One of the interesting aspects of Teams Free and Teams Essentials is that these use a hidden Azure AD tenant. In effect, when using these versions, you connect to a tenant that hides many enterprise features in the GUI.

This means that if an organization wishes to upgrade, it is [relatively straightforward for Microsoft to convert an organization](#) to a full-blown tenant. As you'd expect, because all versions share a common platform, Microsoft updates the code for all versions concurrently. This doesn't mean that Teams Free and Teams Essentials will gain new features over time as it's quite likely that Microsoft will only expose new features if a competitive need exists.

## Industry-Specific Versions

Microsoft has several packaged versions of Teams developed for different verticals. These include:

- [Teams for Healthcare.](#)
- [Teams for Education.](#)
- [Teams for Frontline Workers.](#)
- [Teams for Government.](#)
- [Teams for Nonprofit.](#)
- [Teams for Retail.](#)

Each package includes apps and customizations tailored for the vertical. For example, Teams for Healthcare includes the Virtual Visits app, while Teams for Retail includes apps like Shifts, Approvals, and Walkie-Talkie.

## Teams Architecture

The Teams architecture integrates many facets of Microsoft 365 and exploits functionality from different Microsoft 365 applications and components. Figure 12-1 shows the major components in the Teams architecture, divided into the presentation layer (clients) and the back-end services, which run in Microsoft data centers. Some of the terms used might be unfamiliar to you:

- The current (and original) version of Teams uses [Electron](#) to build the desktop and web versions. Electron is an open-source framework based on Chromium and Node.js developed through GitHub. Electron allows developers to reuse web components to create a desktop GUI. Microsoft designed Teams to be a cross-platform application, so all clients reuse as much code as possible. Much of the Teams code is in TypeScript or Node.js and although Angular was the original framework, the current implementation uses React.js and Edge Chromium to make it easier to share components with other Office applications (as in the Files channel tab with SharePoint Online). The desktop version of Teams is a wrapper (Electron) around the website while the mobile clients combine React Native with platform-specific components like Objective-C (iOS) and Java (Android). In June 2021, Microsoft said that the next version of Teams (2.0) will move away from Electron to use React and the [Edge WebView2 control](#) instead. Office 365 apps like Outlook already use Edge WebView2 for components like the calendar room finder and the Teams meeting add-in.
- **Intelligent Communications** is the media stack used by Teams for functionality such as video and audio conferencing. Skype consumer uses the same stack, which is very different from the old on-premises code base used by Skype for Business Server. Because the Intelligent Communications stack runs in the cloud, changes and evolution occur at a more rapid pace.
- **Experimentation** is a configuration service used for Intelligent Communications components.
- **MRU** means "Most Recently Used," a service used to keep track of files accessed by users through Teams.
- **WAC** means "Web App Companion," more commonly known as the [Office Web Apps Server](#).
- The email service used by Teams allows channels to accept inbound messages. Teams does not support outbound email.
- **AAD** is Azure Active Directory, used by Teams to authenticate tenant and guest users, manage group membership, and for extended services such as group settings and expiration policies, conditional access policies, and so on.

- The **Notification Service** generates notifications for users that appear in their Activity Feed. Applications can post their notifications to the Activity Feed using the [activity feed notification API](#).
- Teams uses **Exchange Online** to access Information Governance services such as the processing of retention and data loss prevention policies for personal chats and channel conversations.
- Microsoft captures a great deal of **Telemetry** data for Teams operations. The data is anonymized.

Teams includes an experimentation layer to support the deployment of different features to targeted tenants. The layer allows Teams to deploy new features to targeted sets of users for testing at different stages of the development process. Users within Microsoft see a very early version of the feature, customers who sign up for testing through the Technology Adoption Program (TAP), or people who choose to enable preview features access features under development. Normal tenants only see the feature when Microsoft makes the software generally available. Chapter 15 discusses the preview channels used by Teams in more detail together with information about how Microsoft links Teams preview with the Current Channel (Preview) for Microsoft 365 apps for enterprise.

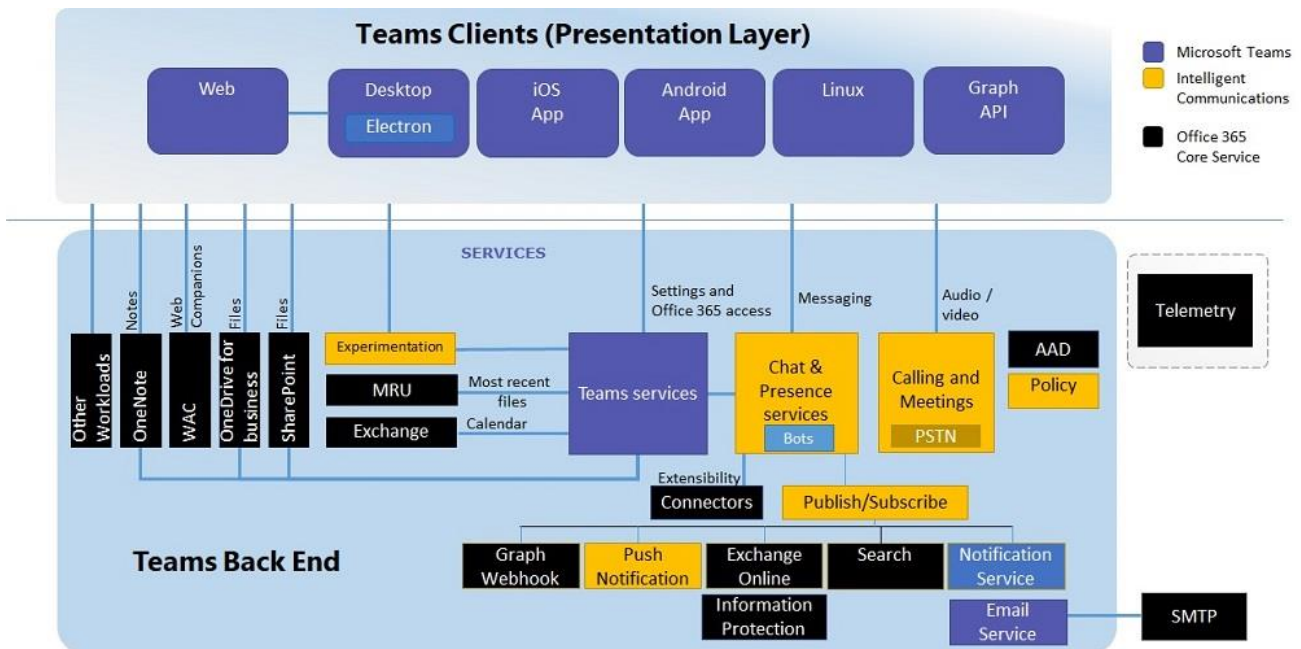


Figure 12-1: Teams Architecture (source: Microsoft)

Because the Teams architecture spans many interconnected pieces, triggers and notifications are used extensively to inform components when they need to do something. For instance, when a user posts a message into a conversation, Teams generates a new thread (for a new conversation) or adds the message to a reply chain (for an existing conversation). The members of the team or private channel hosting the conversation (known as the “roster”) need to know that a new message is available. This happens through direct notifications to desktop and browser clients connected to the Teams service and via platform-specific push notifications to mobile clients. At the same time, the new message synchronizes with clients.

In the background, components receive notifications about new messages to initiate further processing. For example, Teams might need to check a message against one or more data loss prevention policies. For compliance purposes, the Microsoft 365 substrate also captures one or more copies of the message and stores these items in Exchange Online. An update goes to the Activity Feed for each recipient, @mentions in the message might generate notifications in Teams or the operating system, and so on. In other cases, such as adding a new member to a team, result in the capture of audit records in the audit log. In other words, actions within Teams can generate a cascade of further actions across a range of Microsoft 365 services.

## Teams Building Blocks

Figure 12-2 boils Teams down into five major building blocks.

- Teams clients are available in desktop, browser, and mobile versions. Clients share a common code base, meaning that features should be easy to roll out concurrently and that people have access to the same functionality across all clients. However, some differences in functionality are often observed across clients. Sometimes Microsoft deliberately targets the release of a new feature in a specific client. For instance, being able to forward a personal chat to a channel first appeared in the mobile clients.
- Teams consumes a range of services drawn from other applications like Exchange Online and SharePoint Online to avoid recreating any wheels. See the later section covering dependencies on other services.
- Teams has its own set of microservices to manage different elements within its infrastructure. We discuss these microservices in the next section.
- The voice, audio, and telephony components used by Teams share a common infrastructure with the Skype consumer service.
- Teams consumes many Azure services. For performance reasons, the chat service stores recent chats in memory and commits its data to [Azure Cosmos DB](#). Teams stores images inserted in conversations in an Azure media store. Another example is the use of Azure B2B collaboration to support guest user access to Teams.

Microsoft publishes [architectural diagrams for Teams online](#) in PDF and Visio formats.

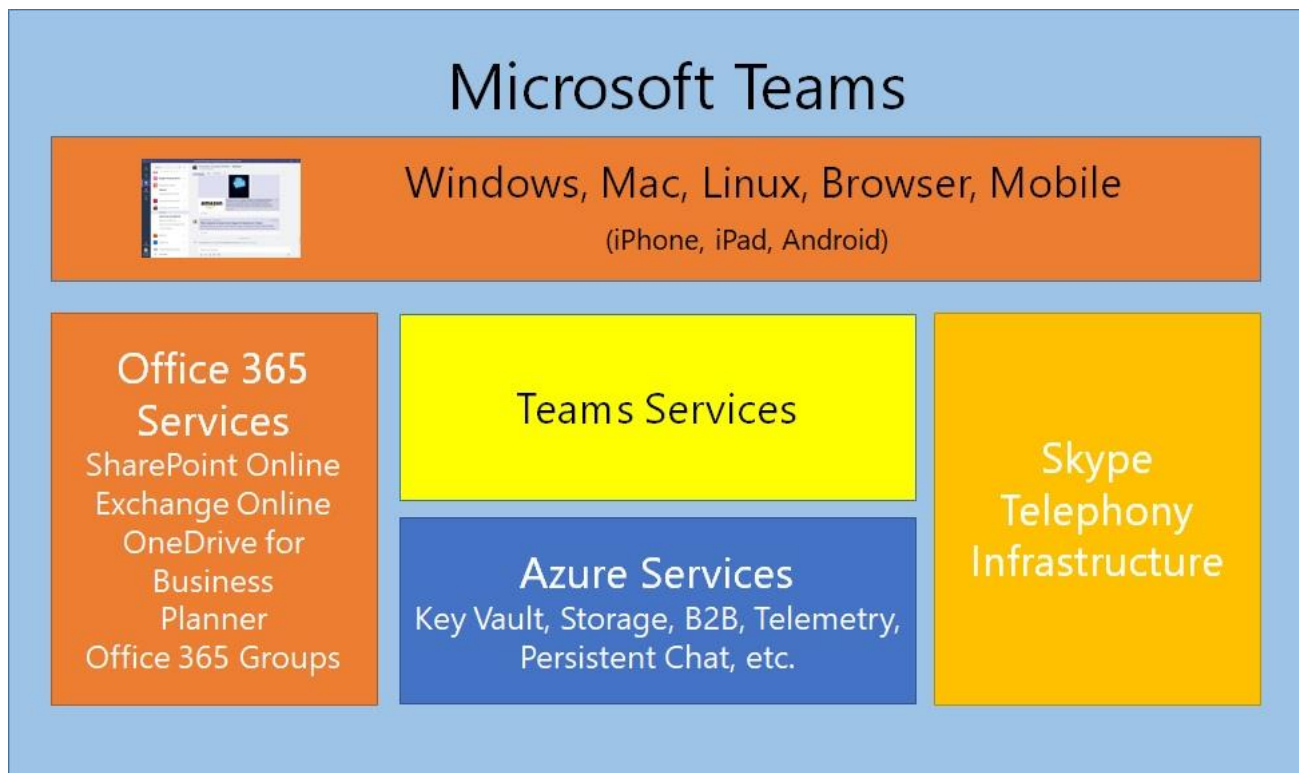


Figure 12-2: Building blocks for Teams

## Teams Microservices

Figure 12-3 shows the set of microservices Teams uses to run various parts of its infrastructure. The purpose of most of the microservices is clear. For example, the compliance microservice generates audit records about activities that occur in Teams, such as the creation or deletion of a channel. It also captures copies of personal and channel conversations for compliance purposes. Like all applications, Teams has some settings that are

individual to it, like the settings that control whether users can remove or edit messages. The configuration microservice manages these settings and works with the Azure AD synchronization microservice to manage other settings that affect Teams, like whether users can create new teams.

Azure AD synchronization also makes sure that changes to the groups that underpin Teams flow from Groups to Teams and Teams to Groups. It is important to ensure that membership and ownership changes are effective, user properties (like photos, job titles, and reporting relationships) become known to Teams, and that Teams executes its side of management actions like the deletion of a group (and its associated team). The expected synchronization interval is less than 15 minutes, but the defined SLA (internal to Microsoft) is 24 hours, so it might take some time before changes made in Azure AD appear in Teams or vice versa. The Azure AD synchronization microservice is also responsible for the recovery of Teams components when an administrator recovers a soft-deleted group (see the Groups chapter for details).

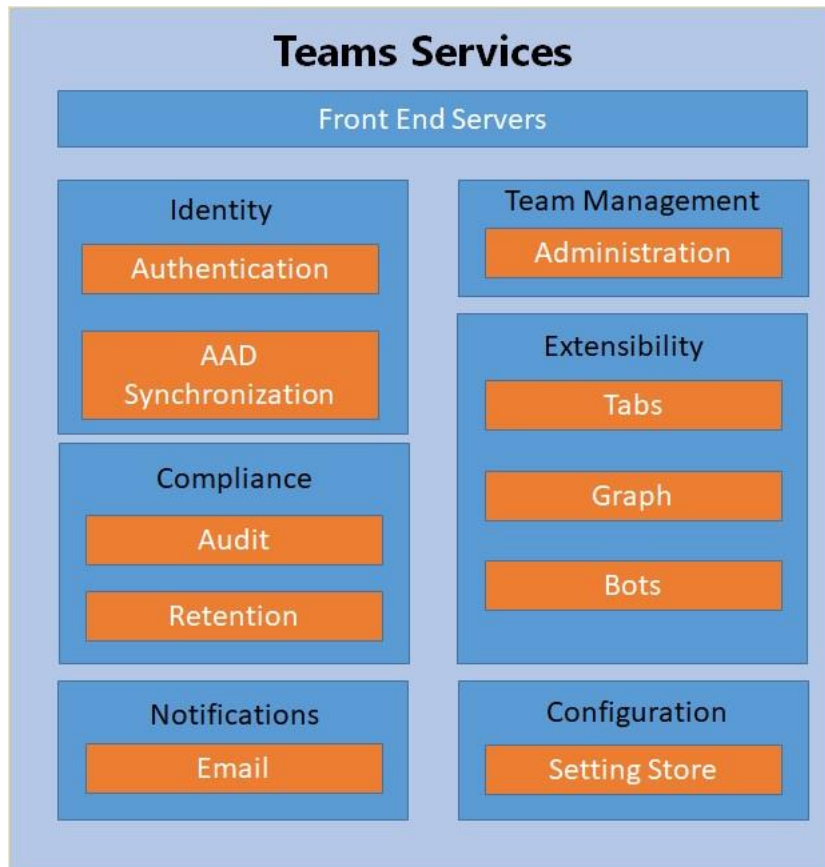


Figure 12-3: Micro-services used by Teams

**The impact of Covid-19 on Teams services:** During the early days of the Covid-19 pandemic, Microsoft experienced a surge in demand for Teams services as people started to work from home. Microsoft saw a dramatic rise in Teams meeting minutes along with knock-on effects on other services like Stream (to store meeting recordings). Coping with the unexpected rise in usage over a short period forced Microsoft to review the Teams architecture and change how services consume resources. The result was successful and Microsoft managed to cope with the growth in demand to become better prepared to scale in the future. For more information, read this [Microsoft blog post](#).

## Teams is a Collection of Apps

Apps form the basic framework of Teams. Some apps are fundamental to Teams, others are optional. Some come from Microsoft, others from third parties, and others are line of business applications unique to an organization. Administrators can control the apps made available to users through app setup and permissions policies (explained in the next chapter).



- **Fundamental Apps.**
  - **Teams:** Access to teams and channels, including the applications, bots, and connectors installed in the channels.
  - **Activity:** The activity feed.
  - **Chat:** Personal and group chats. This also includes federated chat with users in other tenants and Skype consumer users.
  - **Files:** Lists recent files accessed by a user (inside and outside Teams), files used in Teams, downloads, and cloud storage (like Google Drive).
  - **Calendar:** Access to the user's Exchange Online calendar.
  - **Calls:** Access to the Teams calling subsystem to make VOIP calls to other Teams subscribers and external parties (external calls and PSTN connections are available only after buying a calling plan).
- **Microsoft First-Party Apps for Teams:**
  - **Tasks by Planner:** Access to group (Planner) tasks and personal (To Do) tasks.
  - **Yammer Communities:** Access to Yammer communities.
  - **Insights:** Viva Insights for Teams combines insights from Viva Insights, Workplace Analytics, Wellbeing, and third-party data sources.
  - **Stream:** Access to videos managed by the Stream service.
  - **Forms:** Conduct polls (including during meetings) and questionnaires.
  - **Channel Calendar:** Access to meetings scheduled within a channel.
  - **Lists:** Access to Microsoft Lists stored in SharePoint Online.
  - **Approvals:** Basic approval workflow based on Power Automate.
  - **OneNote:** Access to OneNote notebooks.
  - **Praise:** Recognize the achievements of co-workers. The Viva Insights app includes Praise and allows users to track praise they have sent and received over the previous six months.
  - **Shifts:** Front-line worker shift setup and management.
  - **Wiki:** Access to a personal Wiki (like that used for channels and meetings).

The list of first-party apps has grown over time and can be expected to increase further.

## Dependencies on Other Microsoft 365 Components

Along with its microservices, Teams consumes a range of services drawn from other parts of Microsoft 365. These dependencies are:

- **Groups.** A group exists for each team. The group object is in Azure AD and the members of the group are the members of the team.
- **Exchange Online.** The Teams Calendar app displays information synchronized from the calendar folder in the signed-in user's mailbox (the channel calendar app reads information about channel meetings from the group mailbox belonging to the team). The contacts used in the Calls application are in the user's mailbox, as is their profile picture (avatar). The Microsoft 365 substrate captures compliance records for channel conversations in the group mailbox belonging to the team while the compliance records for personal and group chats are in the mailboxes of chat participants. Users with on-premises mailboxes can use Teams, with [some limitations](#), notably that on-premises users must have mailboxes on Exchange 2016 CU3 (it's best to use the latest available version of Exchange Server) servers to create and view meetings. This requirement exists because the Teams middle layer uses Autodiscover V2 to find how to retrieve mailbox and calendar data. Read more about [how Teams retrieves calendar information from Exchange on-premises servers](#) and how to [troubleshoot common issues](#).
- **SharePoint Online:** The Teams Files functionality (for channels) stores documents in a SharePoint team site provisioned during the creation of a new team. Teams users do not need a SharePoint

license (in many cases, users have a SharePoint Online license through one or more of the product licenses assigned to their account) to access Files, but SharePoint Online must be available within the tenant. Teams does not support SharePoint on-premises.

- **OneDrive for Business:** Users require a OneDrive for Business license to be able to share files in personal and group chats (including loop components) and to store recordings of the meetings they organize. OneDrive for Business licenses are in products like Office 365 E3 and E5.
- **Microsoft Forms:** Teams meetings use Forms for meeting polls and quizzes.
- **Planner/To Do:** The Tasks by Planner app integrates tasks created in Planner and To-Do. Team members can access Planner plans through channel tabs.
- **Power Automate:** The Teams Approvals app uses Power Automate for its workflow.
- **Microsoft Bookings:** The Teams [Virtual Visits app](#) depends on Bookings.
- **Whiteboard:** Teams meetings can share whiteboards to allow collaboration around a shared canvas.
- **Yammer:** The Teams Community app gives users access to Yammer communities. Parts of the Microsoft Viva suite use Yammer for conversations and comments.

In addition to these base services, Teams takes advantage of other Microsoft 365 components to extend its functionality as required by channel tabs and apps. Examples of other Microsoft 365 components used by Teams include Stream, Microsoft Information Protection, OneNote, the Office Online apps, connectors, Power BI, and actionable messages.

**Teams and Service Plan Licensing:** Products like Office 365 E3 and Office 365 E5 include access to the service plans for all the Microsoft 365 components listed above, meaning that users can use the individual services as needed. However, administrators can disable selected service plans. For instance, they could decide to disable Forms and Power Automate because the organization considers these services unnecessary. If this happens, users cannot access the standalone services, and they won't be able to use Teams features that depend on those services.

## Teams Limits

Microsoft's service limits for Teams are [documented online](#). Because of its dependencies on other Microsoft 365 components, some of the limits applying to Teams come from those components. For example, while Teams can support hundreds of thousands of teams in a tenant, only 5,000 of those teams can have dynamic membership due to a restriction placed by Azure AD.

Each team has a fully provisioned SharePoint Online team site to hold files, the shared notebook, lists, and the wiki. The storage used by team sites comes from the overall quota available to SharePoint Online in the tenant. Usually, SharePoint manages storage automatically and assigns storage to individual sites as necessary. You can also configure SharePoint for manual storage and [set specific limits on sites](#). If necessary, organizations can buy extra storage to increase the overall amount available to the tenant.

Microsoft has not documented any restrictions on the number of conversations, messages, or graphics that Teams can hold for a tenant in its Azure-based services. Teams does not apply any numeric or storage quotas for messages in a single team or channel.

**Teams mega-tenants:** Some tenants need to support more Azure AD objects than allowed by [the service limits](#). For instance, some education tenants need to go past the limit for groups (500,000). Microsoft has a process to allow customers who need to exceed the documented limits. If you're in this situation, contact Microsoft support and ask how to increase the limits applying to your tenant.

## Teams Storage

We've already mentioned how Teams uses different cloud locations to store its data, such as how SharePoint Online stores the files created through the Files interface. Table 12-2 describes how Teams stores the different

types of data used by its components. The dependencies that exist with other parts of the Microsoft cloud are obvious.

<b>Data</b>	<b>Primary Storage</b>	<b>Other Storage</b>
Messages	The Chat service stores its messages in Azure Cosmos DB.	Compliance records (copies of messages) are in group and personal Exchange Online mailboxes.
Images	Media service (Azure) using blob storage.	Any images referenced in messages are copied into compliance records and stored in Exchange Online.
Files	Personal files are in users' OneDrive for Business accounts. Files shared in channels are in the team's SharePoint document library.	
Voicemail	Personal Exchange mailboxes.	
Recordings	Meeting recordings are in OneDrive for Business (personal meetings) or SharePoint Online (channel meetings).	Links to recordings posted in the meeting chat.
Calendars	Group and personal Exchange mailboxes.	
Contacts	Personal Exchange mailboxes.	
Telemetry	Microsoft data warehouse (inaccessible to customers).	
Meeting resources	Transcripts, attendance, and registration reports are in the Exchange mailbox of the meeting organizer.	

Table 12-2: Where Teams stores its data

The messages making up Teams channel conversations and personal chats form "reply chains" (the original topic and all subsequent replies). Messages for channel conversations remain in the Teams store unless removed by retention policies or the deletion of the team. Unlike channel messages, which the hosting team owns, participants in a personal or group chat collectively "own" the messages in a conversation. Messages remain in the Teams store unless removed by retention policies or following the removal of the accounts of all participants.

## Geographic Location of Teams Data

Like other non-core workloads, Microsoft does not store the Teams-specific data in every data center region. This data includes personal chats and channel conversations, teams and channel metadata for a tenant, and media inserted into chats (data inserted by Teams into SharePoint and Exchange exists in the data center region to which a tenant belongs). The Microsoft 365 substrate captures copies of chats in group and user mailboxes to make them available for searches and eDiscovery. The storage for images and media used in chats (except for Giphy GIFs, which Teams stores as a URL to the original file) runs in an Azure-based Media service deployed alongside the chat service.

Although Teams is not yet available in every data center region, Microsoft plans to deploy the necessary Azure services to support Teams more broadly to satisfy customer data sovereignty concerns. Microsoft documents the [current set of data center locations for Teams data online](#).

Some data accessed through Teams might be in other data center regions. For instance, in a multi-geo situation, the SharePoint team site (holding documents, the shared notebook, and wiki) used by a team running in the same region as the person who created the Microsoft 365 group for the team. The user's personal Teams data such as the compliance records captured for personal and group chats and any OneDrive for Business documents shared in chats are in the satellite region. Remember that Microsoft's commitment to data sovereignty only extends to the data that they manage and does not include third-party

data accessible through Teams. For instance, if you create a tab for Trello, the data manipulated through that tab are in Trello's cloud and might be outside the region.

## Teams Advanced Communications

The Teams Advanced Communication add-on is available for all Office 365 and Microsoft 365 plans. It covers several enterprise capabilities including:

- Customization for Together mode background scenes.
- Deployment of custom Teams policy packages.
- Custom screens for meeting pre-join and lobby screens.
- Advanced real-time telemetry for user communications.

More information on the current capabilities licensed by the add-on [is available online](#).

## The Structure of Teams

A team can have up to 10,000 members. The basic structure for a team is a set of channels that give members a framework to organize conversations (discussions composed of "persistent" chats) and other relevant information, activities, and applications. A channel is a dedicated section within a team used by members to communicate about a topic, like a space to organize a project that the team must manage. All team members can see everything in a channel, so if you want to keep something private, you must do so in a personal chat. Only the people involved in a personal chat can see what goes on there.

Private teams limit access to any team resources to people added as team members, including guest members. Public teams are open to anyone who wishes to join. When someone joins a team, they can work with the resources available to the team. When they leave a team, they lose access to those resources. Although some restrictions do exist for guest members, the principle that all members share equal access to team resources holds true.

## Owners and Members

Teams use Groups to manage team membership, which means that the roles available within a team are owners (who manage the team) and members. An owner is also a member and there should be at least two owners for each team to make sure that someone is always available to handle basic administration such as adding new members. If all the owners leave a team (perhaps because they leave the organization), a tenant administrator can promote a member to become an owner or add a new owner to the team. Unlike Groups, Teams does not support the concept of subscribers who receive copies of any new posts for a conversation via email. If a member wants to see what is happening inside a team, they must open it to join in the conversation.

Microsoft recommends that you should allow Teams users to create new groups. If limited by policy, users cannot create new Teams because of the dependency on Groups. Although such an approach is liberating for users because they can create new teams as and when they like, it creates some management challenges that we will get to later.

A team owner can promote any other team member except a guest to become an owner. To do this, use **Manage team** to display the current membership and then click the down arrow in the role column beside the member's name to select Owner. When someone is a group owner, they automatically become an owner (or administrator) of the associated team and can perform actions such as adding or removing users, adding new channels and connectors, and so on. To remove owner status, reverse the process and change Owner to Member. If you want to remove an owner from a team, you must first demote them to Member status and then go ahead and remove them by clicking the **X** alongside their name (Teams displays the X when you

move the screen cursor over the role column). If an owner removes a user from a team, they also lose their membership in the underlying group. You cannot have a situation where a subset of the members of a group can access the associated team and other members cannot.

Like any other group, a team can have up to a hundred owners. An individual user can create up to 250 teams. However, a global administrator can create as many teams as they need to within a tenant.

## Communications Rosters and Disabled Accounts

The membership of each team comes from its Microsoft 365 group and Azure AD is the directory of record. Internally, Teams uses the concept of communication rosters or people who can communicate with each other, to control access to features. For example, in hybrid environments, you can have disabled AD accounts synchronized with Azure AD. Because these users cannot communicate with the other team members, Teams excludes these accounts from its rosters even though they are present in the underlying group (you can see this by running the *Get-MgGroupMember* or *Get-TeamUser* cmdlets against the group). Organizations often disable the accounts of ex-employees to stop them from accessing applications. Disabled accounts can't sign into Microsoft 365, so Teams suppresses their presence because they can't participate in conversations, share files, or otherwise engage with the rest of the members. The membership of private channels is another example of a roster used by Teams.

Teams also removes blocked Azure AD accounts from the rosters of regular and org-wide teams. A Teams background process scans for blocked accounts and removes them from the rosters (but not from the underlying Microsoft 365 Groups). This process can take some time to happen. Upon the unblocking of the account, Teams restores the account to team rosters. Some [issues exist in this process](#) that you should be aware of.

## The General Channel

Channels divide conversations into logical sections within a team. Every team has a default channel called "General," which is the physical embodiment of the team. You cannot remove or rename the General channel.

The General channel is often where team owners post information about the team's goals and charter and discuss topics such as what those goals should be or how to reach them. Teams used to post messages in the General channel about people joining or leaving the team, or team events such as the creation of new channels. These messages no longer appear in the General channel and are now visible in the channel's information pane, accessible through the **(i)** icon in the channel header. The Office 365 audit log also captures events about member removal and deletion or channel operations.

Even with the removal of system messages, it is still best to avoid having discussions in the General channel. Instead, reserve the General channel for team announcements and important posts. You can do this by editing channel settings to restrict the ability to post new topics to team owners.

## Managing Channels

To support discussions, you should create a set of channels in the team and encourage users to post in those channels rather than the General channel. Before moving the focus of conversations away from the General channel, team owners should think about how to use channels to organize the work of the team. A well-thought-out, well-named, and logical channel structure avoids the potential that the team becomes a confusing mess. When the overall structure of the team is known, you can create the channels and encourage users to think about which channel is best for their posts.

In addition to conversations, channels host tabs, apps, connectors, and bots. These components allow owners to equip teams with the necessary functionality needed by members. You can also mail-enable a channel by asking Teams to generate an SMTP email address for the channel (see Chapter 13). When a channel is email-

enabled, authorized users can send emails to the channel to start conversations. The email capability only extends to inbound communications as Teams does not support the ability to send email as a channel or on behalf of a channel.

Teams supports three types of channels:

- A **regular channel** is available to all team members. This is the default type of channel created by Teams. The General channel is always a regular channel. Each team can have up to 200 regular channels.
- A **private channel** is available to a subset of the team members. A team can support up to 30 private channels, each of which has a dedicated SharePoint site to hold the files for the channel. Like a regular channel, you can add tabs (except Forms, Stream, and Planner) and connectors to private channels. A private channel supports up to 250 members, all of whom must become members of the team owning the private channel before they can join the channel membership.
- Individual teams can host **shared channels**. The membership (or roster) of a shared channel is composed of individual members and teams from the host tenant or external tenants who want to collaborate on a common topic. People who aren't members of the host team can join a shared channel. A team can support up to 200 shared channels, with each channel having a dedicated SharePoint site connected to the home team.

Table 12-3 compares the capabilities available in the three types of channel. The limited support noted for channel tabs, apps, and bots for shared and private channels reflects the need for app developers to support these channels and the authentication models used to limit access to channel membership.

<b>Capability</b>	<b>Standard channel</b>	<b>Private channel</b>	<b>Shared channel</b>
Guest access (Azure AD B2B Collaboration)	Yes	Yes	No
Federated access (Azure AD Direct Connect)	No	No	Yes
Automatic access to all team members	Yes	No	No
Join channel without team membership	No	No	Yes
Share channel with other teams	No	No	Yes
Share channel with the parent team	No	No	Yes
Channel moderation	Yes	No	No
Breakout rooms	Yes	No	No
Email support for channel	Yes	Yes	Yes
Dedicated SharePoint site	No	Yes	Yes
Support for channel meetings	Yes	No	Yes
Support for channel tabs	Yes	Limited	Limited
Support for bots and apps	Yes	Limited	Limited
Tag members in channel conversations	Yes	No	No
Channel analytics	Yes	No	No

Table 12-3: Capability comparison for Teams channel types

Guest members enjoy full access to everything in a team except private and shared channels. They can access a private channel, but only after a channel owner adds their account to the membership of that channel. Shared channels don't support guest accounts as members. However, a guest member with a Microsoft 365 account can join a shared channel if the channel owner sends an invitation to their account in their home tenant. In this scenario, the user joins the shared channel using their account rather than the guest account.

## Channel Type and Controlling Creation

You cannot convert a channel from one type to another. Once you assign a channel type at creation, a channel retains that type until its deletion. If you make a mistake and create a channel with the wrong type, you must delete the channel created in error and create a new channel of the correct type. The teams policy assigned to an account controls its ability to create:

- Org-wide teams (*AllowOrgWideTeamCreation*)
- Private channels (*AllowPrivateChannelCreation*)
- Shared channels (*AllowSharedChannelCreation*)

In addition, policy settings are available to control if a user can share channels with external users or participate in shared channels hosted by other tenants. These settings are available through the Teams admin center or PowerShell. For example:

```
[PS] C:\> Get-CsTeamsChannelsPolicy | Format-List Identity, AllowPrivateChannelCreation,
AllowSharedChannelCreation, AllowChannelSharingToExternalUser,
AllowUserToParticipateInExternalSharedChannel
```

```
Identity : Global
AllowPrivateChannelCreation : False
AllowSharedChannelCreation : True
AllowChannelSharingToExternalUser : True
AllowUserToParticipateInExternalSharedChannel : True
```

## Channel Purpose

Each channel should have a clear purpose. For instance, a team might have a channel for their weekly update calls and others for specific topics, like ideas for a marketing campaign or plans for a trade show. Although the ability to split conversations over many channels gives a wide degree of latitude in the topics supported by a team, the trick is to create just the right number of channels. Creating too many channels runs the risk of confusing users (the “where do I start this conversation” syndrome). Take the example shown in Figure 12-4 (taken from a real-life deployment). How do you think the average user will respond when they see that a team has 114 hidden channels to explore? How will the average person decide what channel is the best destination for their topic? The other side of the coin is that too few channels can result in a mixture of wildly different topics in channels and lead to people losing sight of important conversations. It’s important to achieve a balance, with the preference to go slowly when creating new channels.



114 hidden channels **New**

Figure 12-4: Too much choice in too many channels?

When users find channels with content that interests them, they can follow those channels so that notifications for items posted to the channel appear in their activity feed. Unless you like dealing with hundreds of notifications daily, it can be a bad idea to do this for a busy channel. Instead, users should only choose channels where they want to see notifications of new activity.

## Naming Channels

A channel should have a name and description to inform members what they should expect to discuss in the channel. When composing the name of a channel, you cannot use the special characters (~#%&\*{}+/\:;<>?|'"), but you can include the @ sign. You can include a full stop in a channel name if it is not at the end. You can also include emojis in channel names.

Although you don’t need to enter a description for a channel, it is good practice to do so. Teams displays the name given to a channel no matter what language someone uses. For this reason, in multinational organizations, it is wise to seek channel names that make sense in as many languages as possible. This isn’t

needed for the General channel because Teams displays a translated value for “General” based on the language chosen for the client. For instance, the channel is *Général* in French and *Allgemein* in German.

Except for the General channel, you can edit channel settings to rename a channel if necessary. Until September 2021, Teams didn’t synchronize the new name assigned to a channel with SharePoint Online, meaning that the channel name shown in Teams could differ from what appeared in SharePoint Online. Microsoft then deployed an update to make sure that Teams renames the folder created for the channel in SharePoint Online when it renames a channel. Channels renamed before September 2021 do not synchronize their name with SharePoint Online, so if you want the same name to appear in both places, you need to rename the channel again.

## Showing Channels and Notifications

Some teams have more channels than others and the discussions in some channels are more interesting than others. Users can organize the channels that they find most interesting by making sure that these channels are visible while putting the less interesting channels into the background. For instance, you might want to keep channels where active discussions about your current projects happen in the foreground while pushing channels used to track posts about code updates or similar developments into the background. Use the **Show** option (in the ellipsis menu) to mark a channel to appear in the visible list while the **Hide** option puts it on the hidden list. You can decide to access a channel at any time (with or without putting it back into the visible list). Note that if you create a channel, Teams assumes that you want the channel to be visible and makes it visible.

To help new members get to know the ebb and flow of what happens inside a team when someone joins a team, the application automatically adds the five most popular channels from the team to the visible list, meaning that the channels appear under the team name when you open it. Popular means the busiest channel in terms of conversations. The user can accept the recommended set of channels or organize a custom list by hiding and revealing available channels. If a team has fewer than five channels, Teams automatically adds all existing channels to the shown list for new members. When they create a new channel, owners can mark the channel to show up automatically in the list visible to team members, which is a good way to make members aware of the existence of an important new channel. If an owner forgets to highlight an important channel, they can do this later through the Channels section of team settings by checking the auto-show box. Users can always hide the channel if they decide that they don’t consider the content to be important.

Some channels are inevitably more important than others. For these channels, use the **Channel notifications** option in the channel menu to enable notifications for new posts and replies in the channel. You can choose to have notifications for:

- All new posts.
- All new posts and replies.
- Channel mentions.

Notifications come in two types: banners and the activity feed (you can choose both). A banner is a notification message that pops up to tell the user that someone has posted in the channel. Banner notifications can become quite distracting, especially in busy channels, so it is usually better to have Teams update the activity feed with notifications for these channels. If the notification settings are not set for a channel, Teams uses the user’s **All teams and channels** notification setting (controlled in **Settings** under the user avatar).

## Pinned Channels

Over time, the number of teams and associated channels available to users will inevitably grow. This can create a challenge for users to keep track of the channels where the most important discussions occur as their



activity feed becomes cluttered and they receive too many notifications. The solution is to use the **Pin** option (in the channel menu) to mark selected channels as important.

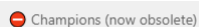
Teams places the set of pinned channels on top of the teams list in the sidebar. Teams lists the set of pinned channels in the order the user pins them to the list and the user can reorder the list by dragging and dropping channels into their preferred order. The name of the team a pinned channel belongs to appears under the channel name, and if the channel has unread messages, Teams highlights this fact by bolding the channel name. The user can remove a channel from the pinned list at any time with the **Unpin** option.

## Removing Obsolete Channels

Channels can become underused or obsolete over time. Team owners can delete a channel at any time. This action removes the channel from the list of active channels displayed to users. Deleted channels enter a holding pattern for 30 days. During this period, team owners can restore the channel through the Channels section of the Manage team dashboard. After 30 days, Teams removes deleted channels and they become irrecoverable. The content in the channel folder in the team's SharePoint Online document library remains unaffected by the channel deletion. There's no way to force the removal of a deleted channel before the 30-day retention period lapses.

If a channel becomes obsolete, there's no way to archive it (set the channel to be read-only). Instead, you can delete it to clean up the team. Before deleting a channel, it's a good idea to mark it as obsolete and pending deletion. This allows users to look through the messages in the channel to see if any valuable information exists that team members should extract and keep (you can't copy or move channel messages to another channel in the same or another team). It's also a good idea to check the contents of apps reached through channel tabs. Then, when you're ready to mark the channel as obsolete, a team owner can:

- Use the Edit this channel option to change the channel name. For instance, you could insert a Windows emoji (accessed using the Windows logo key and period) as a prefix as a visual marker for the channel's status, and then add a suffix to indicate that it is obsolete. For example:



Also, make sure that the *Automatically show this channel in everyone's channel list* checkbox is not set.

- Post an announcement to the channel to tell team members the deletion date for the channel.

When the deletion date comes around, a team owner can delete the channel and begin the 30-day deleted channel recovery period.

**No Way to Reuse Channels:** If a channel becomes disused, you can remove it to clean up the team. However, you cannot then reuse the same name for a new channel in that team. Once a channel name exists within a team, Teams does not allow you to reuse it. Microsoft says that the restriction exists "for information protection scenarios."

## Private Channels

Private channels have a small lock shown against their name to mark them in a team's channel list. You can't convert a regular channel to become a private channel or vice versa. In a multi-geo environment, private channels and their dependent sites are in the same geography as the parent team. Teams administrators can manage the membership of private channels through the Teams admin center. However, Teams administrators cannot access the content of private channels unless they are a member of the channel.

Private channels handle scenarios such as the need to limit access to information to a subset of a team. For example, if you create a team to work on a project, there might be financial or other sensitive aspects of the project restricted to people directly involved with that data. To address the need, you can create a private channel in the team and add those people, including guest users, to the membership of the channel. Only the

channel members can then access the conversations in the private channel and the files in the SharePoint site belonging to the channel.

Channel owners can add members to a private channel using the Teams client. Team owners and administrators can update channel membership with PowerShell. However, team owners can't access the contents of a private channel unless they become a member. Although they can see private channels listed in the Channels tab of the Manage team option in the Teams client, if they're not a member, the only sign of activity in a private channel available to a team owner is through administrative interfaces. For instance, if they use the *Get-SPOSite* cmdlet to examine the properties of a private channel, the *LastContentModifiedDate* property might indicate recent activity in the site belonging to the channel. Audit records for file activities in the channel are available in the audit log.

Because the channel owners manage the membership of a private channel, these channels need at least one owner. If the last channel owner leaves a tenant (for instance, an administrator deletes their account), the Microsoft Teams AadSync background process detects that the channel is now ownerless and selects a channel member at random, and promotes that member to be the channel owner. The user receives a notification of their new status and can then decide if they wish to hand over the responsibility to another channel member by promoting them to be an owner.

The SharePoint Online sites used by private channels inherit their properties and membership from the parent team. A synchronization process updates the site properties when necessary. If an administrator updates the site properties using a SharePoint admin interface, the synchronization process will reverse the change within six hours.

## Limiting Creation of Private Channels

By default, any tenant user can create a private channel. If you want to limit the creation of private channels, go to the Teams section of the Teams admin center, and edit the Global (org-wide default) Teams policy to change the **Create private channels** settings to **Off**. Then create a new Teams policy with Create private channels set to On and assign this policy to the accounts you want to be able to create private channels. The PowerShell *Get-CsTeamsChannelsPolicy* cmdlet returns details of the Teams policies in a tenant (the *AllowPrivateChannelCreation* property is the one you're interested in), while the *Grant-CsTeamsChannelsPolicy* cmdlet assigns an appropriate policy to accounts.

## Shared Channels

Shared channels are part of the Microsoft Teams Connect initiative. Shared channels depend on a new federation mechanism called Azure B2B Direct Connect. External users can present credentials gained by signing into their home organization and have another Microsoft 365 tenant accept their credentials to authenticate access to resources in that tenant. This mechanism eliminates the need for guest accounts as used for team membership. This does not mean that guest accounts are bad or that Microsoft will remove them. Direct Connect is a different way to grant access to tenant resources that comes with its foibles and concerns.

Users do not need to switch tenants to access shared channels hosted by other tenants. Teams clients use the authentication granted through Azure AD Direct Connect to open and display shared channels alongside standard and private channels from the user's home tenant. Figure 12-5 shows the manage channel screen for a shared channel. In the teams list, the set of channels for the selected team includes standard, shared, and private channels. As a further visual indicator to users, when browsing the teams list, teams from other tenants with shared channels display the name of the home tenant under the team name, as in "@Tenant Name."

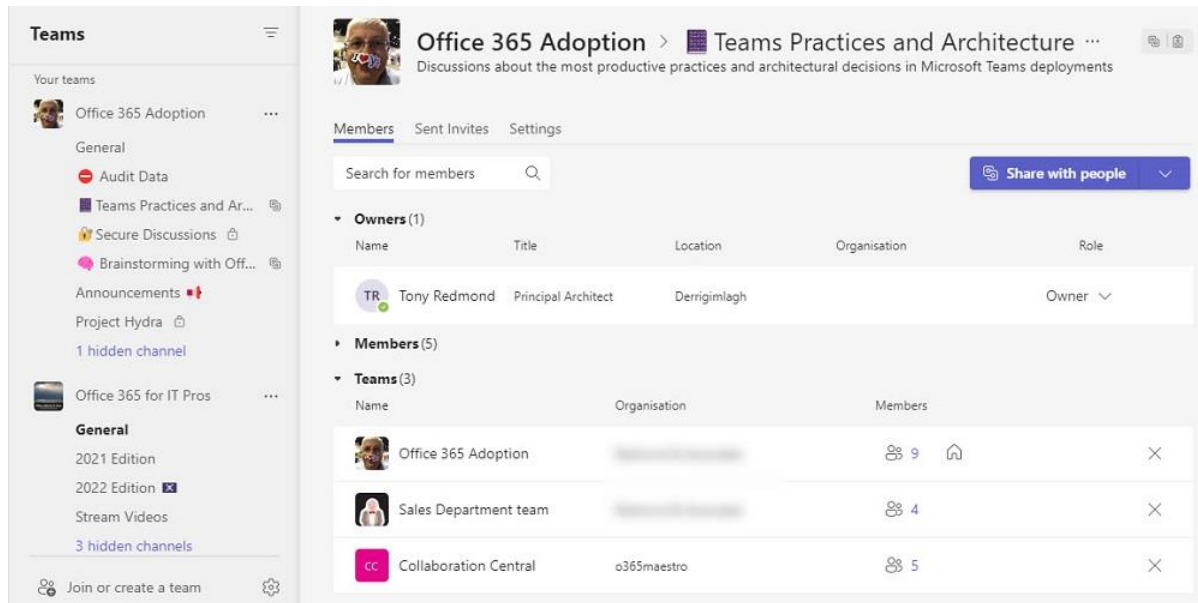


Figure 12-5: Three types of channels are visible in the Office 365 Adoption team

Azure AD cross-tenant access settings control Direct Connect sharing with other organizations. Cross-tenant access works on a mutual trust basis. In other words, the tenant sharing channels must be happy to allow inbound access for external users while their home tenants must allow outbound access (to allow the users to access resources in your tenant). If you allow inbound access to a tenant and that tenant doesn't allow outbound access to your tenant, mutual trust doesn't exist, and adding external users to shared channels isn't possible.

The simplest policy allows inbound access and outbound access from any other Microsoft 365 tenants. However, cross-tenant access settings can limit sharing to a granular level. For instance, only members of a specific group can join shared channels in a specific tenant, or only specific external users or groups (identified by the GUIDs for the objects in the Azure AD of the external organization) can join shared channels in your tenant. Applying granular control to cross-tenant settings is an Azure AD Premium feature. It's usually best to start with a policy that allows flexible sharing and then amend the policy over time as necessary. See the identity chapter for more information about Azure AD cross-tenant access settings.

Like private channels, shared channels have owners who manage the channel membership. Owners, who can only be accounts from the host tenant but do not have to be a member of the host team, can share a channel with:

- People (individual users from inside or outside the tenant). There's no need to add users from the home tenant to the host team. The first time an external user accesses a shared channel in a host tenant, they must give consent to allow the host tenant to access details of their profile like their picture, display name, and email address. This information is available to other channel members in the channel roster and profile cards.
- A team (invited using the email address of an owner of the team). Only tenant members from the team will join the shared channel.
- A team from the home tenant that's owned by the channel owner. Teams lists the available teams, and the channel owner selects which one to add.

To share, a channel owner selects the appropriate action from the Share dropdown menu. In Figure 12-5, the *Share with people* option is visible, this shares with individual users. We can see that the current membership of the shared channel includes three teams, including a team from an external tenant. The membership of a shared channel can contain individual users and teams from multiple tenants.

In the case of team memberships, an individual's access to the shared channel is available only while they remain a member of a team in the shared channel's membership. Channel owners must invite users or teams to share a channel: users or teams cannot browse for shared channels and decide to join on their initiative. You cannot invite a guest account to share a channel. Invitations are only valid when extended to full member accounts in the home tenant or other Microsoft 365 tenants.

When a channel owner extends a sharing invitation to the owner of a team, that person sees the invitation in their activity feed and must respond within 14 days. During this period, the channel owner can see the status of the invitation and how long remains before the invitation becomes void. Once extended, a channel owner cannot revoke an invitation. However, they can reject the acceptance when it comes back from the team owner. Part of the response is the selection of the team whose members will gain access to the channel (this cannot be an org-wide team). When the channel owners receive the response, they can review the team selected by the team owner and decide if they wish to share the channel with the nominated team. A channel owner can see the number of members in the chosen team, but they cannot see any details of the individual members. If necessary, the channel owner who processes the response can reject it. If accepted, Teams validates the membership of the chosen team to ensure that none of the members conflict with the cross-tenant access policy. Teams drops these members, as also happens for any guest members of the team. After Teams completes its check, the team shows up in the membership of the shared channel. A member of a shared channel can leave it at any time, and the owner of a team that shares a channel can remove the entire team from the sharing roster when they wish.

Teams limits members of a shared channel to whatever resources are available to the channel, including a SharePoint team site associated with the channel like a private channel. External members can only see limited directory information for other channel members, but they can chat or call anyone in the channel. Teams adds an External suffix to the display names of external members to indicate their status. In addition, if a shared channel has external members, Teams flags this with a banner saying *This channel is shared with members in other orgs*.

Unlike private channels, which only support Meet Now meetings, shared channels support both Meet Now and scheduled meetings, but only tenant users can schedule meetings. The person who schedules the meeting becomes the meeting organizer, but the meeting is a channel meeting, not a personal meeting. The Teams channel meeting app does not currently support shared channels. Microsoft stores the calendar data for a shared channel in the cloud-only mailbox created to store compliance data for the channel.

## Channel Information Pane

The channel information pane is available by clicking the **(i)** icon. The information pane displays some basic information about the channel (display name and description) plus:

- **Recent contributors:** Members who have recently participated in the channel by posting or replying to the last twenty topics in the channel.
- **All members:** If you select *See all members*, you're brought to the Manage team screen to view the full list of owners, members, and guests.
- **Updates:** This is where Teams posts system messages covering team membership changes and channel updates (Figure 12-6).

When you see *Unknown User* in a system message, it means that Teams can't find the user in Azure AD. Usually, this is because the user's account is deleted. A note that *Microsoft Teams AadSync* has added or removed someone from a team means that the action happened through some other administrative interface, like an update to a Microsoft 365 group using the Exchange Online PowerShell module. Teams learns about changes made to group membership by monitoring a pipeline within Azure AD. Once a change is detected, the *Microsoft Teams AadSync* process replicates the information to Teams to update the team roster.

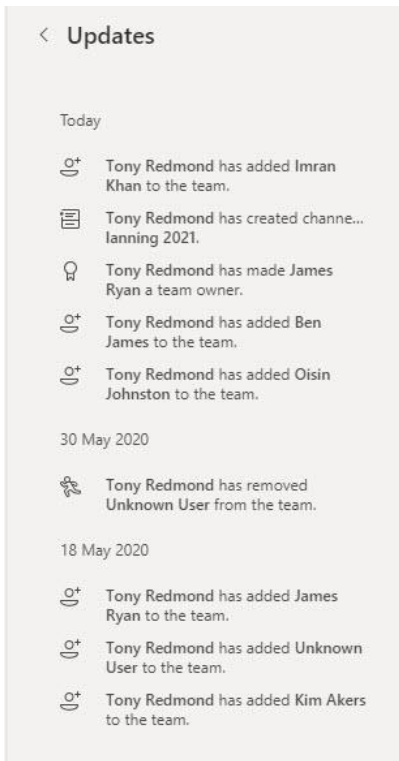


Figure 12-6: Teams system messages in the channel information pane

Any team member, including guests, can pin or unpin a channel topic or reply. No limit exists for the number of pinned messages in a channel. However, only the latest two pinned messages appear in the About section of the information pane.

## Channel Tabs

Users access the content owned by a channel through a set of tabs. A tab is analogous to a browser tab and acts as a link to a resource available within the channel. Some channel tabs are available only to regular channels while others can be used in shared and private channels too. Conversations, Files, and Wiki are default tabs that appear in all channels. Unlike the Files and Conversations tabs, you can remove the Wiki tab from a channel. Examples of other tabs that you can add to a channel include:

- **Planner:** Link to a Microsoft Planner plan owned by the underlying group.
- **SharePoint:** Link to a SharePoint site or folder to which team members have the necessary access rights.
- **OneNote:** Link to sections of the shared OneNote notebook for the underlying group.
- **Power BI:** Link to a Power BI report or workspace.
- **Office documents:** Link to specific Excel, Word, or PowerPoint documents in the SharePoint document library for the team. You can also add a link to a **PDF** document in the same manner.
- **Stream:** Link (URL) to a video file or channel stored in Microsoft Stream.
- **Website:** Link to the URL for any website. For example, you can create a very easy integration with Yammer by creating a link to a Yammer group.
- **Visual Studio:** Link to a Visual Studio Team Services board to allow the team to work on code projects.
- **Third-party apps** such as YouTube, HootSuite, Smartsheet, Polyscribe, Trello, Intercom, Polly, Zendesk, and Asana. You can browse the complete set of available apps through the Teams Store.

Where it makes sense (as in the case of links to individual documents or websites), you can assign your chosen name to the tab. When someone adds a new tab to a channel, Teams highlights the tab with a “New”

icon. The icon stays for seven days after the creation of the tab and then disappears. It also disappears after you access the tab for the first time.

**Access to Office and other Web Apps:** Because of the way it integrates different functionality into a single client, some people compare Teams to Outlook. You can make Teams a launching point for many Office applications by creating website tabs to link to those apps. For example, you could create a tab for OWA by pointing it to <https://outlook.office365.com>, for the user's OneDrive for Business site by linking a tab to <https://tenant-my.sharepoint.com/>, or for To-Do by linking to <https://to-do.microsoft.com/?app>. The general rule is that if users can authenticate themselves with a website, they can access that site through Teams. This approach allows users to work in Teams while having access to all the applications they need to use.

## The Teams Wiki and OneNote

Microsoft's [definition for the Teams wiki](#) is that it is "a smart text editor that doubles as a communication machine, because you can draft, edit, and chat all in one place." The provisioning process for channels creates a tab for the wiki when it creates a channel, but you can rename or remove the tab as needed. Wiki content is in a document library called *Teams Wiki Data* in the SharePoint team site belonging to the team. The first time someone accesses the wiki in a channel, Teams creates a folder in the document library named after the channel and stores MHT (web archive file) files for the wiki pages and sections there. It is possible to create multiple wiki tabs for a channel. In this case, the pages for all the wikis go into the channel folder. The MHT files don't hold any real data. Instead, they are pointers to information stored in hidden SharePoint lists. When you edit the wiki, Teams extracts the data from the list and presents it in the editor.

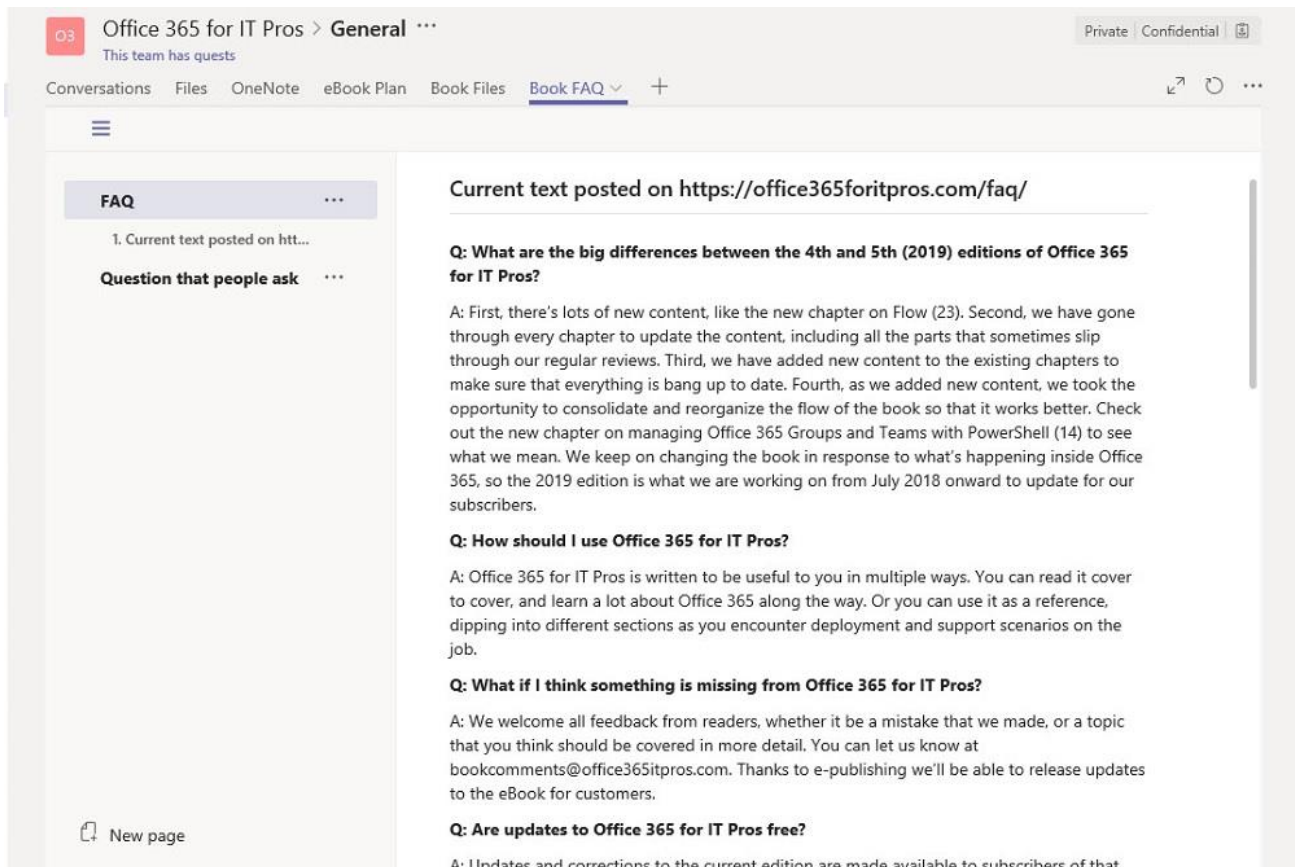


Figure 12-7: Editing content in a Teams wiki (the channel tab is renamed)

The wiki editor (Figure 12-7) is available in the desktop and browser clients and is like the editor used to compose Teams messages. The editor includes the ability to insert formatted text using a small number of predefined styles (paragraph, heading 1, and so on), hyperlinks, graphics, pages, and tables. One obvious

deficiency is the lack of support for search as you can't use Teams search (command box) or SharePoint search to find content stored in wiki pages. Meeting Notes are stored as a Teams wiki file and use the same editor.

## Wiki versus OneNote

The way that the Teams wiki works, especially in terms of notetaking and the page/section structure, often causes customers to ask why Teams includes a component that seems to compete with OneNote. OneNote is Microsoft's flagship notetaking application which supports many high-end features such as inking. You can integrate OneNote into Teams by creating a tab to point to a new or existing notebook. Many experienced deployment consultants recommend that if a tenant has invested heavily in OneNote, it should not try to force the use of yet another notetaking application on users as this approach forces users to change the way they work. It's also the case that people will lose functionality if they change from OneNote as the Teams wiki misses some of the features supported by OneNote. In these circumstances, it is best to remove the wiki tab from channels after creation and replace it with a tab linking to the most appropriate OneNote notebook.

Compared to OneNote, the Teams wiki is not as good for dynamic notetaking; it is better suited for capturing and sharing relatively static text, such as the rules of engagement about the topics discussed in a channel. One way of thinking about the two tools is that a small team working in a channel might use the wiki to sketch out some ideas that they later publish to a wider audience. A shared OneNote notebook might be the final publication vehicle, but it is more likely that a Word document or web page will be.

Because it is part of Teams, the wiki is more integrated with Teams than OneNote is. For example, you can start a conversation within a wiki page and use @ mentions to refer to the team, channel, or team members within wiki content if you need to draw their attention to something. Microsoft says that the wiki will pick up other features over time, such as the ability to send information directly from a channel conversation into the wiki. Whether tight integration is enough to prefer Wiki to OneNote remains an open question.

If you don't want to use the Teams wiki, you can remove it manually or programmatically. An example of how to remove the wiki tab from all channels in a tenant is outlined in [this blog post](#).

# Maintaining Team Membership

Teams does not maintain a central directory of teams (like the Exchange Online Global Address List) to allow end-users to browse to find teams to join. People can find teams to join through suggestions or owners can add them to team memberships, but there's no notion of browsing to find interesting teams. In some way, the lack of a directory is understandable because not all teams are public and some host very sensitive discussions, so no desire exists to make their existence known to all users. Some tenants create and publish a Teams directory. Different approaches to the task are [explored in this article](#).

The usual ways that people join teams include:

- Tenant or Teams administrators add users to team memberships via **admin portals**.
- Administrators can also update team memberships **programmatically** using PowerShell or the Microsoft Graph.
- Team owners can **manage the membership** of their teams, including being able to add guest users if allowed by policy settings.
- **Members of a team** (except guests) can add another tenant user to the membership from the team's [...] menu. If the team is private, an email request to add the new member goes to the team owners for their approval.
- Tenant accounts can **discover and join teams** suggested by Teams based on their membership of other teams and other signals gathered in the Microsoft Graph.
- Tenant accounts can also receive a **code or deeplink** to join a team.

In all cases, when someone joins the membership of a team, they also join the underlying group and can access all the resources available to both the team and the group.

## Suggested Teams

To find new and interesting teams to join, tenant users can browse the set of suggestions that Teams creates for them to join by selecting **Teams** in the navigation bar and then **Join or create a team** in the bottom left-hand corner (guests do not see this choice). Teams uses signals collected in the Microsoft Graph to display a list of suggested teams the user might like to join. The list can contain both public and private teams.

Many suggestions are based on “common interests” shared by the existing membership and the potential new member. In other words, if a public team exists with a membership composed of people who are also members of other teams with the user, a reasonable chance exists that the user will be interested in the topics discussed in that team. Groups perform similar calculations when users browse for groups to join. The computations to suggest teams to users occur behind the scenes and it takes up to a day after its creation or after a team’s access changes from private to public before Teams suggests a team to users.

To join one of the suggested teams, move over the icon for the team and a **Join team** button appears. Click the button and the team appears in the navigation bar for the user to access its resources. Once someone is a team member of a public team, they can add new members by selecting **Manage team** from the ellipsis menu for the team, followed by **Add Member**. If the team is private, requests to add members to the team are routed to team owners for approval.

**Limited Directory Searches:** If a tenant uses Exchange address book policies to set the scope for users to see items in the directory (enabled by the Search by name setting under Teams settings), Teams is no longer able to suggest public teams for people to join. This is because the query against the Microsoft Graph to find teams to suggest doesn’t take any notice of directory scoping, so Teams disables the functionality for everyone in a tenant when address book policies are in use. In addition, the /Join command, which also lists the teams available for someone to join, is disabled. If you deploy information barrier policies (see Chapter 21), scoped directory searches are used with the same effect.

## Joining a Team with a Code

To make things easier for owners to manage teams with hundreds of users, Teams supports the ability of users to join a team with a system-generated code. The idea is that it is easier to supply potential members with a code that they can input to join a team than it is for a team owner to manually add them to the membership. Using a code is also a good way for someone to know that they are joining the right team in a situation where many teams with similar names exist in an organization, such as classes in a school. Team codes work for both public and private teams. The process is as follows:

- A team owner goes to the **Team code** section under team **Settings** and clicks the **Generate** button to have Teams generate a unique seven-character code for the team. The code is a value something like “jny9ota.”
- The team owner publishes the code to potential team members. For example, if the team is open to anyone in the organization, you could publish the code on a website for all to see. On the other hand, if the potential membership comes from a more limited population, you could email the code using a distribution list.
- When someone receives a team code, they can join the team using the normal **Join or create a team** process and input the code into the **Join a team with a code** box, or by typing the /Join command in the command box. For example, /Join jny9ota.
- Teams checks the supplied code and flags an error if it is invalid. If valid, Teams adds the user to the membership and opens the team.



- Team codes are only valid for tenant members. Guest users cannot use codes to join teams.
- The team owner also can remove a code at any time or reset the code. You can't use a removed or reset code to join a team.

Apart from the way that they join a team, members who join with codes are like any other member.

## Joining a Team with a Deeplink

A [deeplink](#) is a hyperlink (URL) used to navigate to some item within Teams. Examples of deeplinks are the URLs generated by Teams when someone fetches a link to a team, channel, tab, or message. For instance, if you go to the ellipsis menu for a message and select **Copy Link**, Teams creates and copies a deeplink to the message to the clipboard. You can paste the link into another message or use it to open Teams at the linked message. The same happens if you use **Get link to team**, **Get link to channel**, or **Copy link to tab** to create a link to these elements.

The format of a deeplink varies depending on its purpose. When you link to something in Teams, the link includes several important elements to allow Teams to find the right data. If we look at the link below, three pieces of information are highlighted:

```
https://teams.microsoft.com/l/message/19:45ac50d1afa2425a80362e94f381b1e7@thread.skype/1523005757318?tenantId=c662313f-14fc-43a2-9a7a-d2e27f4f3478&groupId=ce23aef8-c551-40d4-bdb6-2bcc1eb64626&parentMessageId=1523005757318&teamName=The%20New%20Hydra%20Project%20Team&channelName=General&createdTime=1523005757318
```

These important parts of the deeplink are:

- The reply chain identifier, which identifies the message thread (**1523005757318**). This identifier appears in several places within the link. The identifier is a Unix timestamp (epoch), based on the number of seconds since 1 January 1970. The timestamp is calculated to the millisecond. You can convert the identifier to human readable time with converters such as [the Epoch Converter](#), which reveals that the identifier shown above is Friday, April 6, 2018, at 9:09:17.318AM.
- The tenant identifier, which tells Teams what tenant to access (**c662313f-14fc-43a2-9a7a-d2e27f4f3478**).
- The group identifier, which tells Teams the team that the message belongs to (**ce23aef8-c551-40d4-bdb6-2bcc1eb64626**). The identifier points to the Azure AD group underpinning the team.

The team name and channel name are also included in the link.

If you send a deeplink pointing to an item in a team to another person and they are not a member of that team, they can use it to join the team. The link opens a dialog to ask the person if they want to join the team. If the team is public and the user accepts, they become a member of the team. If the team is private, Teams sends a request to join by email to the team owners, who can then accept or deny the request. In addition to receiving email notifications about requests to join their team, owners can go to the Manage team page to view requests waiting to be processed under the **Pending Requests** tab.

## Starting a Chat with a Deeplink

Applications usually generate and consume deeplinks. However, anyone can create and use a deeplink if they know the format of the hyperlink needed by Teams. For example, let's assume that you want to allow other people in the tenant to start a personal chat to follow up with you after receiving some email. To automate the process, you can include a deeplink in your email signature. When someone reads an email from you and clicks the deeplink, Teams navigates to personal chat and starts a new chat (if they have never chatted with you before) or continues an existing chat.

The form of deeplink that you need to include in the email signature is:

<https://teams.microsoft.com/l/chat/0/0?users=Kim.Akers@office365itpros.com&topicname=Chat>

Anyone clicking this link will start a chat with Kim Akers. Deeplinks to personal chats only work for people inside the same tenant. They don't work for guest users or external (federated) access to users in another tenant.

## Leaving a Team

A user can leave a team at any time by selecting the **Leave the team** option from the ellipsis menu available for the team and any channel in the team. Alternatively, a team owner can remove a member by selecting **Manage team**, moving to the entry for the member, and then clicking the "X" to the far right of the entry.

Guests can also leave teams. However, this action does not remove their Azure AD account because the guest account might be in active use for other purposes, such as membership of a group, another team, or access to some SharePoint or OneDrive for Business documents. If you want to remove the account, you must do this through the Azure or Microsoft 365 portals, or by running the *Remove-MgUser* cmdlet.

## Maintaining Team Membership

Normally, team owners handle membership. If the team is private, the team owners take care of membership by adding and removing users, promoting members to become owners, or demoting owners to become members, which is a necessary first step before you can remove an owner from a team. If the team is public, users can join the team without owner intervention, so owners only need to add or remove guests and decide who are owners. It is also sensible for team owners to check team membership periodically in case someone joins that should not be there.

Tenant administrators and other users assigned with user management permissions can update team membership even if they are not a team owner. A variety of methods is available to do the job:

- **The Teams admin center.** Update the membership of teams.
- **The Azure AD portal.** Update the membership of the Azure Active Directory groups used by Teams.
- **The Microsoft 365 admin center.** Apart from tenant administrators, accounts assigned the User Management Administrator role can update the membership of groups, including those enabled for Teams.
- **The Exchange admin center.** Accounts assigned the Exchange administrator or User Management Administrator role can update the membership of Groups through the EAC.
- **PowerShell.** You can run the *Add-TeamUser* or *Add-UnifiedGroupLinks* cmdlets to update team membership. Given the choice, use the *Add-TeamUser* cmdlet because this mimics what happens when you add a user through a Teams client. See the discussion in Chapter 23 explaining how to use PowerShell to manage Teams. For now, this example shows how to add a user to a team:

```
[PS] C:\> Add-TeamUser -GroupId (Get-UnifiedGroup -Identity "My Microsoft 365 Group").ExternalDirectoryObjectId -User Donald.Vickers@Office365itpros.com
```

Because of the need to synchronize directories, users might experience a short delay before Teams picks up the new membership data and they can access a team.

**Unknown Users:** If you see an "unknown user" listed in team membership, it refers to someone whose account has been removed from the tenant directory or whose account cannot be resolved against the directory, or a glitch in the local Teams cache. Because the Teams client cannot resolve the member against the directory, they are deemed to be unknown.

## Stale Teams

Over time, people usually accumulate membership in multiple teams. They might end up being a member of so many teams that it's difficult to find a specific team in their list of teams (or "teams gallery"). Teams divides the teams a user belongs to into a section called *Your teams*, meaning the teams the user has opted to highlight because they are important to them, and *Hidden teams*. Teams in the hidden section might be as important to the user because of the information they contain, but the user might not want to access them as regularly. By hiding teams and placing them in the *Hidden teams* list, the user focuses on the teams they work with most often.

To help users manage the set of teams they belong to, Teams periodically checks the set of teams available to each user (but not for guest users) to detect those that might be "stale" and could be moved out of view to allow the user to focus on the teams most important to the user. Several tests are used to decide whether a team has become stale. The basic test is that the user has not accessed the team in the last 45 days, but there are some qualifiers. For instance, a team with an unread @mention for the user is never regarded as stale, the team the user last moved from the more list is never regarded as stale, and the same is true for any team marked as a favorite in the last 45 days.

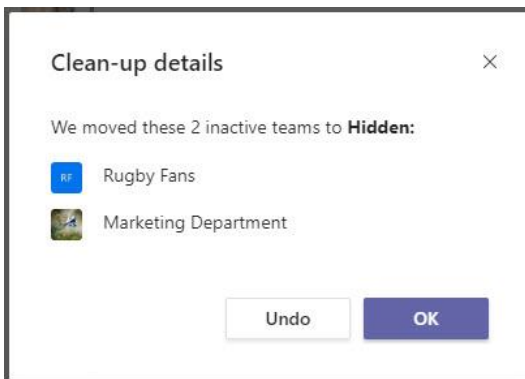


Figure 12-8: Teams shows that it has hidden some teams

When Teams decides that some teams are stale, it notifies the user that these teams have been moved to the *Hidden teams* section (Figure 12-8). The user can undo the move if they wish, which moves the team back into their *Your teams* list.

Users always have access to every team they belong to, but they must navigate to the *Hidden teams* section to find hidden teams when they need to access their contents (alternatively, they can go to the Manage teams section and search for the team name). For ease of finding a team in what can be a very large list, Teams organizes the *Hidden teams* section alphabetically, and users cannot move teams up or down within the list. Users can move a team from the hidden list by selecting the team and choosing the Show option from the ellipsis menu. Likewise, they can move a team into the hidden list by selecting the Hide option from the menu.

## List All My Teams

To see the set of teams that they belong to, a user can click the cogwheel icon at the bottom of the Teams list to expose the **Manage teams** screen. The splits into active and archived teams and lists:

- Team Name. Click on the name to access the General channel for the team.
- Description.
- Membership (Owner or Member).
- People: The number of members of the team.
- Type: Org-wide, Public, or Private.
- Show or hide the team to put it in the *My Teams* (Show) or *Hidden teams* (Hide) lists.

- More options: Manage team, add channel, add members, leave the team, **edit team**, get link to team, **archive team**, or **delete the team**. The bolded options are only available to team owners.

The Analytics option available in this section gives users a view of how active the teams they belong to are. Statistics available for active teams include the number of members, guests, posts, and replies over the last 7, 30, or 90 days.

## Arranging the Teams List

After you join a team, you can move the team within the list of teams to which you belong (*Your teams*). Teams does not display the list alphabetically but in the order that you joined the teams. To reorder the list, simply drag and drop a team into the position that you want. In most cases, people place their most important teams at the top. You can also use the **CTRL+Shift+Up** and **CTRL+Shift+Down** key combinations to move a team up or down within the list (on a Mac, use the Command key).

You cannot change the channel order within a team. The General channel is always first, with the other channels following in alphabetical order.

## Members and the Manage Team Option

The **Manage team** option displays information about a team and is available to members (including guests) and owners. The ability to take management actions, like changing a team setting or adding a channel, is restricted to team owners, tenant administrators, and teams service administrators. Members can see:

- The **membership** list, including their title, location, and role. Users can add a member here. If the team is private, the membership request must be approved by a team owner.
- The **channels** in the team, including the channel name, description, and date of last activity. The ellipsis menu is available to allow the member to get the email address for the channel (if enabled), get a link to the channel, or follow/unfollow the channel. If team settings allow members to add, remove, or restore channels, they can take those actions here.
- The **apps** that are available to the team. If team settings allow members to add or remove apps, they can do so here.
- **Analytics** for the team are available to give an overview of how active the team is.
- **Tags** defined for use in the team. Some tags are available in all teams, but most are specific to a team.

## Teams Messaging

For many organizations, the central focus and purpose of Teams is the ability to connect team members through short, interactive conversations. Or as Microsoft puts it "*communicate in the moment and keep everyone in the know.*" The individual messages that make up conversations are persistent, meaning that they continue to exist unless the author or a team owner decides to remove them. Teams conversations are complete with all the modern ways to brighten discussions (reactions with icons such as "Like" and "Angry", emoji, and memes). Reactions (like, laugh, angry, and so on) are important signals to chat participants (and recorded in the Microsoft Graph) to show the view someone has about a message in a conversation. If desired, you can disable the ability to add images to conversations through Teams messaging policies managed in the Teams admin center.

Conversations in channels are public because they are available to any member of the team. Those in chats are private to chat participants. You can't change a private conversation to be public or vice versa or move a chat from one channel to another, either within a team or between teams. People control the participants in private chats by adding others to or removing them from chats. The mechanics of contribution are identical for public and private teams, but people indeed tend to be more chatty or informal in personal chats and

more formal and perhaps verbose in channel conversations. In either case, participants type their thoughts into the compose box and then post them to the chat or channel. Or just use the thumbs-up reaction to indicate that you've seen and approved of a message without the need to post a more complex response.

Chats are persistent, and people can leave and rejoin the conversation as they wish. Persistent means that conversations are enduring and can be resumed at any time. For performance reasons, Teams caches recent messages in memory. Clients might have to page older messages back from storage before being able to display details of a conversation. Normally, it takes a client just a few seconds to retrieve and show old chats.

If you find that you want to develop a personal chat into a more general discussion, you can either expand the chat up to the participant limit or start a new conversation in an appropriate channel.

## Starting Conversations

Channels divide into topics or conversations. Each conversation has a base note to set the context for the conversation followed by a set of replies. In technical terms, a conversation is a set of messages connected by a common thread identifier, which is how Teams knows what messages belong to a conversation. Replies are ordered within a conversation using the message timestamp. To start a new conversation, click **New conversation** and enter the text for the initial post. Although effective, this isn't a great way to start a conversation. It's much better to give the new topic a subject to inform team members about what you want to discuss. Entering a subject also makes it easier for people to scan a busy channel and find conversations important to them.

To start a new topic and give it a subject, click the Format icon under the Start a new conversation box (the icon looks like a capital A with a pen at its bottom right). This exposes several options to give much more control over how and what is posted. The important options here are:

- **Add a subject:** Give your topic a meaningful subject to clearly show what's going to be discussed here. A good subject is brief, to the point, and clear about the intent of the post.
- **Post as a conversation or announcement:** The default is to start a regular conversation, but you can choose to post an announcement to give your topic more impact (see below).
- **Restrict who can reply:** The default is that all members of the channel can reply, but you can restrict replies to just you and channel moderators. People often do this for announcement posts when organizations want to convey information to team members when the need to debate a matter no longer exists. As explained in the Managing Teams chapter, individual teams can have moderators to restrict who can add topics to channels, but you don't need to enable channel moderation to restrict replies. If specific moderators aren't defined for the channel, team owners serve in their place.
- **Post in multiple channels:** Normally you only want to post a conversation in the channel you're working in, but you can post in up to 49 other channels if you want. Naturally, you can only post in channels that you are a member of. You can't post in channels where posting is restricted to moderators (and you're not a moderator). Teams treats each message separately and, apart from visiting each of the target channels, there's no way to see all the replies for the same message from all the channels where the message is posted. You can include attachments in multi-channel posts. If you upload a document from your workstation, Teams stores it in the channel folder of the SharePoint document library belonging to the current team and shares it from that folder with the teams for the target channels.
- **Use an advanced editor:** Instead of a single line editor, you can use an editor like WordPad. The editor includes the ability to insert tables, code snippets, weblinks, bulleted and numbered lists, and highlight text in different ways (see *Making Your Point* below).

**A use case for multi-channel posts:** A good example of where it's very useful to post in multiple channels is when a company wants to spread the news of an announcement as widely as possible. Create an

announcement post and limit replies to the author and channel moderators to eliminate the need to deal with responses to the individual posts. You still want to encourage communication and responses to the announcement, and a good way to do this is to include a link in the announcement to redirect users to a topic in a channel in an org-wide team where they can post replies.

## Editing Posts

If you make a mistake, you can edit a post (message) to correct the error, but only if the Teams messaging policy applied to your account allows you to edit sent messages and the team settings allow members to edit their messages. In most cases, it's reasonable to allow people to edit their messages. All of us make mistakes and there's nothing as annoying as sending a message and then discovering a spelling or grammatical error.

If you post to multiple channels, you can edit the original post and the edited text will appear in all the channels where the original message was posted. However, you won't be able to update the post in channels owned by teams whose settings don't allow members to edit posts. You can also edit a conversation that is only posted in a single channel to add extra channels and make it a multi-channel post. When this happens, the original topic message is posted to the new channels but any existing replies for that topic are not.

The author of a message can recall a multi-channel post from a channel by editing the post to remove that channel. You can even edit a post and remove the channel the message originated from and the message will disappear from that channel. Replies to recalled messages are not removed from channels. If you delete a multi-channel post, Teams removes it from all channels but leaves any replies intact.

Unlike regular posts, Team owners can't edit or delete multi-channel posts.

## Finishing Conversations

Starting a conversation is easy (but please make sure you start it with a subject), but finishing is more difficult. Structured conversations about a well-defined topic usually conclude, but how many of them are summarized? If a conversation finishes, consider posting an "end of debate" note to let everyone know what was decided and what the next steps will be. Another advantage is that people will then be able to see what happened without having to look through all the messages in the thread.

## Announcements

Many channel conversations begin with someone asking a question or sharing a piece of information they have learned. Announcements are a more formal kind of post where people highlight something important. Teams supports the announcement post type for this purpose. When you compose a new conversation, you can select the post to be an announcement. To mark announcements in the message stream, they have a large heading with a background composed of a graphic or a solid color (Figure 12-9). If you decide to use a graphic, you upload a JPEG or PNG file and then select the area of the file to fit the oblong shape used for the heading. The ideal proportions for the graphic are 914 x 120 pixels, but if your graphic is a different shape, Teams allows you to select part of the graphic. The title of the announcement can be up to 40 characters long.

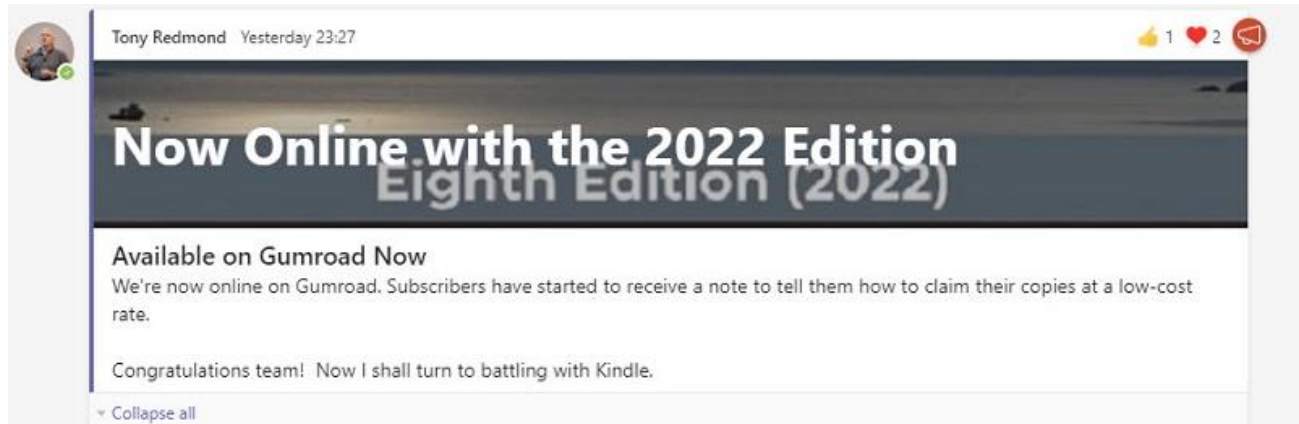


Figure 12-9: An announcement post as seen by readers

The heading attracts the attention of readers, so announcements are a good way to make people aware of important company events or news. Anyone can compose an announcement and if everyone begins to post announcements to a channel, the headings lose their uniqueness. With that in mind, it's a good idea to coach people to post announcements only when they need to highlight something important.

## Keeping Conversations Focused

Team conversations tend to be “high velocity,” meaning that contributions flow quick and fast. Typical interactions are often brief and use simple formatting, just like text messaging. Like any conversation, people are most productive and achieve the best results when everyone obeys some simple guidelines to structure communications.

A common mistake is to create new conversations instead of replying to an existing conversation. This quickly creates a confusing mass of posts that make it difficult for anyone to understand what happens in a conversation. People sometimes assume that Teams is like email and that it does not matter where they post. This is untrue. Email can preserve the context of a conversation by including the text of earlier replies, meaning that a recipient can see how a conversation unfolds by reading those replies. However, Teams conversations do not include earlier replies, so readers can only see the full context of a discussion if people reply to the conversation and do not split off into new conversations.

**Teams Conversation Etiquette:** Left to their own devices, it is easy for users to reduce a channel to a confused mass of incoherent conversations. Just like any other communications medium, it is sensible to advise people so that they understand how to extract full advantage from their contribution. Here are some simple rules to help keep Teams conversations civil and focused:

- **Always create a subject for new conversations:** Teams doesn't force users to create a subject for new conversations. However, you should always add a subject to highlight the commencement of a new thread and tell team members what you want to discuss.
- **Leverage existing threads when possible:** The temptation exists for users to start new conversations to express their ideas when existing conversations are better locations for their thoughts. Ask people to check if their contribution should be a reply to an existing thread instead of a new topic as this will help the channel avoid degenerating into many split and duplicated conversations. Using the search feature to check if a team previously discussed something takes only a few seconds. In addition, when you keep all the posts relating to a topic in a single thread, it maintains the full context of the discussion and makes it easier for people to find the information later.
- **Don't hijack conversations:** Keep conversations to a single topic and resist the temptation to introduce new topics in the middle of a conversation. It doesn't make sense to mix topics in a thread and it makes it harder to search and find information. If a topic is important enough, it deserves a dedicated conversation.

- **Mark important messages.** You can mark a message as being especially important by clicking the exclamation mark when composing the message or using the CTRL+Shift+1 key sequence. Teams then includes the text "Important!" (highlighted in red bolded and uppercased text) in the message and highlights the message with a red and white exclamation tab as a visual indicator that this is a "must-read" message.
- **Use Announcement posts for important news:** Another way to highlight important messages is to post them as announcements, which allows you to add an attractive graphic and titles to grab the attention of team members. Make sure that the graphic you choose is appropriate to the topic. It's just more professional that way.
- **Don't post private information.** Everything in a channel is available to its members. Never post anything in a channel you are unhappy to share with everyone in the membership. If you're unsure, check the membership before posting anything sensitive.
- **Use reactions instead of short replies.** If you receive a question in an email, you might respond with a one-line answer like "OK" or "Go ahead." You can do the same in Teams, but it's usually better to respond with a reaction (for instance, "Like" or "Laugh") to let the author know that you've seen and approved of its contents. Apart from anything else, this also reduces the number of messages in a thread and makes the conversation easier to read. If you don't like the content, the "Sad", "Surprised", or even "Angry" might be a good response. It's sometimes appropriate to react to one of your messages, such as when someone has liked a message to acknowledge what it's saying, and you want to react in turn to show that you've seen their response.
- **Use @mentions and tags intelligently.** Channels can be busy places where it's easy for people to overlook a question to which you want them to respond. Highlight the need for specific individuals to respond by using an @mention or a tag (if a suitable tag is available). And if you expect everyone in a team to respond, use an @Team mention, or even an @channel mention (only those who have the channel visible in their channel lists receive these notifications). Don't use these "reply-all" mentions unless you need to as you don't want to clutter up the activity feed of team members. It's also true that you cannot assume that someone has seen something in a channel conversation unless you @mention them to force the conversation into their activity feed.
- **Be Friendly.** If you @mention someone, Teams automatically inserts their full display name into the text. You can use the backspace key to remove everything but their first name, which seems a little friendlier than something like "Kim Akers (Operations)."
- **Use styles to highlight text.** The in-built styles exist to help users highlight important parts of their conversations. You can add headings, quotations, or code examples.
- **Summarize conversations:** If a conversation comes to a decision, conclude the conversion by writing an end of debate message that summarizes what the decision was and what the next steps will be.
- **Use private chats for personal conversations:** Teams allows people to share their ideas and viewpoints with anyone who can access a team, but sometimes it is best to take a conversation somewhere more private. Personal chats exist for this purpose. Use these chats when small groups need to thrash something out before making a topic public.
- **Don't just say Hello in Chats:** It's polite to let people know why you're contacting them in personal chats. Follow the [guidelines described here](#) and avoid starting a chat with Hello, Are you there, or something similar. Apart from anything else, starting with a meaningful message gives the recipient some context and allows them to respond with something useful immediately instead of waiting for you to describe what you want in other messages.

Users can use the save option in the [...] menu for a conversation to mark a complete conversation as something they wish to bookmark. The set of saved conversations (or bookmarked) is available by selecting **Saved** in the options revealed through the user photo or by typing /Saved in the command box. To remove a conversation from the saved set of conversations, select it and use the Unsave option from the [...] menu.



## Making Your Point

When you need to compose more than a simple sentence or two, click the Format icon (it looks like a capital A with an attached pen) under the compose box. This exposes a more capable editor that supports highlighting, font size and color, bulleted and numbered lists, bolding, italics, underlining, and a small set of styles. A primitive table formatter is also available to insert and populate tables. You can apply basic formatting to text, such as different headings or marking some text as a quote from a source (you can also use the markdown convention and highlight text as a quote by prefixing it with the ">" character). There is also a code style to include code examples in a code box (click the `</>` icon) to help programmers share ideas. As shown in Figure 12-10, you can name a code example and choose the language (in this case, PowerShell) from a list including C++, CSS, HTML, JSON, Markdown (GitHub), Python, Ruby, and TypeScript.

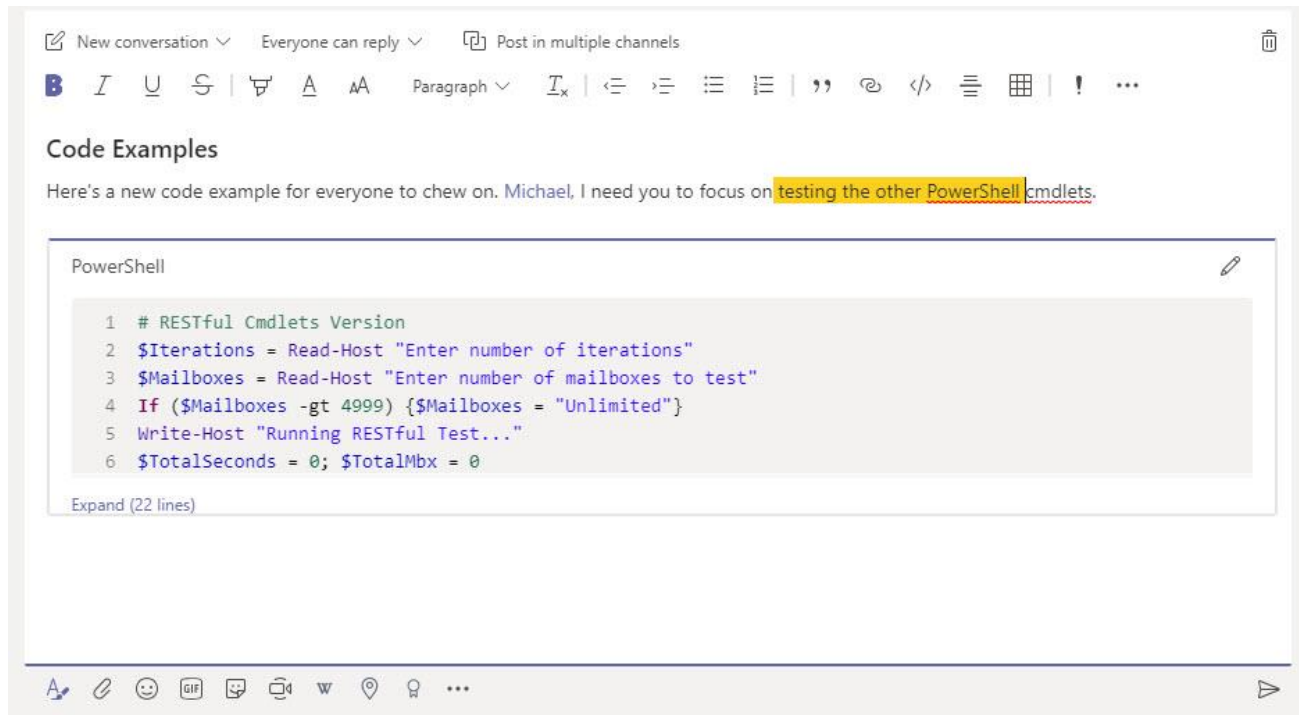


Figure 12-10: Sharing code snippets in a conversation

If you make a mistake in something you write, you can edit the content to fix the error or remove the item if it is no longer valid. Team owners can also exercise editorial control by removing messages from conversations if the need arises, but they cannot edit someone else's message.

The Teams editor does not offer as much control over formatting as a full-featured editor like Word, but if you need to compose something like a complex announcement, you can write it in your preferred editor and cut and paste the text into Teams. Text and basic formatting, including hyperlinks, survive the transition. Specific styles and embedded graphics do not. You cannot copy and paste graphics from other applications directly into a Teams conversation. If you want to include a graphic, copy it to the clipboard and then paste it into Teams. If you have a complex document that you want to share and discuss with others, upload it to the channel where you want to discuss the content. Teams stores these documents under Files, but members can review and discuss the content in the channel.

## Quoting Someone

Often you might want to highlight something another person says in a channel conversation. You can use a quoted reply in a chat, but not in channel conversations. To include someone's text in a reply, do the following:

- Select the text you want to quote and copy it (CTRL + C).

- Start a new reply and make sure that you use the full editor (click the Format button).
- Insert a line. Go back to the top of the text and type Shift and > (right arrow) together. Teams creates a highlighted area into the reply. Position the cursor in the area and paste the text you copied (CTRL + V). Teams inserts the copied text into the highlighted area.
- To insert your comment about the quoted text, move the cursor to the unhighlighted area and start typing.

## Pasting Text into a Conversation

If you compose text outside Teams and then want to include that text in a conversation, you can paste up to the maximum supported message size of 25 KB of text (roughly 2,000 words in English or approximately 12,000 characters) into a conversation. Including tables and other complex structures affects the amount of text you can paste. The maximum size of a message is approximately 28 KB, but this size includes reactions, @ mentions, and other elements. If the message is too long, Teams generates a preview of the message and tells you to shorten the text before it can be posted. You cannot drag and drop a message from one channel to another or from one team to another. If you make a mistake and post to the wrong channel, you must copy the message and repost it in the correct place.

**Spell Checking:** Teams has its own spell-checking dictionary, which it downloads to `%AppData%\Roaming\Microsoft\Teams\ictionaries` (you might find several language files there, one for each language you use Teams with). This dictionary is separate from any other used by Office applications, and you can't add new words (like technical terms) to it.

## Translating Messages

Teams uses the Microsoft Translator service for inline translation of messages in channel conversations and personal chats in the desktop and browser clients (the mobile clients can translate channel messages). In many cases, the members of a team can communicate quite happily in a common language. For larger teams, such as those used to share ideas across an entire company, being able to express a thought in a person's native language might reduce the barrier to collaboration, and that's why Teams supports inline translation for chat and channel messages in the ellipsis menu. The Translate option is available if Teams detects that a message is in a different language to the user's chosen language in Teams settings. To translate text, Teams calls Microsoft Translator to detect the language of the message and then translate it into the user's chosen language. Guest users don't have a preferred language, so the option to translate doesn't appear for them.

Languages are complex and regional accents, local sayings, idioms, jargon, argots, and technical terms often cause problems in translation. Some glitches might still happen even with simple text, but the translated output is usually good enough to convey the sense of a message. At the time of writing, Microsoft Translator supports [90 languages and dialects](#).

Translation can handle large messages, but a limit does exist. Teams supports posts of up to 25 KB. However, the buffer allowed for translations is smaller than this. In practical terms, you can expect to be able to translate messages of up to 1,000 words. The exact number depends on the language and size of the words. If you try to translate a larger post, Teams will tell you that the message is too long to translate.

The ability of a user to translate messages is controlled by the `AllowUserTranslation` setting in the Teams messaging policy assigned to their account. If the setting is `$True` (the default), translation is available. To check the translation setting across all policies, connect to Teams with PowerShell and run this command:

```
[PS] C:\> Get-CsTeamsMessagingPolicy | Format-Table Identity, AllowUserTranslation
```

Identity	AllowUserTranslation
-----	-----
Global	True
Tag:Advanced	True

**Tag:Advanced Users**  
**Tag:Restricted - No Chat**

**True**  
**False**

You can also update the allow translation setting in messaging policies through the Teams admin center, where the setting is called *Allow users to translate messages*. Remember that disabling translation in a messaging policy affects all users covered by that policy.

## Other Multilingual Features

Teams supports localized translations for the @Team and @Channel mentions. For instance, where English users type @Channel to draw attention to a topic in the channel, French users can type either @Channel or @Canal, while German users can type @Kanal. Teams recognizes the English and local language version of the term.

Normally Teams uses the default language configured on a device to know what language to use for spell check. The Teams desktop client for Windows can detect when users switch languages in chats and channel conversations. If a user sends several messages in a different language than their norm, Teams will switch spell checking to that language and prompt the user to confirm that this is OK. Unlike message translation, no data goes to the server. There's no administrative control over this feature.

## The Activity Feed

The Teams Activity Feed (Figure 12-11), is an app accessed through **Activity** in the top left-hand corner of the desktop and browser client. The role of the app is to highlight important items for the attention of a user. As people respond to chats and channel conversations with replies and reactions, notifications appear as banners in the activity feed. Notifications for messages posted by bots or connectors, including email sent to channels, do not show up in the activity feed. Apps can use the [notifications API for the activity feed](#) to post notifications too. For example, the Viva Insights app posts reminders to notify users to begin their virtual commute or note their current state of feeling.

The default view in the activity feed is *Feed*, which reveals all notifications related to the user. For example, someone uses an @mention to ask you a question or to bring something to your attention, or it's time to renew a team that you own. Along with a text snippet for each event, Teams uses different icons as visual clues of why the event appears in the feed. The clues include @mentions, channel mentions, team mentions, replies, and reactions. The number in the red circle beside the Activity Feed "bell" shows how many conversations have recent unseen activity.

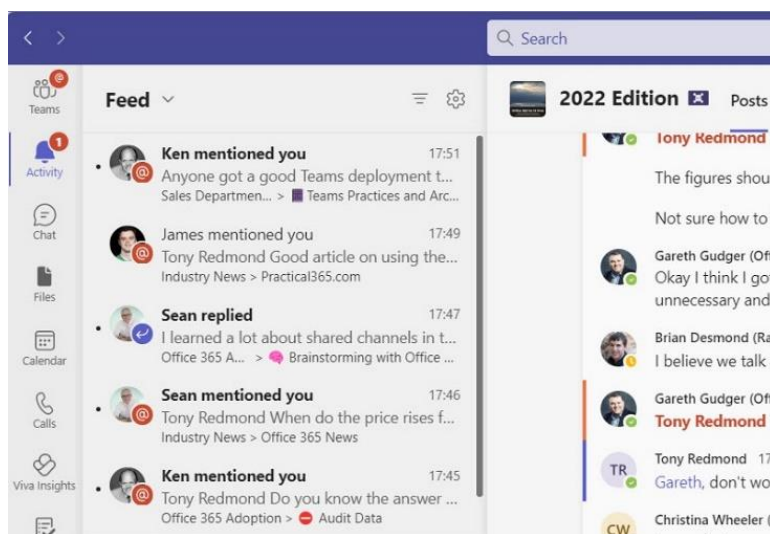


Figure 12-11: The Teams Activity feed

If you are a member of some busy teams, your activity feed is also likely busy. To help find messages in specific categories, you can filter by:

- **Unread:** Show unread notifications. You can also type /unread in the command box to see the unread items in the activity feed.
- **Mentions:** Show notifications where someone @mentions you in a conversation. Typing /mentions in the command box has the same effect.
- **Replies:** Show notifications for replies to your messages.
- **Reactions:** Show notifications for reactions to your messages.
- **Apps:** Show notifications posted by apps.
- **Missed call:** Show notifications for missed calls.
- **Voicemail:** Show notifications for voicemail messages.
- **Trending:** Show notifications for messages considered by Teams to be trending. In other words, messages with large amounts of activity (reads, responses).
- **Suggested:** Show notifications for messages that Teams believes will interest you. Teams derives both Trending and Suggested notifications from signals about user activity gathered in the Microsoft Graph with the idea being to encourage greater user involvement with these messages. An administrator can disable these notifications for the tenant in the Teams settings section of the Teams admin center.

If someone responds to one of your messages or mentions you in a conversation, Teams generates a notification in your activity feed. However, if its author or a team owner later removes a message, you might see an entry in the activity feed that doesn't point to anything. This doesn't happen very often, but it can.

## Refining the Activity Feed

After you've worked with Teams for a little while, you'll probably have a lot of items showing up in the activity feed. In email terms, an overloaded activity feed is like an overloaded inbox – it's hard to know where to start processing items, especially if you've been away from Teams for a while. One strategy often used to impose order on an activity feed is:

- Be very careful about what channels you follow. Unfollow channels that host discussions of lower importance and try to keep the channels you follow to the bare minimum. This will reduce the number of notifications you see.
- Process direct questions first. If someone uses an @mention to refer to you, it means that they want to bring something to your attention, so you should scan these entries and decide which ones need attention.
- Check replies to conversations you participate in next. You've joined these conversations, so you probably want to find out how they end.
- After you've cleared items from important channels, questions, and replies, you can leave other channel conversations until you have time.

The notifications section of the Teams settings app allows users to control notifications for personal and team @ mentions, replies, and likes and reactions. A user can turn these settings off to reduce the volume of notifications posted to the activity feed. They can also use the [...] menu for a notification for a reaction (thumbs-up, heart, etc.) or generated by an app to turn these types of notifications off. Notifications for an app are dealt with separately. In other words, if you want to disable notifications generated by ten different apps, you must find a notification from each app and use the menu option to instruct Teams to turn the notifications off for that app.

## The Importance of @Mentions

To make sure that an item shows up in a user's activity feed, you can address them with an @mention. Four options are available.

- **Direct @mention:** To bring something to the attention of a specific team member, input the @ sign followed by the name of the team member in a personal or group chat or channel conversation. For example, **@Kim Akers**. You can include a list of people, each prefixed by @. Teams checks mentions against a cached set of member and channel names (for channel conversations) and people you communicate with often (for personal chats). If you're in a channel conversation, you can't mention someone who doesn't belong to the team. If in a personal or group chat, you can't mention someone who isn't participating in the chat. Teams is intelligent enough to detect the names of team members as you type text into messages. If Teams finds a match against one or more members is detected, it displays a list of suggested names for you to select and add as a mention in the text.
- Teams also supports **@-less mentions**, which means that you don't need to prefix a name with the @ character. Instead, the client scans text as you write messages to check text against its cache. The initial letter of the member's first name (as in *Kim* for *Kim Akers*) must be capitalized otherwise Teams treats it as just another word. If it detects a match against the membership list, Teams displays a list of matching names. You can then select a name to make it into a mention.
- **@Channel mention:** Instead of the member's name, input the name of the channel. For example, @Budgets. Use a channel mention when you want to highlight something to the team members who include the channel in their channel list. You can also use @Channel (or its local language equivalent, like @Canale in Italian) and Teams will insert the channel name. The @Channel mention is particularly useful for private channels when only a subset of team members have access to content.
- **@Team mention:** Use the team name. For example, @Engineering. Use a team mention when you want to bring something to the attention of everyone in the team. You can also use @Team (or its local language equivalent) and Teams will insert the name of the team.

You can also use the tags defined in a team to refer to subsets of team members. See the section below.

Team owners can control whether members can use channel and team @mentions through team settings. By default, team settings allow users to use mentions in conversations.

**Toast Replies:** When someone sends you a message or replies to you in a personal or group chat (or meeting), Teams sends a "toast" notification to tell you what's happening. You can then respond to that user by entering your reply in the toast, but only with the Windows and Mac clients. This is a great way to send a quick response, but the reply cannot exceed 1,000 characters and can only include text.

When someone involves you in a conversation with a mention, a notification appears in your activity feed. You can click the notification to go to the conversation and pick up the thread. Figure 12-12 shows how Teams uses different icons to highlight conversation threads, including a team mention, personal mention, and an item marked as important.

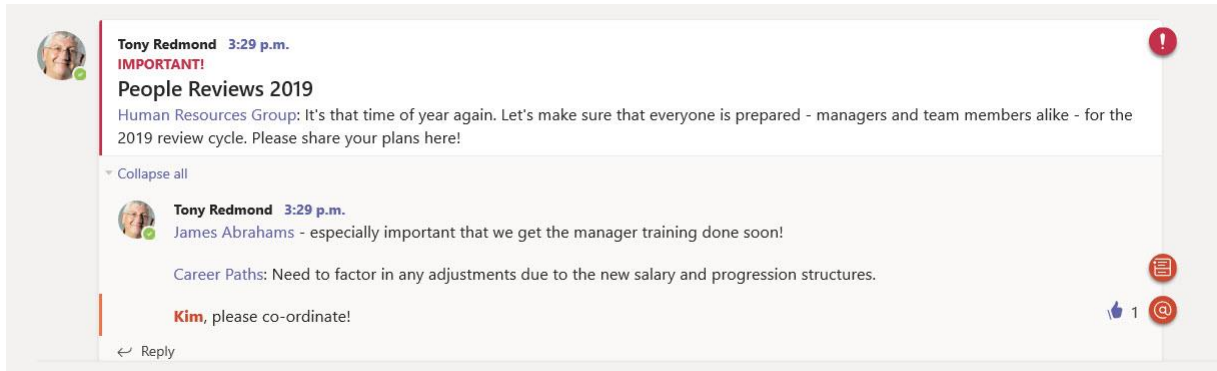


Figure 12-12: Visual indicators in conversations

Unlike most email systems, there is no equivalent of an “all users” distribution list that people can use to notify every account in the organization. If company-wide announcements are important to you, the closest equivalent is to create a team that has everyone as a member and use that team for general announcements. If you use dynamic groups, this is an easy way to keep the membership updated and ensure that new people automatically join the team when an administrator creates their account.

## Tags and Tagging

@Mentions allow users to address messages to an individual member, a channel, or the complete team. Team and channel mentions aren’t very precise because they flag a message to everyone in a team or channel. Heavy use of these mentions can clutter the activity feed of people who don’t need to know about something. Tags fill the gap between individual mentions and channel/team mentions by giving people a way to address a subset of members in a team. A tag has a descriptive name like “Managers” or “Individual Contributors” that team owners can assign to members to allow Teams to know who should receive notifications when someone uses the tag to address a message. Tags are different from hashtags because tags refer to a collection of team members while hashtags refer to a topic, place, or thing.

When you access a team, Teams builds a list of the team members and channel names in memory to allow personal and channel @ mentions. Tags are added to the roster of team members, which means that you can address the members of a tag by prefixing the tag name with an @. Because dynamic teams have non-fixed membership, you can’t use tags with these teams.

### Tag Management

Tag management is available in the Teams admin center in the Tagging section of Teams settings. There you can choose to use the following options to create tags:

- Team owners.
- Team owners and members.
- Not enabled (tagging is unavailable).

You can also allow team owners to override the organization settings (for instance, you might allow team members to create tags across the organization, but some team owners might not want this to happen in their teams). Owners control who can define tags in a team through the **Manage team** option (select Settings, then Tags).

Organizations can define a set of default tags that then become available for use in all teams. For example, you could create a default tag called *Owners* to allow members to address the team owners using @Owners. Or you could have a default tag for *Guests* to address guest members in a team. Other ideas include role-based tags like Programmers, Managers, Support, or other jobs that people commonly have, tweaked for your organization. Defining a default tag isn’t enough to make it useful; members must be assigned to tags (default or team-specific) in a team before the tags can be used to address messages within a team.

Other points to remember about tags include:

- Except for default tags defined at the organization level, tags are specific to a team and not shared between teams.
- A team can define up to 100 tags and each tag can have up to 100 members.
- The name of a tag can be up to 25 characters.
- An individual team member can have up to 100 tags.
- You cannot use tags in messages posted to multiple channels or in personal or group chats.
- Because the membership of private channels is a subset of the main team, tags cannot be used to address messages in private channels.
- Tags cannot be used to address messages in personal, group, or meeting chats.
- You can't use tags as a search filter.

By default, Teams enables Tags to allow team owners to create and assign tags using the **Manage tags** option in the [...] menu (Figure 12-13). Click on a tag to edit the people with the tag (you can also launch a chat with the tagged people from here). The tags defined by the organization are suggested at the bottom of the list. After someone is added to a default tag, it joins the set of tags for the team.

The screenshot shows the 'Tags' management page for a team. At the top, there's a search bar and a 'Create tag' button. Below that, there are two main sections: 'Tags assigned to you (4)' and 'Other tags (3)'. Each section contains a table of tags with columns for Name, Members, Description, and Imported from.

Name	Members	Description	Imported from
Admin	1		...
Editors	2		...
Managers	1		...
Writers	6		...

Name	Members	Description	Imported from
Guest Contributors	2		...
Guests	6		...
Technical Editor	2		...

Figure 12-13: A set of tags defined for a team

Another way of managing tags is through the **Manage team** option. You can view the list of team owners and members and see their assigned tags. You can then assign tags to people and see the set assigned to each tag. If a suitable tag doesn't exist, you can create a new tag and assign it to appropriate members.

Apart from the settings to control tags in the Teams admin center, there's no other management interface to deploy tags or have any insight into how team members use tags. The Microsoft Graph supports APIs to find

the set of tags in a team and return the members of a tag. See [this article](#) for an example of how to create a report using these APIs about tag usage in a tenant.

**Automatic grouping:** Teams generates an automatic grouping (tag) for team owners to allow people to address team owners in channel conversations using *@Team Owners*. You don't need to do anything to maintain this tag and it is available for use in all the regular channels of a team.

## Praising Others

The Praise app allows users to highlight the work of co-workers with a message composed of a badge and some celebratory text. People can praise others in personal chats or channel conversations or using the Viva Insights app. Both tenant and guest users can give and receive praise. The flow of the application is simple:

- Click the Praise icon below the compose box for a new message or reply.
- If using Viva Insights, choose to post the praise in a chat or channel conversation.
- Select the badge from a set provided by Teams. Each badge expresses a tribute to give to a recipient like "Thank You" or "Problem Solver."
- Select the recipients. For personal chats, the recipients must be members of the chat. For channel conversations, the recipients must be members of the team hosting the channel.
- Add some words to show why you're praising the recipients.
- Review and send the message when you're happy with the content.

Teams treats praise messages as a graphical form of an @mention, so the message shows up in the activity feed of the recipients.

## Searching Teams Content

Teams uses Microsoft Search to help users find messages, people, and files. To run a search, type something into the Search box and press return. Because this form of search looks through all the content in chats and channels available to the user, many matches likely result, divided into messages, people, and files. If you do not see what you are looking for, click the Filter icon to reduce the set of results to a more manageable number. Filtering supports preset date ranges like "Last month" or "Yesterday" and can restrict search to a specific team or channel or look for items with attachments. If you look for Files, you can restrict the search to a certain file type (for instance, PowerPoint).

When searching Files, as you enter characters into the search box, Microsoft Search scans for the most relevant hits based on multiple criteria such as the files you recently worked on or shared. These files are suggested as soon as Search finds them. Often, one of the suggested files is the file you're looking for and there's no need to progress to a full-scale search.

### Refining Search Queries

Applying filters to search results can help to find the right information, but it's even better to use precise search queries to find information, which isn't always the case. Most people mimic what they do with Google or another general-purpose search engine and search based on a word or phrase. Teams search is keyword-based and supports some (but not all) of the [KQL \(search\) syntax](#) for Exchange (email) items and used by Office 365 content searches. This means that precise queries can be constructed. For example, by adding a date range.

Let's start with a query to look for chats and channel conversations containing the phrase "Office 365 book" posted in the current month (the uppercased words in these examples are search operators):

*"Office 365 book" AND Sent = "this month"*

We could also use a specific date range:



*"Office 365 book" AND Sent >= "1 June 2020" AND Sent <= "15 June 2020"*

And limit the search further by including the name of the user who posted the message:

*"Office 365 book" AND Sent >= "1 June 2020" AND Sent <= "15 June 2020" AND From:"Tony Redmond"*

We could even refine the original search keywords ("Office 365 book") to say that two further phrases must be within ten words of each other:

*"Office 365 book" AND "book" NEAR(10) "planet" AND Sent >= "1 June 2020" AND Sent <= "15 June 2020" AND From:"Tony Redmond"*

Precise search queries work with the Teams mobile client too.

## Contextual Searches

Teams refers to searches confined to a single chat or channel as "contextual." The CTRL + F (Command + F for macOS) command tells Teams to launch a contextual search within the currently selected chat or channel. The query entered in the command box is used for the search, including precise queries as described above, so the number of items returned by a search is likely much fewer than with a more general search.

Using a contextual search effectively applies a filter, and you can't apply other filters on the results of a contextual search as are available for normal searches (which is a good reason to use precise searches). Also, CTRL + F only searches messages in channels and chats, so the **People** and **Files** results returned by normal searches aren't available.

## Search Suggestions

When a user inputs a term into the command box, Teams looks for the presence of a "/" to see if the user wants to execute a command. If none is found, Teams uses the term to generate suggestions matching the search term and displays them in a drop-down list known as the "suggestions well." The suggestions are organized into Top Hits (which include teams and documents stored in OneDrive for Business or SharePoint Online), People, Group chats, teams, and files. Selecting an item either brings the user to the location (chat or channel) or opens a file. Alternatively, the user can view full search results generated using the search term by pressing enter.

## Removing Messages

Some organizations hate the idea of anyone being able to remove anything they posted electronically (to email, SharePoint, or Teams) and some think this is sensible. After all, you can always have second thoughts and want to remove something written in error. Teams controls the ability of users and team owners to remove messages from conversations at a tenant level and within individual teams.

The **Messaging Policies** section of the Teams admin center includes the tenant-wide controls for message deletion.

- **Allow owners to delete all messages** controls whether team owners can remove messages in any channel in the teams they own.
- **Allow users to delete their messages** controls whether people (including guests) can delete messages that they post.

The intention behind owner control is to allow them to police conversations within a team and be able to remove anything posted that is inappropriate, insulting, or otherwise objectionable. Within a specific team, an owner can select **Manage Team** from the ellipsis menu, and then **Settings** to access the settings for that team. Navigate to **Member Permissions** to select whether:

- Everyone can edit their messages.
- Everyone can delete their messages.

- Owners can delete all messages.

The settings apply to all channels within the team. If the tenant-wide settings do not allow users to remove or edit messages, Teams hides these settings at the team level.

When someone removes a message, Teams flags it with *"This message has been deleted."* Contributions to the conversation before or after the deleted item are unaffected. If someone removes a message from a channel, they can restore it by clicking Undo. Teams does not remove the compliance records for deleted messages. These items persist as evidence that a conversation occurred.

Although you can remove individual messages from personal chats or channel conversations, you cannot remove a complete thread (all the messages that compose a conversation). Attachments pose another issue. When someone includes an attachment in a message, Teams uploads the file to that user's personal OneDrive for Business site (for personal chats) or the channel folder in the SharePoint document library belonging to the team. If the team owner or author later removes a message that has an attachment, Teams removes the message but leaves the attachment intact. I can understand the logic of leaving the file in place, but it might be the case that the attachment holds the objectionable content that you want to eliminate. If so, checking and cleaning up SharePoint and OneDrive content are extra steps in the removal process.

Searching for and removing content from multiple teams is another issue. This might happen if someone posts something inappropriate to one or more teams or people copy content posted to one team into others. To purge the content, you then must find and remove each post individually. Unlike the purge action for content searches, Teams does not have a method to scan all teams in a tenant to find and remove specific content.

## When Members Leave

Team conversations are persistent, so they exist even when people involved in conversations leave the organization. The contributions of removed users to private chats and channel conversations remain in place. The user no longer exists in the membership list (or the tenant directory), which means that other members can no longer @mention the user. Because their account is no longer in the tenant directory, the removed user shows up as an "unknown user" as a participant in private chats. However, Teams keeps their display name for their contributions in the conversation history.

If you restore a user account within 30 days of deletion, they regain their membership in groups and teams and can function as they did before the deletion of their account.

## Personal (1:1) and Group Chats

To this point, we have focused on using team channels as the basis for conversations. Channel conversations are under the control of the organization because channels only exist inside teams. By comparison, personal chats occur when two or more people converse on an ad-hoc or persistent basis. Chats are under the control of the participants, who decide when to start the conversation and who should take part. You cannot "promote" a personal chat to move the conversation to a channel in a team.

Like channel conversations, personal chats are persistent and are always available. The content of personal chats is always confidential to the participants. In other words, no one outside the chat can see what's happening unless they join the chat. Even then, a newly added participant can be restricted to just seeing chats after they join.

Personal chats are either one-to-one (1:1), meaning that two people take part, or group chats, meaning that between three and 250 people participate, including federated and guest participants. Both kinds of chats are accessed through the **Chat** section of the Teams client, which lists the set of active conversations for the last month. This doesn't mean that older chats are dropped. They're merely hidden and ready to be shown if you

restart a conversation. Although you cannot drag and drop chats to arrange them into a preferred order, fast access to important conversations can be achieved by using the **Pin** option in the ellipsis menu. Up to fifteen chats can be pinned at the top of the chats list and you can drag and drop the pinned chats to arrange them in whatever order is best. Within a chat, any participant can choose to pin an individual message to appear at the top of the chat. This is a useful way to highlight important information about a chat such as its purpose. On the downside, any chat participant can unpin the pinned message. Because no one in a chat has more rights over chat contents than other participants, there's no way to lock a pinned message.

Table 12-4 lists the most important differences between channel conversations and personal chats.

	<b>Channel Conversations</b>	<b>Personal/Group Chats</b>
<i>Participants</i>	Open to anyone in the team owning the channel. Shared and private channel conversations are open to channel members.	Open to those invited to join the chat, from 1:1 or group chats with 3 to 250 participants. Teams also supports chat with self.
<i>Purpose</i>	Discussion of ideas, concepts, and announcements.	More detailed discussion (often preliminary) about topics that might later be shared in channels.
<i>Structure</i>	Each conversation is a separate thread.	The conversation in a chat is a single thread.
<i>Notifications</i>	Based on @mentions and channel favorites.	Messages in a chat generate pop-up notifications for the participants.
<i>History</i>	Full history of channel conversations available to all team members.	Previous conversations can be shared with new members, but don't have to be.
<i>Sharing</i>	Files shared in the team's SharePoint site.	Files shared in sharer's OneDrive for Business account or by uploading to the sharer's account.
<i>Calendar</i>	Scheduled and "Meet now" ad-hoc meetings.	On-demand audio and video calls (with screen sharing). Can schedule meetings with chat participants.
<i>Pinning</i>	Pinned channels appear at the top of the channel list. Team members can pin individual messages which appear in the channel information pane.	Users can pin up to 15 chats to appear at the top of their chat list. Anyone in a chat can pin a message to appear at the top of the chat. Anyone can unpin the pinned message.
<i>Apps</i>	Channels can deploy apps in tabs and menus and as bots.	Chats can include apps from the Teams app store and add tabs to invoke other apps.
<i>Search</i>	Use Teams search to find messages in channel conversations.	Use Teams search to find messages in chats or use the Filter function to find specific conversations. A specific filter is available to find chats from meetings.
<i>Read Receipts</i>	Unavailable in channel conversations.	Available for 1:1 chats and group chats with up to 20 participants.
<i>Quoted Replies</i>	Unavailable unless the user cuts and pastes a reply into a new response.	Supported through the reply option in the [...] menu. Teams inserts approximately 200 characters from the chosen message into the reply.
<i>Smart Replies</i>	Not available.	Supported for 1:1 chats with other tenant members.
<i>Loop components</i>	Not supported in channel conversations.	Fully supported in chats with internal and guest users.

Table 12-4: Differences between channel and personal conversations

Tenant participants can record a 1:1 meeting if both use Teams clients and the Teams calling policy assigned to the account of the person who wishes to start the recording permits this action. Recordings don't work when calls involve PSTN lines or federated connections to Skype consumer users. After a recorded chat finishes, the recording is available in the meeting chat.

**Limitations of Large Chats:** Although Teams supports large group chats, to save system resources, some limitations apply to group chats of more than 20 participants:

- Outlook out-of-office replies and Teams status messages are not displayed for participants.
- The indicator showing that someone is typing a message is disabled.
- Video and audio calls can't be started in the chat.
- Sharing of documents (using OneDrive for Business) isn't allowed.
- Read receipts for messages don't work.

Teams disables these features to reduce the strain on the service. Fetching out-of-office information for many participants requires a lot of interaction with Exchange Online mailboxes while tracking who's read a message for the same number consumes many processing cycles.

## Starting Chats

To start a new chat, click **Chat**, right-click, and then **New Chat** (or click the New Chat icon in the menu bar). Teams invites you to start typing a name of an individual or existing group chat in the To: line. You can add the following as chat participants:

- Tenant accounts.
- Guest accounts.
- People in other tenants by entering their User Principal Names (as explained below, these people can only participate in 1:1 chats).
- The name of an existing group. In this context, "group" means the name of an existing group chat rather than a Microsoft 365 group or a distribution list. If you want to chat with a team, you do so within a channel in that team.

If you restart a new chat with a set of people with whom you have already chatted, Teams recognizes that a chat with those people exists and displays the messages from the earlier discussion to allow you to continue from that point.

You can also start a chat or continue a conversation with someone by typing @ and their name in the command box, followed by the message you want to send. This approach works well when you want to send a quick note without switching context away from something else, like a meeting or discussion in a channel. You can't address more than one person through the command box.

## Read Receipts

Read receipts are a well-known concept in email. When someone reads a message marked with a read receipt, the email service creates a notification for the original sender to tell them that the recipient has read the message. Teams uses a different concept. In personal chats and group chats involving up to 20 participants, Teams tracks who has seen a message and marks its status with a small icon to the right of the message. The icon is either a checkmark (the message is *Sent*) or an eye (the message is *Seen*). In group chats, you can see a list of who has read a message by selecting the Read by option in the [...] menu for a message. When everyone in the chat has read the message, the icon changes to Seen.

Read receipts only work if they are enabled for all participants in a chat. Read receipts can be blocked by the messaging policy assigned to user accounts. If enabled by policy, users can turn off read receipts through the Privacy settings for their account. Read receipts don't work for federated chats and they are not supported in channel conversations.

## Tabs in a Personal Chat

In the desktop and browser clients, personal chats have three default tabs:

- **Chat:** shows the messages in the conversation.
- **Files:** access to files shared with chat participants. As explained below, these files are in the OneDrive for Business accounts of the people who share them. Federated participants don't see the Files tab.
- **Activity:** shows recent messages posted by the person to team channels. Both tenant users and guests can see messages through the activity tab.
- **LinkedIn** is a tab to access LinkedIn profile information for the other user in a 1:1 chat. This tab isn't available in group chats.

The **Organization** tab is also visible to tenant users if it is enabled by the Show Organization tab in chats setting. The information exposed through the organization tab is explained later. Group chats do not display the activity tab, nor is the activity tab available for federated chats.

The idea of the activity tab is that it gives you an insight into what the person you're chatting with has been doing recently in teams where you share common membership. The lookback period is approximately two weeks. If you see something interesting in the list of messages, you can click it to go to the channel to see the full conversation. There's no risk that you'll see something you shouldn't because the messages that show up here are from teams that you both belong to.

Any tenant user in a chat can add a tab. For instance, you might decide that you want to discuss a website or a PDF file on a website with the other participants in a chat. To add a tab, click the plus sign in the menu bar to add tabs for:

- Documents that you have shared in the conversation (Word, Excel, PDF, PowerPoint).
- Link to videos stored in Stream.
- Link to a website page.
- Link to Power BI report (such as usage analytics for the tenant).
- Apps authorized by the tenant.

The intention behind adding tabs to a conversation is to give fast access to information being discussed or needed for the conversation. Participants can quickly move from the conversation to a tab and back again without pause. Teams opens objects as soon as the tab is accessed to make it even easier to access the content.

## Sharing Files in Chats with OneDrive for Business

Sharing a file to a personal or group chat is like posting to a channel. If you include a file in a conversation that's stored on your PC, Teams uploads the file to a folder called *Microsoft Teams Chat Files* in your personal OneDrive for Business account and shares the file with the other users in the conversation. On the other hand, if you share a file that's already in OneDrive for Business, the file is left where it is and its access is amended to include chat participants.

Files are shared with chat participants on a direct basis. In other words, each chat participant receives explicit access to the file being shared. If you make a mistake and share the wrong file or share a file with the wrong set of people, you must remove the access permissions from the file in OneDrive for Business as Teams does not support removing a file from a conversation. If Teams detects that the permissions on a shared file do not allow chat participants to access the file, the sharer is asked to adjust the sharing link for the file to include them.

Retrospective sharing isn't currently supported, so if someone joins the conversation after you share a file, they will be able to see that a file was shared with the chat (in the chat history or in its Files view), but you must share the file with the newly-added participant before they can access its content.

## Expanding a 1:1 Personal Chat to be a Group Chat

Group chats involve three or more people. Conversely, 1:1 or personal chats occur between two people. The two participants in a personal chat can be accounts belonging to the tenant, between a tenant account and a guest account, between two guest accounts, or between a tenant account and a Teams user in another tenant. Guest users can also create new personal chats and include other users and guests in the conversation. In fact, a guest can create a chat that only guest users join. However, guests can only communicate with the members of the teams they belong to. A guest can't chat with a tenant user who doesn't share membership of a team with the guest.

Anyone taking part in a group chat can add someone else to the conversation by clicking the **Add people** icon in the top right-hand corner, which you might do to include an expert to answer a question that has come up in the conversation.

When you start a group chat, you can give it a name (Figure 12-14) to let people know what topic you want to discuss. You can also edit a chat name after the conversation starts. It is always good to assign a name to group chats as it makes it easier to find the right chat in the list the next time you want to share a thought. If necessary, anyone taking part in a chat can rename it at any time to reflect the current scope and focus of the conversation. One-to-one chats use the name of the person with whom you chat.

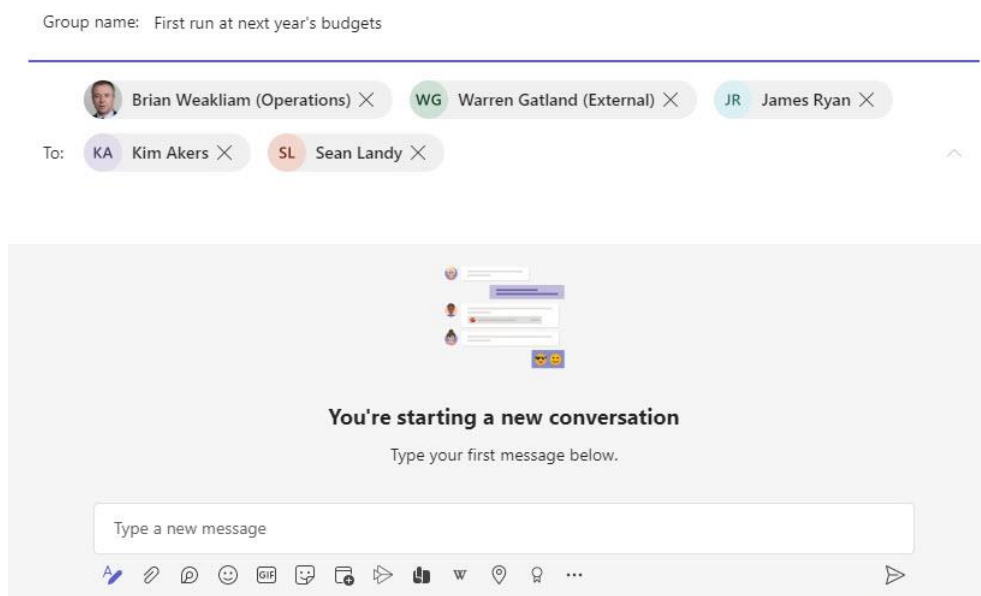


Figure 12-14: Creating a new group chat

You cannot add a guest to a chat unless their account exists in the tenant directory, meaning that they must have previously joined a team. When you add someone to a group chat, you decide whether they can see all or some of the prior message history. For instance, if you bring someone into a conversation to answer a specific question, you might allow them to see messages from the last five days so that they understand the context.

Personal chats behave differently from group chats in that you cannot reveal previous messages with someone you add to a 1:1 chat. The reason is simple. Whatever's discussed in a 1:1 chat is available only to the other person in the conversation. Chats are persistent and so if you could reveal prior messages to someone joining a conversation, you might impinge on the other person's privacy by revealing something that they prefer to keep confidential. An argument can be made that people should control their data and

decide how to share conversations, but in this case, Microsoft decided that it was best to limit the sharing of prior messages in a conversation to group chats. When a 1:1 chat evolves to become a group chat, Teams creates a brand-new chat and those participating in the chat then have full access to anything discussed thereafter.

## Leaving and Muting Chats

If you make a mistake and discover that the wrong set of people are in a group chat, any of the tenant accounts participating in the chat can remove the person who should not be in the conversation. The option to remove users from group chats is controlled by the Teams messaging policy assigned to accounts, so it is possible that some participants can remove people while others cannot. Guest accounts cannot remove other participants from group chats.

To remove someone from a group chat, click on the list of participants, find the person you want to remove, and click the **X** beside their name. Someone removed from a group chat can see the messages sent in the chat up to the point they are removed but once removed, they cannot send or receive any further messages. If you remove someone in error, you can add them back to the chat and they can pick up from where they left off. Remember to allow them to see the full chat history to ensure that they don't miss anything that happened during their removal from the chat.

Participants who do not want to be in a group can leave it at any time by selecting it from the list of personal chats and then **Leave** from the right-click menu (you can't leave a chat when only two people are in the chat). When you leave a chat, you no longer see new messages posted to the chat, but you can still see any message posted up to the time you left. Alternatively, if you simply want to remove a chat from the list of active chats, you can **Mute** the chat (another choice in the right-click menu). The chat is still there, and you can take it up again at any time, but notifications for new content in the chat do not show up in your activity feed.

If you don't want people to be able to remove other participants from group chats, change the **Allow a user to remove users from a group chat** setting to *Off* in the Teams messaging policy assigned to the user accounts you want to restrict.

## Chat with Self

Chat with Self means that a user starts a chat by entering their own name as the chat participant. It is a similar capability to the functionality found in other messaging platforms that allow people to send themselves messages as a form of aide memoire.

There have always been situations where someone ended up as the only person in a chat, such as when everyone else left a group chat either voluntarily or following the deletion of their account, but Teams did not allow users to start a chat with only themselves until June 2022. Unlike the other instances that created a single-person chat, each user can only have one chat with self and they cannot transform the chat into a regular chat by adding other participants.

If someone needs to share information in their chat with self, they can create a Loop component in the chat and make it available by sharing the link to the loop file in OneDrive for Business to allow others to open and interact with the component.

## External Access: Federated Chat Outside the Tenant

External Access or federated chat is the ability to have 1:1 and group chats with Teams users in other tenants, Skype consumer users, or [Teams consumer users](#). This is not the same as guest user access because federated chat makes a restricted set of features available to chat participants. If you make an external user a guest in your tenant, they enjoy extensive access to resources (personal chats, channel conversations, documents,

apps, and so on) in your tenant, including access to private channels. External participants can access whatever's in the chat.

The [following conditions](#) govern external access:

- External access must be enabled for the tenant (managed through Users section of the Teams admin center.)
- Your tenant does not block access to the remote user's domain. Each tenant can construct allow or blocklists to limit the connections users can make. On the other hand, if you don't use these lists, you allow inbound connections (or federation) from any other tenant.
- The remote user's tenant also allows external access.
- The remote tenant does not block your tenant. They might have blocked your tenant or limited the connections the tenant will accept to an allow list.

The [Microsoft 365 admin center includes diagnostics](#) to help resolve difficulties in Teams federation which might give you some hints about why external access to a specific address does not work.

Once everything is set up, users can enter the user principal name (or email or SIP address) for external users into the search box to have Teams search their tenant's directory to find their account. Teams first checks if federation is available between the tenants, and if so, uses the domain name in the address to connect to the other tenant to look up the user. If it finds a match, Teams adds the external user to the chat. For 1:1 or group chats involving an external user, Teams flags that you're using federated chat by including the **External** label at the top of the chat. Teams also flags external participants in the participant list of group chats.

External chat participants use rich native federated chat. In other words, they can compose and send messages like normal chats containing formatted text, links, emojis, stickers, and so on. They can also call or participate in meetings with other chat participants.

## Chats and Calls with Skype Consumer Users

External access also includes communication with Skype consumer users. This connection is natural because Teams and Skype consumer share the same media stack. The feature is disabled for tenants by default, so the first step is to enable the *Users can communicate with Skype users* setting in the **External access** section of the Users settings in the Teams admin center. Once enabled, it takes an hour or so before Teams is ready to connect to Skype.

Teams and Skype consumer users can use chats and VOIP calls to communicate. Because Skype consumer users don't belong to a verifiable organization, when a Skype consumer user reaches out to connect with a Teams user for the first time, the Teams user has the option to accept or block the connection. The Teams user is also able to view the message sent from Skype as an aid to decide whether they want to connect. Both Skype and Teams users find each other using their email addresses (you can't use a Skype ID or phone number to find a consumer user). In the case of Teams, the connection is like finding a Teams user in another tenant. You input the email address into a new chat, Teams won't find the address in the local directory, and you must instruct Teams to search externally to find the user.

Chats between Teams and Skype consumer accounts support only plain-text messages. Unlike federated chat between tenants, you can't use text formatting and emojis. The two platforms do not share presence information, so you can't find out if a Skype consumer user is free or busy.

## Urgent Messages and Priority Notifications

Urgent messages cause Teams to prompt the recipient with priority notifications every two minutes for twenty minutes (or until the message is read) to grab their attention and hopefully respond to the message. The idea is that you can use these messages to signal critical events to recipients, like the arrival of a high-priority



patient at a hospital emergency department. The ability to send urgent messages is controlled by the *AllowPriorityMessages* setting in Teams messaging policies. By default, the setting is **On**, so all users can send urgent messages. If you have people who abuse priority notifications and make all their messages urgent, you can create a new messaging policy with *AllowPriorityMessages* set to **Off** and assign that policy to their accounts.

To create an urgent message, open the Chat window, select the recipient, create a new message (Figure 12-15), and click the exclamation icon to reveal the option to mark the message as urgent. Compose and send the message as normal.

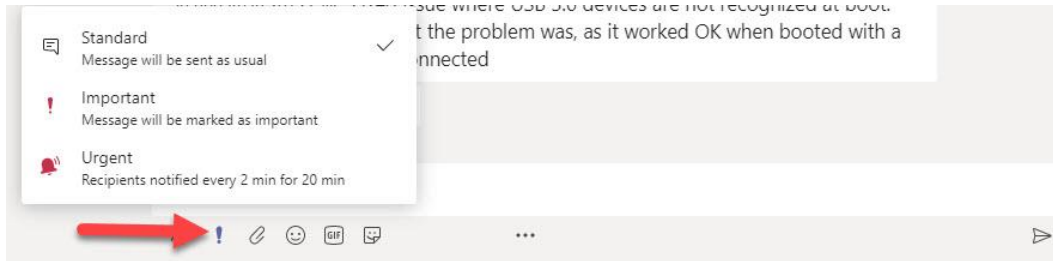


Figure 12-15: Creating an urgent message

Unlike notifications for normal Teams messages, you cannot use a priority notification to create an inline reply. Teams wants you to read the urgent message, so clicking Reply opens the chat with the sender to see the urgent message and complete the priority notification cycle.

Urgent messages are intended for direct communication with an intended recipient and are not supported for channel conversations. However, you can send an urgent message to a group chat so that everyone else in the chat receives a priority notification.

## Loop Components and Teams Chat

Loop (previously live or fluid) components are based on [Microsoft's Fluid Framework](#). Announced at the Ignite conference in November 2021, a new Microsoft 365 app called Microsoft Loop is built with loop components. In addition, loop components are available in apps, with [Teams chat](#) (not channel conversations) being the first. The Teams Windows, macOS, and Linux desktop clients and browser and mobile clients support the following loop components:

- Agenda. Build an agenda for a meeting (we'll see more use of this component when it shows up in Teams meetings).
- Table. Just like a table in Word.
- Bulleted list and Numbered list. Work like any other bulleted and numbered list in a Microsoft word processor.
- Checklist. Write down all the things people need to do. Like a task list, but with no assigned task owners and target dates.
- Paragraph. Free text component that's good for capturing ideas and sharing information like web links.
- Task list. Build a set of tasks or follow-up items and assign tasks and expected completion dates to chat participants.

The major advantage of a loop component is that once sent in a chat message, updates to the content of the component synchronize and appear for chat participants in almost real-time. It's like the way updates appear in Office documents using co-authoring with the notable difference that the updates appear much faster than in Office. According to Microsoft, using a loop component for dynamic collaboration avoids the need for back-and-forward debate in chat because those involved can work out details within the component. All participants see what's happening and it's easier to follow what develops. Figure 12-16 shows a bulleted list loop component with three active users. The component includes an embedded table, and one of the table

cells has a comment. This is a good example of how to use a loop component to have an active discussion in chat.

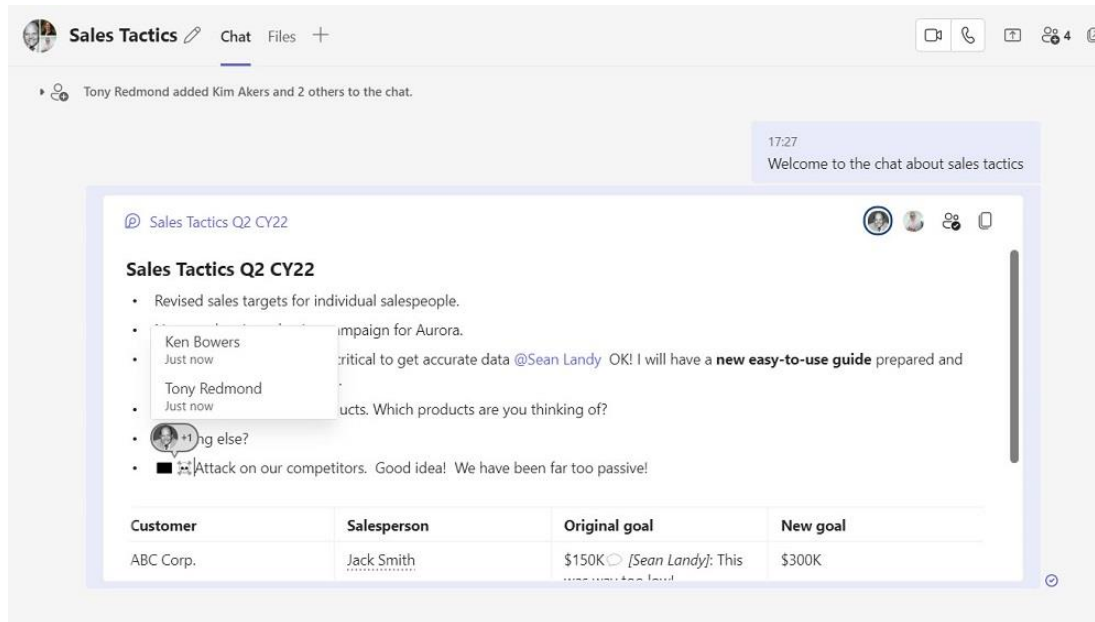


Figure 12-16: A loop component in active use in a Teams chat

Points to note about the use of loop components in Teams chat include:

- A loop component must be the only item in a message (a component can store other loop components). Teams stores the files for loop components in the *Microsoft Teams Chats Files* folder in the sender's OneDrive for Business account (so only those with OneDrive accounts can send loop components). These files have a fluid extension. To make a loop component available to chat participants, Teams shares the file. If necessary, the owner can share the file with others.
- Because of the dependency on OneDrive for Business, guest users can't create new chat messages with loop components. However, they can edit loop components sent in chat by tenant users.
- You can paste text from other applications into loop components. However, the resulting formatting might not be perfect.
- Unlike normal chat messages, translation of content in a loop component is not available.
- Share to Outlook doesn't work for chat messages with loop components.
- DLP for Teams doesn't detect policy violations in loop components.
- Loop components cannot be used in chats with external (federated) participants.
- Microsoft search processes the content of fluid files, meaning that you can include loop components in content searches. However, the copies of the files retrieved by a search can't be opened.
- Compliance records captured for chats containing loop components reference the file stored in OneDrive for Business and don't contain any of the content. As such, these records are useless to features like communications compliance policies which rely on compliance records.

Microsoft will likely resolve some or all these issues over time. We will track progress here.

## Teams Meetings

Teams meetings are a very popular part of the application. In April 2020, Microsoft said that Teams users consumed [2.7 billion meeting minutes on March 31, 2020](#), with 43% of calls involving video. Three months later, [Teams handled 5 billion meeting minutes in a single day](#). Given the growing interest in working from home, Microsoft expects Teams calling to continue to have strong growth in the future.

A Teams meeting is a virtual workspace with an identity (a unique URI) in the Teams meeting service. Although meetings have start and end times, participants can join a meeting immediately after creation, even if the scheduled time is many days away. A participant can join a meeting multiple times with different devices. A meeting has associated objects, like the participant list, including those nominated in different roles (organizer, presenters, and attendees), optional recording, and a chat thread. Collectively, these form the complete meeting. A Teams meeting can last a maximum of 24 hours. Broadly speaking, the meeting (and calling) experience available in the Teams desktop and browser (Chrome and Edge) clients support equivalent features. Other browsers, like Safari, use a simplified meeting experience.

Users create and manage Teams meetings with the Teams **Calendar app** or via an Outlook client. Control over the ability for users to create Teams meetings is set via the meeting policy assigned to accounts. Guest users can't create meetings. Four settings in the general section of the meeting policy assigned to accounts exert control over the kind of meetings users can create:

- Allow the **Outlook** add-in: If the setting is on, Outlook (Windows and Mac desktop clients, Mobile, or OWA) can schedule online Teams meetings (you can [configure settings to make Teams online meetings the default](#)) and Outlook automatically loads the [Teams meeting add-in](#) when the client starts. Modern authentication must be used for the connection between Outlook and Exchange. If these conditions don't exist, Outlook won't load the add-in. Outlook desktop and mobile clients can support several mail accounts in a profile, with one account assigned as the default (used to send outbound email). Outlook add-ins run in the context of the default account, so it's important that the account chosen as the default is the one used to create Teams meetings. Outlook can only create private meetings, so the Teams meeting policy assigned to their account must allow the user to create private meetings (any meeting not published in a channel is private).

The Teams meeting add-in uses the Edge WebView2 component, and it receives updates during the automatic update cycle for the Teams desktop client. It also depends on the .Net framework. The workstation should run the latest version (4.8 at the time of writing). On Windows PCs, you can check the version of the .Net framework with PowerShell. This command returns *True* if 4.8 or above is installed.

```
[PS] C:\> (Get-ItemProperty "HKLM:\SOFTWARE\Microsoft\NET Framework Setup\NDP\v4\Fu11").Release -ge 528040
```

Apart from creating new Teams private meetings, the add-in also allows users to update meeting settings (like who can join without going through the lobby, who can present during the meeting, etc.). The add-in includes a *Meet Now* button to allow users to create impromptu personal meetings. The user must be signed into their home tenant to allow the Meet Now option to work, so if they're connected to a different tenant, they must switch back to use Meet Now. See this post for [more information about the Teams meeting add-in for Outlook for Windows](#). A version of the add-in is available to [allow Google Calendar and Google Workspace users schedule and manage Teams meetings](#).

- Allow **channel meeting** scheduling: If on, the user can create a meeting in any standard or shared channel in teams they belong to by:
  - Selecting a channel when creating the meeting in the Teams calendar app.
  - Using the *Schedule a meeting* button in the compose reply options.
  - Using the *Meet Now* option in the channel header.
  - Creating a new meeting in the channel calendar.

Channel meetings appear as a new topic within the channel to make the meeting available to team members. Chat messages from the meeting also appear in the topic and are visible to all members, including those who do not take part in the meeting. You can add participants to a channel meeting

who don't belong to the team. These attendees can join the meeting but cannot access resources shared through the channel. One advantage of channel meetings is that meeting resources like documents and recordings are in the channel folder of the document library in the team's SharePoint site instead of an individual user's OneDrive for Business account, as is the case for private meetings. This can be important when people leave the organization along with their OneDrive for Business accounts, which might contain documents relating to important meetings. Although it's possible to give someone else access to a user's OneDrive for Business account when they leave to review and recover information, this is a step easy to overlook.

- Allow scheduling **private meetings**: If set, the user can create a private (or personal) meeting. As noted above, a private meeting is a meeting not scheduled in a channel. Like personal chats, you dictate who receives invitations to join the meeting. If a participant cannot join the video or audio part of a meeting, they can still contribute to the conversation. Invitations for private meetings come from the meeting organizer's mailbox. Teams considers meetings created with the Teams add-in for Outlook as private, even if you add a team as an attendee for a meeting. You can invite people outside your organization to a meeting, even if they do not have a guest account in the tenant, by adding their email address as a participant. This is known as an *anonymous join* and generates an invitation to the email address to allow the recipient to join the meeting.
- Allow **Meet Now** in channels: If on, the user can create an ad-hoc meeting. These meetings are like channel meetings in that they occur in a channel. To start an ad-hoc meeting, click the *Meet Now* icon under the compose message box. Team members can join the meeting if they wish. Another policy setting (Meet Now in private meetings) controls whether users can create unscheduled meetings outside channels.

**Exploiting the Virtual Lobby:** The *automatically admit people* setting in the meeting policy assigned to a meeting organizer's account dictates who can join a meeting. Depending on the setting, some participants can join direct by bypassing the virtual lobby while others must remain until the meeting organizer, a co-organizer, or a presenter admits them. Available options include the ability to admit everyone in the organization automatically (this includes guest users) or everyone in the organization and other federated Office 365 tenants. Organizers can update the settings of an individual meeting to fine-tune who waits in the lobby. Options such as "Only me" or "People I invite" allow organizers to exert more precise control over the lobby. For instance, for confidential meetings it's a good idea to update meeting options so that only the organizer or those invited explicitly can join automatically. All other participants must wait in the lobby for admittance, which removes the risk that someone else might be able to join the meeting without being noticed and gain access to confidential information.

Teams meetings support desktop sharing (including the ability to only share a specific window rather than the entire screen), recordings, muting noisy attendees, and transfer control to other participants so that they can run a meeting. Participants can use [Whiteboard as a collaborative space](#) to share and develop ideas during meetings.

If you record a meeting, a Teams bot joins the meeting to capture the audio, video, and screen sharing activity in a feed sent to Azure Media Services. When the recording is turned off (or the meeting ends), the video file is stored in Stream or OneDrive for Business. Any of the meeting participants can play the recording. The owner of the recording can share it to allow others access to the content. The Videos chapter has more information about the management of Teams meeting recordings.

Up to 1,000 participants with full access to all resources can join a Teams personal or channel meeting. This means that attendees can use video and access meeting resources such as the whiteboard, and chat. After the meeting reaches its capacity, Teams can switch automatically to admit view-only attendees. See the later section covering meeting overflows and view-only attendees.

## Sharing Files in Teams Meetings

People accustomed to Outlook often notice a difference in the way files are handled when scheduling meetings with Outlook and Teams. Outlook allows the attachment of files to meeting invitations while Teams does not. The Teams calendar app doesn't display attachments added by Outlook. If you want to share files with meeting participants, the organizer can include links to the files in the meeting invitation or participants can upload them to the Files section of the meeting. If the text isn't too long, the organizer can include it in the body of the meeting invitation.

Any tenant user can share a file before, during, or even after a meeting. If the attachment is a text file and isn't too long, you can also cut and paste the text into the body of the message invitation. Guest users who have OneDrive for Business accounts can also share files once the meeting starts (permissions might need adjustment to allow access to these files). Other participants can share links to files in the meeting chat during the meeting.

Files shared in meetings are in the *Microsoft Teams Chat Files* folder of the sharer's OneDrive for Business account. Teams uses sharing links to control access to files by meeting participants. One thing to note is that the sharing link applied to a file shared in a meeting doesn't give access to people who subsequently join the meeting. If new attendees join, file sharers must update the permissions to allow them access.

## Streaming Meetings and View-Only Attendees

Organizations can update Teams meeting policies to enable the use of Teams Live Events streaming services for overflow attendance at large meetings. When enabled, Teams will allow people to join a meeting after the meeting reaches its capacity. These people are view-only participants and have limited access to meeting content. Up to 20,000 view-only attendees can join a meeting. Microsoft will reduce this limit to 10,000 from July 1, 2022.

View-only attendees can:

- See the video feed for the current presenter or content they share from their desktop.
- Hear the audio feed for other full participants.

They cannot:

- Use the gallery, large gallery, or together mode views to see other participants.
- Interact with other participants through meeting chat, polls, or file sharing.

View-only attendees can join using any Teams client. They cannot join using a Teams Room system or Cloud Video Interop (CVI) services. View-only attendees don't appear in the participant list, which means that the meeting organizer can't remove them from the meeting. They don't appear in the attendance report. Because Teams uses streaming services to deliver the video and audio content to view-only attendees, the content is approximately 30 seconds behind the live meeting.

Control for view-only attendance for meetings is through the *StreamingAttendeeMode* setting in the Teams meeting policy assigned to the organizer. By default, this is False. To enable view-only attendees, update the policy to True. For example:

```
[PS] C:\> Set-CsTeamsMeetingPolicy -Identity "Allow Meeting Recording" -StreamingAttendeeMode Enabled
```

As always, it can take several hours before a policy change becomes effective.

## Using Teams Meeting Policies to Control Features

Teams meeting policies include settings to control aspects like meeting and call recording, transcription, screen sharing, PowerPoint sharing, and use of the whiteboard. The settings of the default meeting policy enable users to create all supported types of meetings within Teams. If necessary, you can create new meeting policies to disable access to some meeting features and assign the policies on a per-user basis. For example, you might decide that only specific users can record meetings. To do this, you would:

- Set *Allow cloud recording* to *Off* in the default Teams meeting policy. This setting blocks the ability of users to record meetings. To enable meeting recordings, assign users a meeting policy with the setting turned to *On*.
- Create a new Teams meeting policy where the *Allow cloud recording* and *Allow transcription* settings are enabled (set to *\$True*). This will allow meeting organizers assigned the policy to record meetings and generate automatic transcripts. The transcription setting is only available in PowerShell, so you will need to run a command like:

```
[PS] C:\> Set-CsTeamsMeetingPolicy -Identity Global -AllowTranscription $True -AllowCloudRecording $True
```

- Assign the new Teams meeting policy to selected accounts.

Although you can assign the new meeting policy to accounts through the Teams admin center, PowerShell is usually the best way to assign a policy if a large set of accounts is involved. This example shows how to fetch a set of mailboxes based on a filter against one of the custom attributes and then use the *Grant-CsTeamsMeetingPolicy* cmdlet to assign a custom Teams meeting policy to each account.

```
[PS] C:\> $Mbx = Get-Mailbox -RecipientTypeDetails UserMailbox -Filter {CustomAttribute1 -eq "Meetings"}
ForEach ($M in $Mbx) {
    Try {
        Grant-CsTeamsMeetingPolicy -PolicyName "Allow meeting recording" -Identity
    $M.UserPrincipalName
        Write-Host $M.DisplayName "is allowed to record Teams meetings" }
    Catch {
        Write-Host "Problem occurred when assigning the Allow meeting recording policy to"
    $M.DisplayName } }
```

By default, Teams meeting policies disable recordings for 1:1 calls. If you want to change this, use the *Set-CsTeamsCallingPolicy* to update the *AllowCloudRecordingForCalls* setting to *\$True*. More information about the settings in Teams meeting policies is discussed in the Managing Teams chapter.

**Teams Meeting Etiquette:** Like any other business activity, some preparation on the part of meeting attendees and a sense of how people should behave in meetings make online gatherings run smoother. Here are some basic points about Teams meeting etiquette:

- **Be on time:** There's nothing worse than people showing up late for meetings, even if the meetings are online.
- **Start the meeting muted and stay muted unless you need to talk:** There's no need to share background noise and other distractions that might be happening around you, like a jet passing overhead. Keep your microphone muted until you have something to say and make sure that your microphone is unmuted a second or so before you speak. Always mute if you need to cough, sneeze, burp, or otherwise do something that others might find unpleasant.
- **Use a camera if your PC has one:** People are more connected and involved in meetings when everyone is visible. If your network connection supports the video load, turn on your PC's camera when you join meetings, after positioning the camera so that your face is visible and head-on. Once the meeting starts, it's acceptable to turn your camera off when listening to a presentation or need to step out to do something. No one needs to see a video feed of an empty seat. Well-lit

rooms are better than dark rooms for video calls as the lighting helps camera performance. Use lighting above or to the side, never behind you. Above all, you'll look better when people can see you.

- **Blur your background:** No one needs to see just how messy your office is. Remove the distraction by blurring your background. Alternatively, you can choose to use a custom background image. If you do, make sure that your chosen image is professional and appropriate.
- **Understand Q&A protocol:** Some presenters like taking questions during a meeting. Others hate the idea because they believe questions are distracting and spoil the flow of their presentation. These presenters prefer to take questions in a structured manner when they're finished. It's good to set ground rules before presentations begin. When asking questions, make sure they're directed to someone as open-ended questions often result in a cacophony of competing voices.
- **Use meeting chat intelligently:** Every meeting includes a group chat. You can use the chat to capture questions to make sure that they are noted and either addressed during the meeting or followed up afterward. Making notes in the meeting chat is also a great way to make attendees aware of resources, such as posting links to supporting material.
- **Raise your hand:** Teams has a virtual raised hand feature that you can use to indicate to other participants that you want to speak. Use it. And presenters should keep an eye on the meeting roster to see if anyone raises their hand.
- **Use reactions to communicate with the presenter:** Reactions like the thumbs-up or like emoticon are good ways to inform a presenter what you think of their material and delivery. They are also excellent ways to avoid a barrage of responses in chats to common presenter queries like "Can you hear me" or "Can you see my screen."
- **Restrain multitasking:** You can multitask during a meeting, but if your video is on people might know that you're not concentrating on the meeting. Is that an impression you want to give?
- **Silence your phone:** If you're in a meeting, you should be involved. Don't disturb the flow of discussions by having your mobile or landline ring during the meeting. And if you do need to take another call, make sure to mute your microphone as it's highly unlikely that the other meeting participants want to listen to your call.
- **Ask before recording:** Teams is good at capturing the audio and video content for a meeting in a recording. Before you record a meeting, make sure that all participants are happy with the idea. And let them know that the recording will be available in OneDrive for Business soon after the meeting ends. The recording and transcript allow participants to check what happened during the meeting. Remember that only tenant users have access to these meeting resources.
- **Add color to the call:** Sharing app content to make your point is good, unless you have some boring PowerPoint to share (or boring Excel). If you're going to share content, make sure it's relevant, accurate, and restrained to just what people need to know.
- **Invest in a good headset:** There's nothing worse when someone in a meeting uses poor audio equipment that forces participants to strain to hear what they're saying. Do everyone a favor and invest in a good headset to use for Teams calls, preferably one that the manufacturer attests will work well with Teams.

A good way to remind people about the etiquette for online Teams meetings is to include some text in calendar invitation messages. As [explained in this post](#), a mail flow rule can insert a disclaimer in calendar messages. The disclaimer text can point to a website or contain some basic rules like those explained above.

## Creating Teams Meetings with Outlook and the Calendar App

If Teams is enabled for an account, the Teams meeting add-in is automatically loaded when Outlook starts. The add-in is available for Outlook desktop (Windows and Mac), OWA, and Outlook mobile to allow users to create and manage private Teams meetings. Each meeting has a link to the online location or thread which is included in meeting notifications and reminders. The link is a [Globally Routable User Agent URI \(GRUU\)](#), a URI formatted to allow Session Initiation Protocol (SIP) clients to connect to an online event. In the case of Teams,

the URI is a deeplink to the location in the Teams infrastructure where the meeting is hosted. Meetings created through the calendar app insert the same kind of URI.

You can think of the online location as the workspace where audio and video feeds come together to instantiate the meeting. It's one of the resources belonging to the meeting along with the participant list, recording, notes, whiteboard, and so on. Once created, the online space is available for any participant to join, even if the starting time for the meeting is long in the future. This facility exists to allow people to prepopulate a meeting with resources, like notes or shared files, before it begins. Likewise, a meeting persists after its formal end time to allow participants to access its resources after the meeting finishes.

To allow meeting participants to navigate to the online workspace, Outlook populates several properties of the calendar event such as *OnlineMeetingConfLink* and *SkypeTeamsMeetingURI* with joining information. Outlook and the Teams calendar app use these properties to recognize the event as an online event and to show the Join button in meeting reminders and other places in the client UI. Clicking the Join button (or the *Join Microsoft Teams Meeting* link in the body of the meeting item) starts the process of joining the meeting, which might involve navigating through a web page to choose how to join and waiting in the meeting lobby for admittance. An example deeplink for a Teams meeting looks like this:

```
https://teams.microsoft.com/l/meetup-join/19%3ameeting_ZjcxMDVhNjYtOTE4OS00ZTFmLTlkOWQtMzQ1MjZjNjc4M2Ni%40thread.v2/0?context=%7b%22id%22%3a%22b662313f-14fc-43a2-9a7a-d2e27f4f3478%22%2c%22oid%22%3a%22eff4cd58-1bb8-4899-94de-795f656b4a18%22%7d
```

All instances of a recurring meeting use the same joining information. In other words, the same online workspace hosts all instances of a recurring meeting.

Users can copy the deeplink for a meeting from the body of the invitation or the Teams calendar app (it's available in several places, including if you right-click on a meeting). People often retrieve the deeplink to send it to other people who might want to join the call. Any client can use the deeplink to join the meeting.

Apart from using the deeplink to connect to a Teams meeting, people can use a combination of meeting identifier and passcode to join. The Teams meeting invitation includes the necessary details, which can be entered into a Teams client or the or the special webpage <https://www.microsoft.com/microsoft-teams/joinmeeting>.

**Tip:** If you edit an Outlook meeting and update the Join Microsoft Teams Meeting link, you can add the string *&webjoin=true* to the end of the link to force users to join the meeting with the browser client instead of getting the option to choose between joining with the browser or desktop client.

## Private Meetings

A private (or personal) meeting is one scheduled by an individual user with other tenant accounts, external users, distribution lists, and Microsoft 365 Groups (for the groups included in the GAL). After entering the invitees, Exchange Online sends a copy of the meeting invitation to each participant, who can accept or decline the invitation. Like Outlook, the Teams calendar app has a scheduling assistant to help find the right time for a call.

When you add a distribution list to an invitation, Exchange Online expands the membership and sends notifications to the individual recipients. Distribution lists can contain recipients like other groups, mail contacts, mail users, and even public folders, so you might not know the full set of users who receive invitations to a meeting. Microsoft 365 groups don't support nested groups and are composed of mailboxes and guests.

Teams meetings appear along with other events in the personal calendars in the mailboxes of participants. The Teams calendar app synchronizes data from the personal calendar. If reminders are set for meetings,



Outlook notifies users when meetings are about to happen. Except for an icon posted alongside a channel name when a meeting is in progress, Teams does not have an external notification mechanism of its own to let people know about active meetings, but they can join a meeting from an Outlook reminder or the events in the Outlook or Teams calendars.

You can also schedule recurring meetings by selecting a value such as Daily, Monthly, or Yearly from the Repeat drop-down list.

## Meeting Recap

A notable feature available to private meetings is “meeting recap,” which is how Teams highlights important resources under the Details tab for the meeting. After a meeting is over, participants can access the meeting recording and transcript, and the organizer can access the attendance report here. Channel meetings and Meet Now meetings support recordings, but don’t generate a transcript.

## Transcription and Recording

Meeting organizers have the option to capture a recording or transcript as records of what occurred during the meeting. Meeting recording is available for all forms of meetings. Organizers, co-organizers, and presenters can initiate recording from the desktop, browser, and mobile clients. Teams stores meeting recordings as MP4 files in OneDrive for Business (personal meetings) or SharePoint Online (channel meetings).

Enterprise Teams licenses include support for live captions, CART captions, captions in other languages, and automatic transcription. This functionality includes access to captions for anonymous meeting participants. If permitted by the meeting policy assigned to the meeting organizer, participants with the organizer, co-organizer, or presenter roles can initiate transcription for private meetings using the Teams desktop client. Microsoft says that they will introduce the feature for channel and Meet Now meetings in the future. To generate live captions and create the transcript, Teams uses an AI-based learning model to listen to the audio feeds of participants and analyze what they say. Originally only available in English, live captions and transcription now support multiple languages including Arabic, Chinese, German, French, Italian, Russian, Spanish, and Dutch with more languages becoming available over time. Support means that the AI can recognize and parse the spoken contributions of people in those languages. The AI can capture but cannot understand participant contributions in other languages, so the output won’t make much sense.

Transcription includes speaker attribution, meaning that Teams includes the name of the speaker alongside their contributions. Attribution makes it easier for people to know who said what during a meeting. If users prefer, they can disable attribution in the *Captions and transcripts* section of Teams client settings. When this happens, Teams inserts generic speaker identifications (like “Speaker 1”) instead of their real names.

After the meeting finishes, Teams generates the final copy of the transcript and makes it available in the meeting recap and through the Recordings & Transcript tab for the meeting. Users can then download the transcript in Word (DOCX) or Video Text Track (VTT) format. Guest users don’t have access to the Teams calendar app, so they can’t access the transcript. Only the meeting organizer or a Teams administrator can remove a transcript.

Pointers to transcript data are in the meeting organizer’s mailbox. However, the transcript texts are not in the mailbox and are therefore not indexed and discoverable. In many cases, organizers download the transcript to a Word document and fix any errors created by the AI to produce a final record of the meeting. If this document is stored in OneDrive for Business or SharePoint Online, or attached to an Exchange Online message, it will be indexed and discoverable.

## Attendance Report

After private meetings finish, Teams generates an attendance report for the meeting organizer. The Teams back end generates the meeting data from the information gathered in the Graph as people join and leave

the call. Teams stores the data for these reports in [a hidden folder of the meeting organizer's Exchange Online mailbox](#). The meeting organizer can view the information through Teams by opening the meeting and going to the attendance tab. Alternatively, they can download the data to a CSV file (a sample attendance CSV file is shown below).

#### Meeting Summary

Total Number of Participants	8
Meeting Title	High Performer Review
Meeting Start Time	15/1/2021, 11:13:08
Meeting End Time	15/1/2021, 12:45:48

Full Name	Join Time	Leave Time	Duration	Email	Role
Sarah James	15/1/2021, 11:22:59	15/1/2021, 12:45:48	1h 22m	Sarah.James@Office365itpros.com	Attendee
James Ryan	15/1/2021, 12:02:24	15/1/2021, 12:37:15	34m 51s	James.Ryan@office365itpros.com	Attendee
Michael Conroy	15/1/2021, 11:22:59	15/1/2021, 12:45:48	1h 22m	Michael.Conroy@Office365itpros.com	Attendee
Tony Redmond	15/1/2021, 11:13:08	15/1/2021, 12:45:48	1h 32m	Tony.Redmond@office365itpros.com	Organizer

Meeting attendees cannot access the attendance report. However, the organizer can download the data and share the CSV file with attendees if they wish.

It's possible that the meeting organizer will restart a meeting several times with each instance having a different attendance. Teams generates a separate version of the attendance report each time a meeting restarts. The organizer can choose which version to view using a drop-down menu in the attendance tab. Downloading the CSV file uses the data for the latest instance of the meeting. However, if the organizer goes to the meeting chat, they can select whichever instance they wish to download that data.

## Lobby Bypass

The lobby is a virtual area where participants wait for admittance to a meeting. The lobby bypass setting for a meeting dictates who can join a meeting without waiting. The default value for lobby bypass is set by the Teams meeting policy assigned to the meeting organizer. The organizer can amend the lobby bypass setting for a meeting afterward by updating the meeting options. The following values are available:

- **Everyone:** All users, including anonymous users, can join the meeting without waiting.
- **People in my organization, trusted organizations, and guests:** Everyone with a tenant or guest account, or those with accounts in federated (trusted) organizations, can join without waiting.
- **People in my organization and guests:** Only people with accounts in the tenant directory can join without waiting.
- **People in my organization:** Only tenant users can join without waiting.
- **Invited people only:** Only those who receive an invitation to the meeting (they are on the attendee list) can join without waiting. It's a good idea to amend meeting settings to disable the ability for users to forward invitations to other people for meetings with restricted attendance. The flag prevents [email clients which can block forwarding of meeting invitations](#) from showing the forward option. However, those who receive invitations via forwarding can still join automatically, which is why it's a good idea to check the participant roster during confidential meetings.
- **Only me:** Only the meeting organizer can join automatically. Everyone else must wait for admittance.

The more sensitive a meeting, the tighter control should apply to lobby bypass.

## Channel Meetings

Channel meetings differ from personal meetings because they are scheduled by someone on behalf of the channel using the calendar app or the channel calendar app. Because it's designed to create personal

meetings, the Teams meeting add-in for Outlook doesn't support the scheduling of channel meetings. When a channel meeting is created, the team owning the channel becomes the meeting owner and the meeting is created in the group calendar of the team's group mailbox instead of the organizer's calendar. The group calendar stores meetings for all channels in a team, and the channel calendar app applies a filter to display only the events belonging to the channel into which the app is installed.

Because they act as the meeting organizer, the original creator of a channel meeting can change its details, such as adding extra attendees, altering the time, or deleting the event. Channel meetings are open to any member of the team the channel belongs to and therefore members don't need specific invitations to attend these meetings. Another way of looking at this is that people don't need to attend the meeting unless they are interested in the topic. It's there for them to attend, but there's no compulsion to be there.

If you don't add team members to the participant (invitation) list for channel meetings, they learn about the meeting by seeing it in the channel. When the meeting is active, Teams posts a new topic in the channel with a join button for members to participate. A camera icon is also shown beside the channel name in the team list. Unless they are added to the attendee list, only members receive invitations to channel meetings if they are subscribed to the group for calendar events. Users can update their subscription settings for a Microsoft 365 group by opening it in OWA and editing the **Manage group mail** settings from the [...] menu. The group is only available in OWA if it is not hidden from Exchange clients. However, most of the groups used by Teams are hidden, so some administrator intervention is often necessary to update group subscriber lists. See this [article for more information](#) about how to use PowerShell to update group subscriptions so that team members receive invitations for channel meetings.

Any member of a team can attend a channel meeting. The organizer can invite people from outside a team to join a channel meeting. These participants can attend the meeting, but they won't be able to participate in the meeting chat, or access shared documents, or the transcript because they don't have access to resources stored in the channel.



Private channels can't host channel meetings. Although they belong to a team, private channels don't have access to its group mailbox, so they can't schedule the meeting in its calendar. Another restriction to remember is that after you schedule a meeting in a channel, you can't edit the meeting to move it to another channel. If you create a meeting in the wrong channel, you must remove the original meeting and reschedule it in the correct location.

**Don't Use Channel Meetings for Sensitive Subjects:** Because any member of a team can attend a channel meeting, it's a bad idea to use channel meetings to discuss sensitive or confidential topics. Doing so creates the risk of inadvertent information disclosure. For instance, any member of the team can view the recording of a channel meeting or access files shared in the meeting. Use personal meetings instead as this gives you the ability to control who attends the meeting and who can access content shared in the meeting.

## Channel Meetings App

The channel meetings app is one of the Microsoft apps automatically available to a tenant. The app can be installed as a tab to allow team members to access meetings scheduled for the channel. Events for all channels are in the group calendar, so the app filters and displays the meetings for the channel into which it is installed. The interaction between users and meetings is like that available in the Teams calendar app. Because guests don't have permission to access the group calendar, the channel meetings tab isn't available to them. Any other member can use the app to schedule or view channel meetings. As with channel meetings scheduled through the Teams calendar app, the person who creates a meeting is its organizer and is the only one who can update meeting settings.

## Meet Now Unscheduled Meetings

The Meet Now feature is a way for users to launch impromptu or unplanned meetings without having to schedule a formal event. *Meet Now* events are launched by clicking  the Teams calendar app (to create a private meeting) or a channel (to create a channel meeting) on the  desktop, browser, and mobile clients. The Teams meeting add-in for Outlook for Windows also includes a Meet Now button to create a private meeting. As with other Teams meetings of the same type, participants must be invited to private meetings while channel meetings are available to the members of the team which owns the channel.

The ability to use Meet Now is controlled by two settings in the Teams meeting policy assigned to user accounts:

- **Allow Meet Now in channels:** Controls if the user can use Meet Now to create a new channel meeting. Applies to all channels in all teams the user belongs to.
- **Allow Meet Now in private meetings:** Controls if the user can use Meet Now to create private meetings.

When using PowerShell to manage policy settings, run the `Set-CsTeamsMeetingPolicy` cmdlet and set `AllowMeetNow` and `AllowPrivateMeetNow` to True (allowed – the default) or False (blocked). In most cases, being able to create meetings on demand is a welcome capability and organizations are happy to let the feature be used. The usual situation where Meet Now is disabled is in education scenarios where organizations often block students from organizing online events.

## Planning Meetings

Teams currently doesn't support a way for a user or administrator to create a specific type of meeting, such as starting with everyone's video turned on to display a corporate background image. Meetings are treated separately, meaning that when you create a meeting, you can configure some meeting settings (like who can present in the meeting), and when you join a meeting, you need to configure your video and audio settings for that meeting.

Some aspects of meeting functionality are controlled by the Teams meeting policy assigned to the organizer's account. For example, the screen sharing mode control in the meeting policy assigned to the organizer determines if desktop and/or window sharing is allowed in a meeting while the allow chat setting controls if participants can chat during the meeting. Other important controls are those that allow anonymous people (dial-in users) to start a meeting (without an authenticated user being present) and which participants can join a meeting automatically (without being forced to wait in the meeting lobby). This is an important control if your tenant commonly organizes meetings with external people because you might want to force organizers to be sure that a meeting is ready to begin before admitting certain types of participants. The available options to restrict the ability to join meetings automatically are described above under "Lobby bypass."

The last option is useful in a policy assigned to organizers who often create meetings where confidential or sensitive information is discussed. Remember that a meeting policy sets the default for meetings created by people; organizers can change settings for created meetings through the Teams meeting add-in for Outlook or the Teams Calendar app to adjust the lobby bypass settings if necessary. For instance, the organizer can decide that only people who receive direct invitations to the meeting (their email address is in the invitee list) can bypass the meeting lobby. Another interesting control available to the organizer is how meeting chat works. They can enable chat (the default), disable chat, or enable it only while the meeting is active.

**Guard the Meeting Link:** As described earlier, clients use the deeplink pointing to the online meeting space to join the meeting and access meeting resources. Anyone with a deeplink for a meeting can join it and might gain automatic admittance if allowed by the meeting policy assigned to the organizer or the lobby bypass settings for the meeting. Uninvited attendees will be able to access the video and audio

feeds, including any information shared during the meeting. In private meetings, they can access the meeting chat. This isn't possible in channel meetings because access is restricted to members of the team owning the channel. Because possession of a deeplink makes this access possible, it's important that the organization considers who can join a meeting automatically and organizers tune lobby bypass settings for confidential meetings to restrict access to those invited.

## Meeting Roles

When planning larger or more structured Teams meetings, it is usually important to define specific roles for some participants. For regular meetings, including webinars, the roles are:

- **Organizer:** The organizer is the person who creates the meeting. The organizer has full control over the meeting, its participants, and meeting settings.
- **Co-organizer:** Optionally, a meeting can have up to ten co-organizers chosen from the tenant users invited to the meeting. Guest accounts can't be co-organizers. Apart from managing breakout rooms, users with this role can manage the other aspects of meeting while it is active. However, they cannot edit the meeting invitation or access the attendance report.
- **Presenter:** This is also an optional role. By making a tenant user or guest user from another tenant a presenter, the meeting organizer allows them to share content like presentations during the meeting. Presenters can also start and stop the meeting recording and manage participants (mute and remove participants, admit people from the lobby).
- **Attendee:** All other invited participants are attendees. These are full participants in the meeting (they can speak and chat). However, they can't share content, present, or manage the meeting in any way.

For Teams Live Events, the roles are:

- The **Producer** is responsible for making sure the presentation works. The producer is also responsible for managing the feeds in the event and monitoring the quality and stability of the event. This role serves the same purpose as the organizer of a regular meeting.
- **Presenters** focus on content delivery.
- The **Moderator** follows questions asked by attendees to answer questions or to ask the speaker when they finish presenting. The moderator can be a person assigned the producer or presenter role.

The big difference between live events and regular meetings is that attendees are more restricted in live events as their role is essentially passive. By comparison, participants in regular team meetings can usually interact with each other unless meeting settings prohibit chat or a meeting organizer, co-organizer, or presenter mutes them.

[This article](#) describes what presenters and attendees can do during meetings. The short version is that presenters have all capabilities, but attendees can only join in the meeting chat (unless this feature is disabled), speak and share videos, and view PowerPoint presentations shared privately.

## Updating Meeting Options

The *Roles that have presenter rights* setting in the Teams meeting policy assigned to the meeting organizer sets the default presenter right for meetings created by that person. The options are:

- Everyone can present.
- Everyone in the organization (including guests) can present.
- Specific people can present.
- Only the organizer and co-organizers can present.

When the meeting is created, Teams sets the presenters for the meeting according to policy. Later, you can modify the roles assigned to meeting participants. For structured meetings, it's common to find that you need to nominate some people to be presenters before the meeting starts. To preassign the presenter role, you:

- Add the people you want to be presenters to the meeting participant list.
- After the meeting is created, you can update meeting options through the calendar app or the Teams meeting add-in.
- Select specific people to be presenters (the other option to nominate presenters are Everyone, people in the organization, and only me).

Once you limit who can present in a meeting, you also restrict people outside your organization from presenting because you can't preassign the presenter role to external participants (including guests). However, once the meeting starts, you can update the role of these participants to allow them to present.

For larger meetings, it's best to restrict the set of presenters to those you expect to speak and use the (virtual) raise hand feature to allow other attendees indicate if they wish to contribute. Presenters can see who's got their hands raised (if multiple people raise their hands, they are listed in chronological order based on when they raised their hands). If people are muted, the presenter can unmute them to allow them to speak and lower their hand in the list.

When a meeting is active, you can change a person's role. For instance, you can make an attendee a presenter and vice versa. Do this by viewing the meeting participants in and selecting *More options* for an individual participant. These controls are not available in channel meetings.

## Galleries and Together Mode

The view of participants displayed in Teams meetings is known as a gallery. Teams creates a tile for each participant to display their video feed. If the user disables their video feed, Teams uses the photo from the user's account or their initials if no photo is available. Galleries have three view options:

- The **default gallery view** shows nine tiles in a 3x3 arrangement. If more than nine participants are in a call, Teams prioritizes those with video turned on and highlights the current speaker. If more than nine video-enabled participants are in a call, Teams shows the most active speakers in the tiles. The default gallery view is available to all clients. Teams meeting recordings use the 3x3 view.
- The **large gallery view** accommodates up to 98 tiles in a 7x7 view (organized in pages) with navigation controls to move between the pages. The standard 3x3 view also supports paging. The desktop and Chromium-based browser clients support the large gallery view.
- **Together mode** is available in meetings with more than five participants. The difference between Together mode and the gallery view is that the gallery shows tiles for participants while Together mode uses artificial intelligence to isolate the head and shoulders from the video feed for each participant before combining the feeds in a background scene. According to Microsoft, active participation increases when large meetings use together mode. Meeting organizers and presenters can select a background to place people into from a set provided by Microsoft, including some suitable for large meetings (like an amphitheater) and others suitable for smaller gatherings (like a boardroom). When they enable together mode, meeting organizers and presenters can choose to *Select Together Mode for Everyone*. This option applies together mode with the selected scene for all participants. New participants join in together mode with the selected scene. Participants can switch to the regular gallery mode if they prefer.

Teams dynamically adjusts the view available to meeting participants based on factors such as the number of attendees, who enables their video and who uses audio only, active speakers during meetings, and when people share content during the meeting. Here's how dynamic view affects what attendees see:

- Tiles for attendees with video feed enabled appear differently to those who only use an audio feed, Audio participants appear in smaller tiles (a reasonable call because a set of initials or a static photo in a circle isn't very visually compelling).

- Content shared in a meeting, like a presentation, app, or whiteboard, is given more space. Depending on how many people are in a meeting, dynamic view resizes attendee tiles to make them more visible.
- Users can pin or spotlight selected attendees to make their cards larger than other participants.
- The together mode view can appear alongside content.
- Users can “dock” the gallery of attendee cards on the top of their screen.

Dynamic view aims to make meetings more visually interesting than the flat gallery views used in the past. The hope is that user attention and engagement will be higher because Teams adjusts the view to concentrate on the most important content in the meeting.

**Customizing Together Mode:** Tenants can [create and package new background scenes for Together Mode in the developer portal](#) (here's a [description of how to build a new scene](#)). This capability requires a Teams advanced communications license.

## Background Effects: Images and Blurring

No one likes distracting other meeting participants with a video feed of a messy office. To help people disguise messy surroundings, they can use a background filter in personal and channel meetings. Background filters work by isolating the image of a user from their video feed and replacing the background with either:

- **Background blur:** The user's surroundings are blurred.
- **Custom background image:** The user's image is merged with a background image. The image can be selected from a standard set provided by Microsoft, or a custom image obtained from another source and uploaded to a workstation or mobile device. In effect, the mechanism works like a video green screen.

If their camera is on, users can select a background filter through the pre-join screen or by using the **Apply background effects** option during a meeting. CTRL + Shift + P is the Windows keyboard combination to invoke this option.

To apply the chosen filter, Teams uses artificial intelligence to work out what part of the image is the user and what forms the background, and then applies the filter to the background. On Windows PCs, the ability to apply background filters depends on the workstation having a post-Haswell chipset with AVX (Advanced Vector Extensions); all recent PCs should have the necessary hardware.

### Background Filters Policy Setting

The ability for users to choose background filters is controlled by the *VideoFiltersMode* setting in the Teams meeting policy assigned to an account. The available values for the *VideoFiltersMode* setting are:

- **NoFilters:** No filters are available.
- **BlurOnly:** Background blur is available (but only if certain hardware conditions are met).
- **BlurAndDefaultBackgrounds:** Background blur and the set of curated background images selected by Microsoft can be used.
- **AllFilters:** All filters are available, and the user can upload and remove custom images. This is the default value for meeting policies.

You can change the background filter setting by updating a Teams meeting policy in the Teams admin center. Alternatively, you can update the *VideoFiltersMode* setting for a meeting policy by running the PowerShell *Set-CsTeamsMeetingPolicy* cmdlet. For example, this command removes the ability to use background filters from any user assigned the *RestrictedFunctionality* meeting policy.

```
[PS] C:\> Set-CsTeamsMeetingPolicy -Identity RestrictedFunctionality -VideoFiltersMode NoFilters
```

Two standard filters designed to improve the appearance of users in video meetings are unaffected by the policy setting. The brightness filter lightens a participant's image in the video feed to make it clearer when the

lighting in the area where they sit is dark. The effect is much like using a ring light. The soft-focus filter smoothens facial wrinkles and lines to make them less distinct. The filter works better for some people than it does for others. These filters are supported on the Teams desktop client for Windows and can be enabled before joining a meeting or while the meeting is in progress.

## Standard and Custom Background Images

The basic Teams license allows the use of two types of background images:

- **Standard images.** A set of curated background images chosen by Microsoft. These images are stored on a content delivery network (CDN) to make them available to all users. If you choose one of the standard images, Teams downloads a copy of the image to the device. On a PC, downloaded images are stored in the %AppData%\Microsoft\Teams\Backgrounds folder. Sometimes users can't reach the CDN to access the standard images. Usually, this is because of a VPN or other network restriction. To test, try and access one of the standard Teams background images, like [the contemporary office scene](#).
- **Custom images.** You can upload your JPEG or PNG images using the Teams client or by copying the files to the following folders.
  - PC: %AppData%\Microsoft\Teams\Backgrounds\Uploads.
  - Mac: /users/<username>/Library/Application Support/Microsoft/Teams/Backgrounds/Uploads (you may need to hold down the Option key before you choose Go from the Finder Menu to get the Library to appear).
  - Mobile clients: The iOS and Android clients allow users to upload their images to local storage.

When you browse the set of available images, Teams shows the standard images first followed by custom images (blur is included as a standard image). The images are ordered alphabetically within their respective folders. Users can remove images they upload by hovering over the image and selecting the Remove option. Standard images cannot be removed. After a user selects a background filter (blurring or image), Teams displays that effect for every meeting they join with video-enabled until the user selects a new filter.

Microsoft publishes suitable images for use with Teams in a [public custom backgrounds gallery](#). Other companies create their custom background images for use with Teams and other conferencing software, and some release the images for public use (images released for Zoom can also be used by Teams). Two examples are the images [made available by IKEA](#) (a major Teams customer) and [Star Wars](#). A collection of free-to-use background images intended for use as [PC wallpapers is available online](#). You can also download and use the images used on the Bing home page using [this PowerShell script](#).

Microsoft's guidelines for custom background images are that images should be a minimum of 360 x 360 pixels and a maximum of 2048 x 2048 pixels. Images at the minimum size will look horrible. Best results are gained when you use images with a 16:9 ratio sized at 1920 x 1080 pixels (the default size used by Microsoft and other sites) with the resolution scaled to produce files of around 1 MB. Files can be larger in both pixels and size as the uploading process will downscale them to fit. Best results are always obtained when you control the sizing process and generate images yourself. When displayed in a video feed, the operating system will adjust the image to fit the screen display. Depending on the shape of the screen, some of the images might not be visible at the edge. To avoid problems, make sure that the important part of an image is in the center.

If you use Teams on several workstations and want to use the same custom background images everywhere, you must load the images onto each workstation. Although they can control the level of access users have to background filters, tenants cannot mandate the use of a standard image.



## Organization Images

Organization or corporate images are a feature of the Teams Advanced Communications add-on. Tenant administrators can upload images in the meeting policies section of the Teams admin center. The images then become available to users with the Teams Advanced Communications add-on and are shown before the Microsoft-curated set of standard images. If a user selects an image, Teams downloads it from the cloud to the same folder used to store custom images and uses it as the meeting background. Teams desktop and mobile clients support organization images. See [this article](#) for more information.

## Noise Suppression

Device settings in user profiles have settings to control how Teams suppresses background noise during a meeting. Available for the Teams Windows and macOS desktop clients and the iOS mobile client, noise suppression uses the CPU to process background noise around the meeting participant using the workstation to eliminate sounds like paper rustling or fans. The higher the level of suppression, the more computer resources are used. You can choose from:

- **Auto:** Teams manages background suppression and tune it up or down depending on the level of ambient noise.
- **Low:** Use this level when music or a low consistent level of noise is present in the background.
- **High:** The maximum level of noise suppression is used. This might be chosen by someone working in a very noisy environment.
- **Off:** No background noise suppression is used.

The setting in a user profile becomes the default for all meetings. If needed, you can select a different setting during a meeting from the Devices option in the meeting menu.

Noise suppression is done on the audio feed from the workstation to the meeting. It does nothing to improve the sound received by headsets or phones. To ensure the best quality sound in meetings, use a Teams-certified device. Teams does not apply noise suppression when a meeting is recorded (suppression is applied to the recording as a whole) or when live captions are used.

## Maintaining Focus During Meetings

A busy Teams user can receive multiple notifications during a meeting to inform them of @mentions, replies posted in chats, and so on. These notifications can become disruptive and distracting during meetings. Two controls are available to help. First, a user can opt to mute notifications during all meetings by selecting this option in the Meetings and Calls section under Notifications in Teams client settings. Second, if they prefer to control notifications on a meeting-by-meeting basis, they can use the *Mute notifications* option (in the [...] menu) during a meeting. This option disables notifications for the current meeting. An *Allow notifications* option restores notifications for the current meeting.

Teams displays the video feeds from individual users on cards in a gallery during meetings. User feedback revealed that the presence of a user's video feed in the set of cards shown in the gallery can distract and become a source of anxiety for some users. To resolve the issue, meeting participants can select their card and then use the Hide for me option from the [...] menu to transform their card into a chevron. Everyone else in the meeting sees the user's video feed as normal. By clicking the chevron, the user can restore their video feed and display it in their gallery.

## Music Mode

Teams music mode leverages [the Satin codec](#) to transmit high-fidelity sound in meetings and calls. When music mode is used, Teams adjusts the sampling rate automatically up to 32 kHz at 128 kbps using the available bandwidth to deliver the best possible sound quality (the transmission of acceptable human voices

requires a much lower bitrate). The lowest bitrate for good quality sound is 48 kbps. Because sound quality is linked to available bandwidth, Microsoft recommends that you use music mode only when connected to wired networks.

In environments like a studio with low background noise and microphone control, users (including guests joining calls in other tenants) can control the audio stream further by turning off noise cancellation and disabling echo cancellation if using closed-back headphones. If using professional microphones with external gain adjustment, you can disable the auto-adjust microphone sensitivity setting. Microsoft recommends that you don't use Bluetooth headsets with music mode.

The important thing to realize about using music mode is that you must enable the feature in Teams settings before you join a call or meeting. You can also decide to enable echo cancellation, noise cancellation, or auto-adjust microphone sensitivity. If you forget to enable music mode before starting a call, you'll have to leave the call, enable music mode, and restart. With music mode enabled, you'll see a music note symbol in the control bar to toggle music mode on and off. You can leave music mode enabled for the entire meeting or turn it off once the music finishes reducing the demand for bandwidth and codec processing. In addition, music mode doesn't suppress background noise as well as regular mode does.

## Using Whiteboard in Teams Meetings

Whiteboard is an application intended to allow people to collaborate by drawing and refining ideas on a digital canvas (a board), or as Microsoft says an *"infinite canvas where imagination has room to grow."*

Whiteboard runs as an Azure service and is enabled by default for all Office 365 enterprise tenants. If you want, you can disable Whiteboard through tenant settings in the Microsoft 365 admin center.

Users can access Whiteboard through the Office menu, which launches the [browser version](#), signing in with their Azure AD credentials to create, update, or remove boards. Versions of Whiteboard are also available in the Microsoft Store and [for iOS](#) and [Android](#).

Whiteboard in Teams can be installed as [an app in a channel tab](#). It is also available in the share tray for Teams meetings. When invoked in a channel tab or meeting, you can choose to open Whiteboard in Teams or use the app (if installed on a Windows PC). If your PC supports digital inking, you can draw with your finger or a digital pen on a board, which works surprisingly well. Whiteboard includes support for note grids, an wide range of colors for sticky notes, ink, and highlighting, along with templates designed to get ideas flowing. The same capabilities are available in the Teams, Windows, or browser clients.

When you create a board for a meeting, it is associated with the meeting and all meeting participants can access the board during the meeting and afterward (the whiteboard is listed as a meeting resource). Because a board is a common resource that everyone connects to, changes made to the board are seen everywhere in real-time. Guest members can interact with boards during a Teams meeting, but full guest access to the Whiteboard service isn't currently available.

Often a discussion centered around a board will conclude that you want to share with other people who weren't in the discussion. A Post to Teams option is available in the Windows app to post a link to the board as a message in a selected channel. The link can be copied from Teams and used to share the board with other people via email or in a personal chat. Clicking the link opens the board in the app (if installed) or the browser. You can only share boards with other tenant users. It's also possible to invite other users to collaborate in Whiteboard through sharing invitations sent via email. The link can be read-only or allow full write access. When the recipient accepts the invitation, Whiteboard adds the board to their list of available boards and opens the board to allow them to contribute immediately. Invited users aren't allowed to delete boards.

If you don't want to send a link to allow access to a board, you can save a static view of a board as a graphic file (PNG format) and include the file in an email, document, or web page.

All boards created in Teams meetings are stored in OneDrive for Business and have the same “Whiteboard meeting” name. To avoid confusion, you can use the app to select a board in the list of boards and edit it to assign a more meaningful name. For now, Whiteboard content is not captured by Teams compliance records, nor is the content indexed and available to content searches and eDiscovery cases.

Teams isn't limited to Whiteboard when it comes to discussing ideas on a digital canvas as [other brainstorming applications](#) are available in the Teams app store.

## Meeting Reactions

Teams meetings allow attendees to react to whatever's happening in the meeting with a set of four emoticons. Like how the virtual hands-up feature works, the attendee selects the emoticon they wish to use (like a laugh or thumbs-up) and sends it. The reaction is visible to others on their attendee card or, if someone is presenting or sharing information at the time, reactions float up from the bottom of the screen.

Some organizations don't like meeting reactions. You can disable the ability of users to send reactions through the *AllowMeetingReaction* setting in the Teams meeting policy assigned to meeting organizers. In this example, we disable reactions for meetings organized by anyone assigned the default meeting policy.

```
Set-CsTeamsMeetingPolicy -Identity "Global" -AllowMeetingReactions $False
```

Note that meeting organizers can override the meeting policy by editing the settings of a meeting to allow reactions.

## Meeting Polls

Meeting organizers and presenters of personal meetings can create simple single-question polls with up to six answers and make them available to attendees before and during meetings. Polls aren't supported for channel meetings or live events. The Microsoft Forms app for Teams is used to create and manage polls and the app must be added to a meeting before polls are available.

Polls can be published to attendees before and during meetings. When published during a meeting, users of Teams desktop and browser clients see the poll questions in a pop-up notification in the middle of the meeting window and can respond there. Teams publishes the poll as a card within the meeting chat, where it can be accessed by Teams mobile clients. Polls can be opened and closed as needed during a meeting and the results are displayed to users as responses come in unless the poll is marked as anonymous. Once complete, poll results can be exported and downloaded in an Excel worksheet.

Meeting polls are stored as a personal form in the user's account and can be accessed to view results through the Forms app (the forms used by Teams polls are read-only in the Forms app). See Chapter 9 of the companion volume to learn more about Forms.

## Teams Webinars

A Teams webinar is a special form of meeting created when a user selects *webinar* from the drop-down *New meeting* menu or if they choose a value for the *Require registration* setting for a meeting. Along with the normal meeting functionality, a webinar has a sign-up page created by the organizer to allow potential internal or external participants to register for the event. The organizer can customize the sign-up page to collect information about attendees such as their name, company, and custom fields that can be used to collect data relevant to the webinar topic. Presenters and other interested parties are invited to the meeting. Those who register through the sign-up page attend the meeting at the regular time using event details (including an .ics file) emailed to them by Teams. Teams keeps registration and attendance reports to allow webinar organizers to compare those who signed up with the eventual attendance.

Settings in the Teams meeting policy assigned to user accounts control who can schedule webinar meetings and if webinars are internal-only or accessible by both internal and external attendees. The settings are configurable by PowerShell and are:

- **AllowMeetingRegistration:** Controls if a user can create a webinar meeting. The default is *True*.
- **WhoCanRegister:** Controls the attendees who can attend a webinar meeting. The default is *EveryoneInCompany*, meaning that internal accounts and guest accounts can attend. If you want to organize public webinars, set the value to *Everyone*. If anonymous join is disabled in the meeting settings for the tenant, anonymous users will be unable to join webinars even if the *WhoCanRegister* setting is *Everyone*.
- **AllowEngagementReport:** Controls if the user can download the [meeting's attendance report](#) and the registration report. Make sure this value is *Enabled* as a big part of running a webinar is knowing about audience acquisition and participation.
- **StreamingAttendeeMode:** Controls if Teams uses overflow capability once a meeting reaches its capacity (1,000 users with full functionality). Set this to *Enabled* to allow up to 20,000 extra view-only attendees to join.

For example, this command enables the necessary settings in a Teams meeting policy:

```
[PS] C:\> Set-CsTeamsMeetingPolicy -Identity "Allow Meeting Recording" -AllowMeetingRegistration $True -WhoCanRegister Everyone -AllowEngagementReport Enabled -StreamingAttendeeMode Enabled
```

Teams generates attendance reports for meetings when the *AllowEngagementReport* setting of the meeting policy assigned to the organizer is *True*. Webinars also have a registration report, controlled by the same settings. The attendance data can also be downloaded as a CSV file, which is especially useful if a meeting organizer wants to do some analysis of meeting attendance.

## Breakout Rooms

Teams breakout rooms allow a Teams meeting to be split into several subordinate meetings (the breakout rooms) linked to the main meeting. The feature is designed to support scenarios like brainstorming sessions, online classes, and corporate events which often start by assembling all the participants to set the goals before dividing into smaller groups to work on specific issues, and then come back together to report findings and make decisions. The basic flow for breakout rooms is:

- A meeting starts as normal. The meeting organizer chooses the breakout rooms option in the meeting control bar to create the number of breakout rooms needed for sub-groups. Additional breakout rooms can be added or removed later, up to the maximum of 50 rooms. To make their purpose clear, breakout rooms can be renamed. For example, a group working on a corporate merger might have breakout rooms for Finance, HR, and Legal. Microsoft says that in the future you'll be able to predefine breakout rooms including room assignments before a meeting starts.
- The meeting organizer assigns meeting participants to the different rooms. This can be done manually or automatically (participants are evenly divided at random among the available rooms). Users can be moved between breakout rooms.
- After assigning participants to breakout rooms, the organizer uses the Start rooms command to allow the participants assigned to each room to begin work. It's possible to open rooms individually if you don't want them all to begin at the same time. A setting controls whether people are moved automatically into their assigned rooms (the default) or receive a prompt to join. Those assigned to a breakout room cannot add other people – this can only be done by the meeting organizer.
- The organizer can set a countdown timer for the rooms. When the timer expires, participants rejoin the main meeting or leave the meeting.

- Participants meet in the breakout rooms and use normal meeting functionality such as chat, app sharing, turn on together mode, and collaborate with a whiteboard. To encourage people to participate, everyone in a breakout room is assigned the presenter role.
- The meeting organizer can visit the breakout rooms to help keep everything on track. When they join a breakout room, the organizer can work with the other participants.
- The meeting organizer can also make announcements to all breakout rooms. For instance, they might send a note to remind people that the breakout rooms will close in five minutes and that someone should be nominated to present findings. Announcements are posted to the meeting chat in each open breakout room. If participants in a room need to contact the meeting organizer, they can send an @mention message in the chat.
- To bring the meeting back together again, the organizer closes the breakout rooms (or the countdown timer expires). After a short delay, the participants from the breakout rooms rejoin the main meeting. If necessary, the organizer can reopen a breakout room to allow people to restart discussions. Attendees cannot close breakout rooms.
- After wrapping everything up with the complete set of participants, the organizer ends the meeting.
- Separate meeting chats and notes are kept for each room and the main meeting. Separate recordings and transcripts can be captured for each breakout. Access to the information shared or generated in a breakout room is limited to the participants in that room. For instance, if a file is shared in the Finance breakout room, the permissions on the file uploaded to the sharer's OneDrive account are restricted to the people in the breakout room at that time. In the future, Microsoft says that it will be possible to share information more easily from a breakout room with the main meeting.

Managing breakout rooms depends on the pop-out meeting and chat experience, so the desktop client must be used by meeting organizers. Participants can use the desktop, browser, or mobile client. The ability to use breakout rooms in meetings depends on several settings in the Teams meeting policy assigned to the meeting organizer. See [this article](#) for more information.

It's not always the case that those who schedule meetings are the people who run the meetings, and it's also possible that a meeting creator might not be available when the meeting happens. To avoid the obvious issue that the meeting organizer is the only person initially allowed to manage breakout rooms, Microsoft says that it will be possible to assign multiple organizers in the future.

## Real-Time Streaming from Teams Meetings

Teams personal (not channel) meetings support real-time streaming of content to streaming platforms like YouTube via the Real-Time Media Protocol (RTMP). The essential steps are:

- Create a suitable video feed within a meeting, using background images and other effects (like the brightening filter).
- Load the Custom Streaming app into the meeting.
- Get an RTMP URI and an RTMP key from the target streaming platform to identify the content coming from Teams.
- Connect Teams to the target streaming platform.
- Deliver the content and terminate it from Teams or the streaming platform when it's complete.
- Perform any post-production you want on the video before making it available long-term on the streaming platform.

The process is quick and easy and can be accomplished using regular PC hardware. Naturally, you can increase the quality of the video output by adding a better camera, lights, and microphones. It's a great way of making events like webinars and product announcements available from Teams to a public audience. See [this article](#) for more information.

## Teams Live Events

A live event is a structured form of meeting intended for large-scale information dissemination such as company announcements, product launches, training, and so on. The ability to present is limited to those assigned the presenter role; attendees are limited to listening and interacting through moderated Q&A. Organizers can create events with live captions and subtitles, including [the ability to translate spoken words into captions shown in multiple languages](#). Live events support recordings. Those invited to events can play back recordings for up to 180 days after the event finishes. Recordings are not stored in Stream, but they can be [downloaded by event producers](#) and then uploaded to Stream.

A tenant can run up to fifteen concurrent live events, each of which can last up to four hours (sixteen hours until December 31, 2022). Stream and Yammer can also host Live events.

**Large Events:** Normally, Live Events support up to 20,000 attendees. However, with the support of the [Microsoft 365 Live Events Assistance](#) team, an event can extend to support up to 100,000 attendees.

Two types of Teams Live Events are available: quick start and external encoder. Anyone equipped with a PC and webcam can create and run a quick start live event with audio, video, screen sharing, application sharing, and QA. These events, also called “produced with Teams,” are limited in terms of the video quality and the type of information presented during the broadcast (for example, no overlaid graphics), but the output is more than good enough for many topics.

Teams Live Events are created in the Teams calendar app in a similar way to a personal or channel meeting. Instead of *Schedule meeting* (the default), select *Live event* from the **New meeting** drop-down menu. The person who creates a live event is called the organizer or producer. Event organizers must meet the following criteria:

- Their account is assigned an enterprise license (E3 or E5 or academic/government equivalent).
- The assigned license must include Teams and Stream. The need for a Teams license is obvious. Behind the scenes, Teams uses Stream for all the necessary video processing including storage, caption and transcript generation, and delivery of content in formats suitable for different devices.
- Their mailbox must be hosted in Exchange Online.

Setting up a Teams live event is straightforward. You can limit event attendance to:

- Specific people and groups.
- Organization-wide (anyone who can sign into your tenant can attend, including guest accounts).
- Public. Anyone can attend and no sign-in is required. These events are designed for public briefings and presentations.

When creating a live event, the organizer decides to produce the event using Teams, meaning that the content for the event comes from presenter workstations, or using an external app or device. The former option is sufficient for internal meetings while the latter uses an external encoder (see below) to create more of a TV studio-like event, complete with special effects. The event recording is automatically available to the organizer and presenters. It is then up to the organizer to make the recording available to attendees via a website or another repository. You can choose to add captions in the spoken language translated in up to six other languages. Other options control the generation of an attendee report and if attendees can ask questions during the event. Q&A is not like the interaction chat used in other Teams meetings; attendee questions go through a moderation process before they become public (visible to event attendees).

### Notifying Event Participants

Presenters automatically receive calendar invitations generated when the event is created. These links contain the special links to enable the presenter role during the event. Sometimes event organizers forget to send out notifications to their intended audience to tell them how to join the event. To do this, copy the deeplink for

the event and paste the link into an email (use a standard non-Teams meeting invitation to have the event inserted into recipient calendars). The deeplink looks something like this:

```
https://teams.microsoft.com/l/meetup-join/19%3ameeting_Mzc4M2E0MzgtM2ZINS00M2UwLTlkMmQtYTBhYjk5NjMzMzMzNmMy%40thread.v2/0?context=%7b%22Tid%22%3a%22b662313f-14fc-43a2-9a7a-d2e27f4f3478%22%2c%22Oid%22%3a%22eff4cd58-1bb8-4899-94de-795f656b4a18%22%2c%22IsBroadcastMeeting%22%3atrue%7d
```

The reason for sending the deeplink generated for the live event via email or a calendar invitation is simple: distributing the deeplink like this avoids the possibility of including two sets of join information in the invitation (one for the live event, the other for a regular Teams meeting).

## Creating Pass-through Deeplinks

When a deeplink is used by an external participant to join a live event, they see a screen to allow them to sign in or join anonymously. If you organize many public events, you can [edit the deeplink to bring participants directly into the event](#). The edit must be accurate to generate a working link and involves scanning and replacing the following strings:

- "%3a" with ":" (colon – four times).
- "%40" with "@" (at sign – once).
- "%2c" with "," (comma – twice).

In addition, add the string `"setting=enableDirectPublicBroadcastJoin:true"` to the end of the link. The result is a link that works well on mobile devices. If you want it to work with desktop and browser clients, you must also replace "l" in the original link with "/\_#/l/". Taking the deeplink shown above as an example, when the edits are made, the output for desktop and browser clients is:

```
https://teams.microsoft.com/_#/l/meetup-join/19:meeting_Mzc4M2E0MzgtM2ZINS00M2UwLTlkMmQtYTBhYjk5NjMzMzMzNmMy@thread.v2/0?context=%7B%22Tid%22:%22b662313f-14fc-43a2-9a7a-d2e27f4f3478%22,%22Oid%22:%22eff4cd58-1bb8-4899-94de-795f656b4a18%22,%22IsBroadcastMeeting%22:true%7d&anon=true&anon=true&setting=enableDirectPublicBroadcastJoin:true
```

## Live Events Policy Settings

Teams supports anonymous access to live events, but only if this is allowed by the live events policy assigned to the person who organizes the event. If the event permissions allow, anonymous attendees can access the event recording. Live events policies are managed in the Teams admin center.

The settings controlling public access and disable recording for download are both off by default. To enable these settings, go into Live events policies in the Teams Admin Center and change the global policy or create a new policy with these settings enabled and assign them to specific users. You can also control these settings in PowerShell using the `Set-CsTeamsMeetingBroadcastPolicy` cmdlet. For example:

```
[PS] C:\> Set-CsTeamsMeetingBroadcastPolicy -Identity Global -BroadcastAttendeeVisibility Everyone -BroadcastRecordingMode UserOverride
```

## Live Event Presenters

Unlike regular meetings, where everyone can speak, share their video feed, and chat, only people assigned the organizer and presenter roles can speak, present information, and be visible during live events. Presenters can be:

- Tenant accounts.
- Guest and federated users.

- Anonymous external accounts. These are accounts without an Azure AD or Microsoft (MSA) account. If you want to use anonymous presenters for an event, you must set the [Allow external presenters](#) toggle when creating the event.

Attendees can ask questions, but only through a moderated Q&A facility. Because attendees have limited functionality in live events, the number of participants is much higher than for normal Teams meetings. The structured nature of Live Events means that some preparation is necessary to ensure the delivery of the best possible event. Steps to ensure a smooth event include:

- Having a dress rehearsal to make sure presenters understand the flow of the event and how they transition to speak. Presenters must join the event by signing into the host tenant. If they are a guest in that tenant, the presenter must switch to the host tenant before they join; if not, they will join as an attendee.
- Projecting an “Event will start soon” slide complete with information about the event and some music from five minutes before the event is scheduled to begin. This will reassure attendees that they’ve connected to the right event and will be able to hear the proceedings.
- Beginning the event with “house rules” to let attendees know how to ask questions and how presenters will respond to questions. During the event, someone should monitor incoming questions and note those that a presenter should answer. Moderators can reply to other questions straightway while other topics might need later follow-up.

**Note:** Microsoft has temporarily increased the [limit for attendees at a live event](#) to 20,000 with up to 50 simultaneous events per tenant. Each event can last up to 16 hours. The increased limits will expire on December 31, 2022, after which a Teams advanced communications license will be required to host live events with more than 10,000 participants or events lasting longer than four hours.

## Recordings

Organizers can record live events for later playback. The recording is stored in Stream and accessible for up to 180 days after the event, which makes it convenient for people to replay an event later and listen to specific parts of the discussion. Event organizers can download the recording and upload it to Stream if they want to keep it for a longer period. Live events can also be live-streamed to platforms like YouTube and Facebook using an encoder like [OBS Studio](#) (an open-source encoder).

## Teams Live Events External Encoder

Largescale live events which need output of the highest quality are usually carefully-planned productions created with studio-quality recording, camera, and broadcast facilities. These events often involve external encoder software to connect to studio production equipment.

The hardware or software-based encoder must support streaming to a Real-time Messaging Protocol (RTMP) service. Supported encoders include Haivision KB, OBS Studio, Wirecast, and XSplit Broadcaster (here is a [link to the encoders](#) tested by Microsoft and a [write-up](#) by MVP Luca Vitali on how to set up external encoder integration with OBS Studio).

To schedule a live event with an external encoder, create the meeting the same way as for a quick event. On the settings page, you select external encoder. Note that external encoder events are only consumed via Microsoft Stream, which is why public viewing will be unavailable for external encoder meetings. Users can view the streamed content in Stream or via the meeting invite in Teams.

# Teams Calling

Teams is the preferred communications solution for Microsoft 365 and supports two methods to enable users to make, receive, and transfer calls to landline and mobile phones connected to public networks (PSTN). You



can buy Phone system licenses with a [calling plan](#) (various plans are available in different countries for national and international calls), in which case the phone numbers assigned to users come from a pool assigned and managed by Microsoft. Alternatively, you can use [Direct Routing](#) (or direct connect) to link Teams to your existing connection to the PSTN network via a Session Border Controller (SBC). The technology for Teams calling comes from the Microsoft Phone System (originally called Cloud PBX).

Users must have a Microsoft calling plan to place calls to PSTN numbers. Accounts without a calling plan can call other Teams users using VOIP, including those in other tenants. The Teams Calling and Devices chapter contains detailed information about how to deploy and manage the foundational elements for voice and audio meetings, including calling inside and outside the organization.

Teams Devices and Calling are covered in a separate chapter.

## End to End Encryption for Teams Calls

Teams secures VOIP traffic with Transport Layer Security (TLS) and Secure Real-time Transport Protocol (SRTP). The combination is sufficient to secure most voice traffic, but if users need an extra degree of protection for ultra-confidential calls, the Teams Windows and Mac desktop clients can use end-to-end encryption (E2EE). In this scenario, the workstations at both sides of the conversation agree on an encryption key that they use to secure the call. Because of the additional layer of encryption, some Teams calling features don't work, including call transfer, recordings, and live captions. However, E2EE calls support chat.

By default, the Teams enhanced encryption policy for the tenant disables E2EE calling. To enable E2EE, an administrator must:

- Update the default policy to allow users to make E2EE calls, or:
- Define a new Teams enhanced encryption policy that allows E2EE calls and assign the policy to specific users.

In this example, we update the default Teams enhanced encryption policy to allow users to make E2EE calls.

```
[PS] C:\> Set-CsTeamsEnhancedEncryptionPolicy -Identity Global -CallingEndToEndEncryptionEnabledType DisabledUserOverride
```

Setting the policy to *DisabledUserOverride* allows Teams to display an option to use E2EE calls in the Privacy section of client settings.

End-to-end encrypted calls 

One-on-one Teams calls are end-to-end encrypted if both participants turn on this setting. Some features won't be available, including recording and transcription. [Learn more](#)

Directory self both users in a call enable E2EE, their call will be in E2EE mode, they'll know that this is happening by checking the shield icon at the top left-hand corner of the call screen. If it has a lock, E2EE is in force. Clicking the shield icon reveals a set of five 4-digit codes which should be the same for both users.

## Viewing Organizational Information

If the **Show organization tab for users** (in the Teams settings section in the Teams admin center) is **On**, users can view the organizational information about tenant users by:

- Searching for their name, and then selecting the organization view.
- Typing /org followed by their name in the command box.
- Hovering over their people card.

Like other Microsoft 365 applications, the people card reveals some personal information about the user such as their title, phone number, and if set, their email out of office message and their Teams status message. The people card also has links to a set of tasks:

- Switch to a personal (1:1) chat.
- Send a quick message to the person (the message is listed in your personal chat).
- Email the person. In this case, Teams launches the default email program configured for the workstation (for example, Outlook) with a new message addressed to the person using their email address defined in their Microsoft 365 account.
- Start an audio call with the person.
- Start a video chat with the person.
- View their organization information.

The organizational information about a person (Figure 12-17) comes from the reporting relationships, job title, and other information held in Azure AD. The information is only as good and as accurate or complete as it is in the directory, so it's a good idea to invest effort in populating and maintaining the directory.

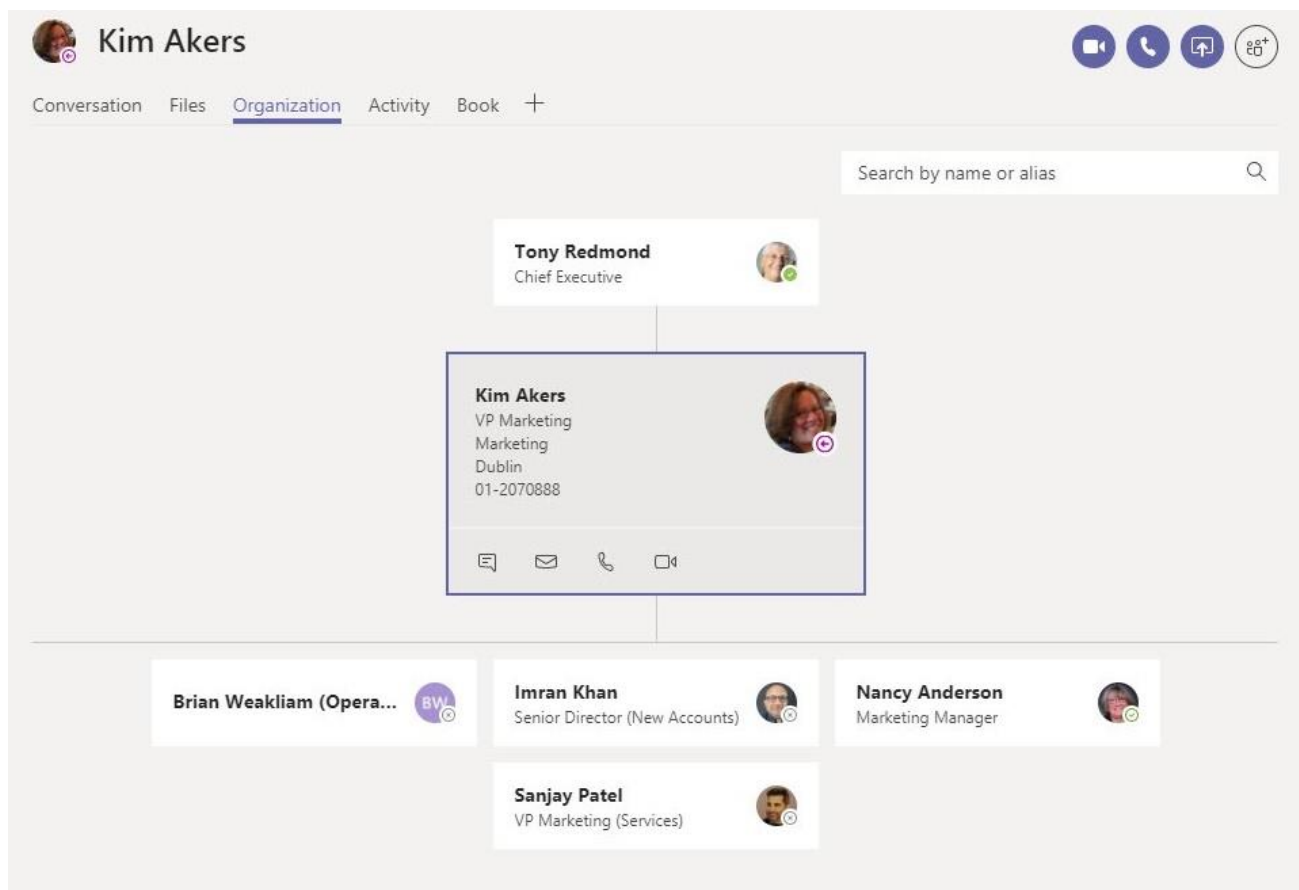


Figure 12-17: Viewing organizational information about a user

If incorrect or incomplete information about management relationships exists in Azure AD, it is impossible to get much insight into the organization through Teams. An inconsistent directory also makes it much harder to use features like dynamic Teams which depend on queries run against the directory. To ensure that up-to-date information is available to end-users, some companies use their HR system as the directory of record and have a direct feed between it and Azure AD, with updates in reporting relationships, job titles, phone numbers, and so on being transmitted periodically, while others use third-party products like [Hyperfish](#) to find and fix problems in Azure AD. For instance, you might find inconsistent use of account properties like city or state, and province. Another frequent problem is missing values for properties like department or office.

You can edit the contact information section of a user's account through the Microsoft 365 admin center to update properties like their job title, department, office, phone number, address, including a user's reporting relationship. You can also update this information through the:

- **Azure AD portal:** Select the user's account and update the job title and manager properties in the *Job Info* section. Update contact information (phone number, address) in the *Contact Info* section.
- **Exchange admin center:** Select the user's mailbox, then Organization, and update their manager in the *Manage Organization Information* section. You can also update their contact information in the *Manage Contact Information* section.

Changes made in EAC synchronize to Azure AD and are then available to Teams. Guests do not have access to organizational information. They can see someone's title, work address, and phone number, but the org chart icon does not appear on the people card.

When you view information about a person, you can also see details of personal chats you have had with them (Conversation), if you share files from your OneDrive for Business site with them in a 1:1 chat (Files – documents shared in group chats do not show up here), and recent posts from them to channels in teams to which you belong (Activity). Guests can only see the Conversation and Activity tabs. Another way of navigating the organization structure is to use the WhoBot, discussed in the section on bots later. WhoBot can track connections between users in a tenant, but it still depends on an accurately populated Azure AD to understand the organizational structure.

## Presence and Status

Teams keeps two indicators about someone to help users understand the current situation for that person.

**User presence** shows their availability to connect and appears in the user's people card and avatar in places like the activity feed. **User status** is a free-text message configurable by the user for Teams to display to others.

### User Presence

User presence is set in two ways:

- **User-configured:** The user picks a status and sets it in the Teams client. The presence values are Available, Busy, Do Not Disturb, Be Right Back, Appear Away, and Appear Offline. Users set their presence status by entering a command in the command bar (for instance, **/dnd** is the shorthand code used to set your status to *Do Not Disturb*) or by clicking their picture in the top right-hand corner and setting their presence there. Federation allows guest users to see user presence. If Teams cannot determine someone's online presence, Teams shows their presence status as "Offline." Users can also set their presence to Offline if they want to appear unavailable and uncontactable to other people. Teams supports the ability to set a duration for a presence, meaning that you configure a presence to last for a period from 30 minutes to a future custom date (for instance, after you return from an extended vacation). When the presence duration lapses, Teams reverts to an app-configured presence. You can't use the Available status when you set a duration.
- **App-configured:** A presence status chosen by a user always has priority and Teams uses this status whenever possible. If the user doesn't set their presence status, Teams tries to determine a presence based on their activity. For example, if someone is in a meeting or a call, Teams knows this and can change their presence status to reflect this activity. Teams synchronizes out-of-office information and calendar data with the user's Exchange mailbox and uses this information to figure out their status and current presence. For instance, if the user blocks some time in their calendar for an appointment or meeting and marks the time as busy, Teams knows that they are unavailable during their period. Time blocks booked by Viva Insights for "focus time" cause Teams to automatically set a user's

presence to Do Not Disturb. In addition to synchronizing calendar data, when a Teams client is signed in, it uses the [Graph presence subscription API](#) to detect calendar updates to take new events into account when calculating the user's presence.

Teams publishes presence changes immediately (or very soon afterward) following an update. When someone's presence is set to Do Not Disturb, Teams does not deliver notifications for normal messages and @mentions, while continuing to deliver notifications for urgent messages or those from people on the user's priority access list (managed in the **Privacy** section of **Settings**).

See [this page for more information](#) about the icons displayed by Teams to reflect a user's presence status. Of course, just because Teams reports someone as being available doesn't mean that they are. They might just be asleep!

## User Status Message and the Profile Card

Teams shows a user's profile card (Figure 12-18) when users hover over their photo or avatar in chats and conversations. The profile card contains:

- User information: Name, job title, department.
- Current presence status as calculated by Teams.
- User status (see below).
- Local time for the user and the time difference between the user and the person viewing the profile card. This information comes from the user's calendar time zone setting. Local time information is available for guest accounts if a federated organization sharing policy to share free/busy calendar information exists with their home tenant. Users can update their calendar time zone through Outlook settings, or administrators can run the *Set-MailboxCalendarConfiguration* cmdlet. See [this article](#) for more information.
- Primary SMTP email address.
- Fixed and mobile phone numbers.
- Office.
- Direct reports and manager.

This information will not be accurate if the data held for user accounts in Azure AD is inaccurate.

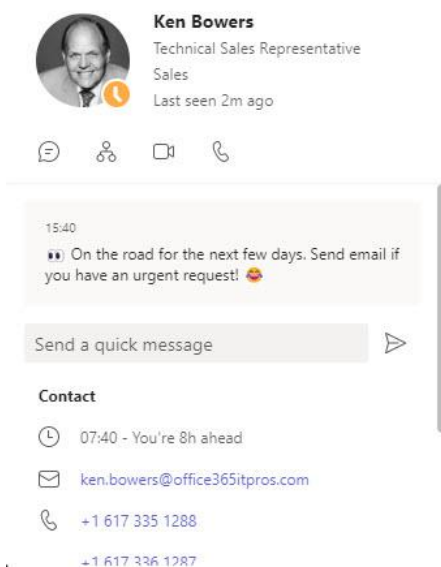


Figure 12-18: The user profile card holds lots of information

In addition to the information extracted from Azure AD, users can publish their personal status message. This is a free-form 280-character value status message intended to give other people a more descriptive insight

into their status. To update the status message in the desktop and browser clients, click the user photo in the top right-hand corner and select **Set status message**. Plain text and emojis are usable in a status message. Mobile clients support setting the status message through the user menu (top left-hand corner). When a status message is available, Teams displays it under the user's photo.

Like the presence status, a user status has a lifetime, which can be Today (until midnight), 1 hour, 4 hours, This week, Custom, or Never (the status never expires). After its lifetime expires, Teams removes the status message. Users can choose for Teams to display their status when someone sends them a message or @mentions them in a message. In these cases, Teams displays the status message over the compose message box. It also displays a reminder about the status message and its lifetime in chat.

In Figure 12-18, the user's presence is set to *Appear Away*, and their status gives some information about their ability to respond to messages. The people card also shows when the user was last active in Teams.

If an Exchange out of office notification is set for a mailbox, Teams includes this information in the people card unless the user has already explicitly set a status. If the user sets an out of office notification first, it takes precedence until the user clears the notification in Exchange. On the other hand, if the user sets their status in Teams first, subsequently setting an out of office notification in Exchange won't overwrite the status. Guest users can see the status message for other users in host tenants, so it's not a good idea to include anything confidential in a status message. Users can set the Exchange out of office notification from the Teams desktop and browser clients and Teams will update the auto-reply configuration for their mailbox to make sure that all clients use the same settings.

Guests can also set a status message, and they can set a different message in each tenant where they have an account. The status message set by federated contacts in other tenants aren't available during chats.

## Files: Linking Teams and SharePoint Online

The process to create a new team includes provisioning a new SharePoint Online team site containing a document library with a folder for the General channel. The purpose of the folder is to hold documents posted to the General channel. The creation of a new channel also creates a new folder of the same name in the document library. Users access the contents of the document library through the **Files** channel tab. This is one of two default tabs created for every channel that team owners or administrators cannot remove or rename.

Accessing the Files tab opens the folder holding channel files. Team members can upload documents to the folder or drag and drop files from File Explorer or the Mac Finder to move items into the folder. Naturally, any files created in the folder through the SharePoint browser interface or using Office apps are also available. Users can edit Word, Excel, and PowerPoint documents within Teams or open the file with the Online version of the app. The editing experience is similar but opening the file with Office Online or the desktop application allows the user to continue to do something else with Teams during the edit.

The Files tab supports navigation to sub-folders belonging to the channel, but if users want to move to another folder in the document library, they can click **Open in SharePoint** (available in the Files tab or the [...] menu for Posts) to open the site with SharePoint's browser interface. They can then access other folders in the document library, such as the folders belonging to other channels, or, if you team-enable an existing group, folders previously created in the default document library. It's important to understand that Teams does not integrate other document libraries in its navigation. If a site spans multiple document libraries that you'd like to access from Teams, you need to create individual channel tabs pointing to the root folder of those libraries (see below).

Another example of where you might create a channel tab pointing to a SharePoint resource is where you want to expose the site Recycle Bin to users to allow them to recover items deleted in error.

## Using Tabs to Access SharePoint Resources

To enable access to other SharePoint resources from within Teams, you can add tabs to bring users directly to a specific folder in a document library, a list, or a page in the SharePoint site belonging to the team. Here's how:

- Select the channel you want to make SharePoint resources available to team members, and then click the plus sign (+) to invoke the **Add a tab** dialog. Now select the **Document library** tab to add a link to a document library or the **SharePoint** tab to link to a published page or list in the site.
- If you select the Document library tab, you can now select one of the **Relevant sites** (essentially, a list of the sites that you accessed recently) and then a document library from that site; or select **Use a SharePoint link** and input the URL to the library you want to make available through the tab. The easiest way to get the URL is to open the target location in a browser and copy the URL from there to Teams. Make sure that team members have the necessary permissions to open the target.
- Give a name to the tab to help users understand its purpose (you can rename the tab later if necessary) and then **Save**.
- If you select the SharePoint tab, Teams retrieves a set of available pages to choose from. Select the page in the team site to display (for example, Home or News) and **Save** the setting. The tab will take the name of the chosen page (you can rename the tab if you want to).

An example of linking to a specific page is when you want to publish news articles to Teams. You can bring news articles into Teams as cards created by the News connector, but if you link to the News page for the site, you see all the news items posted, including the web parts (images, etc.) used for each item, and can comment on an item. In other words, you can interact with the page instead of just seeing a static snapshot of the content.

### Linking to Resources in Another Site

As noted above, the SharePoint tab is a good way to bring people to the News page for the site. However, what happens if you want a tab to open resources in another SharePoint site that doesn't belong to the team? Teams does not have an out-of-the-box way to do this, but a workaround exists through the website tab. This tab accepts and displays any valid URL, so it is possible to input a URL to take users to any SharePoint resource that they can access. All you need to do is format the URL to tell Teams what to do. For example, let's say that you have a URL pointing to an important news item:

*<https://office365itpros.com.sharepoint.com/sites/BlogsAndProjects/SitePages/Microsoft-Launches-New-Teams-Exploratory-Experience.aspx>*

The format of the URL to input into the website tab is:

*[https://office365itpros.sharepoint.com/sites/BlogsAndProjects/\\_layouts/15/teamslogon.aspx?spfx=true&dest=/sites/BlogsAndProjects/SitePages/Microsoft-Launches-New-Teams-Exploratory-Experience.aspx](https://office365itpros.sharepoint.com/sites/BlogsAndProjects/_layouts/15/teamslogon.aspx?spfx=true&dest=/sites/BlogsAndProjects/SitePages/Microsoft-Launches-New-Teams-Exploratory-Experience.aspx)*

This URL construct takes care of any authentication necessary to access the site hosting the content (assuming the user has access to that site). This is a workaround and it's not guaranteed to work every time (the best solution would be if Microsoft delivered a tab that could access any SharePoint page in the tenant). If you encounter problems, consider the [solution proposed here](#). It takes more work, but it might be just the trick for you.

## The Files Channel Tab

The Files Channel tab is a standard tab available for all channels to enable access to the files stored in the folder belonging to a channel in the team's SharePoint Online document library. Figure 12-19 shows the source Word documents for this book. The original version of the Files channel tab presented a simplified

view of documents when compared to the view available in the SharePoint browser interface. Microsoft has made steady progress towards the goal of giving the Files tab parity with the browser interface, and the two are much closer now, with support included for features such as column formatting and customization. By default, Teams uses the OneDrive file viewers to open and work with files. These viewers support over 300 file formats. For Office documents, users can configure clients to open files as follows:

- Desktop: Use Teams (the viewers), Office Online, or desktop apps. Only Office 16 or later is supported.
- Browser: Use Teams or Office Online.

The Files channel tab lets users open, move, copy, rename, download, and delete files to work with documents without the need to open the SharePoint browser client. You can also check out documents. Co-authoring is supported because Teams supports both the [WOPI](#) and [FSSHTTP](#) protocols.

Some limitations exist in the Files tab compared to the SharePoint interface. For example, you can't access the version history for a document, so you can't restore a previous version. You can't add a Flow or an alert to a document, and you can't apply retention labels to documents or access other custom properties. Finally, you can't share a document. In these cases, you're forced to open the library in SharePoint to perform these actions. The lack of label support in the Files tab might mean that some documents do not get the labels they need for retention or compliance purposes. If this is a concern, make sure that users receive advice about how to apply retention labels (see Chapter 17) to individual files. Alternatively, if the organization has Office 365 E5 licenses, it can deploy auto-label policies to detect and label files containing sensitive information.

Tenant users can synchronize files in the SharePoint document libraries used by Teams just like any other SharePoint site. The situation is different for guest users, where a OneDrive for Business feature called [B2B Sync](#) must be used to allow guest accounts to synchronize documents.

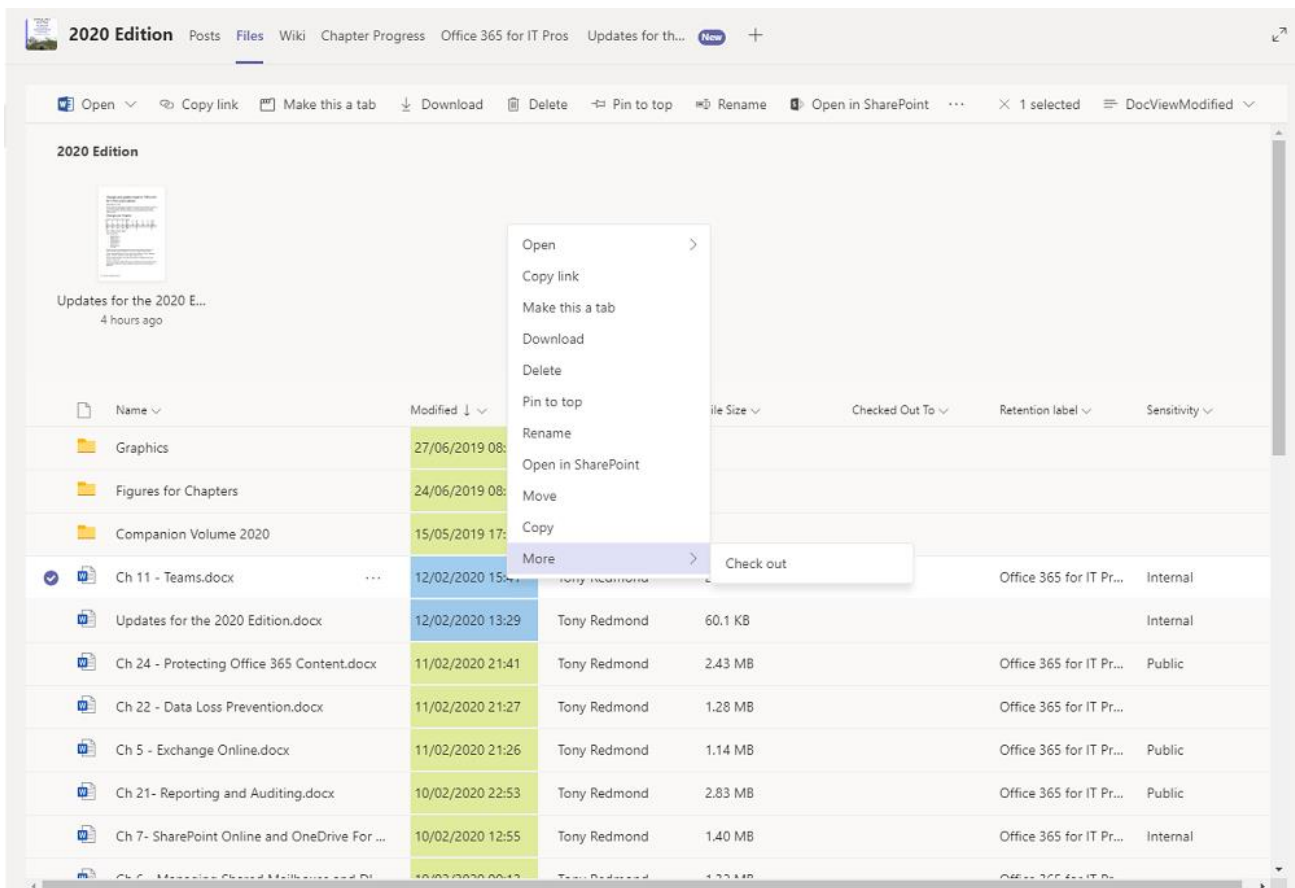


Figure 12-19: The Files channel tab displays files in a folder in a SharePoint document library

## SharePoint Site Membership

You should let Teams manage the membership of the SharePoint sites connected to teams. If you are a SharePoint expert and understand the consequences of making changes to site membership, feel free to do so. However, remember that if something doesn't work afterward, you will have to support what you've done.

Another thing to consider is that Microsoft might synchronize site membership between Teams and SharePoint in the future. This happens for private channels today where every four hours (approximately) Teams synchronizes membership of the private channel to SharePoint and overwrites the existing site membership. In other words, if you change the site membership in between synchronizations, Teams will reset the membership the next time synchronization occurs.

## Sharing Links

The **Copy link** option creates a sharing link to the folder belonging to a channel or a selected document. The organization's sharing policy for SharePoint Online dictates which types of sharing links are available for use. See Chapter 8 for more information about sharing links.

To use a link, click **Copy Link** and then adjust the link settings to control who can use the link, if they can edit the file, and so on. When the link is ready, click **Copy**. The sharing link is then available in the clipboard. You can then paste the link into a chat or email and send it to the intended recipients.

## Tabs to Link to Specific Documents

Apart from tabs connected to SharePoint sites, team owners can add tabs for fast access to a selected document. Open the SharePoint library with Teams, select the document, and then use the **Make this a tab** option. Teams creates the tab to point to the document. Alternatively, add a tab and then select the file type (Word, PowerPoint, etc.), and then select the document you want from the document library belonging to the team (you cannot select a file from another document library). When you create a tab for a file, selecting the tab invokes an Office viewer to display the content of the file. You can also edit the file in either the online app or by launching the desktop app if it is available on the workstation.

## Files in the Navigation Rail

The **Files** link in the left-hand navigation rail exposes a different set of files to the Files tab for a channel. Instead, the default (Recent) view for this link lists the recent files worked on by the user stored in SharePoint sites and their OneDrive for Business site. Some other shortcuts appear under Files:

- **Recent** is a list of files in SharePoint document libraries belonging to teams that the user recently accessed. Guests do not see this link.
- **Microsoft Teams** shows a list of files, including email messages, recently loaded into teams to which the user belongs.
- The **OneDrive** link exposes a list of folders and files stored in the user's OneDrive for Business site. Guests do not see this link because they do not have a OneDrive site in the tenant.
- **Downloads** shows the files in the Downloads folder on the device.
- If you connect a **cloud storage service** to Teams, like Google Drive, a shortcut to that service appears here.

## OneNote

When you create a tab in a channel to access OneNote, you have the choice to create a new notebook, browse to find existing notebooks to link to the tab, or paste [a notebook link](#) to bring users to a specific notebook, or a section, page, or paragraph within a notebook.



## Sharing Files

When someone shares files in a public conversation, Teams captures copies of the files in the folder of the group document library for the channel. However, when someone shares files in a private chat, Teams first uploads the file to a folder called “Microsoft Teams Chat Files” (under the Files folder) in the owner’s OneDrive for Business site and then shares it with the other chat participants.

## SharePoint News Connector

SharePoint news items are a special form of web page created in a SharePoint site. The SharePoint News Connector creates notifications about news items in a channel after they are posted. You can only create notifications for news items posted to the SharePoint site belonging to a team. News items posted to other sites are ignored. After the notification messages appear in the channel, users can click on a message to go to the full news item in the SharePoint site or use the item to begin a conversation.

To create a connector, select **Connectors** in the ellipsis menu for the channel where you want the notifications for news items to appear, search for **SharePoint News**, and then click **Configure**. Teams then creates the webhook to check for news items posted to the site. Click **Save** to set up the link.

## Email Folders

If people send an email to a channel, the email travels through a connector. SharePoint Online captures copies of the messages sent through the connector in sub-folders of the channel folder, created as needed on a month-by-month basis. For instance, messages sent to a channel in July 2022 are in the *EmailMessages\_7\_2022* sub-folder of the channel folder.

## Cloud Storage

If allowed by the Teams settings for the tenant, the Files tab supports access to different cloud services such as Dropbox, Box, Egnyte, Citrix ShareFile, and Google Drive (including personal drives) that the tenant might use instead of OneDrive for Business or SharePoint. The process of connecting to a cloud service means that you provide credentials to connect to an account in the service and use that account to select a folder in the storage accessible to the account. Once configured, a link to the cloud service appears in the list of files. Users must be able to authenticate with the service and access the files in the chosen storage location. Once connected, you can view, edit, upload, and remove files stored in the cloud storage. You can also start a new conversation about a file in cloud storage, just like you can discuss files stored in SharePoint.

## Teams for Frontline Workers

Teams for Frontline Workers is a Teams-based app built by Microsoft to replace the old StaffHub app. Designed around the premise that frontline staff such as sales associates and construction workers are often organized into scheduled periods of work activity, the app is also known as Shifts. In addition to the app, the [Shifts Graph API](#) is available for organizations to integrate Shifts into existing workforce management tools. Neither the app nor the API is available in the free version of Teams.

All Teams clients support the Shifts app. To access the app, click the More apps (...) button in the Teams navigation menu and look for the app in the Teams app store. It is also possible to assign a [Teams app policy](#) to selected users so that Shifts shows up in the app bar (along the side in the desktop and browser clients, at the bottom for mobile clients).

Inside Shifts, team owners can select a team to host a schedule and begin creating shifts (periods of work) and worker categories (called groups) to organize the shifts. Like a calendar, you can view shifts by day, week, or

month and see which of the members are already scheduled and those who are still available. Because the schedule is built around the people defined by team membership, a team can only support a single schedule. If you need to maintain several schedules, you must create a team for each schedule.

After creating a schedule, team owners and members permitted to manage schedules can assign members of the team to shifts and note absences such as scheduled vacations (Figure 12-20). To make things easier to build a schedule, you can copy the data for an existing shift from a previous period. If you add someone to a shift who isn't already a member of the team, the app adds them to the underlying group. Although Shifts allows you to create a schedule for a dynamic team, you shouldn't use these teams with Shifts. Two problems are apparent. First, if the query against Azure AD changes the team membership, some shift assignments might be affected. Second, if you try to add a new user to the team and assign them some time slots, Shifts accepts the assignments, but the person is not added to the membership and so can never see the schedule.

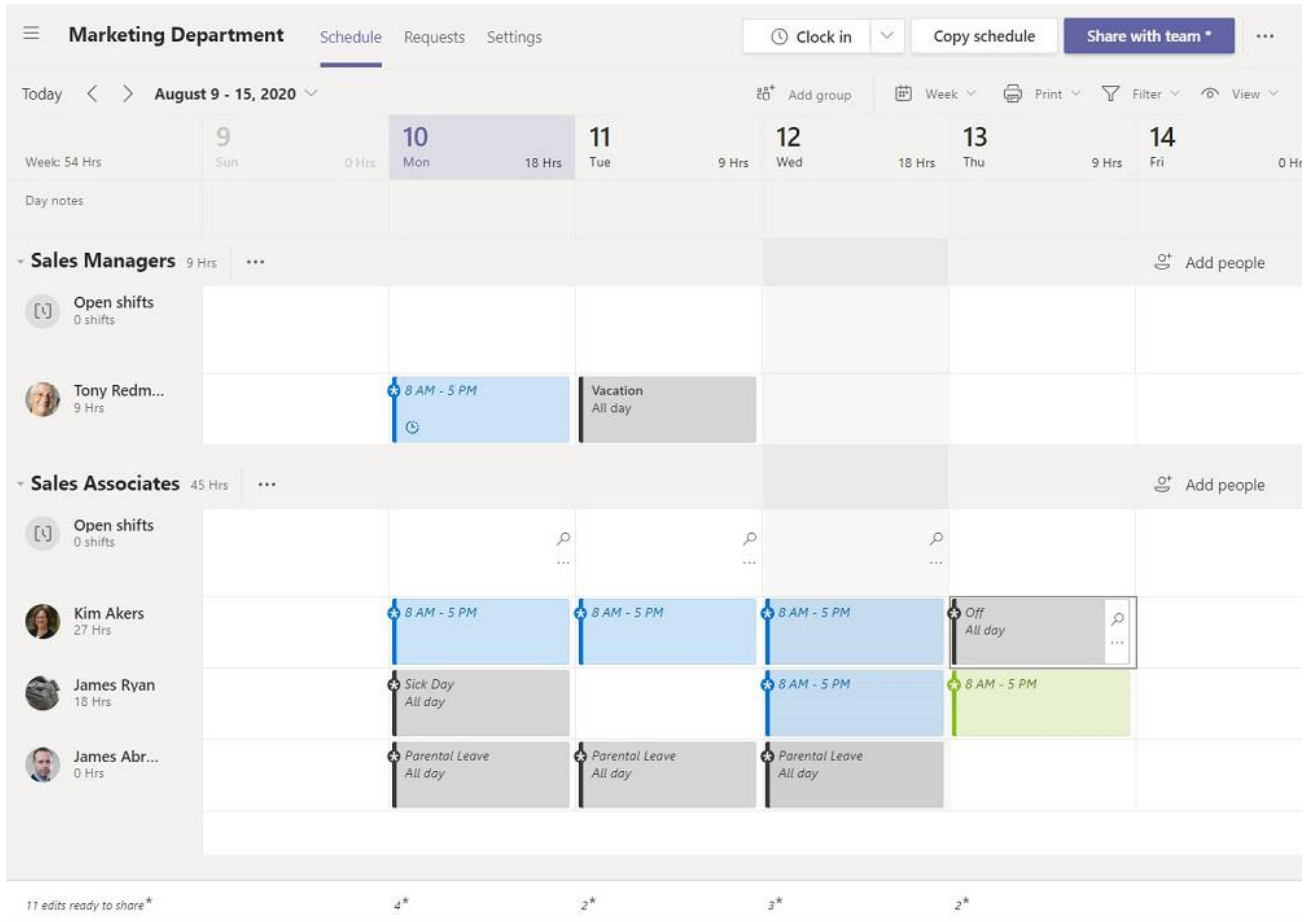


Figure 12-20: Building out a shift with the Shifts app

Once the shift is populated with assignments, you use the **Share with team** button to publish the information to the team. Publication can be to the complete team or just those who are affected by any changes made to the schedule since it was last shared. Sharing notifies team members that details of a new schedule are available. Members then open the Shifts app to review their schedule and request any adjustments they might want to make. For example, someone might want to ask for time off, or swap one of their scheduled shifts with someone else. These requests go to the team owner, who can approve or deny the requests.

## Can Teams Replace Email?

Although messaging is an important part of Teams, it's obvious that other parts of the app are equally if not more important to different people. Meetings and calls, for instance, became the most important part of

Teams for many organizations during the Covid-19 pandemic, while the development of Teams as an app platform has been remarkable. Nevertheless, some proponents of Teams focus on messaging and believe that you can reduce the use of email within an organization by replacing email with chat and channel conversations. Among the advantages cited for this approach are:

- Conversations and documents exist in the team rather than in multiple mailboxes. Thus, users have just one place to find information. Conversations are in channels while documents (including attachments in emails sent to the channel) exist in Files (a SharePoint document library).
- Conversations do not fork. It is easy for a small subset of recipients to begin a separate conversation after receiving email. This does not usually happen inside a channel because the conversation stays there and is visible to all team members.
- The volume of email declines when some traffic moves to Teams and email evolves to serve different purposes. Internal conversations stay in Teams while email handles external communications that Teams cannot handle (because no way exists to send outbound email from Teams). If necessary, people can initiate internal conversations by sending emails to a channel and continue the discussion in the channel from that point.
- Group chats are a good way to get together small teams of people to dissect and understand a problem before bringing a decision to a wider audience.

The advantages its supporters envisage for Teams are aspirational, and no guarantee exists that everyone is willing to change the habits of a lifetime and move from their preferred email client to Teams overnight, even as Teams adopts some of the characteristics of email in features such as message moderation and auto-replies. It's also fair to say that it is as easy to create a chaotic mess of Teams conversations spread across a sprawl of channels as it is to create an unordered messy inbox. Much of the success in any transition to a new technology comes from the effort put into user training and support. Left alone, users will find their bad habits for Teams.

It's also important to consider how different email clients affect the transition. People often refer to their client and how they use it instead of the email application. Their view of email comes from how they use their client(s). Outlook desktop is a great case in point as many people have built their working habits around how Outlook functions. People who use simpler email clients, like the Windows Mail client or a POP3 client, are perhaps less likely to be quite so attached to their client.

## Teams Messaging Can Work Better than Email

There's no doubt that Teams messaging (chat or channel conversations) can be a better tool to drive issues forward and get things done than email is. Take the situation shown in Figure 12-21 where a channel conversation has 269 replies and there are 207 members in the team. If the communication happened over email and every one replied as they did to the conversation, 55,890 (1 topic + 269 replies x 207 recipients) messages would circulate among the team. Not only would people have to spend a lot of time deleting some of the thread messages unread, but they also must cope with out-of-office notifications. Finally, consider the length of the final message and the impossibility of navigating through all the appended replies, auto-signatures, and other junk to get to the gist of the conversation. A well-organized threaded conversation in one place is a much better host for fast-paced to-and-fro communications.

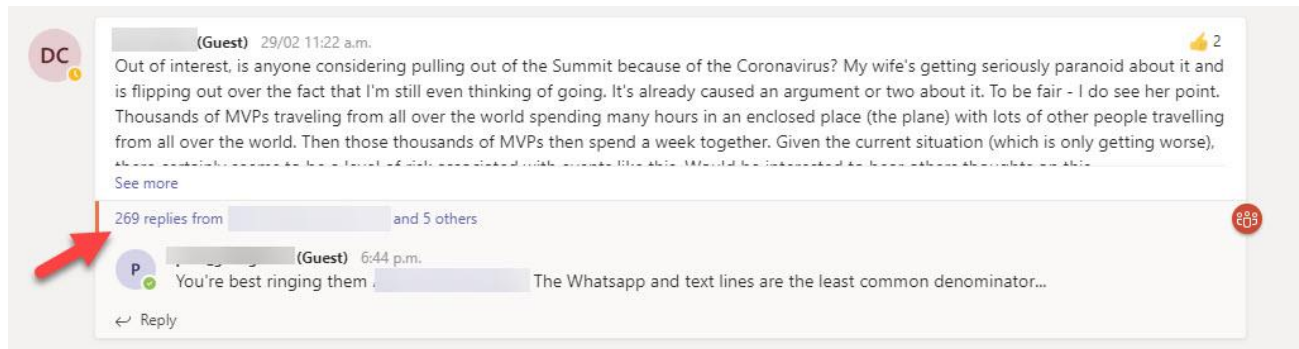


Figure 12-21: A Teams channel conversation with many replies

Channel conversations can be as chaotic as a cluttered inbox, especially when multiple team members contribute to different conversations over a short period. It can be very easy to lose sight of something important in a flurry of messages. It is also important to understand that Teams is very much an internal-facing collaboration mechanism. Even with guest users joining in, the data associated with Teams always remains in the home tenant. Compare this to the way that email serves as the lingua franca for the internet where anyone with an email server can use it to communicate and collaborate with anyone else. Another important factor to consider is the added functionality available in email, such as the ability to protect messages and attachments with rights management and encryption, not to mention features to make it easier for users to organize messages, like rules, categories, flagging items for follow-up, and so on. These features take years to develop and deploy, and Teams is still very much a young application.

Those who champion Teams should consider that all applications have their flaws. Among those seen in Teams are:

- Chaotic, badly organized conversations composed of short, incomplete, and shallow sentences that fail to form fully thought-out ideas. A conversation might have many replies from many individuals and be devoid of analysis and insight.
- Conversations that veer away from the subject at hand because they are hijacked by people who fail to understand the topic.
- Conversations can be dominated by the speedy, the quick-witted, those who have a point of view that they want (and need) to express, and those who suffer from verbal diarrhea. Some members of a team might choose not to contribute to a fast-paced conversation that, although open to them, is taking place between others.
- Conversations about anything that comes into peoples' heads that mask the real work.
- Conversations about topics driven by people who have no authority to speak on a subject or no accountability for an outcome. It is truly amazing how easy it is to waste time in such discussions.
- A rush to decision based on a feeling that chats should develop quickly.
- The fear of losing out if you fail to keep pace with all the conversations that happen in all the channels in all the teams you belong to. Apart from flagging a conversation as important and favoring and following channels, it's hard to prioritize the firehouse of chats daily, let alone sort out what might have occurred when you return from a vacation.

Email can suffer from some of the same problems, but because your inbox is under your control, it can be easier to sort, ignore, prioritize, and respond to what flows into the inbox.

Some organizations will find the transition from inward-circulating email to Teams easy and some will decide that they are best off staying with email and favor Outlook Groups as their collaboration platform instead, especially when regulations or legal requirements dictate the use of some of the compliance functionality presently missing in Teams. Interestingly, it can be easier for users of other collaboration platforms, like Lotus Notes, to move to Teams than to break habits formed around Outlook.

Usage reports, such as those available in the Microsoft 365 admin center or the Teams admin center, can help you understand if any email traffic moves to Teams. However, when you examine usage data, remember that some traffic handled by Teams might replace conversations previously carried in other apps. It is also important to measure after the first burst of activity subsides and people start to use Teams consistently. You will then know whether overall traffic is static or growing and what modality is most popular with users.

## Delegation of Access

Unlike many email systems, Teams does not support the concept of delegation of access for messaging, a feature that underpins common email functionality such as the ability for someone to:

- Send messages on behalf of or as (impersonate) another user.
- Schedule meetings on behalf of another user.
- Manage another user's workspace (their mailbox).

Teams is very different from email, so there's no reason why Teams should slavishly copy functionality from email systems across to its methods of messaging and collaboration. Nevertheless, the desire to move some volume of communications from email to personal chat and channel conversations plus the growth in Teams online meetings creates some need to support the manager-assistant scenario.

There's no good answer for delegation of access in Teams today, only workarounds. For more information, see [this article](#). Teams Calling does support delegation for call handling.

## Other Email Strengths

Other areas of functionality where email is ahead of Teams include:

- Printing. Teams doesn't include the ability to print messages or the calendar.
- Single focus: If someone has guest accounts in multiple tenants, they must switch to each tenant to check chats, channels, and apps. Compare this to the single inbox for all new emails, no matter where they originate.
- Secure communications: Chats and channel conversations are excellent methods for internal communications that extend to some external communication through guests and federated chat. Teams messaging does not include the secure (protected) functionality commonly found in emails such as S/MIME or other encryption (like sensitivity labels).
- Complex messages: Email is better suited to composing longer and more complex messages of the type used for formal business communications.
- Calendar: Outlook's calendar is more functional than the version available in Teams.
- History: Email allows users to keep all messages if desired in online or offline locations.

Different people will assign different levels of importance to each of the points listed above. The important thing is to help users decide when it is best to use Teams and when to use email. Although messaging might seem to be a common thread, they are very different tools.

## Can an Organization Cope Without Email?

Overall, it is hard to see Teams replacing email anytime soon. Teams is a great vehicle for communication, but if it wasn't available, organizations could continue to work with people inside and outside the company because other communication media exist, including email. If an organization decided to remove email, it could continue to communicate internally using Teams, Yammer, and other applications, but its ability to communicate externally with partners, clients, and other companies would be significantly and fundamentally reduced (in some cases, eliminated). With [over 3.9 billion users worldwide](#), email is ubiquitous, secure, and backed by a range of well-understood and widely-deployed standards. Moving on from such a foundation will be difficult.

# Teams and Email Interaction

Because email isn't going away anytime soon, Teams includes several methods for the two modalities to interact:

- **Share to Teams** uses an Outlook add-in to send a message to a Teams channel or chat (including the ability to create a new chat. Share to Teams works with Outlook for Windows (Microsoft 365 apps for Enterprise), Outlook for Mac, and OWA. It isn't available on Outlook mobile.
- **Share to Outlook** is a Teams client option that calls OWA to create and send a message containing a copy of a channel or chat conversation. The message can go to any email recipient, including distribution lists and external recipients.
- **Chat with Teams** and **Reply with IM** are Outlook desktop options available when Teams is the registered chat application for Windows. Teams launches a pop-out chat window connected to the author of the selected message. Chat participants can include Teams users from other domains.
- **Reply to Teams Missed Activity Mail** gives users who receive missed activity notifications the ability to respond to conversations in Teams using [Outlook actionable messages](#).
- **Email-enabled channels** have special email addresses to allow the delivery of messages through a connector to become channel conversations. Organizations can restrict who can send email to an email-enabled channel. See the Managing Teams chapter for more information.
- **Drag and Drop from Outlook desktop** allows users to drag and drop a message (and any attachments) into a Teams channel conversation. The feature doesn't work with chat.

The Share to Teams and Share to Outlook features depend on cloud mailboxes, so users with on-premises Exchange server mailboxes or guest users cannot use the feature.

## Share to Teams

The Outlook desktop client installs the Share to Teams add-in automatically when a user signs into the Teams desktop or Teams browser client. The add-in uses single sign-in to allow Outlook or OWA to post messages into target Teams channels, personal chats, or group chats. To share a message, click the Share to Teams icon in the Outlook menu bar. Outlook checks the connection to the user's home tenant (if the user switches to be a guest to another tenant, they need to switch back to use Share to Teams) and for the presence of the Teams desktop client. If the Teams desktop client is available, Outlook opens a window in the Teams client and loads the current message into a form to allow the user to add addressees (Figure 12-22). After someone uses the feature a few times, Teams learns about the people and locations commonly used for sharing and suggests these as potential sharing targets.

Valid sharing targets are:

- Individual users: Outlook posts the message to a private chat.
- Several users: Outlook posts the message to a group chat.
- Channels in Teams: This includes private channels the user belongs to. Posting to multiple channels is not possible.

To add context to the post in Teams, the user can insert some cover text to explain the message to the channel members. If the message has attachments, you can choose to include or omit these files. When the message is ready for posting, click **Share**.

Behind the scenes, Outlook shares messages through the same mechanism as used to post a message to a channel via email. Depending on the recipients, sharing creates a new topic in the target channel or a new message in a personal or group chat and posts the message. When Outlook posts a message to a channel, Teams saves the message and any attachments in the *Email Messages* folder for the channel in the SharePoint document library belonging to the team. Messages posted to a private or group chat are in the Microsoft

Teams Chat Files folder of the sender's OneDrive for Business account. Teams shares these messages with the group participants to allow them access.

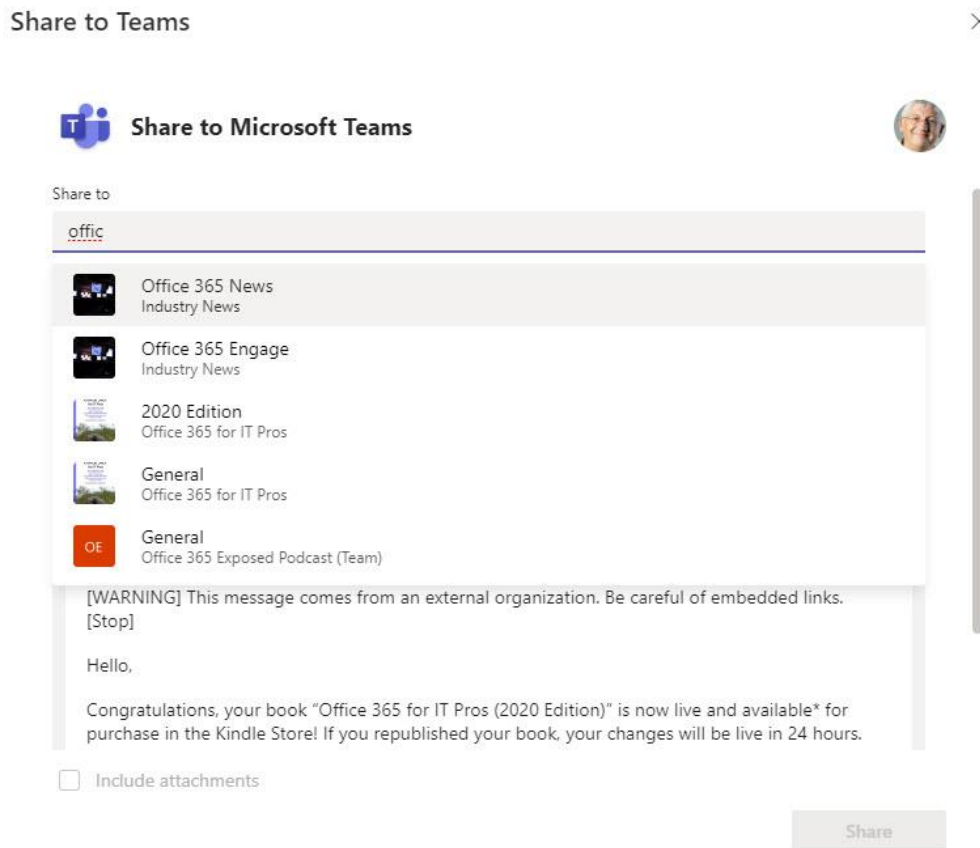


Figure 12-22: Selecting recipients to share an Outlook message with Teams

Share to Teams does not support messages encrypted with Microsoft Information Protection sensitivity labels, Office 365 message encryption, or S/MIME.

If you don't want users to use the add-in, you can disable it using PowerShell by running the *Disable-App* cmdlet. This code finds the identifier for the Share to Teams add-in and uses it to disable the add-in for a mailbox. If necessary, to disable the add-in for multiple accounts, use an input array of mailboxes and pipe the objects to the *Disable-App* cmdlet:

```
[PS] C:\> $ShareToTeamsApp = (Get-App | ? {$_.DisplayName -eq "Share to Teams"}).AppId
Disable-App -Identity $ShareToTeamsApp -Mailbox Chris.Bishop@office365itpros.com -Confirm:$False
```

## Share to Outlook

Sharing to Outlook means *"forward a conversation from a Teams conversation or personal chat to one or more email recipients."* To share, select **Share to Outlook** from the [...] menu, which only appears if the user mailbox is in Exchange Online. Guest accounts can't use Share to Outlook because they don't have a mailbox that they can sign into.

When you forward a personal chat, Teams shares the selected message rather than the entire chat. For a channel conversation, Teams extracts a copy of the complete conversation (this can be slow for a conversation with many replies). Teams loads the HTML content into an OWA message compose form for the user to add address information and make any changes necessary to the message body. If a user's mailbox is not enabled for OWA, Teams won't be able to call the OWA message compose form.

You can add any valid email address, and you can choose to send the message from any mailbox for which you have permission, including shared mailboxes and Microsoft 365 Groups. Because an OWA component creates and sends the message (unsent messages are in the Drafts folder), OWA features like support for sensitivity labels are available. And of course, email clients can print the messages sent from Teams, thus solving the lack of printing support in Teams clients.

Messages forwarded from Teams travel through the Exchange Online transport pipeline and are subject to any mail flow rules enforced there.

## Reply to Teams Missed Activity Mail

Depending on a user's missed activity email settings, Teams sends periodic email reminders to make sure that users know what's happening in a chat or channel conversation, like being @mentioned. Teams missed activity notifications are actionable messages which include two action buttons:

- The **Go to conversation** button uses an embedded deeplink to open Teams positioned ready to reply to the message, much like users can join a Teams meeting from an Outlook calendar event.
- The **Reply** button opens a dialog to allow the user to reply to the Teams conversation without leaving Outlook. Being able to compose and send an inline reply to a conversation without having to switch context from email is especially valuable in terms of helping users to maintain focus. Actionable messages support only simple text replies (no attachments, emoticons, reactions, or text formatting). One of the nice things about this feature is that you can reply from email even when you have not signed into Teams in the tenant that hosts the conversation.

Outlook desktop, OWA, and Outlook mobile support actionable messages for Teams responses. Email clients that can't process the JSON content within the message which defines the actions and the transactions with Teams to perform responses replace the two action buttons with a **Reply in Teams** button. This button works like the **Go to conversation** button in that it uses the deeplink in the message to open the Teams client. In addition, if someone is using Outlook desktop and is not signed into their home tenant with the Teams desktop client, Outlook offers only the Reply in Teams option. To restore full functionality, the user must switch back to their home tenant and then restart Outlook desktop.

A tenant can block actionable messages by setting *SmtplibActionableMessagesEnabled* in the Exchange Online organization configuration to False. By default, this value is True, meaning that applications like Outlook can interact with actionable messages. To block this capability, run the *Set-OrganizationConfig* cmdlet:

```
[PS] C:\> Set-OrganizationConfig -SmtplibActionableMessagesEnabled $False
```

The block affects messages generated by all Microsoft 365 workloads, including Teams and Yammer.

## Reply with IM

The idea behind the Reply with IM feature is simple. You receive an email in Outlook and instead of having endless rounds of to-and-fro replies, you continue the conversation in an instant messaging platform that's more suitable for an interactive debate. The Reply with IM option is in the [...] menu of Outlook's read message window or in the Respond section of the Outlook menu bar (Microsoft has a habit of moving these options around). Reply with IM launches a conversation with the sender while Reply All with IM includes all the recipients in the conversation.

To use the feature with Teams, a user must be:

- Configured in [TeamsOnly mode](#). The value of the registry key *HKCU\Software\IM Providers\DefaultIMApp* should be "Teams." This value is set when you choose to register Teams as the chat app for Office in Teams settings. You can [update the registry using this script](#).



- Signed into the Teams tenant hosting the users you want to chat. In other words, if you want to chat with people from your home tenant, make sure that you sign in there.

There are some details to remember when using Reply with IM:

- If an existing chat with the recipients exists, Teams will use that. Otherwise, Teams creates a new chat.
- Teams doesn't take the message subject and use it to name the chat, even when it creates a new chat. Apart from the recipients, Teams copies nothing from the message into the chat, so you must cut and paste information from the message body into the chat if you want to provide context for the conversation.
- Reply with IM supports federated (external access) chat with Teams users from other Microsoft 365 domains. However, it does not support chat with Teams consumer users.
- If one of the message recipients is blocked for chats by Teams, you won't be able to send messages to the chat.
- If you are signed in as a guest to a Teams tenant where an external recipient is homed, Reply with IM can launch a conversation with that person.
- Rather bizarrely, if a shared mailbox is in message recipients, Teams includes the shared mailbox in the chat (you can clean things up by removing the shared mailbox from the chat).
- If the message recipients contain a group, Teams drops the group when it starts the chat.

Despite some oddities, Reply with IM works well and is a useful feature to have when a chat is a more appropriate host for a conversation than keeping it in email.

## Drag and Drop from Outlook Desktop

Outlook for Windows supports drag and drop of a message and any attachments from any folder to a Teams channel conversation. You can't drag and drop a message to a personal or group chat and the feature isn't available in OWA or Outlook for Mac.

To get the message into Teams, Outlook uploads a copy of the message (as a .msg file) into the channel folder in the SharePoint site belonging to the target team and creates a link to the email in the Teams message. The user can then add extra context for the message, just like they would for any other attachment shared in a channel before posting. Users can also drag and drop messages from Outlook to the Files channel tab. This action uploads the message to SharePoint without creating a message in the channel.

SharePoint Online stores messages as .msg files. These are copies that preserve the structure of the messages. The viewers included in Teams and SharePoint Online only display the text of the message and don't support access to any attachments. If users want to see the attachments, they must download a copy of the .msg file and open it with Outlook.

Although Outlook can upload messages protected with sensitivity labels (or S/MIME or another protection mechanism) to Teams, users can't read the content unless they download the message and open it with Outlook. When this happens, Outlook can call the protection mechanism applied to the message to access its content. For example, if a sensitivity label with encryption protected the message, Outlook can obtain the necessary use license, check if the user has the necessary rights to view the content, and if so, decrypt and display the message.

Another way of handling protected email is to copy the decrypted text from Outlook and paste it into a Teams message. If you want to include the message header to show recipients, forward the message to someone (but don't send it) and copy the text inserted into the forwarded copy. Any attachments (which will also be protected) must be downloaded and posted to Teams separately.

# Chapter 13: Managing Teams

*Tony Redmond*

## Keeping Teams in Good Shape

Applications require management if they are not to fall into a state of utter chaos over time. Teams is no different. Here we discuss the tools available to manage Teams from the policies managed through the Teams admin center to day-to-day management tasks like creating and removing teams. We also cover guest user access to Teams and the compliance and auditing infrastructure that helps to manage the content stored in Teams.

## Creating a Deployment Plan for Teams

Before describing different aspects of Teams management, it's worth listing the areas to consider in a deployment plan. Here are some things to consider:

- What **training** will you deliver to users? How will you keep that training updated to match what users see in the Teams clients?
- What training is available for administrators and team owners? How will you keep the training up-to-date with new developments in Teams and Microsoft 365?
- Are the default settings for messaging, calling, and meeting **policies** appropriate, and if not, what changes are necessary? What method will the organization use to assign policies and what policies will different sets of users receive?
- Will you allow all users to **create new teams**, or will the organization impose control over team creation? If so, how can users request the creation of new teams, and who will approve their creation?
- Will you impose a **naming convention** for Teams? If so, what is an appropriate and useful naming structure?
- Are **org-wide** or similar large-scale Teams needed? If yes, who will manage these teams, and what is their purpose?
- Will you use the **Groups expiration policy** to control the retention of inactive Teams? This feature requires Azure AD Premium licenses, but it can be very helpful when managing large numbers of teams. Another policy to consider is the ownerless groups policy.
- Will Teams **affect how you use other applications** like SharePoint Online or Exchange Online? For instance, if you use a hybrid Exchange organization, do you need to upgrade on-premises servers or should you move some mailboxes to Exchange Online?
- Will you allow **external access** to Teams? If yes, will you use the Azure B2B Collaboration (guest accounts) or Azure AD B2B Direct Connect (for shared channels)? How will you manage external access to confidential information held in teams? Will you require owners to validate the need for continued external access to their teams periodically?
- Do you need to migrate **meetings and voice applications** from other platforms? Will this affect the devices you use in meeting rooms or on user desktops? Do you need to invest in calling plans to allow users to make PSTN calls with Teams?
- How does Teams affect the **data governance framework** for the company? Will Teams be within the scope of data lifecycle policies like Data Loss Prevention, Retention, and communications compliance?
- What **first-party apps** will be used with Teams (Tasks, SharePoint, Insights, OneNote, etc.) and how will users be trained to use these apps effectively?

- What **third-party apps** will be used alongside Teams and by whom? How is access controlled to these apps and how do they feature in the data governance framework?
- Will **bots and connectors** be used with Teams? If yes, what purpose will these components serve?
- Will the organization invest in the development of apps, bots, and channel tabs for Teams? If yes, who will receive training and what training will they receive?

It's a long list that you should tailor to meet the needs of your organization. It's also important to realize that some of the defaults Microsoft chooses for Teams (such as allowing everyone to create new Teams) are there to accelerate the adoption of Teams instead of making long-term management easier. This is especially true for larger organizations where the number of Teams is more than a single administrator can remember. With that thought in mind, let's consider the details of how to approach the management of Teams.

## Teams Management

The [Teams admin center](#) (Figure 13-1) is the focus of team and policy management. Because the admin center is under active development, some of the settings described here might not be the same as presented in your tenant. The next chapter covers the settings for Teams Voice and the Phone system, including aspects such as calling plans, resource accounts, and holidays.

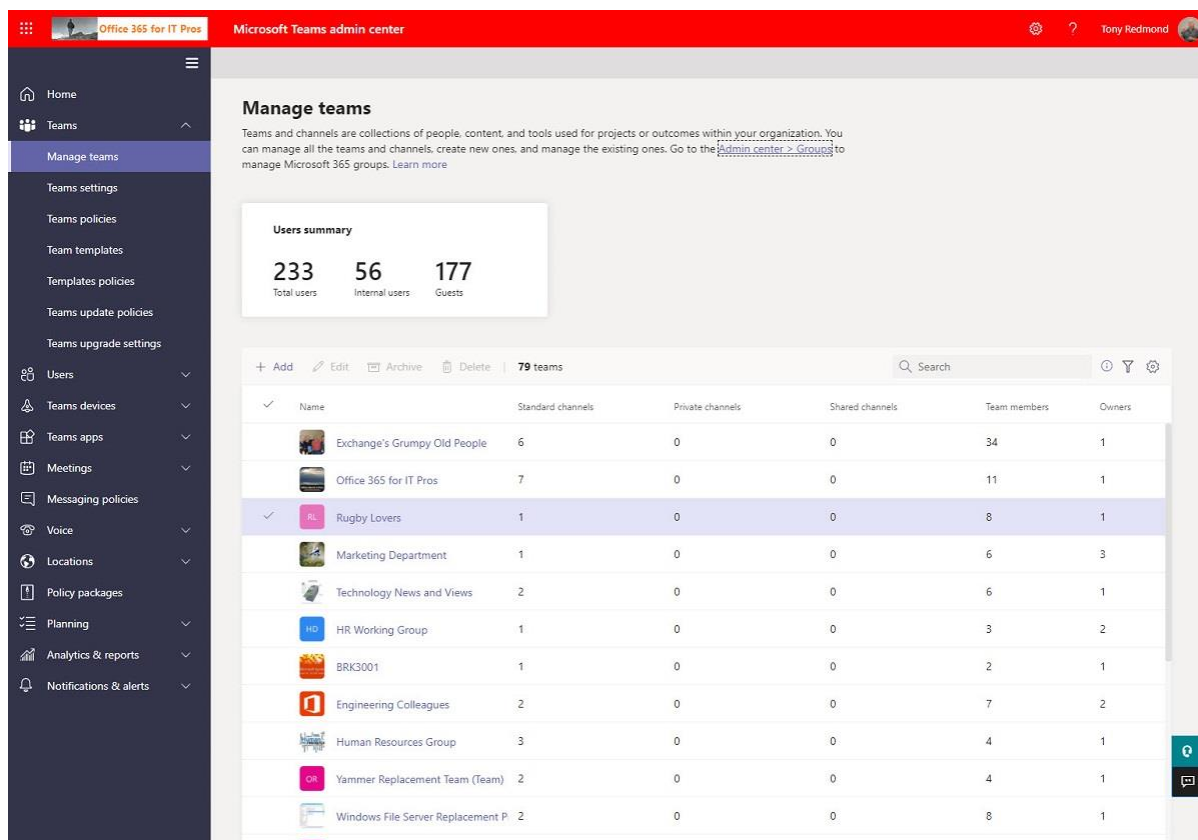


Figure 13-1: The Teams admin center

Not only does the admin center allow tenants to manage the creation and maintenance of individual teams, but it also supports the management of a wide range of user and organization policies to control how different aspects of Teams work, such as messaging, meetings, and guest access. As explained later, while tenant administrators have full access to all the settings and policies, a set of Teams administrative roles is available for assignment to users to grant them the ability to manage specific parts of Teams. See Chapter 4 for more information on these roles.

**Teams in the Microsoft 365 admin center:** If you look at the Groups section of the Microsoft 365 admin center, you can apply the *Groups with Teams* filter to see the list of Groups in the tenant that are teams-enabled (a small Teams icon also appears in the Teams status column). You can edit details of teams such as description, or membership through the Microsoft 365 admin center, but you need to use the Teams admin center or PowerShell to update team-specific settings.

The **Admin** app is available in the Teams app store. Microsoft created the app to help administrators of small to medium Microsoft 365 tenants manage Teams more easily. In these tenants, administrators often take care of Microsoft 365 as a whole, whereas in larger enterprises administrators might be more specialized. Also, in large enterprises, it's common to see the principle of account separation at work, where people use permissioned administrative accounts to manage Microsoft 365 and personal accounts to interact with Exchange Online, SharePoint Online, OneDrive for Business, and Teams. The app allows administrators to manage user accounts, teams, licenses, and some organization-wide settings for Teams. Although the Admin app is convenient for a limited set of tasks, management through the Microsoft 365 admin center and Teams admin center is often more effective.

## How Teams Global Policies Work

Microsoft creates a default policy called *Global (org-wide default)* for most of the policies used to manage teams. The purpose of a default policy is to dictate how a certain aspect of Teams works, like messaging, meetings, or app setup.

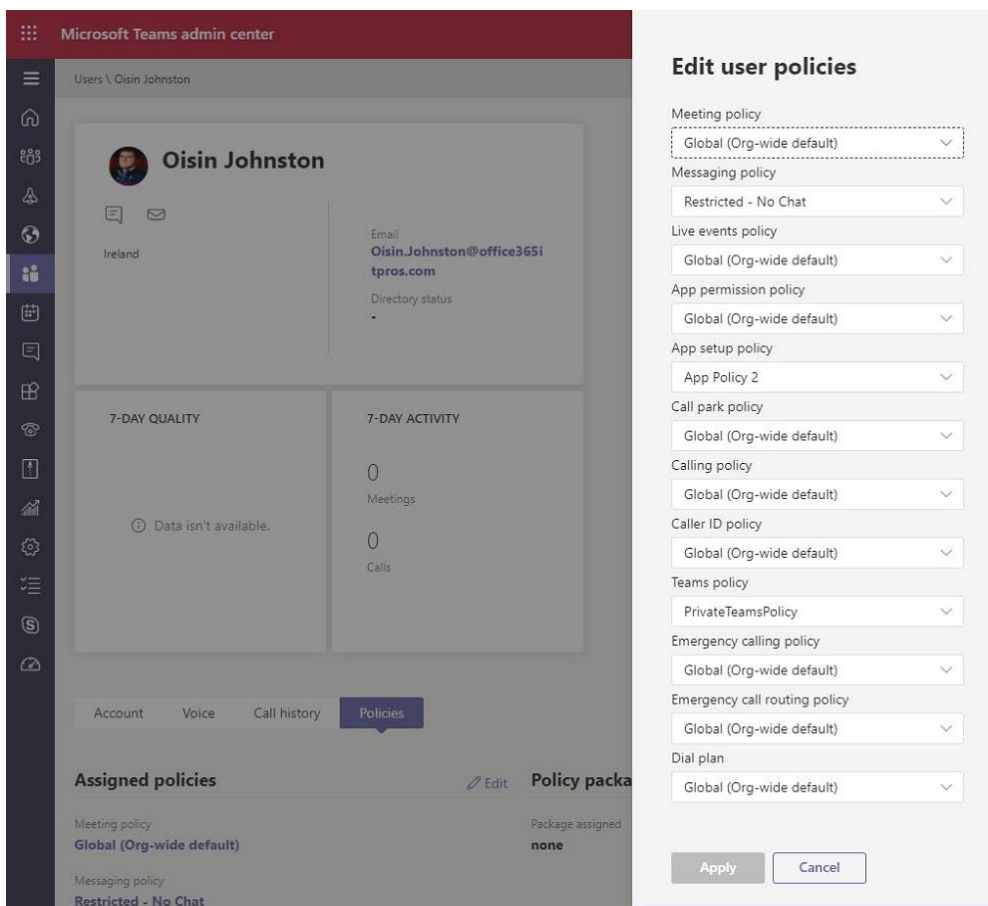


Figure 13-2: Editing the policies assigned to a Teams user

The default policies apply to all users until the organization creates a custom version of the policy and assigns that policy to some or all accounts. This mechanism allows a tenant to manage functionality on a user-by-user basis. Although you can edit the default policies to reflect the way that you want Teams to work for your organization, some administrators prefer to leave the default policies intact and create custom versions of

Teams policies instead. They take this approach because they believe this gives them more control and avoids problems if Microsoft changes something in a default policy in the future. However, if you use custom policies, remember that you must apply those policies to new accounts after creation. This is one reason to depend on the default policies whenever possible.

You can manage Teams with default policies applicable to all users or create policies with different settings and assign them to specific users using the Teams admin center or PowerShell, including batch policy assignments to process large sets of users. Individual accounts have a mixture of default and custom policies. In Figure 13-2 we see a user account assigned nine default policies and three custom policies.

Changes made to user policies need to synchronize with Teams clients before they are effective. Normally, this process takes an hour or so, but it can take longer.

## Microsoft Management of Default Policies

Microsoft manages default versions of Teams policies to allow them to easily change policy settings to reflect to support the introduction of new functionality. A tenant uses the Microsoft versions of default policies until the first time an admin views or updates a default policy through the Teams admin center or PowerShell.

When this happens, Teams copies that default policy to the tenant to allow the tenant to use a custom version of the default policy. From that point forward, any change Microsoft makes to the global policy will not affect the tenant because the custom version of the policy takes precedence.

As an example, Microsoft updated the default meeting policy to force external users to wait in the lobby before admittance to Teams meetings. They did this by changing the *AutoAdmittedUsers* setting to "EveryoneInCompany" and the *AllowPSTNUsersToBypassLobby* setting to False. The net effect was to allow only tenant users to join meetings without going through the lobby. Microsoft updated its copy of the default meeting policy, and the change was picked up by tenants that had never accessed the default meeting policy. A tenant-specific version of a default policy always takes precedence over Microsoft's version. Therefore, Microsoft's update did not affect tenants who had their copy of the default meeting policy and any customizations applied by a tenant remained in force.

## Teams Policies

Teams policies (which sound much more comprehensive than they are) control the enabling or disabling of complete features such as private and shared channels. The default policy enables all features to make functionality available to users immediately after Microsoft releases updated software. If you prefer a more phased approach that takes training and help desk preparation into account, you can disable features in the Teams policy until the organization is prepared for their introduction. Microsoft often exposes the setting to control a new feature sometime before releasing the feature to users, which means that administrators have the time to assess if they want to make the feature available.

To assign a policy to an account, choose the account, and then select the Policies tab. You can then edit the set of assigned policies. If the organization uses policy packages for policy assignment, you can select a policy package to apply to the account.

## Messaging Policies

New users automatically receive the tenant's default messaging policy. The messaging policy controls the actions a user can take when working with Teams chat and channel messages and can override team-specific settings. If you don't want to use the settings in the global policy, you can change them or can create new messaging policies through the *Messaging policies* section of the Teams admin center, which also includes the ability to assign different policies to specific users. For example, you might want to restrict the ability of some users to send voice messages. To achieve the goal, define a new messaging policy that disables voice memos and assign it to accounts you want to restrict.

A messaging policy includes the following settings:

- **Owners can delete sent messages:** Allow team owners to remove all messages from channel conversations.
- **Delete sent messages:** Allow users to delete messages that they sent in channel conversations.
- **Edit sent messages:** Allow users to edit the text of messages they sent to channel conversations.
- **Chat:** If you disable this feature, Teams hides the Chat app from the navigation bar and limits users to channel conversations. Removing chat takes a big piece of functionality from Teams (see below) and is not recommended.
- **Read Receipts:** This setting is enabled for everyone, turned off for everyone, or user-controlled. If user-controlled, users can choose to use read receipts in 1:1 and group chats (read receipts are limited to group chats with up to 20 participants).
- **Translate messages:** Controls if users can translate messages from their entered language to the language of the user.
- **Use Giphy in conversations:** Allow users to add [animated GIF files](#) to conversations. Clients make connections to the Giphy service to fetch available files. The choice to use GIFs links to the content rating, which can be set to "Strict" (nothing offensive to young audiences), "Moderate" (parental guidance needed), and "Allow all content" (all bets are off).
- **Use Memes in conversations:** Allow users to add the [memes](#) available to Teams to personal or channel conversations.
- **Use Stickers in conversations:** Allow users to add the stickers available to Teams to channel or private conversations.
- **Voice message creation:** Controls if users can create voice messages in chats, and channel conversations.
- **Allow URL preview:** By default, clients create previews of the content pointed to by URLs. Turn this setting to *Off* if you don't want Teams to do this.
- **Allow immersive reader for viewing messages:** By default, people can view messages in Microsoft Immersive Reader. When the tool is used, Teams displays message text in a full-screen window. The text is enlarged to make it easier to read and the reader can also choose to have the text read for them in a male or female voice.
- **Remove users from a group chat:** Controls if users can remove another person from a group chat.
- **Send urgent messages using priority notifications:** Controls if users can send priority messages. The number of these messages that a user can send is limited by the license assigned to their account.
- **Create voice messages:** Controls if users can post voice messages to chats and channels, chats only, or not at all.
- **On mobile devices, display favorite channels above recent chats:** If set, this setting forces the iOS and Android clients to display favorite channels on the top of the screen.
- **Suggested replies:** Teams can use AI to analyze the context of a message and generate up to three appropriate responses. This feature is available only in 1:1 chats.
- **Chat permission role:** This setting applies when the organization is configured for supervised chat, usually in an educational environment.

Organizations will communications compliance policies can use the **Report a concern** setting (for PowerShell, this is the `AllowCommunicationComplianceEndUserReporting` setting) to allow users to report messages they believe to violate corporate policies. See Chapter 21 for details about communications compliance policies. By default, the setting is True.

No control is available over the ability of users to insert emojis into channel conversations or personal chats.

## Disabling Chat

Disabling Chat through a messaging policy removes access to the Chat app throughout Teams, but this approach is not recommended. Disabling the Chat app means that users won't be able to:

- Collaborate with other Teams users on an ad-hoc basis, including both personal (1:1) and group chats. Losing the ability to use group chats means that people won't be able to discuss issues and resolve points among a small number of people before bringing the issues for a wider general review in a channel conversation.
- Use federated connections to Teams users in other tenants or Skype consumer users.

Guest users don't have accounts, so they cannot be assigned Teams policies, so if you want to stop guests from using chat between themselves or with tenant users, you must disable Chat in the Guest access section in the Users settings.

Although some feel that it is a good thing to remove chat to force users to conduct conversations in channels, the experience of Teams deployments indicates that users are more likely to seek other options for personal conversations outside Teams, such as WhatsApp. Moving these conversations outside Teams impacts the effectiveness of the organization's compliance policy because chats are then not captured and available for eDiscovery.

## Disabling Meeting Chat

Meeting chats are different from personal chats. A meeting chat is associated with a Teams meeting and is governed by the meeting policy assigned to the participants. The *Allow chat in meetings* setting is usually enabled to permit users to join the meeting chat. If disabled, users cannot chat with other participants during the meeting. This is a big loss because chat is an excellent way for participants to post questions and responses during presentations, or for meeting presenters to share additional information (such as the URL to a web page) with participants.

## Meeting Policies

The meeting policy assigned to a user account controls what that user can do in meetings. The General policy is the default, but you can create as many other policies as you like in the *Meeting policies* section under *Meetings*. Some settings apply to meeting organizers (like who can present in a meeting) and allow organizers to change how meetings work through meeting options. Other settings apply on a per-user basis to control the functionality available to individual accounts when they join meetings.

The settings in a meeting policy are in the following groups:

- **General:**
  - **Allow Meet Now:** Users can create ad-hoc private meetings within a channel.
  - **Allow the Outlook add-in:** Allow users to schedule Teams meetings via the Outlook add-in for Teams.
  - **Allowing channel meeting scheduling:** A channel meeting is available to anyone who can access the channel. For public teams, this means that anyone in the tenant can join the meeting.
  - **Allow scheduling private meetings:** Users can create private meetings, meaning that only invited attendees can access the meeting.
- **Audio and Video:**
  - **Allow transcription:** Allow users to turn on automatic transcription generation for a meeting.
  - **Allow cloud recording:** Allow Teams to record meetings. Teams stores MP4 files for meeting recordings in OneDrive for Business and makes the recordings available to meeting attendees after meetings finish. A separate setting in calling policies controls if users can record 1:1 calls.

The *AllowCloudRecordingForCalls* setting is Off (*\$False*) by default and can only be enabled (set to *\$True*) in PowerShell by running the *Set-CsTeamsCallingPolicy* cmdlet.

- **Allow IP video:** The default is to allow video meetings. If you want to restrict meetings to audio, move this switch to off. Individual users have the choice to restrict their participation to audio when they join a meeting and might choose to do so if the network connection is poor, or they simply don't want other participants to see them in all their glory.
- **Allow NDI streaming:** Allow the user to employ NDI technology where the video for each participant becomes a discrete video source processible by video production software to create an output.
- **Media bit rate:** The default is 50000 (50 MB). Do not change this value unless you have good reason to do so.
- **Content Sharing:**
  - **Screen sharing mode:** Teams allows users to share content from their PC to meeting attendees on the entire screen or as a single application. If you want to prevent this, disable the setting. The default is "Entire Screen."
  - **Allow participant to give or request control:** Allow a meeting participant to give control of a meeting to another participant, or request control if they don't already have control.
  - **Allow an external participant to give or request control:** An external participant is a guest user or an anonymous user who joins a meeting. The setting is *On* for the default policy, but *Off* for other meeting policies that you create.
  - **Allow PowerPoint sharing:** Allow meeting participants to share PowerPoint presentations. Sharing a PowerPoint presentation is an alternative to sharing a complete desktop and consumes less bandwidth when network resources are scarce. See [this page](#) for more information.
  - **Allow whiteboard:** Allow meeting participants to use the whiteboard app.
  - **Allow shared notes:** Allow meeting participants to share notes during the meeting.
  - **Select video filters:** Allow users to select from no filters, background blur, curated background images, or custom images they upload.
- **Participants and Guests:**
  - **Let anonymous users start a meeting:** This setting allows anonymous users (people unknown to the organization) to start a meeting if they are the first to join. The default is *Off*.
  - **Roles that have presenter rights in meetings:** This setting controls who can act in the presenter role during a meeting. The usual value is *Everyone, but user can override*, which means that meeting organizers decide who can present.
  - **Automatically admit people:** This setting controls whether people can join meetings automatically. The default is *Everyone in your organization*, which means that tenant users can join meetings without going through the meeting lobby. You can also set it to *Everyone*, which then allows external users to join without having to wait for admittance in the lobby.
  - **Allow dial-in users to bypass the lobby:** Allow meeting participants who join via phone to enter the meeting without pausing for admittance in the lobby.
  - **Allow meet now in private meetings:** Allow users in personal chats to meet now and create a video or audio meeting.
  - **Enable live captions:** If enabled, Teams generates captions for the conversation.
  - **Allow chat in meetings:** If enabled, meeting participants can chat during the meeting. You can enable or disable chat for all participants or allow all participants to send messages except those who join anonymously. Blocking the ability of anonymous participants to send messages doesn't prevent them from reading the meeting chat.
  - **Meeting reactions:** Set the toggle on to allow participants and guests to use reactions (like applause) during meetings.



## Updating Meeting Policies with PowerShell

In some cases, Microsoft introduces meeting policy settings to support new functionality that is managed through PowerShell. This is usually an interim situation while Microsoft updates the Teams admin center UI. The `Set-CsTeamsMeetingPolicy` cmdlet updates meeting policy settings. This example enables settings in the default policy to:

- Allow meeting chat for everyone except anonymous participants.
- Allow the download of meeting engagement reports.
- Sets the languages used to generate meeting invitations to US English and German (user language preferences cannot override this setting because it's applied on the server when [Teams generates the joining instructions for meetings](#)).

The Global policy applies to Teams users unless their accounts are assigned other meeting policies:

```
[PS] C:\> Set-CsTeamsMeetingPolicy -Identity Global -MeetingChatEnabledType EnabledExceptAnonymous -AllowEngagementReport "Enabled" -MeetingInviteLanguages "en-US,de-DE"
```

## Meeting Customization Policies

Meeting customization policies control features that require users to have the [Teams advanced communications add-on](#). A customization policy allows an organization to define one or more lobby images for meeting organizers to choose from when they set up a meeting. People in the lobby then see the selected image, which is often a corporate logo. The other feature controlled by a meeting customization policy is the ability to upload corporate images for people to use as backgrounds during meetings. Corporate images appear before Microsoft's default set of meeting backgrounds and the custom images uploaded by a user.

## Custom Teams Meeting Invitations

Organizations can customize the emails generated for Teams meeting invitations. Open the Meetings section of the Teams admin center and navigate to *Meeting settings*. You can now customize four components used by Teams to create email invitations for meetings:

- **Logo URL** points to a network-accessible JPEG or PNG file that Teams inserts into meeting details. [Microsoft's documentation](#) recommends using an image that's no more than 188 pixels wide by 30 pixels tall. I have used considerably larger images, but it's best to stay somewhat close to the recommended size to avoid bloating Teams meeting invitations with unnecessary graphic content.
- **Legal URL** points to a web page containing information users need to know about Teams meetings. For instance, the page might explain why Teams displays reminders when an organizer records a meeting.
- **Help URL** points to a web page containing information about Teams meetings. For example, the page could explain that Teams meeting recordings are subject to an automatic expiration policy.
- The **footer** is free-form text containing whatever words of wisdom seem appropriate.

After completing the meeting settings, you can test what the customized body of a Teams meeting invitation looks like by clicking the *Preview invite* button. The view presented is only an approximation of what users see. The actual format depends on the client in use, but it's close enough to understand what people might see.

If meeting organizers update the details of a meeting (like setting a new time), the meeting invitation sent by Teams does not include the organization logo.

## Feedback Policy

By default, users can submit feedback to Microsoft using the **Give feedback** option under Help. Microsoft can also prompt users to take periodic surveys to tell Microsoft how they're getting on with Teams. Many

organizations dislike the idea of users providing direct feedback as they consider this to be a function of the organization after gathering opinions across the entire user base. In these circumstances, you can disable the ability of users to give feedback or participate in surveys by:

- Amend the default (global) feedback policy. This approach has the advantage that it covers every user, including new accounts. If you decide to allow some users to submit feedback, you can create and assign a feedback policy that allows these options.
- Assign the Disabled feedback policy provided by Teams.
- Create and assign a new custom feedback policy to disable the options. The advantage of this approach is that the tenant controls the policy settings. The downside is that you must remember to assign the custom policy to new accounts after their addition.

These actions must be performed with PowerShell because the Teams admin center doesn't include the ability to manage feedback policies. To create a new feedback policy, run the *New-CsTeamsFeedbackPolicy* cmdlet:

```
[PS] C:\> New-CsTeamsFeedbackPolicy -Identity "Tenant Disabled Feedback" -UserInitiatedMode Disabled -ReceiveSurveysMode Disabled
```

*UserInitiatedMode* controls the visibility of the **Give feedback** option while *ReceiveSurveysMode* controls if users see the Microsoft surveys. Unfortunately, bulk policy assignment (which we will cover soon) does not support feedback policies, so you must assign the new policy individually using the *Grant-CsTeamsFeedbackPolicy* cmdlet. These commands read in a set of user principal names stored in a CSV file, assign a policy to each user, and check that the target accounts have the expected feedback policy.

```
[PS] C:\> Users = Import-CSV c:\temp\Users.csv
# Assign the Disabled feedback policy to all users loaded from the CSV file
ForEach ($User in $Users) { Grant-CsTeamsFeedbackPolicy -Identity $User.UserPrincipalName -Policy "Tenant Disabled Feedback" }
# Check that the assignment works
Get-CsOnlineUser | ? {$_.TeamsFeedbackPolicy -eq "Tenant Disabled Feedback"} | Format-Table DisplayName, TeamsFeedbackPolicy
```

## Assigning Policies

Several methods exist to assign Teams policies to user accounts, including:

- Individual assignment by updating the policies section of user accounts in the Teams admin center.
- Bulk editing by selecting a set of users in the Teams admin center and assigning the same policies to the selected accounts.
- Running PowerShell cmdlets like *Grant-CsTeamsMeetingPolicy* to update individual or multiple accounts. Companies often configure Teams policies for users as part of their account provisioning process.
- Using a bulk policy assignment to assign a policy to up to 5,000 accounts. This is a good method to assign the same policy to many accounts at one time and works by defining a set of accounts to receive a policy and submitting a batch job to perform the assignments using the *New-CsBatchPolicyAssignmentOperation* cmdlet.
- Group policy assignment. This is an effective way to assign a set of policies to a large population of target accounts. The membership of a distribution list, security group, or Microsoft 365 group defines the target accounts. Through the Policy packages section of the Teams admin center, the group is associated with a policy package, a set of up to ten different Teams policies including meeting, message, app setup, voice routing, and calling. The association of the policy package with the group causes Teams to assign the individual policies specified in the package to the user accounts defined in the group. Microsoft includes a set of predefined policy packages with Teams. The predefined policy packages are unchangeable, but the details of the policies contained in the predefined packages are editable. Organizations can define custom policy packages using tenant-specific policies if they have

Teams advanced communications add-on licenses for every account assigned policies through these packages.

Direct policy assignment occurs when an administrator assigns a policy to an account by editing the account properties. Direct assignments take precedence over assignments inherited through group membership. In addition, an account might inherit policies through multiple groups. When this happens, the priority order for the policy package dictates which policy becomes effective.

Given the range of options available to manage Teams policies, most companies should be able to find suitable methods to assign policies to individuals or sets of users.

## Bulk Policy Assignment

Although it's easy to write PowerShell scripts to process policy assignments for users, using bulk assignment jobs is the most efficient way to assign policies to large groups of users. Here's an example of the bulk application of a policy (in this case, a meeting policy) to a set of users. First, we find a set of mailboxes based on a value in a custom attribute and extract the user principal name for each account. The data is stored in an array variable.

```
[PS] C:\> $Users = Get-ExoMailbox -RecipientTypeDetails UserMailbox -Filter {CustomAttribute14 -eq "Code22"} -ResultSize Unlimited | Select -ExpandProperty UserPrincipalName
```

Although a maximum of 5,000 accounts can be included in a single operation, it's wise to divide processing up across smaller batches and run a single batch at a time. This approach allows you to take corrective action if errors are encountered for some reason. Another way of managing the set of users for a bulk assignment is to use a CSV to collect the account information. If you do, make sure that the CSV file only contains user principal names. When you're ready to import the data, create an array and populate it with the user principal names from the CSV. For instance, this snippet creates a suitable variable and populates it with the imported data:

```
[PS] C:\> $TargetUsers = Import-CSV c:\temp\UsersToProcess.CSV
$Users = [System.Collections.Generic.List[Object]]::new()
ForEach ($U in $TargetUsers) { $Users.Add($U.UserPrincipalName) }
```

Next, create a policy assignment operation linking the policy to apply and the set of users to process:

```
[PS] C:\> New-CsBatchPolicyAssignmentOperation -PolicyType TeamsMeetingPolicy -PolicyName
RestrictedFunctionality -Identity $Users -OperationName "Teams Meeting Policy Assignment"
2d5c6c07-bc19-4c6c-a359-2f6530b3a652
```

As you can see, Teams responds with an operation identifier (a GUID). This can be used to find the status of the operation:

```
[PS] C:\> Get-CsBatchPolicyAssignmentOperation -OperationId 2d5c6c07-bc19-4c6c-a359-2f6530b3a652 |
Format-List

OperationId      : 2d5c6c07-bc19-4c6c-a359-2f6530b3a652
OperationName    : Teams Meeting Policy Assignment
OverallStatus    : NotStarted
CreatedTime      : 3 Jun 2020 17:13:46
CreatedBy        : 53f08764-07d4-418c-8403-a737a8fac7b3
CompletedTime    :
CompletedCount   : 0
ErrorCount       : 0
PendingCount     : 4
```

We can see that the background processor has not yet started to process the operation. It can take up to an hour before processing starts. We can determine which account submitted the operation by fetching its object identifier from the operation properties and running the Get-MgUser cmdlet to return the display name of the account:

```
[PS] C:\> Get-MgUser -UserId (Get-CsBatchPolicyAssignmentOperation -OperationId 2d5c6c07-bc19-4c6c-a359-2f6530b3a652).CreatedBy | Select DisplayName
```

```
DisplayName
```

```
-----
```

```
Global Tenant Administrator
```

Eventually, the operation will finish, and its status will change to *Completed*. If the *ErrorCount* is non-zero, we can check the state of each account processed with:

```
[PS] C:\> Get-CsBatchPolicyAssignmentOperation -OperationId 2d5c6c07-bc19-4c6c-a359-2f6530b3a652 | Select -ExpandProperty UserState
```

Id	Result	State
--	-----	-----
Ben.James@Office365itpros.com	Success	Completed
Imran.Khan@office365itpros.com	Success	Completed
John.Adams@office365itpros.com	Unknown error occurred.	Completed
Rene.Artois@office365itpros.com	Success	Completed

In this instance, a problem happened with the John Adams account. You'll need to check the account to see what might have caused the problem. It might just be a temporary glitch (or the account was already assigned the policy) and you can go ahead and make the change to the account manually through the Teams admin center. The "user not found" error is common and usually happens when a mistake is made with a property like a user principal name and the batch processor cannot find the account to update.

To discover who has been assigned a policy, you can use the *Get-CsOnlineUser* cmdlet. For example, this command reports the accounts assigned a Teams meeting policy called 'RestrictedFunctionality':

```
[PS] C:\> Get-CsOnlineUser -Filter {TeamsMeetingPolicy -eq 'RestrictedFunctionality'} | Select UserPrincipalName
```

To report users assigned a default policy, check for a null value in the meeting policy:

```
[PS] C:\> Get-CsOnlineUser -Filter {TeamsMeetingPolicy -eq $Null} | Select UserPrincipalName
```

**Reporting Teams Policy Assignments:** After you're finished assigning policies to Teams users, you might like to generate a report showing what policies apply to what accounts. Teams doesn't include a report of this nature, but it's easy to generate with PowerShell. [This article](#) explains how to write a script to create a Teams policy report output in both HTML and CSV formats.

## User Settings

User settings cover the management of users, settings for guest access, and controls for external access (federation). User management deals with the assignment of Teams policies to user accounts.

### External Access

**External access** defines how users communicate with people outside the tenant. The settings are:

- **Teams and Skype for Business users in external organizations** – This setting needs to be *Allow all external domains* (open federation – the default) or *Allow only specific external domains* (limited federation) before people can communicate with users in other organizations. See the later section covering how external access works.
- **Skype users** – Turn *On* to allow Teams to communicate with Skype users.

In addition, you can define a list of domains to block or allow for external communication. If you only want users to communicate with people in specific domains, define those domains here, or block the domains that you don't want people to communicate with. Leave the list blank if you are happy for people to communicate with users in any other Teams organization.

## Guest Access Settings

**Guest access** settings control what a guest user can do within Teams. The settings include enabling or disabling guest access for the tenant, allowing guests to use personal chat, and settings to control the interaction guests have with messages like the ability to edit or delete sent messages. Settings are available to control aspects of guest participation in meetings, such as allowing guests to use the Meet Now feature to create impromptu meetings. Microsoft enables guest access for Teams for new tenants.

## Teams Settings

**Teams settings** are organization-wide settings to control how users interact with Teams for aspects not covered by messaging and meeting policies. These settings include:

- **Notification and Feeds:** Define if you want Teams to add suggested items in users' activity feeds.
- **Tagging:** Define who can manage tags, suggested tags that appear in all teams, and whether users can create custom tags.
- **Email integration:** Settings to control if people can send email to channels (see later section later) using special email addresses generated by Teams for channels. If you enable email integration, you can also confine the ability of people to send email to channels to an allowed list of SMTP domains.
- **Files:** Settings to control how users can share files, including whether they can share files using third-party cloud storage solutions (ShareFile, Box, DropBox, and Google Drive).
- **Organization:** Turn the organization tab on or off. If enabled, the tab allows Teams users (but not guests) to see details of the organizational hierarchy represented by the manager-employee connections stored in Azure AD.
- The **Devices** settings control how Room Systems (like a Surface Hub) and IP phones work in meetings. The settings are:
  - Require a secondary form of authentication (the default is no access).
  - Set content PIN. The default is that a PIN is needed for outside scheduled meetings.
  - Surface hub accounts can send messages. Some devices such as Surface Hub use device accounts to interact with real users through email, IM, and calls. By default, this setting is True, which allows devices to send IM messages.
- **Search by Name** defines whether directory searches performed by Teams users are limited by Exchange address book policies (ABPs). An address book policy limits the visibility of a user to specific sections of the overall corporate address book, such as only the users in a certain department or country. Policies are often used to limit communications between different groups like students and teachers or traders and financial advisors. The default for the directory search setting is Off, meaning that no restrictions are placed on users. If set to On, the user is limited to whatever address book policy is defined for their mailbox (if no policy is defined, they can see and communicate with everyone in the directory). For example, they can only start a personal chat with another user who is visible to them according to the address book policy. In addition, if you use address book policies to limit the directory search for users, Teams is no longer able to suggest teams for them to join (and disables the /Join command). Limiting directory searches for Teams is a prerequisite for Information Barriers (discussed later). See Chapter 3 in the Companion Volume for more information about creating and using address book policies.
- **Safety and Communications:** This setting controls if supervised chat is enabled for the organization. The most common use of supervised chat is in education settings where a teacher or responsible adult oversees chats between students.

## Planning

The Planning section of the admin center covers:

- Teams advisor: This is a wizard-based guide to help organizations who do not currently use Teams roll out different areas of functionality. See [this page](#) for more information. Among the topics covered by the advisor are:
  - Chat, messaging, and apps.
  - Meetings and conferencing.
- Network planner: See the Calling and Devices chapter for more information about planning your network for Teams and the Microsoft Phone system (cloud voice).

## Licensing Teams

The license necessary to access Teams is bundled in many Microsoft 365 offerings, with the notable exception being standalone subscriptions, like Exchange Online Plan 2. These users can't even join teams as guest users because Microsoft 365 considers them to be part of the same organization that hosts the teams and therefore aren't guests.

You can selectively enable or disable Teams on a per-user basis by editing a user's account in the Microsoft 365 admin center and switching their Teams license on or off. If you need to enable or disable Teams for many accounts, it's easier to do this with PowerShell (see the license management section in the PowerShell chapter). To ensure that the widest range of functionality is available to users, make sure that you enable Exchange Online and SharePoint Online for anyone using Teams. Users who do not have a license for SharePoint Online cannot use OneDrive for Business to share files with users in personal or group chats.

### Teams Licenses for Guest Users

Guest users for Teams are licensed based on five guest users per licensed tenant user. In other words, if your tenant has 100 licensed users, it can support up to 500 guest users.

If you don't want a guest user to access Teams in your tenant any longer, you should remove their account from Azure AD. Removing the account has the side-effect of removing their access from any other application that depends on Azure B2B Collaboration, such as Planner. It also removes any sharing access the account has been granted to documents in SharePoint Online or OneDrive for Business libraries.

## Access for On-Premises Users

If a tenant uses AAD Connect to synchronize accounts with an on-premises deployment, users can create and access Teams even if their mailboxes are on Exchange on-premises servers. These users can take part in conversations, group chats, and calls and set up new tabs. However, they cannot create or change the connectors associated with a tab, nor can they create or access meetings or even change their profile picture.

## Creating Teams

Before anyone can create a new team, they must be able to create a new group. Some tenants allow any user to create a group while others control group creation to a limited set of authorized users. The settings to control who can create new groups (and teams) are in the Azure AD policy for Groups. Chapter 11 includes a full discussion of how to manage policy settings using PowerShell. If the policy restricts the ability of users to create new groups, it should also include a pointer to a group holding a list of people who can create new groups. If your account is on that list, you can go ahead and create a new team. If not, you will see an error informing you that *"Your IT department has disabled this Microsoft Teams feature for you"* together with a suggestion to contact the IT department for help.

The recommended methods to create new teams are:

- Teams app (desktop, web, or mobile).

- Teams admin center.
- *New-Team* cmdlet in the Teams PowerShell module.
- Teams Graph API. This includes methods such as “teamifying” a SharePoint Online site (create a new team for an existing site using the SharePoint browser UI).

The advantage of these methods is that they explicitly mark a group as intended for use by Teams. In technical terms, this means that several properties of the group are set for Teams. These include:

- The group is hidden from Exchange clients like Outlook and OWA (*HiddenFromExchangeClientsEnabled* is *\$True*).
- The group is hidden from Exchange address lists like the GAL and OAB (*HiddenFromAddressListsEnabled* is *\$True*).
- New members receive a welcome message for Teams instead of Microsoft 365 Groups.
- Members are not subscribed to receive calendar updates for the group. This means that team members do not receive invitations to meetings scheduled in the group calendar (channel meetings). Some organizations prefer that team members receive invitations. If this is the case, you can update the *AlwaysSubscribeMembersToCalendarEvents* property for the group to *\$True*.

You can also create a new group and add members to it through the Azure AD admin center, OWA, the *New-UnifiedGroup* cmdlet, or any of the other methods available for group creation, and then team-enable the group. If you take this approach, consider updating the group settings listed above to make the group work like those created by Teams.

## Automate Teams Creation

Many organizations restrict the creation of teams and require users to go through a process to seek approval for a new team and use methods like a Power Automate flow or PowerShell script to automate the collection of details for the new team, gaining approval, and creating the team if approved. Microsoft has a sample Power Apps application (called [request-a-team](#)) to automate the provisioning of new teams. You can [download its source code from GitHub](#). Whatever method you use, remember to:

- Update team properties like hidden from address lists, and hidden from Exchange clients, just like a team created by Teams would be.
- Add a description so that people will remember why the team exists.
- Set the privacy of the new team to Public or Private.
- Assign a classification or sensitivity label to the team (assigning the label will set the privacy).
- Ensure that the new team has at least one owner.
- Populate the initial membership.

## Creating a Team in the Teams Client

To create a new team in the desktop or browser client, click **Teams** in the navigation bar and then **Join or Create a Team** at the bottom of the screen. You can now choose to create a new team or join one of the teams suggested by Teams. To create a new team, you can:

- Create a new team from scratch.
- Use an existing group or team as a template.
- Use a template created by Microsoft.

You can also team-enable an existing modern SharePoint Online site (one with a Microsoft 365 group) from the SharePoint browser interface.

The Teams mobile client only allows you to create a new team from scratch. You can't create an org-wide team or team-enable an existing group using the Teams mobile client.

## Creating a New Team from Scratch

Creating a team from scratch means that you have maximum control over its characteristics. The first step is to choose what kind of team to create. A team can be:

- **Private:** The team owners and tenant administrators control membership and decide who can join the team. This kind of team isn't discoverable by users.
- **Public:** Anyone in the organization can join the team.
- **Org-wide:** Every licensed account in the organization automatically becomes a member of the team. Like other teams, org-wide teams can support up to 10,000 members. If your tenant exceeds the membership threshold or five org-wide teams already exist, you won't have the choice to create org-wide teams and should consider an alternative for org-wide communications (Yammer, SharePoint, etc.). Only tenant global administrators can create org-wide teams.

The basic details for the new team are:

- **Team Name:** This is the display name for the team (and the underlying group) and it is what appears in the list of teams in the Teams client and the Exchange GAL (if you decide to reveal Teams to Exchange). It is always best when you give a team a meaningful name that conveys its purpose. If your organization uses a group naming policy, Teams shows the effect of the policy by displaying the name of the team after it applies the policy when it creates the new team. The group naming policy does not affect tenant administrators, so teams created by administrators keep the display name as entered.
- **Description:** This is a free-form text description of why the team exists and what the members of the team use it to do. Users see the team description when they browse teams, so it is important to put some thought into the description to make it accurate, interesting, and easily digestible. Get right to the point and put the essential information like the purpose of the group at the start and make sure that the first 80 characters convey the essence of the group because this is the amount of text displayed in search interfaces. Put less essential information at the end of the description, such as who created the team, the contact name for the team, and so on.

To indicate the level of confidentiality of the information contained in a team, it has a classification. This can be set by assigning a:

- **Sensitivity label:** if the tenant uses sensitivity labels for container management (see below), the team inherits the privacy level (public or private) and guest access (on or off) from the settings in the chosen label.
- **Classification:** Classifications are simple text description strings defined in the Azure AD policy for Groups. Classifications do not affect team settings.

Teams clients do not check that a group, distribution list, or team of the same display name already exists in the tenant, so it is quite possible to create multiple teams with the same name. This is one of the good reasons why some tenants restrict team creation to a small set of people, who have the responsibility to check that the display name does not clash before they create a team.

After entering the properties, click **Next** to continue. Teams creates the group object in Azure AD and Microsoft 365 provisions the resources used by the team, such as the SharePoint team site.

## Sensitivity Labels and Teams

If the tenant enables sensitivity labels for container management, Teams retrieves settings from sensitivity labels when creating new teams or when editing the properties of existing teams. When a team has a sensitivity label, the team uses the name of the label as its classification and applies the settings in the label



for privacy and guest access. If the team has any private or shared channels, the channels and the SharePoint sites belonging to those channels inherit the label assigned to the parent team.

If an administrator or group owner updates the sensitivity label assigned to a group or SharePoint site linked to a team, the updated label assignment synchronizes to Teams to apply the settings from the new label. The synchronization of label assignments from other workloads can take up to 24 hours.

Currently, the only way to assign a sensitivity label to a team is via a client. The Teams PowerShell module does not support sensitivity labels (cmdlets in the Groups and SharePoint modules can assign labels, which then synchronize with Teams) and you can't assign a sensitivity label to a team using the Graph API or when creating a team from a template.

Assigning a sensitivity label to a team does not affect the information stored in the team. See the Information Protection chapter for more information about sensitivity labels and container management.

## Create a Team from an Existing Team or Group

If you don't want to create a team from scratch, you can choose to create a new team based on the settings of an existing team (a template) or you can team-enable a group.

### Using an Existing Team as a Template

The idea behind creating a team using an existing team as a template is that if you are in the position where you need to create multiple teams with approximately the same structure, it's easier to create one team to use as a template and set the team up in a form that you want to replicate to other teams. You can only use a team as a template when you have access to that team (you're a member).

When you create a new team based on a template, you can replicate the membership (including guests), settings, channels, apps, and tabs from the template team. No content is copied across in terms of conversations or files, and some tabs will need to complete a setup process before they can be used. For instance, if the template team includes a tab pointing to Planner in a channel, the channel and the tab are copied to the new team, but because the plan doesn't exist for the channel, you must start Planner and create the plan.

### Enabling Existing Groups for Teams

In many ways, team-enabling a group is an interesting decision because it switches the focus for conversations in the group away from email to channel conversations. There is no way to migrate existing conversations from a group into channels, so the switch is a one-time all-in operation. However, the files stored in the group document library remain available to the team.

Whether you create a team from scratch or team-enable an existing group, you end up in the same place with a team-enabled group. The difference between the two methods is that team-enabling an existing group automatically makes the team available to the current group members because Teams and Groups share the same membership.

Groups can only be team-enabled by their owners. To begin, a group owner goes through the normal process to create a new team, but instead of creating a new team, they select **Create from an existing Microsoft 365 group or team** and then choose **Microsoft 365 group**. This option only appears when the user is an owner of one or more Outlook-based groups that are not yet team-enabled (Yammer-based groups are unsupported). Click the link to expose the set of groups owned by the user that are not team-enabled. (Figure 13-3). As the UI only shows only a small number of groups at a time, it might be necessary to scroll through the list to find the group that you want to upgrade. Among the list, you will find groups that are hidden from Exchange address lists (because Teams does not use address lists). However, groups with hidden membership do not appear. Select the group that you want to enable for Teams and click Choose team to begin the

process to populate the team resources and properties for the group. You cannot team-enable several groups at one time.

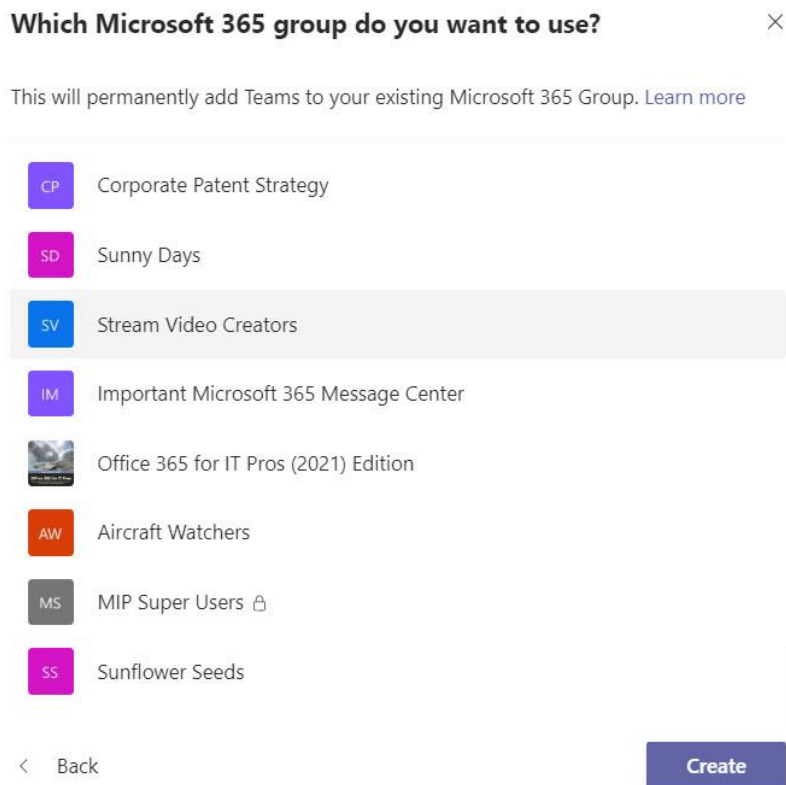


Figure 13-3: Selecting an existing group to create a new team

A group owner must decide to enable the group for Teams. There is no way for a user to search for a group that is not team-enabled and ask its owner to upgrade it. A group owner does not have to be permitted to create new groups to be able to team-enable an existing group. After you team-enable a group, you should consider hiding it from Exchange clients to force users to use Teams for conversations (see below).

## Creating Teams from Templates

Team templates are prepopulated structures created by Microsoft or the organization to make it easy to create teams to do a specific job. Each template consists of a set of channels, tabs, and apps automatically added to teams created using the template. Some of the 13 out-of-the-box templates are for general use (like project management); others (like Organize a Store) are for a particular industry. Among the templates published by Microsoft are:

- Adopt Office 365.
- Manage a project.
- Manage an event.
- Onboard employees.
- Organize help desk.
- Coordinate incident response.

Read [this documentation](#) for more information about creating teams from templates.

As an example of what happens when you create a team using a template, if you use the Manage a Project template, Teams adds four channels and two apps to the new team. The channels are General, Announcements, Resources, and Planning while the apps are OneNote and Wiki. During the creation process, you can rename the channels to make them more appropriate for the new team and add a sensitivity label to

reflect the importance of the information you expect the new team to contain. After creating the team, you still need to add team members, add other apps and channels (including private channels), post a welcome note, and take whatever other actions are necessary to build out the complete team.

Template management is in the Teams section of the Teams admin center. You can add, edit, or remove templates to meet the needs of the organization, including the ability to copy (duplicate) one of Microsoft's policies to create a customized version more suitable for your organization. Custom templates created by a tenant appear before the standard templates from Microsoft when shown to people creating new teams.

[PowerShell cmdlets to manage team templates](#) are also available.

## Template Policies

Template policies control the set of templates visible to users when they create new teams. The default global policy makes all templates visible. It's a good idea to consider creating custom team templates to deliver a more refined view to users. For instance, you could create a template that shows only custom templates and omits all the standard templates created by Microsoft. Template policies are managed in the Teams section of the Teams admin center. No PowerShell cmdlets are available to manage these policies.

Creating a new template policy is simple. Navigate to the Teams section in the Teams admin center and select Teams templates. Create the policy and give it a name and description. Then select the set of templates you want to hide from the available set. Click the *Hide* button to make the choice effective and then save the policy. The next step is to assign the policy to users individually or by using a bulk policy assignment job (the policy type is *TeamsTemplatePermissionPolicy*). Give the policy a couple of hours to become effective and then test it by asking one of the users assigned the policy to create a new team. They should then see the set of templates dictated by the policy.

## Adding Team Members

### Add members to Infrastructure and Technology Plans 2020 (Team)

Start typing a name, distribution list, or mail enabled security group to add to your team.

Figure 13-4: Adding members to a team

After Teams creates the new team, you can add members (Figure 13-4) and specify whether the new member is an owner or just a normal member. You can add individual users or use a distribution list, mail-enabled security group, or group (or another team) as a source, in which case Teams expands the membership of the selected group and adds each user individually to the team. Teams excludes unsupported recipient types found in these groups (such as a mail-enabled public folder) and will not try to add them as members. If you

use any form of a group to add members to a team, remember that any later changes to the membership of the source group will not replicate to update the membership of the team.

As described later, if the tenant supports guest accounts, you can add external people as guest members. If team settings block guest membership, Teams won't show you guest accounts that already exist in the tenant directory or allow you to input the email address of an external user to create a new guest account. Finally, if your tenant uses information barrier policies to stop different sets of users from communicating with each other, you won't be able to add someone to a team if their presence causes a policy violation. To automate processing, any of the supported member types can be added to (or removed from) a team with PowerShell.

## Updating Team Rosters

Internally, Teams refers to the membership of a team as its "roster." If you add someone to a team through a Teams client, the roster is refreshed immediately in that client. Changes to membership can be made elsewhere by another Microsoft 365 application and synchronized to Teams. For example:

- A team membership is updated with a Teams client. The update is synchronized across Teams and to Azure AD.
- A team membership is updated using PowerShell, the Graph API, or another administrative interface. Usually, the change is written to the underlying Microsoft 365 group.
- A team membership is updated with a non-Teams client like OWA. The change is made available to Teams by synchronization between the workload and Azure AD.

Behind the scenes, Teams uses a background process called *Roster sync* to implement changes made in Azure AD in clients. When a user connects to a team with the desktop client, a check is made to see if any changes are waiting in Azure AD. If some are, the Roster sync process downloads the changes and updates the local cache.

The internal SLA for synchronization between Teams and Azure AD is 24 hours, and at worst it can take this long before a change made to group membership by another workload is effective across Teams. Usually, changes are available much sooner than the SLA limit.

**Teams with Duplicate Names:** As noted above, Teams will let you create a team with a display name that duplicates another team. This will confuse users who belong to those teams as they will never be quite sure which team they should use. From a technical perspective, Teams is quite happy to have duplicate display names because behind the scenes the teams have different aliases and names. If you get into this unhappy circumstance, you can change the display name of one of the affected teams by editing the team properties (Edit team) or by running the *Set-UnifiedGroup* cmdlet to update the *DisplayName* property for the underlying group. For example:

```
[PS] C:\> Set-UnifiedGroup -Identity BudgetPlanning -DisplayName "Budget Planning Group"
```

Renaming the team fixes the immediate issue of the duplicate display name. It will not rename the SharePoint team site. If you want to rename the SharePoint site, do so in the SharePoint Admin Center or use the *Start-SPOSiteRename* cmdlet.

## Creating an Org-Wide Team

If your tenant has fewer than 10,000 accounts (the limits are smaller for GCC and GCC/High tenants), you can create an org-wide team. A tenant can have up to five org-wide teams. Org-wide teams are intended for tenant-wide communications without the need for an administrator or team owner to manually add all the employees to the team membership, including the need to check for new employees and add them periodically with PowerShell. The advantage of using PowerShell to manage the membership of an org-wide team is that you have more flexibility in managing the membership; the downside is that you must check and

update the membership with new users and remove those who leave the organization periodically. Using an org-wide team avoids this work.

Larger tenants can consider using:

- **Dynamic Teams** to support discussions for different sections of the organization. For example, you might have a team for each department or each country.
- **Yammer** for company-wide communications and collaboration. Yammer can easily scale up to handle very large organizations with hundreds of thousands of users.

To create an org-wide team, choose **Join or create a team** as usual, then **Build a team from scratch**, and then select **Org-wide** from the list of options. The choice to create an org-wide team is only exposed to tenant (global) administrators and when you create an org-wide team, Teams adds all the global admins as team owners. It then adds all “active users” as members. You know the membership is managed by Teams because the Manage team screen has a banner saying, “*Team members will be automatically added and removed to reflect your Active Directory*”.

Automatic membership should exclude accounts without valid Teams licenses as well as guest users. However, sometimes accounts that should not be included in an org-wide team turn up in the automatically generated membership. Among accounts that should be checked against the team membership are:

- Service accounts (if they are assigned a license).
- Mailboxes used for purposes such as DLP incident reports.
- Accounts that don’t have a valid license where the Teams license is disabled.

Although Microsoft has fixed the bugs at the root of automatic membership, it is still good practice to check the membership after creating an org-wide team to remove anything that doesn’t belong. Once you remove someone from an org-wide team, the service remembers the deletion and won’t try to add that account back to the membership. Likewise, if you find that an account that should be in the membership has been omitted for some reason, you can add them manually. As with any team with a large membership, after you create an org-wide team, consider imposing some order on discussions from the start by updating team settings to restrict channel creation, the ability to post to the General channel, and to use @team mentions (because these generate notifications to everyone).

On an ongoing basis, employees leave and join the company and people lose or gain Teams licenses. When someone leaves the company and their account is removed, their membership in the team is also removed. To handle new joiners and people who gain or lose Teams licenses, a background process running on Teams clients scans the accounts in the tenant periodically (expect new users to appear within a few minutes) and adds or removes the user as needed.

Unlike normal teams, members can’t choose to leave an org-wide team. This limitation is the same as exists for dynamic teams, but unlike dynamic teams, an org-wide team doesn’t use an Azure AD query to calculate its membership. Instead, the background process manages membership and the Azure AD group for the team has an “assigned” membership, so you don’t need to buy Azure AD P1 licenses for all the members.

## Adding New Employees to Org-Wide Teams

Soon after you create an account for a new employee, if the account has a license for Teams, the account is added to the membership of org-wide teams. If your company provisions Microsoft 365 accounts for new employees in advance of their joining date as part of an onboarding process, you might not want this to happen because you don’t want other employees to know that someone is joining the company. In this case, you can either:

- Wait for the employee to join the company and create their account at that point.
- Create the account for the new employee but assign dummy information for the display name and primary SMTP address. For example, you could assign “New Employee” as the display name so that

other employees see that “New Employee.” has joined. The reason why to assign a dummy SMTP address is that users can click on “New Employee” to see more information from their people card. The SMTP address usually contains the first and last name of a person, so you don’t want to expose that information on the people card.

You then update the mailbox display name, name, and SMTP address after the new employee comes on board. If you use a dummy email address, you also need to update the account’s User Principal Name to be the user’s real sign-in address. These updates are easily done using PowerShell. For example:

```
[PS] C:\> Set-Mailbox -Identity NewEmployeexxx20 -DisplayName "Jake Adams" -WindowsEmailAddress "Jake.Adams@Office365itpros.com" -Alias "Jake.Adams" -Name "Jake Adams"
[PS] C:\> Update-MgUser -UserId (Get-Mailbox -Identity Jake.Adams).ExternalDirectoryObjectId -UserPrincipalName Jake.Adams@office365itpros.com
```


You don’t need to worry about keeping the dummy email address assigned to the account because it will never have been used to send an email. Any messages delivered to the mailbox will be waiting for the new employee.

If this arrangement doesn’t work, consider using all-employee teams whose membership is updated manually. It is easy to script additions and removals of employees from membership as part of the HR onboarding or leaving processes. Another workaround is to create accounts for new users in a disabled state and only enable the accounts when people join the organization. Teams ignores disabled accounts when it builds the membership of org-wide teams.

## Updating an Existing Team to be an Org-wide Team

Any existing team with private or public access can be converted by a tenant administrator to be an org-wide team to gain the benefit of automatic membership management. To make a team into an org-wide team, use the **Edit team** feature to change the privacy setting to be org-wide. When you save the setting, Teams updates the membership with all valid accounts. Any users not included in the automatic membership remain in place, including guest users. You can also change an org-wide team to be a private or public team using the same approach, and in this case, the existing membership stays in place, but the automatic background refresh of membership is disabled.

# Dynamic Teams



 This team has membership settings that prevent you from adding or removing members.

You cannot create a team with dynamic membership using a Teams client. Instead, you create a new dynamic Microsoft 365 group through the Azure AD admin center or PowerShell, together with the query to populate the membership. Dynamic groups can have membership rules based on account attributes (such as “everyone in New York City”) or use the membership of other groups, including other Microsoft 365 groups and distribution lists.

When you’re happy that the membership rule for the Azure AD group results in the expected membership, you can go ahead and select the group when creating a new team as described earlier. Teams respects the membership calculated by Azure AD and if you examine the membership (using **Manage team**), you see the same membership there as you see in the Azure AD admin center. However, if you view team membership you see a banner saying that you cannot add or remove members. In addition, the normal **Add Member** button is not shown.

Team owners are static and not computed automatically. You can add team owners to the dynamic group through the Teams client, the Azure AD admin center, or with PowerShell. You can change a member in the

dynamic set to be an owner. However, if you demote an owner to become a member and they are not in the set computed by the query, they lose their membership of the team.

You can also use a dynamic group as the source to add members to another team. In this case, just like a regular distribution list or group, Teams reads the present membership and adds them to the team. And just like when you use other types of groups to populate membership for a team, this is a one-time operation and anyone who joins the original team thereafter does not automatically join the other team.

Dynamic teams depend on the dynamic groups feature in Azure AD. Microsoft considers this to be a premium feature. Every member who comes within the scope of a query used for a dynamic team must have an Azure AD Premium license.

## Hiding Teams from Exchange Online

Teams uses the Microsoft 365 Groups membership service. And because each team has all the resources of a Microsoft 365 group, it has a group mailbox that can function as an Outlook group with email-based conversations. Users therefore could use email or Teams for their conversations. To remove any potential for confusion, Microsoft hides the existence of groups created by Teams from Exchange-based clients (OWA, Outlook desktop, and Outlook mobile).

Following the creation of a new team, Exchange Online sets the *HiddenFromExchangeClientsEnabled* and *HiddenFromAddressListsEnabled* properties of the underlying Microsoft 365 group to *\$True*. The new group is perfect for Teams but remains invisible to Exchange clients and does not appear in address lists like the GAL. In addition, because the focus of communication for the group is Teams communications, the group's subscriber list is not populated with tenant users (guest members are added). Any email sent to the group is delivered to the group inbox, but copies are not distributed to local tenant members (as explained in the Groups chapter, you can set an auto-reply for team-enabled groups to inform email senders that the group doesn't use email). Email from external senders is blocked.

Unfortunately, the flags to hide teams from Exchange Online are not set when an administrator creates a new team-enabled group using an administrative interface like the Teams admin center, Microsoft 365 admin center, Azure AD admin center, SharePoint admin center, or APIs like *New-UnifiedGroup* and the Microsoft Graph API. Administrators have full control over group settings and can update the settings as they wish. The net result is that team-enabled groups can be visible to Exchange Online clients. In addition, older team-enabled groups might not have *HiddenFromExchangeClientsEnabled* set to *\$True* because Microsoft did not update groups retrospectively following the introduction of the setting.

To check if any team-enabled groups are visible to Exchange, fetch the set of team-enabled groups and filter out the hidden groups. For example:

```
[PS] C:\> [array]$Groups = Get-UnifiedGroup -Filter {ResourceProvisioningOptions -eq "Team"}  
-ResultSize Unlimited  
[array]$Groups = $Groups | ? {$_.HiddenFromExchangeClientsEnabled -eq $False}  
$Groups | Format-Table DisplayName
```

To hide a team-enabled group from Exchange clients, update the properties with the *Set-UnifiedGroup* cmdlet. For example:

```
[PS] C:\> Set-UnifiedGroup -Identity Team1 -HiddenFromExchangeClientsEnabled:$True  
-HiddenFromAddressListsEnabled:$True
```

Setting *HiddenFromExchangeClientsEnabled* to *\$True* should also set *HiddenFromAddressListsEnabled* to *\$True*, but it's good to be specific. You can [download a script to find and set the flags for team-enabled groups from our GitHub repository](#).

Exchange clients periodically refresh their list of groups from the server, so it might take between 15 minutes and an hour before a hidden group disappears from a client. When you hide a group from Exchange, you also remove the group from users' Favorites lists for Outlook clients. Another side effect is that because Teams has no user-accessible directory, once a group disappears from address lists, users won't be able to browse groups in the GAL and so won't know of the existence of public teams that they might want to join (or private teams, for that matter). Also, remember that most Outlook desktop clients depend on the OAB and that it takes some time for Exchange to publish the removal of groups in OAB updates that must then be downloaded and applied by clients. See Chapter 3 in the companion volume for information about OAB updates.

## Using a Team-Enabled Group as a Distribution List

Before rushing to hide the groups used by Teams from Exchange clients, remember that Teams uses the Groups service to manage team members. A group is a mail-enabled object, meaning that people can send emails to the group. A group also functions as a distribution list and it can be useful to use the group for that purpose, even when you conduct most discussions inside Teams. For example, you can't send an encrypted message to Teams, but you can send a protected email to the group and the members will receive and be able to read the content in the message (if they subscribe to the group).

It's easier to decide to hide a team when its communications are exclusively internal. Making the call for a team with guest members is harder. When you have a team with many guest members and want to send something important and time-critical to those members, you could post a message in Teams, but each guest then needs to switch to your tenant to read the message. An email to the group might be quicker and more effective if you need fast action.

Even when a team is hidden from Exchange clients, you can send messages to the group by using its SMTP address fetched through the Microsoft 365 admin center or PowerShell. For example:

```
[PS] C:\> Get-UnifiedGroup -Identity MyOffice365Group | Select PrimarySmtpAddress
```

Administrators can find this address relatively easily, but users might not be so lucky.

If you create a team and decide that you want to keep its ability to act as a distribution list, you can reveal it to Exchange by setting the property to `$False`:

```
[PS] C:\> Set-UnifiedGroup -Identity MyTeam -HiddenFromExchangeClientsEnabled:$False
```

Exchange Online delivers messages sent to a group to the Inbox folder of the group mailbox. Local team members do not receive copies of messages because they are not part of the group's subscriber list. This is by design because Teams uses a different type of conversation. For this reason, when you create a new team, Groups sets the `AutoSubscribeNewMembers` property for the group to `False` and does not populate the subscriber list.

Normally, no issue arises because the group's subscriber list is empty. That is, unless you want to be able to use the team to distribute email to all or some of the members by using the team in the same way as a distribution list. To do this, you must add the members who you want to receive emails (or calendar requests) sent to the team to the group subscriber list. You can only do this through PowerShell. For example, this code adds two users (who must already be members of the group) to the subscriber list.

```
[PS] C:\> Add-UnifiedGroupLinks -Identity SeniorLeaders -LinkType Subscribers -Links Jack.Healy@Office365itpros.com, Marc.Vigneau@office365itpros.com
```



# Deleting (and Restoring) Channels and Teams

By default, any member can remove a channel or restore a deleted channel, but you can restrict this ability to team owners by updating team settings and unchecking the “allow members to delete and restore channels” setting under **Member Permissions**. Not only will this prevent accidents (unless a team owner makes a mistake), but it also stops disaffected employees who might be on the way out of the company doing something silly on the way to the exit (they might not know that the channel can be recovered). Like channel creation, Teams generates an audit record whenever someone removes a channel and writes a notification into the General channel. Teams also generates an audit record for the deletion. This information tells you who removed the channel, which might be cold comfort when you consider how much information might be lost in a deleted channel.

## Removing a Channel

You cannot delete the General channel for a team, but you can remove any of the other channels using the desktop, browser, or mobile client. When someone deletes a channel, Teams puts the channel into a soft-deleted state and starts a 30-day countdown, after which the messages and metadata for the channel become irrecoverable. During the 30-day grace period, you can recover the channel by selecting **Manage Team** and then **Channels**. Any deleted channels are listed under the set of active channels. To recover a deleted channel, select it and then click **Restore**. Teams then restores the channel to its original configuration to make its conversations available to users. One point to note is that the restored channel does not regain its previous status in the list of channels visible to the user, so users must mark the channel if they want it to show in their activity feed.

If the 30-day retention period for a deleted channel lapses, you can still recover messages from a deleted channel if the team is subject to a retention hold. To do this, run a content search to find the copies of the items held for compliance purposes in the group mailbox and export the found items to a PST, ZIP file, or individual messages. Although you cannot reassemble the recovered items into the threads they had in the channel, you will be able to email them to a channel or cut and paste content from the recovered items into Teams.

Teams does not remove the folder and files belonging to a deleted channel from the SharePoint document library, including the OneNote sections associated with the deleted channel. If you want to remove the folder and files, you must do so through SharePoint. One good reason why Teams does not remove the SharePoint content is that some or all the content in the folder associated with the channel might come under the scope of a retention policy applied to the site or that individual documents might be assigned retention labels that prevent their removal.

It's also possible that other channel tabs might be associated with content like a plan or form that will remain after the channel deletion. For this reason, it's a good idea to note what tabs exist for a channel and investigate what content is accessed via the tabs to build up a full picture of what's associated with the channel so that you can decide what to remove and what should be left before you go ahead with a channel deletion.

Remember that a team can have a maximum of 200 channels. If a team is at the limit and you remove some channels, the 30-day period must lapse before a Teams background process removes the channels permanently. When this process completes, new channels can be created for the team.

## Removing a Team

A team owner (using the Teams client) or a tenant administrator (using an administrative interface, like PowerShell) can remove a team. When this happens, Azure AD removes the underlying group and all its

resources including its SharePoint team site. For this reason, the owner must confirm that the deletion of the team should go ahead before Azure AD executes the command. This is important because once Azure AD removes the group object all the team data become inaccessible.

Generally, if you want to remove a team, do this from a Teams client as this action makes the team unavailable to all members at once. After removing the team, Teams synchronizes the deletion of the group object to Azure AD. In turn, the deletion command ripples across workloads through the normal directory synchronization process to instruct applications to remove any resources they manage for the group. It can take up to 30 minutes before all the resources belonging to a team are completely removed.

In most cases, unless you need to remove a team, it is better to leave it in place as a disused team. To do this, remove all members from the team except a single owner and update its display name to include a sign that the group is not in active use. This approach keeps the team in a state where you can revive it if needed by simply assigning new owners and members to its membership.

If you remove a team accidentally, you can recover its group (and the team) using the Microsoft 365 admin center or PowerShell, providing you do so within the 30-day retention period for deleted groups. When Azure AD recovers a team belonging to a soft-deleted group, the recovery process puts all content back in place into channels and personal chats and restores any connectors. Any email addresses belonging to channels in the deleted team are reactivated. Although the object for a recovered team is in the directory soon after a recovery begins, it often takes up to 24 hours before all the resources for a team are reconnected. At this point, Teams makes the recovered team available to users.

Although you can team-enable an existing group, no method exists to remove a team from a group, which you might want to do to “reset” a team by removing all channels, including the default channel, removing any tabs, bots, and connectors, and reducing the membership to a single owner. The only way to reset a team is to remove the group (and all its resources) and recreate it from scratch.

## Channel Moderation

You might want to reserve a channel for specific posts such as group announcements. The General channel has a setting to limit new posts to group owners. Moderation serves much the same purpose for channels other than General by controlling who can post new topics and replies to the channel.

Moderation is supported for both public and private teams. To enable moderation, select the channel you want to control and then **Manage channel** from the [...] menu. You can then turn moderation on or off for the channel (the default is Off). Even if moderation is disabled, you still have the option to restrict the creation of new topics (posts) to any member of the team or everyone except guests. After enabling moderation for a channel, the next step is to decide who the moderators should be (Figure 13-5). By default, all team owners are moderators, but you can select a different set of owners and members to act as moderators. Click **Manage** to change the set of moderators. Up to 100 moderators can be defined for a channel.

After moderation is enabled, members who aren't moderators cannot create new conversations in the channel and will see an informational banner saying, “only channel moderators can post in this channel.” In channels where moderation is enabled, it's a good idea to create an “Anything Goes” topic where people can discuss anything they like, including appealing to the channel moderators to create a new topic to discuss something specific.

You can also control if members can reply to posts and if automated processes (bots and connectors) can submit messages to the channel. For example, you could use this setting in situations where the channel hosts automated notifications about builds performed by Visual Studio.

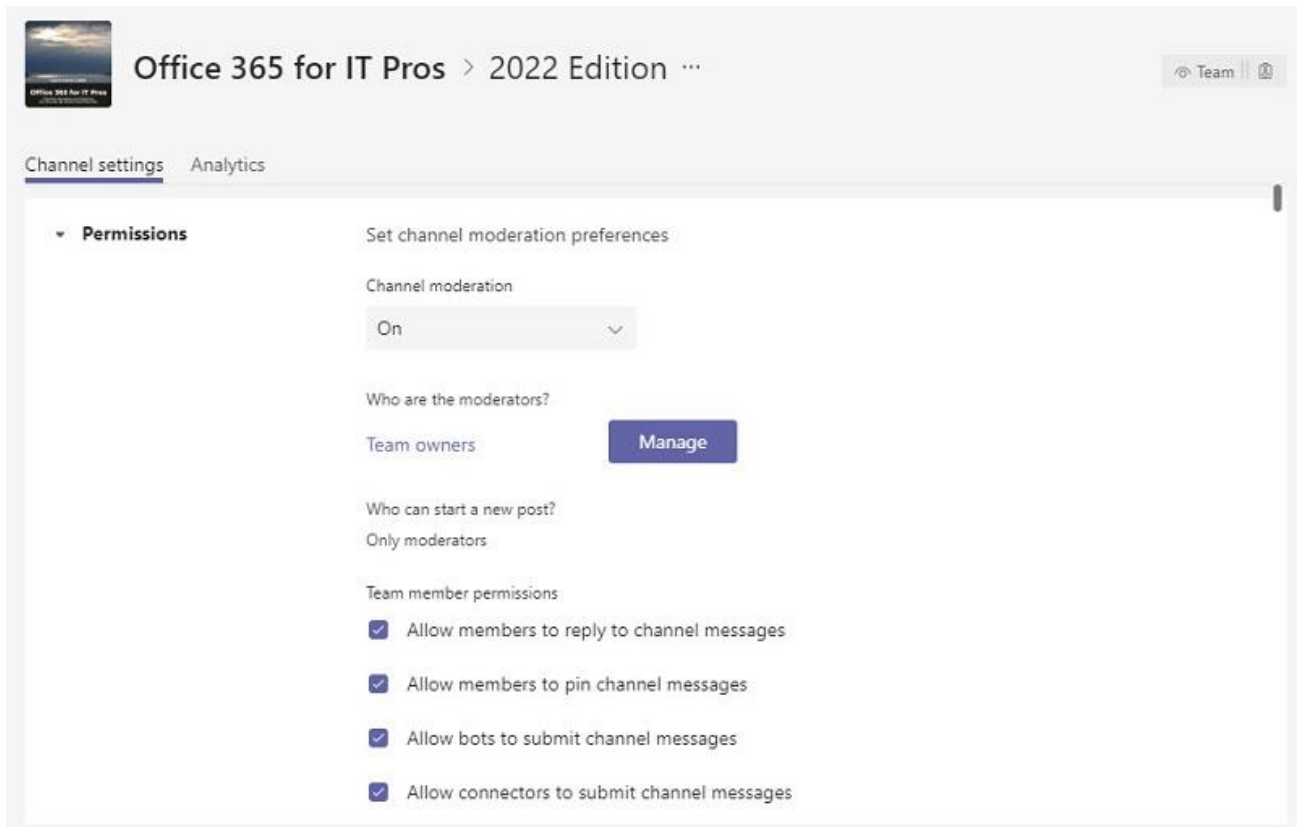


Figure 13-5: Setting up moderation for a channel

You can't update moderation settings for a channel with PowerShell, but you can [with the Graph API](#).

## Managing Settings for a Team

Team owners can manage team settings through the ellipsis [...] menu found beside the team name. Select **Manage Team** from the menu to access the available settings:

- **Members:** List existing team members (including guests), change roles (from Member to Owner and vice versa), remove members from the team, or add new members. When you add or remove members to a team, Teams updates the group membership.
- **Pending Requests.** Process any waiting requests to join the team. This tab is only visible for private teams.
- **Channels:** Manage channels for the team. You can also gain insight into the level of activity for the channels as you can see the number of members who have accessed each channel and the date of the last activity (Figure 13-6). As mentioned earlier, if a channel is important and should be brought to the attention of team members, a team owner can check the auto-show box to force the inclusion of the channel in each member's *Your teams* list.
- **Settings:** Settings to control the team picture; change the discoverability for a private team; assign permissions to users and guests; allow team members to use @mentions to reference other members, teams, and channels; and use stickers and memes (including the uploading of new memes). You can also upload a graphic file (of up to 4 MB) to add a team picture. A file measuring 640 x 420 pixels works well. Under **Member permissions**, you control the actions users can take within a team, such as whether they can add or remove channels, tabs, apps, or connectors; restore deleted channels; whether owners and members can delete messages; and restrict who can post messages to the General channel to team owners instead of anyone. For large teams, you can allow anyone to post but warn them that all members will see the message. If your tenant deploys a group expiration policy,

you'll see the expiry date for the team and be able to renew it here. Settings also include the ability to define **Tags**, a way to address subsets of members (see Chapter 12).

- **Analytics:** View usage information for the team and individual channels. See the later section about reporting usage of teams.
- **Apps:** Add or remove an app for the team. Apps include first-party apps like Forms, Planner, OneNote, SharePoint, and Stream as well as third-party apps like the Hipmunk bot installed in the team.

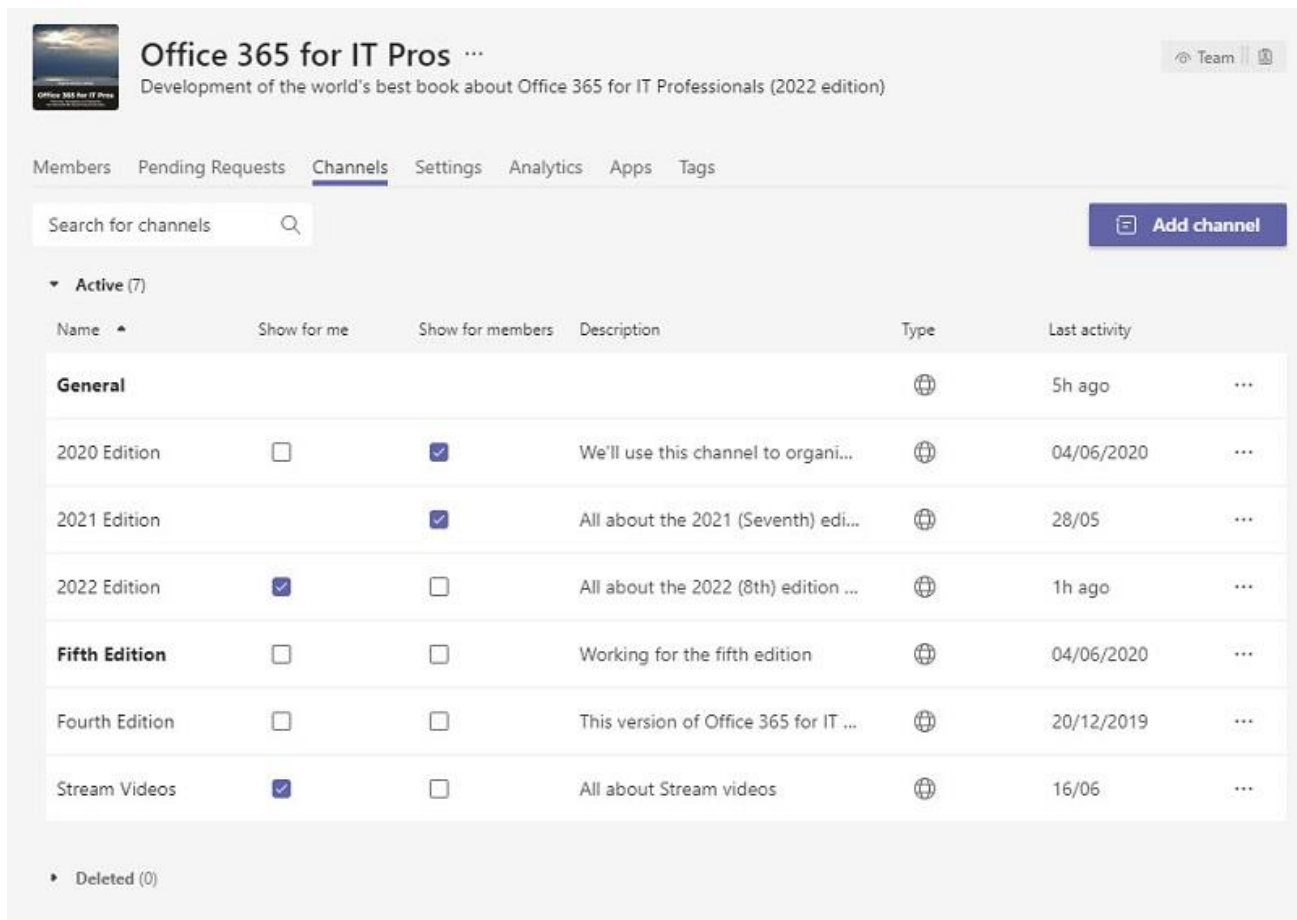


Figure 13-6: Managing the set of channels in a team

The ellipsis menu has choices to **Delete the team**, which also removes the underlying group and any other resources belonging to the group, and **Edit team**, which you can use to amend:

- The team name (display name). When you rename a team, you change the display name for the group. The change, therefore, affects applications like Groups and Planner. Because of the need to synchronize directories and clients, the name change might take some time before it is fully effective.
- Team description.
- Privacy (change the team from Public or Private or vice versa).
- Sensitivity (or classification).

These settings are group-wide and apply to all workloads using the group. Updating the settings is the equivalent of running the *Set-UnifiedGroup* cmdlet to change group properties. For example:

```
[PS] C:\> Set-UnifiedGroup -Identity HydraProjectTeam -DisplayName "The New Hydra Project Team"
-Notes "A very nice group of people that is team-enabled to get stuff done" -AccessType Public
-Classification Confidential
```

As explained in the Information Protection chapter, if you use sensitivity labels to control group settings for privacy and classification, those settings come from the label assigned to the group. The only way to update the settings is to change the sensitivity label to one which imports the settings you want to use.

## System Messages for Membership Changes

Teams posts system messages about updates to team settings, including those made with PowerShell, in the information pane. If you add or delete members for a team, system messages tell team members about the change. If someone uses a Graph API call (including the Teams clients and the *Add-TeamUser* cmdlet) to add or remove a member, the system message includes the display names of the member and the person (owner or administrator) who took the action. Changes made through the other administrative interfaces, such as the *Add-UnifiedGroupLinks* and *Remove-UnifiedGroupLinks* cmdlets (Exchange Online module) or *Add-MgGroupMember* and *Remove-MgGroupMember* cmdlets (Microsoft Graph PowerShell SDK), don't include the name of the person who took the action. It takes some time before Teams posts the system messages to inform members about changes. Changes that occur through other workloads, like Exchange Online or SharePoint Online, update Azure AD first. Subsequently, a background process called Microsoft Teams Aad Sync, synchronizes the changes to Teams. If you want membership changes to be effective immediately, you should manage team membership through the Teams admin center or Teams client rather than using other administrative interfaces.

## Guest Access for Teams

The mechanics of guest user access for Teams function like those for Microsoft 365 Groups. Team owners add guests to teams, joining invitations go to the email address identifying the guests, who then redeem the invitation to complete the process and join the team. As part of the Azure B2B collaboration process, Teams creates Azure AD guest accounts to allow guests to authenticate their access to resources. The Microsoft Invitations service generates the invitations and Teams manages the redemption process, including the verification of the guest account. You can invite anyone to be a guest if they have a valid email address.

Users can accept invitations to join other organizations as a guest from any client. The desktop and browser clients support other features through the Accounts section of Teams settings, including the ability to:

- Decline an invitation received to join a team in another organization.
- Leave an organization by removing their guest account. This option takes the user to the Azure AD [Organizations page](#) to choose the tenant to leave.
- Mute notifications from another organization.
- Hide or show organizations where the user has a guest account from the drop-down list displayed by Teams when the user selects a target organization for switching.

After a guest joins a team in a tenant, they don't have to accept invitations to join other teams through the links contained in the emailed notifications. Instead, the Teams apps detect that the guest is now part of other teams and perform an in-app redemption to add those teams to the list available to the guest.

A user of the free Teams version can be a guest in a tenant that uses the enterprise version and can switch between free and enterprise tenants. The same is true for users belonging to enterprise tenants, who can switch to guest accounts in free tenants.

## Enabling Guest Access for Teams

Because Teams consumes multiple Microsoft cloud services, different settings in the services combine to control guest access. In order of priority, these are:

1. **Azure AD.** The first hurdle is to make sure that Azure AD allows group owners and members to invite guests. In the **External Identities** section of [the Azure AD admin center](#), the **External collaboration settings** control external access to Azure AD. Under Guest invite settings, the default setting is *"member users and users assigned to specific admin roles can invite guest users including guests with member permissions."* If the setting is more restrictive, for instance, the control is set to *"only users assigned to specific admin roles can invite guest users,"* any attempt by users who don't hold an Azure AD administrative role such as Groups administrator fails.
2. **Groups.** In the Settings-section of the Microsoft 365 Admin Center, select Org Settings, then Microsoft 365 Groups, and make sure that *Let group members outside the organization access group content* and *Let group owners add people outside the organization to groups* are both On. When set, you can invite guests to join Teams. If you have specific teams that discuss sensitive information, you can block the ability of the owner to add guests to those teams by either editing the properties of the teams or by assigning a sensitivity label to the teams which prohibit guest membership. For more information, see the discussion about how to block guests for specific groups in Chapter 11.
3. **Teams.** Go to the Teams admin center and then the **Guest access** setting under the **Users** section to make sure that the **Allow guest access in Teams** slider is set to *On* (the default for this setting). Other settings in this policy control the functionality available to guests, such as if they can edit their sent messages. Guest users don't need a license to access Teams. You can also enable guest access to Teams with PowerShell by running the [Set-CsTeamsClientConfiguration](#) cmdlet and setting the *AllowGuestUser* setting to *\$True*.
4. **SharePoint.** In the SharePoint Admin Center, under External Sharing, set *Let users share SharePoint Online and OneDrive for Business content with people outside the organization* to On. This enables guest access to the Files in SharePoint document libraries used by Teams and to files in users' OneDrive sites shared in personal chats.

With all settings turned on, the full spectrum for guest access is supported for all teams in the tenant.

Teams shared channels do not use guest accounts. Instead, users connect to shared channels in other tenants using Azure AD B2B Direct Connect. See the Identities chapter for more information on this topic.

## Adding Guests to Teams

With all the correct settings in place, team owners can add guests. It is important to understand that the group and its associated team share a common membership list, including guests. Therefore, if you add a guest to the group, the guest gains access to the team and vice versa. Sometimes a small delay happens between adding or removing a guest from the membership and that action showing up when viewing membership, but background synchronization processes make sure that any addition or removal of guests applies across both Teams and Groups.

To add a guest member, select **Manage team** and then **Add Member**, or use **Add members** from the ellipsis menu. You then input the email address of the new member (Figure 13-7). In this case, the email address entered for the guest was John.Smith@contoso.com. Left as is, Teams uses John.Smith (Guest) as the display name seen by other team members. The (Guest) suffix is a language-specific string added by Teams to mark someone as external. The string is translated when displayed by Teams clients and notifications – for example, if a guest user runs Teams in Spanish and you receive an email notification for something they post, you'll see (Invitado) after their name. You cannot change the guest suffix because that is a visual reminder to other members that a guest is an external person, but you can amend the other part of the display name by clicking the pencil mark beside the name to edit guest information (Figure 13-7). You can now update the display name to whatever you want. For instance, by adding a company name so that the display name is "John Smith (Contoso)" to inform other members what organization the guest is from.

### Add members to "Corporate Business Development"

Start typing a name, distribution list or security group to add to your team. You can also add people outside your organisation as guests by typing their email addresses. Note: Guests need a work or school account in Office 365. [Learn more about who can join as a guest](#)

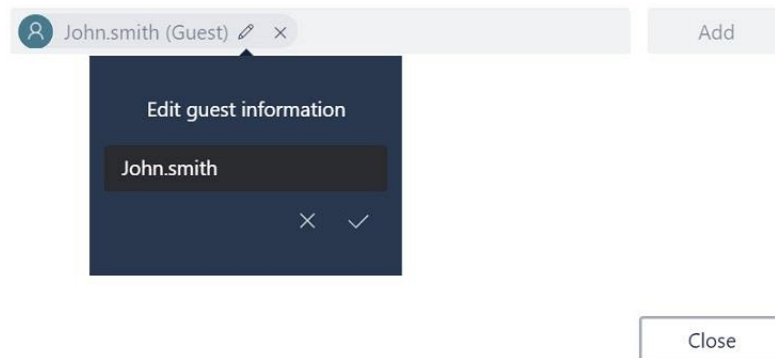


Figure 13-7: Adding a guest to a team

When the information for the guest member is complete, press **Add**. Teams checks whether a guest account for this address already exists in the tenant directory and adds it to the membership if found. If not, Azure AD creates a new guest account. Teams then adds the account to the team membership and generates an invitation for the user to redeem. The invitation has all the information necessary for the recipient to know about the team they are joining, including a GUID to find the right tenant, another GUID for the team, and a token request to redeem the invitation to join the team. The recipient redeems the invitation by clicking the **Open Microsoft Teams** link in the message, which then invokes the necessary flow to confirm the user's details, update the Azure AD guest account for the user with credentials to allow them to connect to the tenant, and access the team. Azure AD logs the redemption in a "redeem external user invite" audit record.

To give a visual clue to tenant users that they should be careful about the information shared externally, once you add a guest to a team, Teams displays the number of guests in the team in the top right-hand corner of the conversations pane.

**Updating Guest Information:** Do not worry if you forget to update the display name for a guest when you add them as you can always update it afterward. You can edit the contact information for the guest through the Microsoft 365 admin center, update the guest account properties in the Azure AD admin center, or run the *Update-MgUser* cmdlet. For example, this code updates the display name for the guest account created for John.Smith@contoso.com bypassing the user principal name as the object identifier. It also updates the user's job title and city as Teams displays these properties when you view a team. Finally, we update the telephone and mobile numbers for the user so that they appear when someone looks at their contact card.

```
[PS] C:\> Update-MgUser -User John.Smith_contoso.com#EXT#office365itpros.onmicrosoft.com
-DisplayName "John Smith (Contoso)" -JobTitle "Account Manager" -City "NY" -BusinessPhones "+1
617 551 6531" -MobilePhone "+1 650 561 4136"
```

It's also a good idea to upload photos for guest accounts. You can either upload an actual photo (the nice approach) or use a default photo to mark the guest as an external person. In either case, uploading a photo for a guest is done by editing their account in the Azure AD admin center or by running the *Set-MgUserPhotoContent* cmdlet. See the Groups chapter for more information. Teams must synchronize with Azure AD before the updated account information shows up in the Teams clients.

Every guest has an Azure AD account. An administrator can amend settings for the guest user account (like their photo), but because guests don't have mailboxes in the tenant, they can't set an out of office notification or any other feature which depends on a mailbox.

## Multi-Factor Authentication for Guests

Like the other Microsoft 365 applications, Teams supports multi-factor authentication (MFA). If your tenant implements an Azure conditional access policy to require MFA, you can include guest members in the scope of the policy to force them to use MFA to connect to Teams, even if their home tenant does not require MFA for connection to services. The logic here is that a tenant always controls its resources and can therefore dictate the level of authentication required to access those resources.

To ensure that a conditional access policy applies to guest users, choose *All Guests and External Users* as the target group for the policy to apply to.

## Blocking Guests from Specific Domains

If you do not want team owners to add guests from specific domains, you can create a block list using the Azure B2B Collaboration policy. The same policy can also create an allow list to restrict guest user access to specific domains, but a tenant can only support either an allow or a block list. When a policy is in place, team owners cannot add guests from blocked domains. However, any guests from blocked domains who are already members of teams continue as before and that guest can be added to other teams because Teams only uses the deny list when adding new guest accounts. It could be the case that the account was added by another application (to share a SharePoint document, for instance) that has nothing to do with Teams.

If you want to revoke membership for guests belonging to blocked domains, you must use PowerShell to search the membership lists for all Microsoft 365 Groups and remove any guest members found that belong to those domains. An example of how to remove a user whom you wish to block is in Chapter 11. If you want to ensure that no guest users are added to specific teams, disable the ability for team owners to add guests to those teams with a sensitivity label or by updating the policy setting for the group as described in Chapter 11.

## Switching Between Tenants

When a user with a guest account wants to access teams in your domain, they must “switch” Teams client by signing into your tenant using their guest account. Switching to work as a guest into another tenant means that an account has a limited view of the Teams environment within that tenant when compared to what they can do in their home tenant. Guests can only see teams to which they belong and cannot browse to join other public teams; they cannot see organizational information about other members, nor can they create new teams or add apps to channels. Because guests cannot access the directory, they cannot update the picture for their account, or any other setting related to their account such as the display name. In addition, they cannot browse the directory to find people to chat with. Instead, guests enter the email address of the person they are looking for and Teams checks the directory and creates a chat if that person exists.

On the other hand, guests can join in private and public conversations, take part in video and audio chats and access files from the team SharePoint site (or the sites used by private channels, if their account is added to the membership of private channels). Another point to note is that guests can only access applications available to the team if they have the right credentials and the application supports guest user access. It is as if they signed into your tenant with more restricted access than a normal tenant user has, which is exactly what you want.

Assuming you have an account in at least one other tenant where Teams is active, switching to an account in a different tenant is simple. In the desktop and browser clients, click the name of the tenant you’re currently connected to in the title bar to reveal the set of known tenants in your Teams profile. A tenant is listed only if your account is added to at least one team in that tenant. If your access expires for a tenant, you’ll be asked to reauthenticate before you can connect. For mobile clients, click the hamburger menu in the top left-hand corner and select the target tenant from the list shown at the bottom of the screen. In Figure 13-8, the



desktop client shows that eight host tenants (and the home tenant) are available for switching. One tenant has a warning sign, showing that I need to reauthenticate before I can connect.

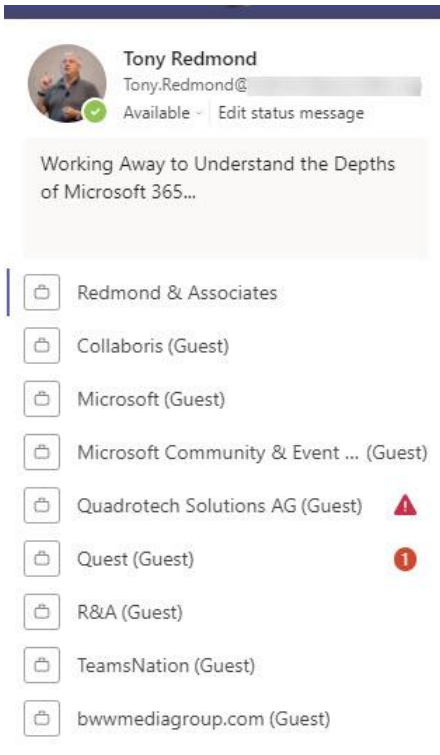


Figure 13-8: Selecting a host tenant to switch to

When you select a different tenant, Teams signs out of the active tenant, signs into the target tenant, and then displays the teams that the user belongs to in that tenant. The sign-in process ensures that the user picks up the rights and permissions they enjoy within the tenant; when they are a guest, they have those rights in that tenant. When they return their host tenant, they can interact with Teams with all the rights of that tenant account.

Generally, users do not see events for other tenants in their activity feed (the exception is when you participate in a call hosted by another tenant – chat notifications for the meeting do appear in the activity feed). Although Teams flags activity with indicators in your avatar, the lack of notifications can be a problem if you need to keep an eye on an active conversation about an evolving situation as you must then sign into the tenant hosting the relevant team to check what’s happening. Teams does not have a good solution for concurrent access to multiple tenants, but you can consider using the [approach described in this article](#) to create “wrappers” for each tenant where you use Teams. You can then move from one tenant to another while keeping access to Teams in the other tenants.

## Personal Account

Users can add a personal account to their Teams profile. A personal account is a Microsoft Services account with a valid email address. When someone uses Teams with a personal account, they switch away from Teams enterprise to Teams personal. Teams personal has different functionality to Teams enterprise and two separate clients (windows) are active after the user signs into Teams personal. The user can move from one client to the other as they like.

## Removing Guests

You remove guests from a team in the same way as you remove a tenant user. Select **Manage team** from the menu and then **Remove** for the user in the member list. Alternatively, you can run the *Remove-TeamUser*

cmdlet to remove the user from the group. For example, this command removes the guest Joan.Smith@contoso.com from the Industry News team.

```
[PS] C:\> $TeamId = (Get-UnifiedGroup -Identity "Industry News").ExternalDirectoryObjectId
[PS] C:\> Remove-TeamUser -GroupId $TeamId -User "Joan.Smith@contoso.com"
```

Removing an external member from a team does not remove the guest account from Azure AD. If you want to remove an external user from all teams in a tenant, you can remove their guest account from Azure AD using the Azure portal, the Microsoft 365 admin center, or by running the *Remove-MgUser* cmdlet. Removing the account removes membership from all teams and groups and removes any sharing permission the user has for SharePoint and OneDrive for Business files in the tenant.

## Email Integration for Teams Channels

The Teams architecture chapter discusses the various methods available to bridge the gap between email and Teams. The *Email Integration* settings for the tenant (managed through Teams settings in the Teams settings section of the Teams admin center) control if channels can receive SMTP addresses to allow them to accept emails. This is a useful way for people to introduce a new topic into a channel using information contained in emails. To retrieve the email address of a channel, use the desktop or web client to click the ellipsis menu for the channel, select **Get email address**, and then **Copy** (Figure 13-9) to copy the address to the clipboard. If an email address does not already exist for the channel, Teams creates one. You can then use the address to send a message to the channel. Any team member can retrieve the email address for a channel, but guest members can't force Teams to create an email address for a channel that doesn't already have one.

### Get email address

See advanced settings for more options.

xecutive Retreat - Project Condor <af9af07f.office365itpros.com@emea.teams.ms>

 Remove email address

Close

Copy

Figure 13-9: Retrieving the email address for a channel

By default, anyone inside or outside the tenant can use the email address to send a message to the channel. Two levels of restrictions are available.

- On an organization level, tenant administrators can block all channels from receiving email or define a **list of domains allowed** to send email to channels. Any user from these domains can send an email to a channel in the tenant, including private channels. If you create an accepted domains list, make sure to include all the domains you wish to allow to communicate with Teams via email. Note that you cannot specify a wildcard domain like \*.onmicrosoft.com.
- On a channel level, a team owner can use the **Advanced Settings** option to restrict the acceptance of messages to **only members of the team**.

You can also use PowerShell to update the set of accepted domains for the tenant. In this example, we pass the set of accepted domains as a string with each domain separated by a semi-colon (don't leave any spaces). The *Set-CsTeamsClientConfiguration* updates the configuration, which we can then check with *Get-CsTeamsClientConfiguration*.

```
[PS] C:\> [string]$SenderDomains = "Microsoft.com;Office365itpros.com;contoso.com"
Set-CsTeamsClientConfiguration -RestrictedSenderList $SenderDomains -AllowEmailIntoChannel $True
```

It can take a couple of hours before the updated list of accepted domains becomes active. When delivery restrictions are in place, Teams rejects emails sent to a channel from senders not on the allowed list, and the sender receives a message from Teams. The text of the notification is something like this:

**Delivery has failed to these recipients or groups:**

*1e2d0eb3.office365itpros.com@emea.teams.ms*

*The administrator has restricted permissions to send emails to this channel.*

Team members don't receive copies of emails delivered to a channel. If you want people to receive a copy of a message sent to a channel, add their email addresses to the message.

## Managing Channel Email Addresses

When Teams creates an email address for a channel, it also creates a mailbox in a special tenant dedicated to Teams email processing and assigns the email address to the mailbox. The channel's email address is in the form `<unique-identifier>.tenant-domain@region.teams.ms`. Tenants cannot change the format of the address or the domain it uses to make it more human-friendly.

For example, the email address `b400aa20.tenant-smtpaddress@emea.teams.ms` belongs to a tenant in the EMEA region. The SMTP domain included in the email address generated for the channel uses the domain defined in the email address policy for the tenant. If a separate email address policy exists for Microsoft 365 groups, Teams uses that domain.

After creating a mailbox to handle inbound email for the channel, Teams creates a connector (see below) to link the mailbox to the channel so that any email arriving in the mailbox flows through to the channel. Dedicated infrastructure exists to host Teams mailboxes in every Microsoft data center region supported by Teams. It's worth noting here that because Teams uses a connector to bring messages into a channel, any information barrier policies used by the organization don't apply. This means that someone otherwise blocked from communicating with team members can send emails to them via the channel.

Any team member (except guests) can use the **Remove email address** option to nullify the email address for a channel. Allowing users to remove the email address used for inbound communications is in line with the general rule that everyone shares equal access to team resources, but it is easy to see how this could lead to problems if people rely on the address to communicate with the channel. It is easy to reconnect email access for a channel by using the **Get email address** option again, but Teams does not restore the old address to the channel and generates a new address instead.

Teams captures audit records when team members generate or remove channel email addresses. Here's an example of [how to interrogate the audit log to find the relevant records](#). Teams doesn't expose email addresses in the properties reported for a channel in the `Get-TeamChannel` PowerShell cmdlet. You can't generate a report of Teams channels with email addresses using pure PowerShell, but you can with [a combination of PowerShell and Graph API calls](#).

## Posting to a Channel from Exchange Online

To make it easier for people to post emails to a channel, you can create an Exchange Online mail contact with the channel's email address and give the mail contact a suitable display name. The mail contact will appear in the Exchange GAL and users can add it to messages as easily as any other email address, including [using the mail contact in a distribution list](#). If someone changes the email address for the channel, you'll have to update the mail contact because no synchronization exists for these objects between Exchange and Teams.

Messages sent to the email address of the Microsoft 365 group used by a team go to the group mailbox. Members who subscribe to the group receive copies, but Exchange will not deliver a copy of the message to a

channel in the team unless you add the channel's email address as an external member of the team. For example, let's assume that you want any email sent to the Corporate Business Development group to show up in Teams and you create a channel called Email Communications for this purpose. You can then get the email address for the channel and add it to the group as a new guest member with OWA. Exchange Online delivers a copy of any messages sent to the group thereafter to the channel. Guest user accounts do not appear in the Exchange GAL.

**Use BCC to Stop Mail Storms:** The email address for a Teams channel functions in the same way as any other email address in that you can use it as a message recipient. However, there's a very good reason for always using BCC when sending emails to Teams. The purpose of a Teams email address is to allow users to start new discussions in a channel with content that already exists in emails. You do not want the channel to take part in a lively back-and-forth email conversation because every interaction shows up as a new thread in the channel and can confuse team members. Use BCC to address email for a channel and then develop the conversation forward in the channel.

## Message Hygiene for Inbound Email to Teams

Because the mailboxes used by Teams are part of the Microsoft 365 infrastructure, inbound messages go through Exchange Online Protection (EOP). Although EOP will stop the delivery of malware to Teams, the special nature of the configuration used by Teams means that inbound mail does not go through the same processing as the email stream for a regular customer tenant. For example, if you license Office 365 E5 and configure the Microsoft Defender for Office 365 policy to enable Safe Attachments and Dynamic Delivery, you expect that EOP processes all attachments in this manner. However, this will not happen for messages routed through Teams. On the other hand, if you configure Microsoft Defender for SharePoint, when Teams captures the message files in the SharePoint document library (see below), Microsoft Defender processes any attachments at that point.

Teams will not deliver a message to a channel if the message has more than 20 file attachments or more than 50 inline images. Exchange Online sends a non-delivery-notification to the sender whenever Teams rejects a message.

## SharePoint Captures Email Sent to Teams Channels

When a channel receives an email, the text of the message appears as a new topic in the channel. If the text exceeds the limit for a contribution to a chat (25 KB, or roughly 20,000 characters including spaces), you see some text but need to download the "original email" to view the full content. You cannot reply to an inbound message because Teams has no other access to email except its ability to receive inbound messages. Any attempt to reply to an inbound email generates the response that your reply will only be visible in Teams.

When Teams receives a message via email, it strips any attachment and stores the file in the channel folder created in the SharePoint Online document library for the team (a separate folder exists for each channel in the team). Originally, Teams held copies of received messages in a single *Email Messages* folder. To improve performance, Teams creates a separate folder per month named after the month. For instance, the folder used for February 2021 is **EmailMessages\_2\_2021**.

In addition to attachments, Teams captures a copy of the message as an *.eml* file and keeps the file in the same location. Keeping copies of messages and attachments received by Teams in SharePoint means that Microsoft Search can index the information to make it available for eDiscovery. Users can download a copy of the original message from email conversations posted to the channel.

Because copies of messages sent to channels end up in SharePoint Online, each delivery of a message to a channel creates an *Uploaded File* audit record. You'll see that the user noted in the audit record is

*app@sharepoint*, a background process that performs many management operations for SharePoint. This process copies the message from Teams to SharePoint.

# Teams and Compliance

Teams stores chat and channel messages in a data store hosted by Azure Cosmos DB. The information in the Teams data store is not directly accessible to the data governance framework. It is obvious that many interesting conversations occur in Teams and that the information discussed in these conversations might be of interest to eDiscovery investigations. The same fear that eDiscovery investigations might not have access to all relevant information exists when employees use third-party chat networks like Slack.

Despite the lack of access to the native Teams data stores, Teams has other mechanisms to support many of the data governance technologies available in Microsoft 365, including:

- **eDiscovery:** The Microsoft 365 substrate captures and indexes compliance records so that content searches can find, preview, and export these items. The substrate captures compliance records for Teams chats and conversations, meetings, adaptive card content, and call data records.
- **Retention:** Teams supports the application of retention policies against messages sent to Teams channels and private/group chats.
- **Communications Compliance:** Teams supports the capture of messages matched against classifiers identifying offensive, threatening, or other problematic interactions. Administrators can remove messages that violate policies to make them inaccessible to anyone but the sender.
- **Data Loss Prevention:** Teams supports the checking of personal chats and channel messages to detect the presence of sensitive information types.
- **High-value audit events:** If your tenant has the necessary licenses (Office 365 E5 is one example), Teams generates audit events for meetings to capture details of the meetings and their participants. See [this article](#) for more information.

Retention policies govern the preservation or removal of data as required by the organization. By default, Teams does not remove any conversation or other information belonging to a channel or chat. This data remains in the Teams data stores or the underlying application (for example, Planner) unless a user or team owner takes action to remove the data. If you want to ensure the retention of Teams data for a defined period or removed after a certain period, you must deploy retention policies to cover Teams messages (chat and channel conversations) and any other repository used by Teams, such as SharePoint Online. We'll discuss how to use retention policies to manage Teams data later. A more general discussion of retention policies to cover other Microsoft 365 data is in Chapter 17.

## Capturing Teams Compliance Data

To ensure that Teams conversations are available for eDiscovery, the Microsoft 365 substrate captures copies of the messages sent in personal chats and channel conversations as items in Exchange Online mailboxes, including messages emailed to channels. These mail items are compliance records. They are imperfect copies of the complete Teams message data as they contain the information needed for eDiscovery and other compliance functionality, such as communications compliance policy checks. The substrate stores the compliance records in the **TeamsMessagesData** folder in the non-IPM part of user and group mailboxes. You can see how many compliance records exist in these mailboxes with PowerShell (see [this article](#)). Because the compliance messages exist in Exchange Online mailboxes, Microsoft Search indexes their content and metadata to make the items available for eDiscovery.

As people communicate in chats and channel conversations, the Microsoft 365 substrate captures each message as a separate compliance record. A complete conversation between multiple people might involve thirty or more contributions, each of which exists as a separate compliance record. The compliance records

are mail items, so they have a sender (the author), a recipient (the group mailbox for channel conversations or user mailboxes for personal chats), and a timestamp for when Teams created the item. Like other mail items, you can examine the properties and content of Teams compliance items using the MFCMAPI program. If the author or a team owner edits a message, the substrate captures compliance records for the original message and every edited version of the message.

The Microsoft 365 substrate captures compliance records as follows:

- **Personal and group chats:** The substrate creates compliance records in the *TeamsMessagesData* folder of the mailboxes of the participating users. For example, if John and Pat have a chat, the substrate creates compliance records for the chat in both their mailboxes. If a group chat involves ten people, compliance records for every message in the chat exist in all ten mailboxes. Compliance records for chats in private meetings are also in user mailboxes.
- **Regular channel conversations:** The substrate captures compliance for messages posted to channels in the group mailbox belonging to the team. Compliance records for conversations in all channels in a team are intermingled.
- **Private channel conversations:** The substrate captures compliance records for these messages in the user mailboxes of the channel members.
- **Shared channel conversations:** The substrate captures compliance records for these messages in a special cloud-only mailbox created for the shared channel. This mailbox is inaccessible to regular email clients and is also known as a *SubstrateGroup* mailbox.
- **Meetings and Calls.** The substrate generates compliance records as the calling infrastructure creates records of meetings and calls. Meetings include scheduled private meetings and ad-hoc meetings associated with a channel. Group chats that include more than two people are also in this category. Calls cover one-to-one calls. The substrate captures compliance records when users from your tenant participate in meeting chats hosted by another tenant.
- **Meeting artifacts:** These artifacts include meeting attendance reports, transcripts, whiteboards shared during meetings, and recordings. As described below, Teams captures meeting and webinar information in Microsoft Lists stored in the meeting organizer's OneDrive for Business account. This data is available for eDiscovery. Transcripts generated from Teams meeting recordings are in the OneDrive for Business account of the person (usually the organizer) who initiates the recording. OneDrive hides this data from users. A background process indexes the spoken words from the transcripts against meeting recordings to allow concurrent playback of video and text. The text (spoken words) is available to SharePoint Search but is not yet available for eDiscovery. See the Managing Video chapter for more information about Teams meeting recordings.
- **Adaptive card data:** First-party and third-party apps integrated with Teams can use adaptive cards as part of their user interface. Another category is cards representing data from network sources created in channels using the inbound webhook connector. The substrate captures compliance items for adaptive cards. If an app creates a card in a personal or group chat, the substrate creates compliance records in the mailboxes of all those involved in the chat. Compliance records for cards created in channel conversations are in the group mailbox of the team which owns the channel.

Compliance records are available very soon after users or apps create messages in Teams (usually a matter of seconds, never more than a few minutes). The substrate does not capture records for draft messages. These messages remain in the local client cache until the user eventually sends the times, at which point the substrate creates a compliance record.

**Warning:** The compliance records created for voice memos generated in personal chats by Teams mobile clients do not contain any metadata or other information necessary for inclusion in content indexes. No attempt is made to transcribe the voice message into a form that can be indexed, so these memos cannot

be found by content searches. Because some Microsoft 365 workloads do not support eDiscovery, other challenges for Teams compliance capture include:

- Planner tasks.
- Whiteboards in meetings.
- Meeting transcripts (available for Microsoft Search, but not yet for eDiscovery).
- Data belonging to third-party apps stored outside Microsoft 365.
- Loop components in Teams chat. Compliance records exist, but they don't contain any of the content for the loop component.

## Capture of Webinar Event Data

The Microsoft 365 substrate captures attendance information for Teams meetings as items in the hidden `93c8660e-1330-4e40-8fda-fd27f9eafe10AttendanceReportV3Collection` folder in the non-IPM part of the meeting organizer's mailbox. The data in these items is sufficient to display the attendance report for a regular meeting.

Special processing occurs to make information about meetings configured as webinars available for eDiscovery. For a webinar, Teams creates three lists in the meeting organizer's OneDrive for Business account. These lists are:

- **Event:** Stores event information such as its start and end time and webinar description and title. The *ThreadId* for the webinar is stored in this list. The webinar title and description can be edited in the list but the information created by Teams for the meeting cannot.
- **Questionnaire:** Stores the attendance records for individual webinar attendees. The information about attendee details (like name and email address) can be edited in the list but information relating to the Teams meeting (like its URI) cannot.
- **Speakers:** Stores details of the speakers such as their names and bios. This information can be edited in the list.

The lists used to hold webinar data are not usually visible to users unless they navigate to the lists section of their OneDrive account (Site Settings, then Site Libraries and Lists). They can then retrieve the URL to a list and use the URL to access the data in a browser.

Information held in webinar lists (including any updates) is available for searching through SharePoint Search or Microsoft 365 eDiscovery. See [this article](#) for more information.

If your organization has Office 365 E5 or Microsoft 365 E5 compliance licenses, Teams generates audit events for meetings (*MeetingDetail*) and participants (*MeetingParticipantDetail*). See the section on advanced auditing in the reporting and audit chapter.

## Compliance and Communications Compliance Policies

Teams supports communications compliance policies. When a user or team comes within the scope of a policy, copies of their messages which violate communications policies are captured in a special mailbox and kept there until reviewed. The messages generated for review are like those captured for compliance purposes, but they go through a different process. During the review process, messages identified as violations can be blocked so that recipients can no longer see these messages. The sender can still see the message, but it is highlighted as being blocked due to a policy violation. See Chapter 21 for more information about communication compliance policies.

## Compliance Records for Private Channels

Two challenges exist for compliance records collected for private channels:

- *Compliance records for private channel messages are in member mailboxes:* Private channels don't have a group mailbox. For this reason, Exchange Online stores the compliance records captured for

conversations in the private channel in the personal mailboxes of channel members like the way that it stores records for group chats. From a content search perspective, it means that the substrate captures more copies of compliance records (up to 250 copies of each message posted to a private channel). Another way of looking at this is that you'll find a record if you add a single member of a private channel to a content search. The compliance records in private mailboxes contain a mixture of records for private chats, group chats, and private channel conversations. The items for private channel messages have different MAPI property values to chat messages. For instance, property 0x0DE001F is *MicrosoftTeams* for chat messages but *MicrosoftTeamsChannelMessages* for private channel messages.

Because compliance records for private channel conversations are in user mailboxes, retention policies must apply special processing to find and deal with compliance records created for private channels. For this reason, retention policies to process private channel messages are separate from other retention policies. By contrast, the substrate stores compliance records captured for Teams shared channels in cloud-only mailboxes (like those used for hybrid and guest users), meaning that regular retention processing for Teams channel messages covers these items.

- *eDiscovery searches might not include some SharePoint sites*: Private and shared channels use special hidden SharePoint sites to store their documents. If you include teams in content searches, the search only includes documents in the standard SharePoint sites used by the teams. To ensure searches process documents stored in the special sites used by private and shared channels, you must add the URLs for the channel sites to the content search locations.

This code returns the URLs for sites belonging to Teams private channels. If you exclude the filter against the set of sites returned by the *Get-SPOSite* cmdlet, the set includes both private and shared channels.

```
[PS] C:\> [array]$Sites = Get-SPOSite -Template "TeamChannel#1" -Limit All | ? {$_.TeamsChannelType -eq "PrivateChannel"}

ForEach ($Site in $Sites) {
    $SPOSite = Get-SPOSite -Identity $Site.url -Detailed
    $Group = Get-UnifiedGroup -Identity $SPOSite.RelatedGroupID.Guid
    Write-Host "Team" $Group.DisplayName "owns private channel site" $Site.URL}
```

Earlier iterations of private channels used the *TeamChannel#0* template. Microsoft switched the template in mid-2021. However, it's possible that you might find some sites with the old template.

## Teams Compliance Record Structure

Like other mail items stored in an Exchange Online mailbox, a Teams compliance record is composed of a set of MAPI properties that can be viewed using a MAPI editor like [MFCMAPI](#). The properties include:

- *SkypeInternalIds*: a list of GUIDs for each of the participants (other than the sender) in the conversation. Teams uses GUIDs internally to avoid problems with user display and principal names, both of which can change over time due to marriage, divorce, or other circumstances. Teams resolves the GUIDs for display and stores the display names in the *PR\_DISPLAY\_TO* property.
- *ParentMessageId*: The reply identifier for the message thread (also in the *LinkId* property).
- *CreatedDateTime*: Date and time of the message.
- *PR\_BODY*: Text of the message (also captured in HTML format in *PR\_HTML*).
- *PR\_SUBJECT*: The subject of the message (also contains the channel name).
- *PR\_SENDER\_NAME*: The display name of the sender. The GUID for the sender is found in the *FromSkypeInternalId* property.

Compliance records don't currently capture the channel name for channel conversations, which makes it more difficult to track down the original message within its host channel.



Compliance records captured for meetings hosted by other tenants contain slightly different information. To resolve the GUIDs used to record participants in the compliance records, remove the “&orgid:” prefix from the value and run the *Get-MgUser* cmdlet. For example:

```
[PS] C:\> Get-MgUser -UserId c7745bc8-6f5c-4d45-82c7-fee55f384985

DisplayName  UserPrincipalName
-----
Ståle Hansen stale.hansen_cloudway.no#EXT#@office365itpros.onmicrosoft.com
```

## What Teams Compliance Records Capture

It is important to understand that the compliance records captured by the Microsoft 365 substrate in user and group mailboxes are only copies of Teams messages. They are not the real messages (which remain stored in Azure Cosmos DB), nor are compliance records perfect replicas of those messages. Instead, the copies captured by the substrate are mail items that Microsoft Search indexes to make discoverable. If investigators find some issues in an eDiscovery case, they can access the real content in Teams to understand the full context of messages (the investigators will need access to the teams where the content exists to allow this to happen).

When the substrate captures copies of Teams messages, some transformation happens to create the compliance records. Not everything found in a Teams message ends up in the mail item. Because the copies created in Exchange Online are incomplete versions of the original data in Teams, you can't compare the copies generated by this process to email journaling. Elements copied to the mail item include:

- Links to any embedded emojis, stickers, inline images, and GIFs.
- Tables.
- Embedded deeplinks to other Teams messages.
- Sharing links to files in SharePoint Online document libraries.
- For channel messages, the compliance item records the subject of the message (if available) as is the name of the team holding the message. For personal chats, the compliance item captures the names of the people involved in the conversation.
- Code snippets in the body of messages. People can use code snippets to disguise conversations that they want to hide. Compliance records capture code snippets in Teams messages as .DAT attachments. The attachments store the HTML-formatted text for the code snippet.

However, Teams compliance records do not capture:

- Reactions (for example, a like, heart, or smile) to messages. In an eDiscovery context, reactions can be important signs that certain individuals have seen a conversation in the same way that changing the read status of an email from “unread” tells you that someone opened the message. According to [Microsoft 365 roadmap item 65130](#), Microsoft Purview Advanced eDiscovery will support the discovery of message reactions from August 2022.
- Recordings of audio messages sent from mobile clients.

Compliance records created for adaptive cards are in the form of mail items with one or more attachments. The attachments hold the adaptive card content. In addition, compliance records captured for *praise* messages only have the text of the praise and don't include the graphic, and compliance messages for messages with quoted text include the text but not the formatting marking the text as a quote.

## Versions of Compliance Records

If a retention policy is in place for Teams, the substrate captures changes made by a user to a message in the *Versions* sub-folder of Recoverable Items in the target mailbox. Thus, a content search might uncover several different versions of the same message, with the last version stored in the *TeamsMessagesData* folder being the content of the last update applied to the message (the earlier versions are in Recoverable Items).

## Compliance Records for Hybrid and Guest Users

Not every person who interacts with Teams has an Exchange Online mailbox in which the Microsoft 365 substrate can create and store compliance records. To get around the problem, the substrate creates special cloud-only Exchange Online mailboxes to store compliance records for:

- Hybrid users with Exchange on-premises mailboxes.
- Federated connections, such as external access to chat with Teams users in other domains or Skype consumer users.
- Guest accounts.

The substrate provisions a cloud-only mailbox (otherwise known as a “phantom” or “shard” mailbox) the first time it needs to create a compliance record for an external user. Cloud-based mailboxes are in the same data center region as the host tenant. The substrate uses the same approach to store compliance records for Yammer and Planner. Users cannot sign into these mailboxes, nor are they used for sending or receiving emails. They exist purely for storage purposes.

For example, if a Teams user starts a federated (external access) chat with a Skype consumer user, the substrate creates a mailbox called *Skypeld@teams.microsoft.com* (where *Skypeld* is the user’s Skype identifier) to store the compliance records.

On-premises users must be part of a hybrid organization that synchronizes their accounts with Azure AD using AAD Connect. The mail user accounts created for hybrid mailboxes must have both Teams licenses and Exchange Online P1 licenses. If a hybrid mailbox moves from on-premises to Exchange Online, the transfer process moves the compliance records from the phantom mailbox into the user’s new cloud mailbox.

Although Teams supports retention policies, those policies do not apply to cloud-based mailboxes. Likewise, you cannot apply retention holds to these mailboxes.

## Teams Compliance Outside the Tenant

When people collaborate in Teams, they can work with external users from outside their tenant. In this scenario, the rule is that the home tenant has all the compliance data. In other words, if an external user posts a message to a channel in your tenant, the substrate captures the compliance record in your tenant and no trace exists of their contribution within their home tenant. The reverse is also true: a tenant administrator has no oversight into what users from their tenant do inside other tenants.

Usually, this isn’t an issue unless compliance administrators need to access information about what users do in other tenants. For example, an investigation needs to know if someone discussed a confidential issue with people they should not have. A search can scan the Microsoft 365 repositories in the user’s home tenant, but any communications with the relevant individual in a channel (regular or private as a guest, or in a shared channel using their home account) remains invisible to eDiscovery performed in the home tenant. To find evidence, an arrangement must be made with the administrators of the tenant where the conversation occurred.

Administrators usually want to know when people are accessing data outside the tenant. To see signs of cross-tenant activity, you can check the Azure AD sign-in logs. This code uses the Microsoft Graph PowerShell SDK to find the sign-in log entries not related to your tenant. A filter extracts the records for shared channels.

```
[PS] C:\> Connect-MgGraph -Scopes "AuditLog.Read.All, Directory.Read.All"
Select-MgProfile -Name "beta"
$TenantId = (Get-MgOrganization).Id
Get-MgAuditLogSignIn -Filter "ResourceTenantId ne '$TenantID'" -All | ? {$_.CrossTenantAccessType -
eq "b2bDirectConnect" -and $_.AppDisplayName -eq "Microsoft Teams"} | Format-List CreatedDateTime,
UserDisplayName, ResourceDisplayName, AppDisplayName, ResourceTenantId

CreatedDateTime      : 20/03/2022 00:04:51
UserDisplayName      : Tony Redmond
```

```
ResourceDisplayName : Office 365 SharePoint Online
AppDisplayName      : Microsoft Teams
ResourceTenantId   : 22e90715-3da6-4a78-9ec6-b3282389492b
```

You can download a script to [find and analyze Azure AD sign-in records from GitHub](#). A similar approach is possible to find records for guest access (the filter used is `CrossTenantAccessType -eq "b2bCollaboration"`). See the PowerShell chapter for more information about using the Microsoft Graph PowerShell SDK.

The [Get-MSIDCrossTenantAccessActivity](#) function from the [MS Identity tools PowerShell module](#) is a good way of monitoring inbound and outbound cross-tenant connections for an organization.

## Call Records

Compliance records captured for meetings and calls depend on the Call Record Processing service. In telephony terms, the compliance records captured for calls and meetings are known as call detail records or CDRs. The substrate creates CDRs in the mailboxes of all call participants in the same hidden folder used to store the compliance records for messages. It can take up to eight hours before CDRs are available for searching. This usually is not a problem because eDiscovery searches normally happen after an event occurs.

CDRs do not include audio or video events in meetings. If you want this, record the meeting, and Teams will store the audio and video feeds for the meeting in an MP4 file in OneDrive for Business. CDRs capture textual details of a meeting or call such as:

- The start and end time of the meeting or call, and its overall duration.
- Notes of when each participant joined and left the meeting. Participants can join through VOIP or PSTN, and as anonymous, federated, and guest users. Teams assigns an identifier for anonymous joiners in the form `teamsvisitor:13c3a4132a584b918b9c6528b6dabef9`. It isn't possible to translate these identifiers into email addresses or other more identifiable addresses.
- Calls to voicemail.
- Missed or unanswered calls.
- Call transfers (which are represented as two separate calls).

An example of a CDR is:

```
Start Time (UTC): 6/20/2020 3:54:56 PM
End Time (UTC): 6/20/2020 5:05:28 PM
Duration: 01:10:31.9773921
```

```
[6/20/2020 4:21:17 PM (UTC)] teamsvisitor:d1cce625a5ee4a29911a7a954a89f1ee joined.
[6/20/2020 4:23:29 PM (UTC)] teamsvisitor:d1cce625a5ee4a29911a7a954a89f1ee left.
[6/20/2020 3:59:38 PM (UTC)] John.Hubbard@office365itpros.com joined.
[6/20/2020 5:05:27 PM (UTC)] John.Hubbard@office365itpros.com left.
```

Because CDRs are captured as email items, the person who starts a call or schedules a meeting is recorded as the sender, and participants are recorded as message recipients. The first part of the item's subject captures the kind of call or meeting while the remainder is a unique identifier for the event in the Teams media stack.

For example:

- Meeting (ScheduledMeeting).
- Meeting (RecurringMeeting).
- Call (Completed).

Microsoft plans to capture more details for calls such as screen and app sharing in the future.

## Storage of Teams Compliance Records

For both group and personal mailboxes, the substrate stores compliance records captured for Teams conversations in a special hidden folder called `TeamsMessagesData` in the non-IPM (system) part of user and group mailboxes. Because the folder holds compliance data that users should not be able to interfere with,

clients like OWA and Outlook desktop do not expose the folder in their GUI. If you need to examine compliance records, you can open the folder in a personal mailbox (but not a group mailbox) to examine individual items with a program like MFCMAPI. Exchange Online doesn't charge the storage used by Teams compliance records against user mailbox quotas.

Depending on load, the delay before a compliance record appears in the mailbox varies from a few seconds to a few minutes. Once captured in a mailbox, Exchange Online automatically indexes the compliance records to make them discoverable by content searches. To see how many Teams compliance records are in a user or group or user mailbox, run the *Get-ExoMailboxFolderStatistics* cmdlet to report the items held in the *TeamsMessagesData* folder:

```
[PS] C:\> Get-ExoMailboxFolderStatistics -Identity "Office 365 Adoption" -FolderScope NonIPMRoot
-IncludeOldestAndNewestItems | ? {$_.FolderType -eq "TeamsMessagesData"} | Format-Table Name,
ItemsInFolder, NewestItemReceivedDate, OldestItemReceivedDate
```

Name	ItemsInFolder	NewestItemReceivedDate	OldestItemReceivedDate
TeamsMessagesData	4227	02/02/2022 16:10:41	11/03/2017 15:41:34

The reason why we use the *IncludeOldestAndNewestItems* switch with the command is to force Exchange Online to return information about the oldest and newest items found in the folder. This is interesting information if you want to know if retention policies are removing items as their retention period expires.

In addition to compliance records captured for chats and conversations, Teams captures copies of email messages sent to a channel in a sub-folder called "Email Messages" of the channel folder in the SharePoint document library belonging to the underlying group. The messages are individual indexed files discoverable by content searches. Any files uploaded to the SharePoint document libraries used by Teams or to users' OneDrive for Business sites come under the usual compliance controls deployed to those workloads.

## Searching Teams Compliance Records

Content searches for Teams messages do not process Teams data stored in its Azure-based data services. Instead, searches depend on the Teams compliance records held in Exchange as those records are the data indexed and available for search. You do not have to do anything special to include compliance records in content searches because these items are searched whenever you include the mailboxes (group or user) or sites used by Teams in content searches. The sole exception is for users who have their mailboxes on on-premises Exchange servers. The content of these mailboxes is unavailable to content searches, so they cannot be included. In addition, you cannot apply a hold to content stored in on-premises mailboxes.

Figure 13-10 shows how Teams compliance records appear in the preview of items found by a content search. In this case, the compliance record captures details of a message posted to a channel. We know this by examining the item properties, where we can find:

- **From:** The display name and email address of the user who posted the message.
- **To:** The name of the team and the email address of the group mailbox used by the team or the names and email addresses of the participants in a personal chat.
- **Send date:** The date and time in UTC format when the sender posted the item.

As noted above, Teams compliance records for channel and personal conversations use the message item type of IM rather than the normal "Email" used for mailbox items. Note that messages created in channels by Office connectors also use the IM message type. Compliance records for meetings have the "Meeting" message type while personal calls use the "Call" message type.



## Contoso finance search samples

Subject/Title	Date	Sender/Author
No Subject	May 24, 2021 12:18...	Tony Redmond
<input checked="" type="checkbox"/> No Subject	Oct 17, 2019 10:08 ...	Tony Redmond
No Subject	May 24, 2021 12:19...	Kim Akers
Contoso Results	May 24, 2021 12:17...	Tony Redmond

Subject line: [Redacted]

Source

From: Tony Redmond <Tony.Redmond@...>  
 To: Corporate Acquisition Planning 2020 <CorporateAcquisitionPlanning2020@office365itpros.com>  
 Send Date: 17/10/2019 09:08:49 (UTC)  
[Download Original Item](#)

We can't do much if we don't know how well Contoso is doing in terms of their results. Does anyone have a snapshot to share?

Figure 13-10: Teams compliance records in a preview sample generated by a content search

**No Team or Channel Names Displayed for Compliance Items:** In the past, content searches displayed the names of the team and channel in the Subject field of items returned by searches. This information was useful in terms of finding the exact location of an item, but it's no longer displayed. The team name and email address are now the only clues as to the team where an item is found.

## Refine Searches for Specific Teams Items

The parameters for the content search shown in Figure 13-10 scans Exchange Online mailboxes to find any Teams compliance record matching the keyword. However, if it asked to search mailboxes without further qualification, the search will also find other non-Teams messages. To restrict search results to Teams content, we add the IM message kind condition to the search. This will find chat and channel conversations. Other search refinements are:

- Message kind set to *MicrosoftTeams* to find chat and channel conversations and call detail records.
- Keyword: *Itemclass:IPM.AppointmentSnapshot.SkypeTeams.Meeting* to find call detail records for meetings.
- Keyword: *Itemclass:IPM.AppointmentSnapshot.SkypeTeams.Call* to find call detail records for calls.
- Keyword: *Itemclass:IPM.AppointmentSnapshot.SkypeTeams.Meeting OR Itemclass:IPM.AppointmentSnapshot.SkypeTeams.Call* to find all call detail records.

Note: when you create a content search, make sure to set the "add app content for on-premises" checkbox to allow the search to include compliance records generated by hybrid and guest users.

## Searching Other Microsoft 365 Workloads for Teams Content

Remember that Teams content exists in many other places than channel conversations and chats. To do a complete search, you might need to include:

- OneDrive for Business accounts to find files shared in personal and group chats.
- SharePoint Online sites for files shared in channel conversations.

Some information used by Teams will remain unavailable to content searches until applications support indexing. Whiteboards shared during meetings are a good example of information unavailable to content searches.

## Reassembling Conversations

The way that the Microsoft 365 substrate captures individual compliance records for each message posted to a chat or channel conversation makes it more difficult to reconstruct a full conversation from start to finish. Normally, when eDiscovery investigators review information, they want to see the full context to understand who said what to whom and how a conversation developed. A full message thread gives context by recording the different interactions of participants in a conversation. Email can do this too by including the text of prior replies in a message.

The Microsoft 365 Advanced eDiscovery and Communications Compliance services are both able to reconstruct a Teams conversation from discovered messages and present the conversation in the same manner as it appears in a Teams client. This isn't possible with Core eDiscovery or simple content searches as these operations return individual compliance records instead of complete conversations. However, it is possible to use a manual process to construct a complete conversation from the compliance records captured by Teams. To do this, you must:

- Search for and find the compliance records for all contributions to the conversation. You won't find all the compliance records belonging to a conversation unless you use its reply chain identifier as a search keyword. For example, you could search for 1529919175913 to find all items in the conversation with that reply chain identifier.
- Because items from all channels are in the same mailbox, be careful not to mix items from different channels together.
- Arrange the items in date order using the timestamp. Some commercial products combine compliance records to form threads when they present search results.

You can export the information found in Teams by a content search to PST files or as individual items, and then give the PST or items to external investigators for their review. See Chapter 18 for more information about how to build and run content searches and export items found by the searches.

If a Teams meeting is recorded and a transcript is produced, Microsoft Search can look for information in the transcript. Content searches can also find videos based on transcript text, but the content search UI doesn't display the transcript. The found items are the recordings (MP4 files) rather than the transcript. If you download an item, the compliance center creates a file without an extension. You can rename the file to give it an MP4 extension and it will then play as normal. Microsoft expects to correct these issues and deliver full capability to search transcript text via Microsoft 365 eDiscovery in 2022.

## Searching Hybrid and Guest for Teams Compliance Records

If a guest user sends a message to a tenant user in a chat, the substrate captures two copies of the message: one in the cloud-only mailbox created for the guest, the other in the tenant user's mailbox. Searches executed through the Microsoft 365 Compliance Center automatically include Teams messages sent by hybrid and guest users using the copies kept in user or group mailboxes but do not scan the mailboxes belonging to these users.

Two ways exist to search compliance items stored in guest mailboxes. You can:

1. Amend search settings to set the checkbox to include *"Add app content for on-premises users."* The checkbox covers the cloud-based mailboxes created for hybrid, federated, and guest users. This is the default setting for new content searches.
2. Use PowerShell to update the search criteria to include app mailboxes.

Content searches find Teams compliance records in cloud-based mailboxes when two parameters are set to *\$True*:

- **AllowNotFoundExchangeLocationsEnabled** controls if the search covers Exchange Online mailboxes that cannot be verified. These are the cloud-only mailboxes used by hybrid and guest accounts.
- **IncludeUserAppContent** controls if app content is searched. App content means the cloud-based mailboxes used by guest and hybrid users.

You can set these parameters for a new content search with the *New-ComplianceSearch* cmdlet or apply them to an existing search with the *Set-ComplianceSearch* cmdlet.

For example, these commands create a search including cloud-only mailboxes and then start the search.

```
[PS] C:\> New-ComplianceSearch -Name "Teams Chat Scan" -Description "Search for Teams Chat Information about Finance" -IncludeUserAppContent $True -AllowNotFoundExchangeLocationsEnabled $True -ExchangeLocation All -ContentMatchQuery "Finance AND Kind:MicrosoftTeams"
```

```
Start-ComplianceSearch -Identity "Teams Chat Scan"
```

After the search completes, you can use the content search GUI in the Microsoft 365 Compliance Center to preview search results, refine the search keywords and qualifiers, and export results for further investigation.

## Teams and Retention Policies

Teams messages are persistent and remain in the Azure Cosmos DB data store until removed by user action or through retention processing. Channel messages belong to the channel and don't disappear following the removal of the authors from Microsoft 365 (if Teams can't resolve the user who posted a message because their account no longer exists in Azure AD, it displays the message as posted by an "unknown user"). Chat participants collectively own the messages in the chat. A participant can delete their link to a chat message, but the message will remain available to the other participants.

Retention policies can remove messages from Teams. An organization can create retention policies to process Teams chats and/or channel conversations. These policies are separate from the retention policies used to control other Office 365 data such as documents and email and instead operate against the compliance records stored in user and group mailboxes. Retention policies for email do not process the contents of the *TeamsMessagesData* folder.

A background retention assistant job processes retention policies for non-Exchange workloads, including the compliance records created by Teams activity and subject to Teams retention policies. The assistant uses the settings in Teams retention policies to remove compliance records when the items expire. When the assistant removes compliance records, the Microsoft 365 substrate synchronizes the deletions back to the Teams chat service, which then removes the relevant messages from its store. Later, Teams clients connect to the Teams service and synchronize their local cache to complete the removal cycle.

If an in-place hold or litigation hold applies to some group or personal mailboxes which include compliance records, those items come under the scope of the hold. The substrate captures any attempt to remove or edit a compliance record by keeping the copies of the removed or edited item in the *\Purges* or *\DiscoveryHolds* sub-folders under Recoverable Items in the mailbox. However, Teams removes the items from its message store in Azure Cosmos DB.

When you create a retention policy for Teams, you can choose to keep messages for the chosen retention period (a minimum of one day) or remove items after the retention period elapses. What you cannot do is give users the ability to mark specific messages to force retention policies to remove those items sooner or keep them longer. SharePoint and Exchange support this kind of flexibility through retention labels or mailbox personal tags.

# Auditing Teams

In addition to capturing compliance records for individual and channel conversations, Teams generates audit records for many user and administrative operations. The most common audit record is to record each time a user, including guests, signs in. Teams captures an audit record for a user sign-in for every hour in a session. This is because the access token expires after an hour. After Azure AD renews the access token, the user signs in again. The user does not notice the sign-in happening as token renewal and reconnection happens in the background. The audit record tells you when the user signed into Teams and some information about what client they used, but it does not capture details of which team or channel they accessed.

Among the administrative activities that Teams captures audit records for are:

- Team creation and deletion.
- Channel creation and deletion.
- Users added or removed from teams.
- Settings changed for the organization, team, or channel.
- Tab addition and deletion (for a channel).
- Connector addition, deletion, and updates.

For example, audit events are captured for the creation and removal of teams and the creation of channels within teams. For instance, when you use a Teams client to create a new team, the audit log receives an *"add group"* event when Azure AD creates the new group followed by some *"update group"* events when it populates the properties of the new group. Finally, you see a *"TeamCreated"* event for the new team. Audit events are also recorded for any alteration of team settings. However, a *"TeamCreated"* audit event is not captured when you enable an existing group for Teams.

The audit records for the creation of teams and channels capture the names given to the new teams and channels and who created the object. However, the records for tab creation only tell you that someone created a tab. No information is captured about the content the tab links to. For example, if you create a tab to link to a YouTube video, the audit record tells you that the tab type is *"extension"* but not that it links to a video or what the video is (apart from third-party apps like YouTube, the extension tab type covers Microsoft applications like Planner and Stream too). No information is recorded about the content either when someone updates a tab. The lack of data is sometimes explained by the need to protect user privacy, and sometimes because tenants and Microsoft need to define what information Teams should capture (and why) for audit purposes.

Teams periodically uploads its audit data to the audit log. You can interrogate audit records using the audit log search in the Microsoft Purview compliance portal, by running the *Search-UnifiedAuditLog* cmdlet, or with third-party security products. Techniques for using these tools are explained in Chapter 21.

## Teams and the Groups Expiration Policy

Removing old Teams when their period of usefulness expires is part of compliance planning. If your tenant uses an Azure AD group expiration policy (Chapter 11) and the group used by a team comes within the scope of the policy, an extra section called **Team Expiration** (Figure 13-11) is visible to team owners under the Settings tab to inform owners about when the team needs renewal. A team owner can click the **Renew Now** button to extend the lifetime of the team for whatever period is set in the policy (normally a year or two). Renewal can happen at any time.



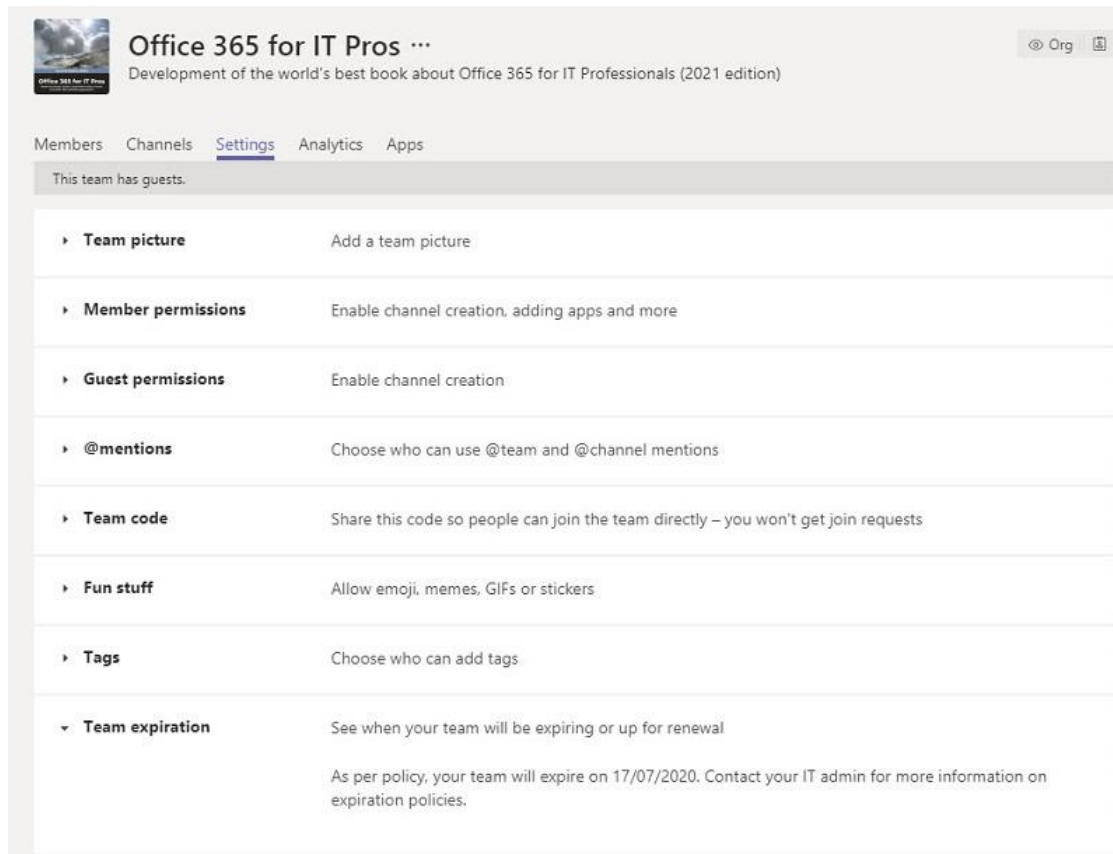


Figure 13-11: Showing the expiry date for a team

When a team is within a month of its expiry date, Teams highlights its potential expiration with a warning posted to the activity feed of the team owners. In addition, a warning triangle appears beside the team name in the navigation pane (but only to team owners). Hovering over the warning triangle reveals the expiry date and the choice to **Renew team** appears in the ellipsis menu. Because the Groups expiration policy auto-renews groups if they are active (and posting topics and replies in team channels is good evidence of activity), you shouldn't see warnings unless the assessment against the policy settings considers the team to be inactive. This can happen for teams used to host connectors where the number of human posts is low. If you see a warning, go to the Manage team options for the team and renew it there.

If a team reaches its expiry date and isn't renewed, Azure AD soft-deletes the group and all associated resources, including the team, and removes the ability of users to access those resources. An administrator can restore the team at any time within 30 days of its being soft-deleted. Once the 30-day period elapses, Azure AD permanently removes the group and all its associated resources, and the group becomes irrecoverable.

## Archiving Teams

Expiring a team removes it from the tenant. Archiving a team is another way of dealing with teams that are no longer active. When you archive a team, you make the elements controlled by Teams (like channel conversations and the wiki) read-only. Users can access messages in an archived team, but they cannot post new messages, edit messages, or remove messages from a channel. In addition, members can access files in the document library belonging to the team, but they cannot upload new documents or remove files from the library, and the link to open the document library in the SharePoint browser interface is not available. When users open an archived team, they see notices to tell them that they can no longer post to the team. Teams also displays an icon (a closed drawer) alongside the names of archived teams in the team list). The idea is

that the team remains available to its membership to allow users to continue accessing information while not being able to add to that information. You can update the membership of an archived team to add or remove members, including guests, or promote members to be owners.

Three methods are available to archive a team:

- In the desktop or browser client, select Teams, then the **Manage** cogwheel icon at the bottom of the list of teams to display the set of teams you belong to, divided into active teams and archived teams. This option is available to all team members, including guest accounts. However, only team owners, tenant administrators, and Teams service administrators see the option to **Archive team** in the ellipsis menu for the team (Figure 13-12). To restore an archived team and make it read-write again, select it in the list of archived teams and then choose **Restore team** from the ellipsis menu.
- In the Teams admin center, find the team in the **Manage teams** section and then select the **Archive** option. Only those assigned a Teams administration role can archive a team using this method.
- Use the Teams PowerShell module to archive a team. For example, to see the set of archived teams:

```
[PS] C:\> Get-Team -Archived $True
```

To archive a team, run the *Set-TeamArchivedState* cmdlet and pass the identifier to the team to archive. The *SetSPOSiteReadOnlyForMembers* setting controls if the SharePoint site belonging to the team is set to read-only.

```
[PS] C:\> Set-TeamArchivedState -GroupId $GroupId -Archived $True  
-SetSPOSiteReadOnlyForMembers $True
```

To reverse the process, set the *Archived* switch to *False*. You can only update the *SetSPOSiteReadOnlyForMembers* setting if it was set when archiving the team:

```
[PS] C:\> Set-TeamArchivedState -GroupId $GroupId -Archived $False  
-SetSPOSiteReadOnlyForMembers $False
```

Before archiving a team, it's important to check the status of any private or shared channels in the team. Administrators or team owners don't have access to the content of these channels unless they are channel members. A private or shared channel can be very active without the knowledge of the team owner, and if the team is archived at this point, team members (including external and guest members) won't be able to post new content. The Teams desktop client doesn't display any details about private and shared channels in a team. This information is available in the Teams admin center and can be retrieved by querying the team channel configuration with PowerShell using the *Get-TeamChannel* cmdlet.

After archiving a team, Teams moves it into the set of hidden teams displayed at the bottom of the teams list. Removing a team from the active set effectively makes the archived team invisible to users unless they go looking for it by opening the hidden set or by using the *Manage teams* option to find the team in the archived section. If you restore an archived team, Teams makes it writeable again but leaves the restored team in the set of hidden teams.

When archiving a team, you can choose to make its SharePoint site read-only for team members (this also sets the wiki to be read-only). Team owners can continue to upload and update content, but team members who access the site after the team is archived have restricted options because Teams adjusts the site permissions for team members to remove their write access. For example, members cannot upload files to the library, rename or remove files, update document details, assign retention or sensitivity labels, and so on. They can still synchronize the library and download files. The sites belonging to private and shared channels also become read-only (even for channel owners), and in these instances, SharePoint displays a banner to inform channel members of the site's read-only state.

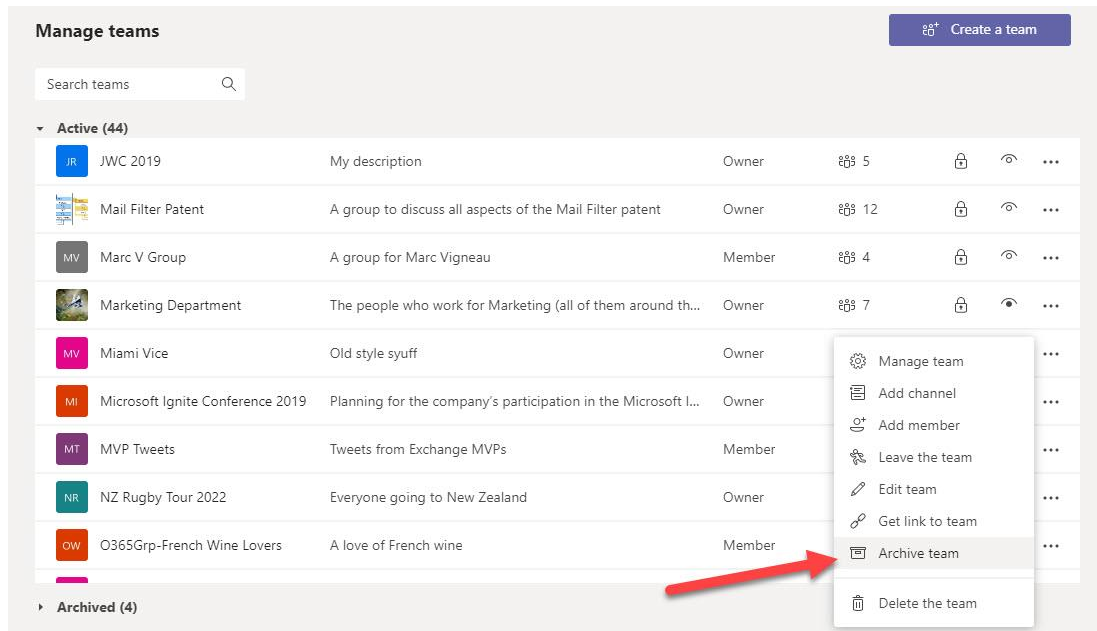


Figure 13-12: The option to archive a team

Setting read-only access to the conversations and files belonging to a team is an effective way of putting it into an archive status, but a big selling point for Teams is its ability to be an integration point for other third-party applications. To make archive status fully effective, every application connected to Teams must understand when a team is archived. SharePoint does this, but the other connected apps don't, which means that team members can continue to have read-write access to other apps like Planner, OneNote, and apps added to the team as tabs, and bots. This makes sense for third-party apps as they might be shared across multiple teams.

Archiving a team does not stop it from expiring if it is within the scope of the group expiration policy. The team still exists, albeit in a read-only state; expiration kicks in once the renewal time of the underlying group is reached and if the group is not renewed, Azure AD soft-deletes the archived team.

The mechanism used to archive teams is reasonable. Its biggest advantage is that members continue to enjoy access to channel conversations. A different approach is outlined in the PowerShell chapter showing how to archive teams by removing everyone but a single owner from group membership. The advantage of this approach is that access to all team resources is removed from previous members, which is what you might want to achieve.

## Reporting Teams Usage

Microsoft 365 includes several methods to gain an insight into the activity levels of Teams:

- The standard usage reports available in the Microsoft 365 admin center include a Teams usage report. The data can be viewed in 7-day, 30-day, 90-day, or 180-day snapshots. The data is available for the overall tenant, per-team, and per-user level (in a sortable table) and is exportable to a CSV file.
- The **Analytics and Reports** section of the Teams admin center includes a set of usage reports. The Teams admin center doesn't support sorting of the report detail, but you can export the information to a CSV file to analyze and report the data as you wish. Reports are available for the last 7, 30, or 90 days, except the PSTN data, which is available for up to 28 days. The same data is available for device usage and user activity in the Microsoft 365 admin center. The reports include:
  - Teams user activity (number of channel messages, chat messages, 1:1 calls, video and audit minutes, and date of last activity).
  - Teams device usage (Windows, Mac, iOS, Android).

- Teams usage (active users, guest members, active channels, number of messages, team privacy setting).
- Teams Live Event usage (event, start time, organizer, presenter, producer, and views).
- PSTN blocked users.
- PSTN minute and SMS pools.
- PSTN and SMS usage.
- Information protection license report. This has nothing to do with Microsoft 365 Information Protection and instead deals with if users have the [necessary license for apps to receive subscriptions](#) to the endpoints tracking changes to messages and chats.
- Apps usage (number of active users per app, number of teams where the apps are used).
- Analytics for individual teams are available through the Manage team option in the Teams desktop and browser clients (Figure 13-13). The data reveals activity for the selected team over a 7, 30, or 90-day period. The information is available to any member of a team, including guests, and is intended to give an overall view of the activity within a team together with some other information such as the breakdown of membership into active and inactive members and tenant users and guests. You can view data for all channels or select the data for an individual channel (including private channels). The data available for a channel covers the number of topics and replies posted. Team analytics track active users over time, the split between tenant users and guests and members and owners, the number of apps installed in the team, and meeting activity. The data doesn't include messages posted to channels by Office connectors (like an RSS feed) or those that come in through email. Analytics for an individual channel is also available through the Manage channel menu. Teams doesn't include a facility to export analytics data from a client, but a Graph API is [available and can be used to generate downloadable data](#). The report also lists the top five inactive channels in the team. Typically, the channels listed have had no activity in over a month, but this depends on the number of channels in the team.
- The [Microsoft 365 usage analytics content pack](#) for Power BI includes a Teams usage report. Like the usage data for all workloads available in the content pack, the Teams data can be up to a month or so behind, depending on when Microsoft 365 last refreshed the data. The content pack exists to track usage over a sustained period and that is how to use it.

Usage data is collected and collated by background processes and stored in the [Microsoft Graph](#), so the information available for analysis through any of these options is at least two days old. The analytics and reports available in the Teams client and the admin centers present data in different ways, so you might observe some minor inconsistencies when comparing the data viewed in one interface against another. Remember that third-party reporting alternatives are available (see Chapter 21) that usually offer more powerful and flexible analysis and reporting functionality than is available in Microsoft 365.

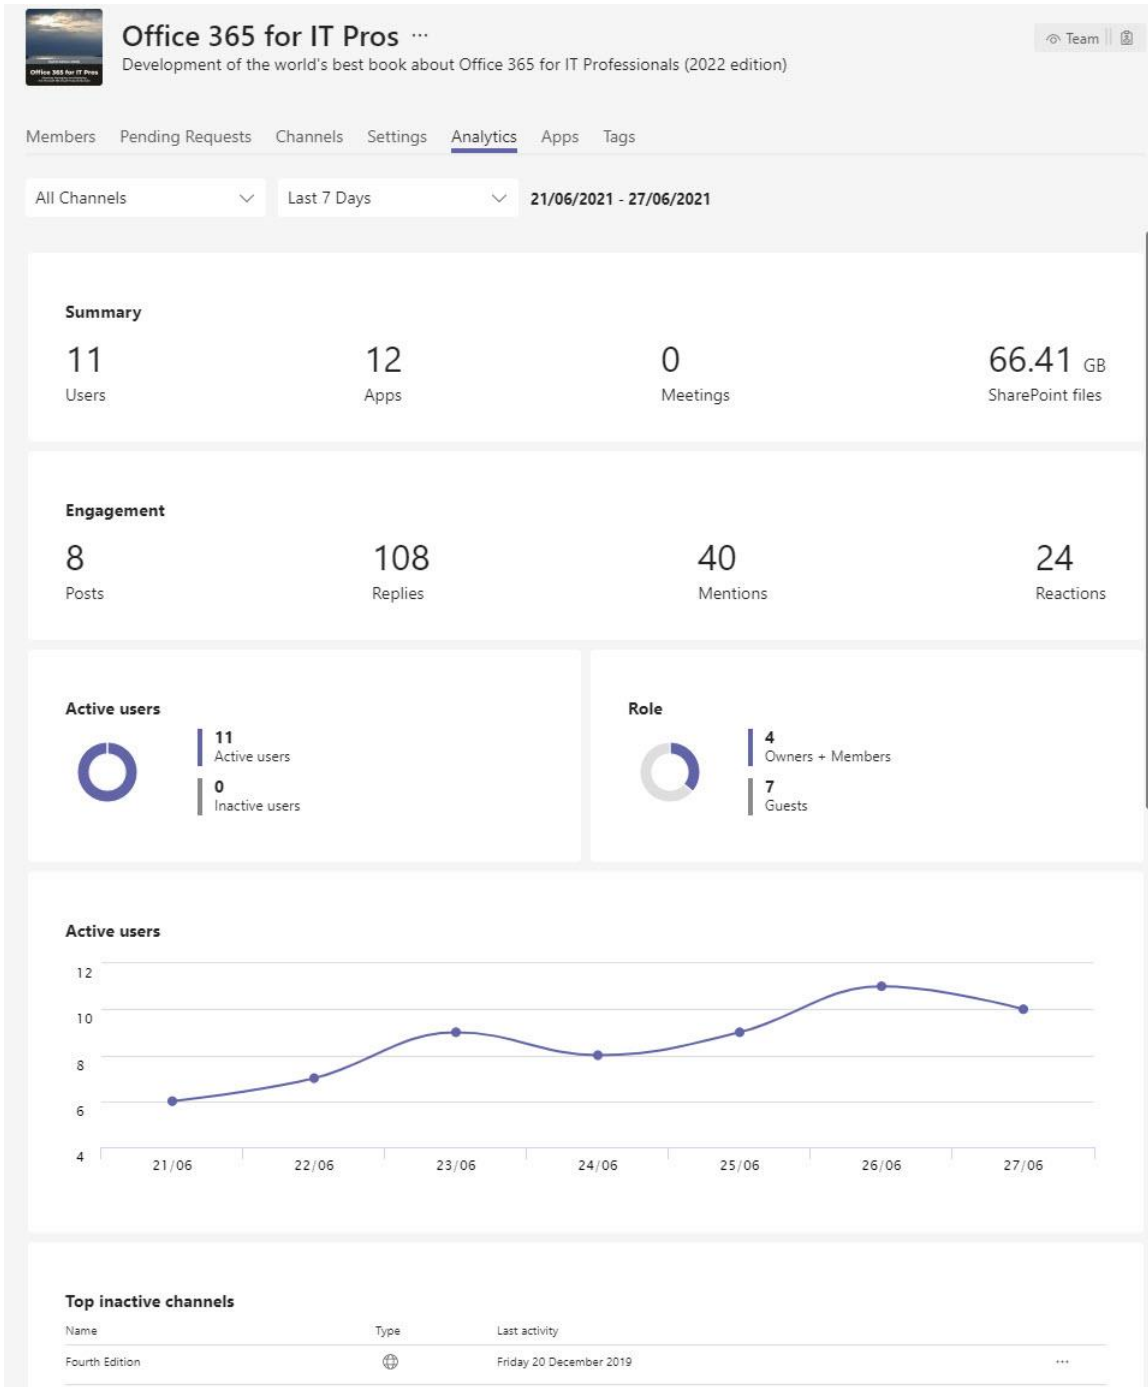


Figure 13-13: Analytics for an individual team (all channels)

## Extending Teams

Out-of-the-box, you can extend Teams by adding tabs, apps, bots, and connectors to channels. Bots fit into conversations to allow users to interact via chat with software assistants, usually to answer questions about specific topics. Connectors set up links between channels and network data sources so that information flows from the data source and appears in the channel like user contributions.

Microsoft has a [developer program for Teams](#) focusing on different aspects of integration like:

- Extending the tabs available for a channel to incorporate access to an application.
- Adding a new bot to answer questions posed by team members.

- Creating a connector to a network data source.

MVP Tom Morgan publishes [Building and Development Apps and Bots for Microsoft Teams](#). It's a good starting point for app developers who want to build apps using the Teams platform. Given the attractiveness of the platform and the hundreds of millions of active Teams users, it is no surprise that the number of Teams apps increases over time. As of mid-June 2022, over 1,600 apps are available in the Teams app store, and the number of apps increases by approximately 80 per month.

The [Microsoft 365 app certification program](#) helps customers to understand what apps they might like to run in Teams by publishing developer-provided information about the apps. Three levels are available:

- **Publisher verification:** The app developer has a Microsoft developer network identity. The app supports modern authentication and is capable of multi-tenant activity. This is the entry-level participation in certification.
- **Publisher attestation:** The app developer completes a questionnaire covering security, data handling, and compliance.
- **Microsoft 365 certification:** Instead of the app developer reporting details of their app, third-party assessors audit the assertions to validate that the app meets Microsoft standards for security and compliance. The process occurs annually, and details gathered during the audit are available online. The audit information is available through the Microsoft 365 certification link for the app. Abode Sign is an example of a Microsoft 365 certified app.

Regretfully, many apps (including some created by Microsoft) are uncertified and few app developers go through the full Microsoft 365 certification process. This might be a question of resources or a lack of perceived value in achieving certification. However, the audit information gathered by the attestation procedure is valuable and worthwhile knowledge for customers interested in the data accessed by apps.

## Managing Teams Apps

Not every organization is willing to allow users free rein over the apps they can install and use with Teams. To exert control over the apps available to users, three mechanisms are available in the **Teams Apps** section of the Teams admin center:

- **Manage apps:** Define which apps are available in the tenant. By default, any app published in the Teams app store is available to a tenant. Some apps might be inappropriate or not very useful for an organization, so tenant administrators can block individual apps here. Blocked apps can't be installed by users and won't be displayed in the app navigation bar if included in an app setup policy. To decide whether to allow an app, you can check its properties:
  - *Publisher:* The organization responsible for creating and maintaining the app.
  - *Version:* The current version of the software.
  - *Categories:* For example, Productivity or Business Management. The publisher chooses these categories as general guidance for the type of solution an app is.
  - *Certification:* The highest level of certification is [Microsoft 365 certified app](#), which means that Microsoft has reviewed and approved the app against a set of security, compliance, and data handling standards.
  - *Licenses:* ISVs can generate Teams apps that support per-user licensing (monthly or annually) or in-app purchases. The Plans and Pricing tab for the app tells administrators what licensing is available. Some apps offer free versions that users can upgrade to a premium version with a license.
  - *Capabilities:* Where in Teams the app can be used. If *Team*, the app can be installed into a team channel. Other categories include *Personal*, meaning that a user can install the app for

their personal use, and *Group chats*, meaning that the app can be installed to be shared by participants in a group chat.

- *App Id*: Each app gets a unique identifier (GUID) during the publication process to make the app available in the Teams app store. The identifier is the same for all tenants.

To block an app, select it from the app inventory and move the App status slider from *On* to *Off*.

If an app is scoped for installation into a team (normally as a channel tab), the administrator can install it into selected teams.

- **App Setup policies**: Controls the set of apps displayed in the Teams app navigation bar. Setup policies also control if users can pin apps or upload custom apps. See the section below.
- **App Permission policies**: Apps come from multiple sources and not every app is suitable for every user. App permission policies allow tenants to control the set of allowed apps that individual users can install in Teams.

The Manage apps page also includes **org-wide app settings** for apps. These settings are tenant settings to:

- **Allow third-party apps**: If *On*, third-party apps can be installed by users. Turning this setting to *Off* prevents users from installing third-party apps and limits them to apps provided by Microsoft.
- **Allow new third-party apps published to the store by default**: If *On*, any new third-party apps published to the Teams app store are visible in the tenant's Teams app store and are available to users if the app permission policy assigned to their account allows third-party apps. If *Off*, new third-party apps do not appear in the app store.
- **Allow interaction with custom apps**: Custom apps are those developed for your organization. If this control is *On*, users can install and access custom apps. If *Off*, they cannot (this setting also disables [outgoing webhooks](#)). This setting is *Off* by default for GCC tenants.

Microsoft creates default app setup and app permission policies (both called Global (Org-wide default)). These policies are assigned automatically to Teams users and stay in place unless an administrator updates the policy assignment for an account. You can amend the default policies or create new policies to assign to individual users or sets of users to match organizational requirements.

## Customizable Apps

Depending on the manifest settings published by an app's developers, the settings of an app can be customized for an organization. For instance, you can replace the app's icons with versions that include corporate branding. When an app is customizable, it is marked as such in the Teams admin center and an administrator can update the following settings:

- **Short name**: a 30-character app name.
- **Short description**: an 80-character description of what the app does. You can also add a full description. I've used descriptions of over a thousand words.
- **Privacy policy URL**: You can use this link to point to an appropriate page on the organization's website where users can find information about the company's privacy policy.
- **Website URL**: This is usually the URL for the main landing page of your company's website.
- **Terms of use URL**: Typically points to a web page where users find information about the terms of use for the application. Normally provided by the app developer.
- **Color and outline icons**: The color image must be smaller than 1,000 KB and must be a PNG file and should be 192 x 192 pixels. The outline icon is 32 x 32 pixels. The color app shows up in the Teams app store. Teams uses the outline app to display in the app navigation bar.
- **Accent color**: A hex code defining the background upon which Teams displays the app's color icon. Here's a [website](#) to help find the code for the background color.

Some experimentation might be needed to identify the best icons and colors.

## Applications and Permissions

Administrators (or users, if access is only needed to that user's information) must grant consent to applications based on the Graph API to use the required permissions before they can access data such as teams, channels, and messages. Although some apps do not need to access tenant data, many third-party and LOB apps interact with data from Teams or other workloads which they access using the Microsoft Graph APIs. Before they can access any tenant data, an app must be granted consent by a tenant administrator. Consent is managed through the Permissions tab of app properties, divided into two sections:

- **Org-wide permissions:** Allow access to org-wide data, such as all sites, users, or teams in the tenant.
- **Teams Resource-specific-consent (RSC) permissions:** These are permissions granted to access [certain types of Teams data](#) (like channel settings or the tabs available in a team) scoped only to the team in which an app is installed. Some [Azure AD configuration](#) is necessary to allow RSC to work. Trello, Zoho CRM, and Template Chooser are examples of apps that support RSC. Not all Graph-based applications support resource-specific consent permissions. The ability to use resource-specific consent permissions is enabled by default at the tenant level. A tenant can enable or disable the feature by blocking team owners from granting consent to apps. To do this, go to the Enterprise application section of the Azure AD admin center and set **Group owner consent for apps accessing data to Off under User consent settings**.

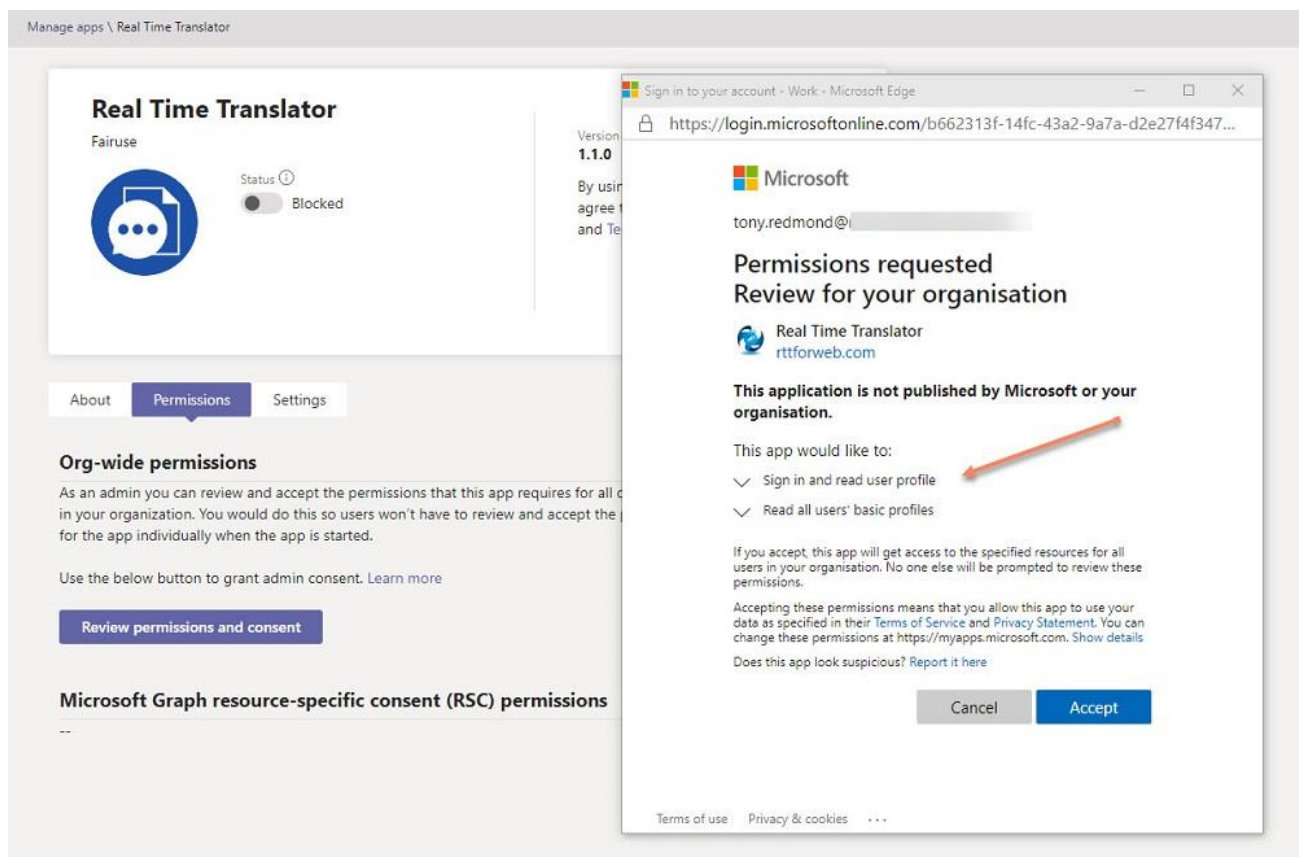


Figure 13-14: Granting consent to permissions requested by an app

When an administrator grants consent to the permissions requested by an app, they do so on behalf of the tenant. This has the advantage that individual team owners won't be asked to grant consent when they install an app into their team. Often, team owners don't have the necessary administrative rights to grant the necessary permission, which causes a delay in making the app available to users. It's also the case that team



owners might be unaware of the danger of granting consent, especially org-wide permissions, to an app published by a company they don't know.

Figure 13-14 shows the screen used to display the permissions requested by an app during an administrator review. In this case, the app is requesting two org-wide permissions needed to read user profiles. To grant consent, the administrator clicks **Accept**. Teams then updates the app's service principal in Azure AD. App consent can be withdrawn at any time by editing the app details through the [Enterprise applications blade](#) in the Azure AD admin center.

Upon receiving consent, the permissions granted to the app are registered in Azure AD and Microsoft Graph allows the app to access data within the scope of the permissions. See [this article](#) for more information about app permissions.

## App Templates

To illustrate the capabilities of the platform, Microsoft has created a set of production-ready apps for Teams. This is both a learning tool and a way to accelerate app development for Teams. The idea is that organizations can browse [the available set](#) to find an app that's similar in concept (or an exact match) for what they want to do and then download the code base and deployment scripts for the chosen app from GitHub before tailoring the code to meet their requirements. The GitHub repository for each app template also includes documentation. Among the available templates are:

- Ask Away: A bot to conduct question and answer sessions within Teams.
- Building Access: Administration of access to buildings.
- Company Communicator: Create and send messages to multiple teams or large numbers of users over chat.
- Employee Ideas: Allow employees to submit ideas.
- Expert Finder: A bot to find someone in the organization based on their skill set (depends on data registered in Azure AD).

## Teams App Setup Policies

The Teams app navigation bar contains a set of pinned apps and an ellipsis [...] menu to allow quick access to other apps. The apps in the navigation bar are controlled by the Teams app setup policy assigned to the user account. [App setup policies](#) are managed through the **Teams Apps** section of the Teams admin center. An app setup policy describes the set of apps to show in the app navigation rail (left-hand side in desktop and browser clients, bottom of the mobile client) and the order in which the apps appear. An app setup policy also controls if the users assigned the policy can pin apps or upload custom apps.

A typical app setup policy includes all or some of the default Teams apps (like Chat and Teams) together with other apps from the Teams app store or the organization's app catalog. The idea is that you can build sets of apps appropriate for different kinds of users and assign them via policy. The default Teams app setup policy (called Global) is used unless another policy is assigned. Teams also installs an out-of-the-box policy for front-line workers in each tenant.

In Figure 13-15, we see that the set of core Teams apps are rearranged so that Teams is at the top of the list in the app bar. Planner is a non-standard app that has been pinned to the bottom of the list. The user can select additional apps to pin to the app navigation rail from the Apps menu. Any of the core Teams apps unpinned from the navigation bar remain accessible in the Apps menu.

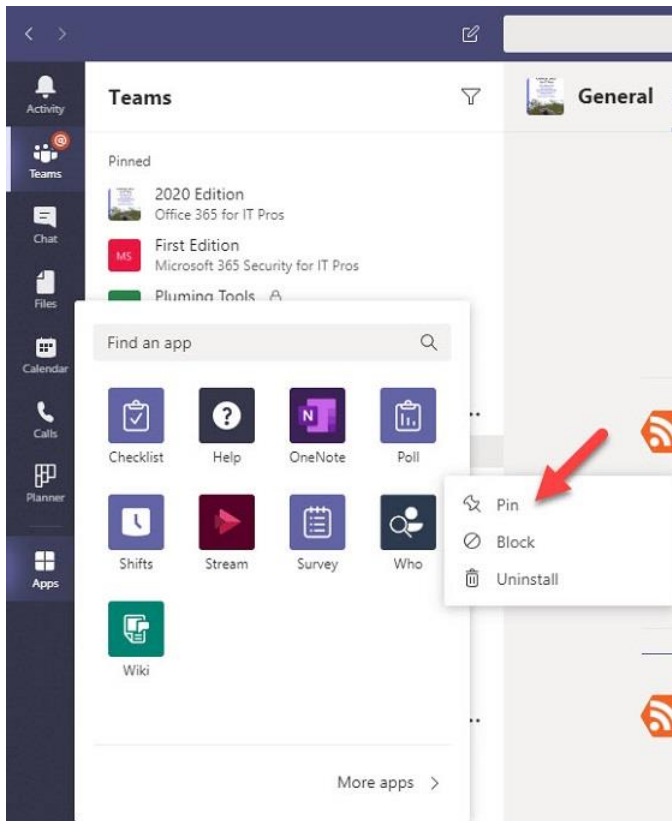


Figure 13-15: Selecting an app to pin to the app bar in the left rail of the Teams client

The *Allow user pinning* setting in an app setup policy controls if users can add (pin) or remove (unpin) apps to the app bar. By default, the setting is On. Update the policy setting to Off if you don't want users to be able to select apps for the app bar. Clients won't implement the new settings until they refresh their copies of policies, which can take an hour or so. Guest users don't have app setup policies and are limited to the apps they can access (Teams, Chat, Activity, and Files).

If an administrator changes the app setup policy assigned to an account, Teams notifies the user that their settings have been updated and that some of their pinned apps might have been moved.

## Analyzing App Usage

The Apps usage report in the Reports section of the Teams admin center helps administrators to understand the effectiveness of Teams app setup policies. For instance, if you see that an app displayed in the navigation pane isn't being used, it's a prompt to either make people aware of the app or to consider replacing the app in the pane. Likewise, if you see an app in heavy use that isn't shown in the navigation pane, perhaps it's time to include it there. The most effective way of analyzing app usage is to download the data to Excel and review it there. The apps with the heaviest usage are likely to be Teams (channels and conversations), Activity (the feed), Files (SharePoint Online), and other standard Microsoft apps found in all Teams installations.

## Assigning App Setup Policies

App setup policies can be assigned individually to accounts in the Teams admin center or via PowerShell. It's usually more convenient to use PowerShell to assign a new policy to multiple accounts. In this example, the set of users for the marketing department is selected and an app policy designed for that department is assigned to each account.

```
[PS] C:\> $Users = (Get-CSOnlineUser -Filter {Department -eq 'Marketing'})
Foreach ($U in $Users) {
    Write-Host "Assigning Teams App Setup Policy for the Marketing department to" $U.DisplayName
```

```
Grant-CsTeamsAppSetupPolicy -PolicyName "Marketing" -Identity $U.UserPrincipalName }
```

# Teams App Permission Policies

App permission policies control the set of apps available to users. Each policy can give access to a separate set of apps. After you assign an app permission policy to a user, they can install any of the apps covered by the policy. An app permission policy can't override a block set in the org-wide app settings. Management of app permission policies is through the Teams apps section of the Teams admin center.

If you want to allow access to different apps, you can customize the set of apps defined in the global app permission policy or create a new app permission policy and assign it to selected accounts. An app permission policy covers three types of apps:

- Microsoft apps.
- Third-party apps. Some of these apps support the purchase of applicable licenses through the Teams admin center.
- Custom or tenant apps (apps published and owned by the organization).

For each category, you can decide to:

- **Allow all apps.** Users can install and use any app of the type published in the Teams app store.
- **Allow specific apps and block all others:** The administrator selects the apps that users can install and use. Teams doesn't allow users to install blocked apps in the Teams app store.
- **Block specific apps and allow all others:** The administrator blocks selected apps available in the Teams app store and makes them unavailable to users.
- **Block all apps:** Users cannot install apps of this type.

When you restrict the set of apps with an app permission policy, Teams filters the set of apps, bots, and connectors it displays to users assigned that policy. Due to caching, it can take up to a day before Teams clients respond to a change in the set of apps allowed to users or a change in the policy assigned to an account.

Users can request administrators to allow them to install blocked apps. Administrators don't receive notifications of these requests. Instead, they must check for waiting requests in the Teams admin center by looking for apps with open requests from users. After finding a request, the administrator can unblock the application and add it to the relevant app permission policy to allow the user to install the app. After clients download the updated policy, users can install the unblocked app.

Once a user installs an app in a chat, the app becomes available for use in any other chat without the need to reinstall the app from scratch.

## Assigning App Permission Policies with PowerShell

The *Get-CsTeamsAppPermissionPolicy* cmdlet lists the Teams app permission policies available in the tenant. Use the *Grant-CsTeamsAppPermissionPolicy* to assign a policy to a user. For example:

```
[PS] C:\> Grant-CsTeamsAppPermissionPolicy -PolicyName "Unrestricted App Access" -Identity Kim.Akers@office365itpros.com
```

To assign an app permission policy to the members of a team, we need to retrieve the members of the team and then assign the policy. Here's how to do it using the *Get-Team* and *Get-TeamUser* cmdlets from the Teams PowerShell module followed by a call to *Grant-CsTeamsAppPermissionPolicy* to assign the policy to the individual members:

```
[PS] C:\> $HRGroup = Get-Team -DisplayName "Human Resources Group"
$TeamUsers = Get-TeamUser -GroupId $HRGroup.GroupId -Role Member
```

```
$TeamUsers | ForEach-Object { Grant-CsTeamsAppPermissionPolicy -PolicyName "HR App Policy" -Identity $_.User }
```

# Teams App Store

The Teams app store is the place where users go to discover the set of apps available to them. Apps blocked by the tenant (in Manage apps) are not shown. Users reach the app store by opening the app selection panel from the app (left-hand) rail and then selecting the [...] option. The More apps link leads to the app store (Figure 13-16) where they can browse the set of available apps permitted for installation by the organization. Depending on the app type, it might be installable in various places. The target varies with app type. For instance, a connector usually installs into a channel while other apps might show up in a channel tab. See [this link](#) for more information.

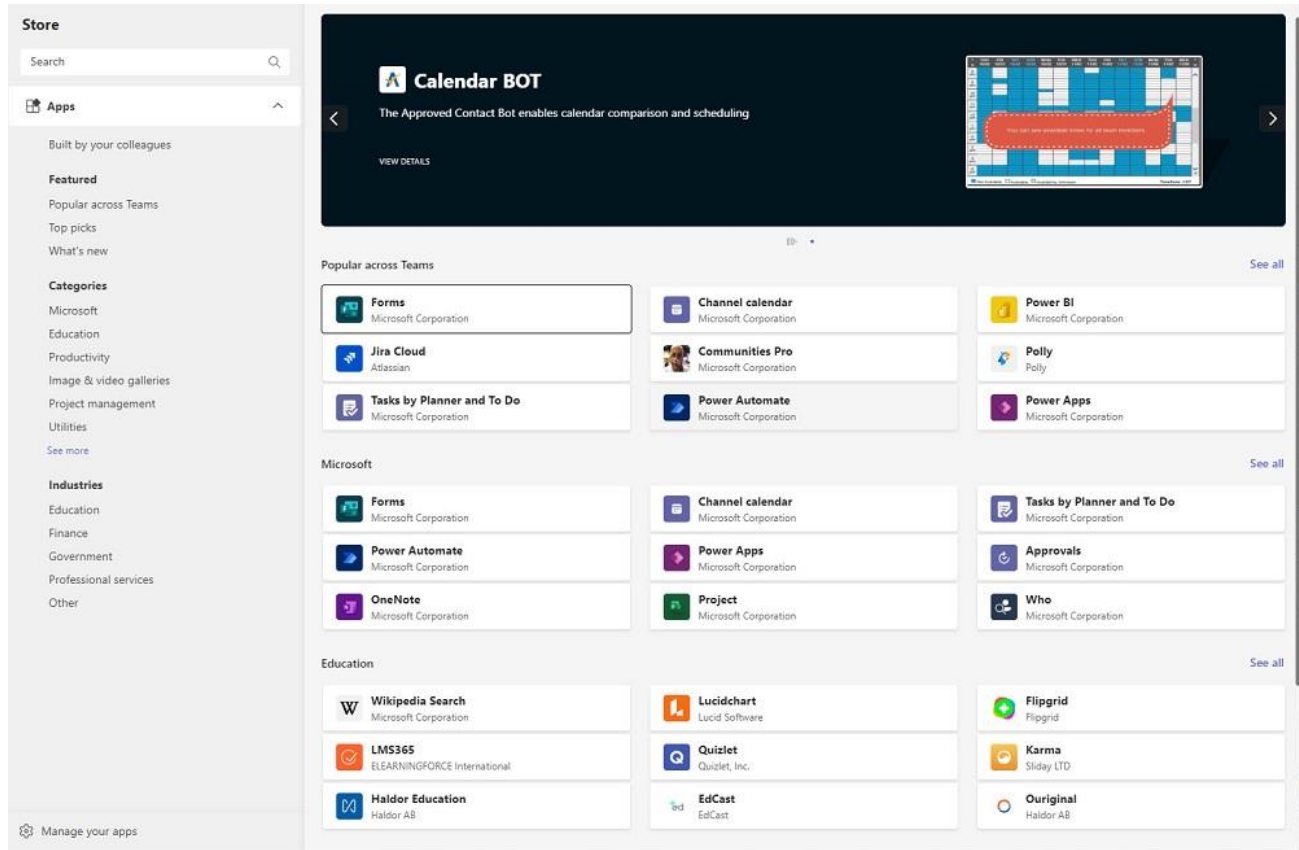


Figure 13-16: The Teams App Store

Organizations can [customize the Teams app store](#) by adding their logo and corporate colors.

## Teams and Bots

Another way of extending the usefulness of a team is to connect it to one or more bots. A bot is a virtual software assistant. In the context of Teams, a bot is an application that accepts questions about specific topics from team members and responds, all within a conversation like those conducted with humans. The Microsoft [Bot Framework](#) assists developers to build their bots for use with Teams. A sample [Teams application written in C#](#) is available too. Another example of how to create a bot for use with Teams based on an [Azure knowledge base](#) is also available. After creating a bot, to enable it for a team, go to the Teams Apps Store and click **Bots**. You can then browse the set of available bots and decide which bot to install into a channel within a team.

The Who bot is one of the standard bots included with Teams. Before you can use the bot, you must install it from the Apps Store. Once installed, the Who bot uses the organizational information included from Azure AD and signals gathered in the Microsoft Graph to answer questions such as “Who is Kim Akers” or “Who knows about collaboration.” The answer to the first question comes from the manager-employee data in Azure AD; the answer to the second is from information in Teams chats. In both cases, to find answers, the Who bot executes searches on behalf of the user. To find experts, the bot scans all the teams to which the user belongs to find messages relating to the topic. The bot then performs some post-processing to score the results, using @mentions as important signals. The bot ranks each user by the number of messages about the topic they post, with their ranking increased for each @mention they receive. The idea here is that if someone mentions another person in a message about a topic, they do so because they think that person is an expert and might know the answer. Messages that include multiple @mentions receive a lower scoring boost because these mentions are often used as notify people of something they should know about.

If you do not want users to interact with bots, you can disable access to bots on a user-specific basis by configuring and assigning an app permission policy.

## Office Connectors and Teams

Office connectors create a link between a network data source and an Office application. In Teams, the destination is a channel within a team. Connectors work much the same way with Teams as they do with Groups. The results fetched through the connector show up as “cards” created as conversations within the target channel. The cards do not hold the full content of an item fetched from a source like an RSS feed or the Office 365 Service Management API. Instead, they hold enough text to let users decide whether they want to discuss the content – or click the link to explore the full content. Note that only users with an Exchange Online license can create a connector for a channel.

You can use connectors with teams in many ways. For instance, you could have a marketing team that needs to keep an eye on information posted for corporate communications. This would be easy to configure within a channel called “Corporate External Communications” inside the “Marketing Group” team. Another example is [how to use Google Analytics](#) to track basic usage patterns for visits of the web client to Teams. Many other different out-of-the-box connectors are available, including the Incoming Webhook connector, which you can use to bring data from custom data sources into a team or group. The PowerShell chapter includes an example of how to use the Incoming Webhook connector for this purpose. The example code explained there also supports Teams.

Because the cards created by connectors appear as messages in channels, it is logical that you add connectors to specific channels. To add a connector, click the ellipsis menu to the right of the channel name you want to use and select “Connectors.” You can then browse through the gallery of available connectors and select which one you want to use to bring network data into the channel. Some of the connectors need credentials before they can fetch content, so you might need to have a username and password to configure the connector.

When someone creates or removes a connector for a channel, Teams notifies members of the fact by posting a system message with details about the connector in that channel. After the connector downloads content and creates messages in the channel, users can comment on them like any other conversation.

## Disabling Connectors

Connectors are part of the default and external apps supported by Teams. As mentioned above, you can enable or disable individual apps through tenant-wide settings. If the app is a connector, it then becomes available or unavailable for inclusion in a channel.

If you want to disable the ability to add any connector to a team, you must disable the feature for the underlying group by running a command like:

```
[PS] C:\> Set-UnifiedGroup -Identity "Marketing Team" -ConnectorsEnabled:$False
```

Setting the *ConnectorsEnabled* property for a group disables the ability to add new connectors in any channel in the team. Any attempt to add a connector results in the error:

```
Connectors have been turned off for this mailbox by the admin.
```

Because of the need to synchronize between Teams and Azure AD, it might take a little time to apply or remove a block. Unfortunately, the implementation of the block within Teams is not particularly graceful as the choice to add Connectors persists in the menu even when a block is in place for the group (or the organization). To block the creation of connectors for all Teams in the tenant, run the following PowerShell command to update the organization configuration:

```
[PS] C:\> Set-OrganizationConfig -ConnectorsEnabled:$False
```

Existing connectors in use before the implementation of a block continue to run as before. If you want to remove connectors, you must do so manually. Teams captures audit records for the addition or removal of connectors in the audit log. The audit records hold details of the type of connector (for example, RSS) but not the source.

## Teams Approvals

The Teams Approvals app delivers easy-to-use and simple workflow processing. The app has links to Power Automate, but Teams hides the details of that connection from users. Users can:

- Use the basic approval form to seek approval from reviewers for a variety of requests from time off to trips.
- Use a variation of the [basic form connected to an eSignature package](#) like Adobe Sign or DocuSign to seek approval for a document. In this case, Teams gathers information about the document for approval and the reviewers and passes the details to the eSignature provider for processing. As the document goes through the approval process, the eSignature provider sends status information back to Teams for the user to track.
- Use approval templates created by the organization. An approval template deals with a specific form of approval, such as a form to request authority to discount a customer order. Administrators create custom approval templates from scratch or use sample templates as a base. Approval templates have a scope defining who can use them. The scope can be org-wide, confined to a named set of people, or a team (in which case, any team member can create approval requests using the template). Figure 13-17 shows the creation of a Teams approval request based on a custom template.

An approver can view and approve requests generated by the Teams Approvals app in the Power Automate admin center.

Initially, approvals with e-signatures could only be processed on the Teams desktop client. From August 2022, the capability is available on the Teams mobile clients too.

The screenshot shows a mobile-style interface for creating a Teams approval request. At the top, there's a header 'Technical or Copy Edit Request' with a sub-header 'This approval template is used by an author to request a technical or copy ...'. Below this is a text input field for 'Name of request' containing 'Copy Edit Teams Architecture'. The 'Approvers' section lists two people: Kim Akers and Vasil Michev (Technical Guru). There's a toggle for 'Require a response from all approvers' which is currently turned off. Under 'Edit type', 'Copy edit (grammar, layout, house style)' is selected. The 'Editing requirements' field contains the text 'Please copy edit the attached chapter'. At the bottom, there are 'Back' and 'Send' buttons.

Figure 13-17: Creating a Teams approval request using a custom template

## Power Platform and Teams

Microsoft is keen that people use the Power Platform to build Teams apps on a low-code or no-code basis. To help people start, several resources are available:

- [Dataverse](#) for data storage. Dataverse organizes data into tables of columns and rows. A set of standard tables are available, and apps can create custom tables if needed. The idea is to have a cloud-based, easy-to-access data service that apps can use without building custom data stores.
- Sample applications such as [Milestones](#) and [Bulletins](#). Milestones is an app to keep track of projects. Bulletins are a way to publish information to people within an organization. Source code is included to enable customization of the apps to meet the needs of organizations. Although off-the-shelf alternatives exist within Microsoft 365 like Project, Planner, and Communications sites, it still might be interesting to use custom versions of these apps to meet specific needs within the business.
- Documentation for [integrating Power Apps and Teams](#).

See Chapter 22 for more information about the Power Platform.

## Debugging Teams Clients

If something goes wrong and you experience a problem with the Teams desktop client, the approach to get started again usually follows these steps:

1. For desktop and mobile clients, sign out of Teams and then sign in again. This removes any lurking problems with authentication and credentials and solves many connectivity and cache synchronization issues. For browser clients, use a private session to connect to Teams. If this works, clear any browser cookies associated with Teams and office.com and try with a normal session again.
2. If the desktop client is still failing, stop all the Teams processes. On Windows, use Task Manager to find the processes and end them. Restart the client to see if the problem disappears.
3. If that doesn't work and the desktop client continues to have problems, clear the Teams cache. On Windows workstations, to speed up performance, Teams caches copies of server data in

`%UserProfile\AppData\Roaming\Microsoft\Teams` folder. Occasionally, the cached data is out of sync with the server or holds some corrupted data. Both conditions cause problems for the client. Before removing the cache, stop the Teams client (all processes). Then remove the complete Teams folder. When you restart Teams, the client rebuilds the cache so that its content is fresh. Although it's not recommended to blow away the Teams cache as the first step in troubleshooting, it can solve problems where other steps fail.

4. If you cannot resolve the problem, you should file a service request and give the diagnostics and client logs (see below) to Microsoft to help them understand where the root cause might lie.

If you find that you can't switch to another tenant, you may have expired credentials. To solve the problem, sign out of Teams and sign back in again. After you sign in, if you then try to switch to another tenant, Teams will prompt for credentials to allow the client to connect to that tenant.

## Teams Client and Diagnostic Logs

As summarized in Table 13-1, Teams can generate several diagnostic logs.

<b>Type</b>	<b>Clients</b>	<b>Contains</b>
Debug	Browser, Windows, Mac, Linux	Text file read from the bottom up to capture client activity such as login, connection, call, and conversation events. Also includes information about the environment such as the client version.
Desktop	Windows, Mac, Linux	Also known as the bootstrapper log. This text file notes details of communications between Teams and the underlying platform.
Media	Windows, Mac, Linux	Contains information about audio, video, and screen sharing in Teams meetings.

Table 13-1: Teams diagnostic logs

To generate the debug log, use the CTRL-Shift-Alt-1 key combination on Windows and Linux, and Command-Option-Shift 1 on a Mac. On all platforms, Teams creates the diagnostics log in the Downloads folder on the workstation and creates a name by combining *MSTeams Diagnostics Log* with the current timestamp.

Use the following commands to get the desktop log:

- Windows and Linux: click the system tray and find the Teams icon. Click the icon and select **Get Logs** from the right-click menu.
- macOS: Choose **Get Logs** from the Help pull-down menu.

By default, Teams turns off media logging. When needed, enable media logging through the General section of the Settings app and restart Teams. The locations of the various files created by media logging are noted in [this support article](#).

Utilities such as [Fiddler](#) and [Charles Proxy](#) are useful tools to collect a browser trace recording the communications between a Teams client and the back-end services. Make sure that you install the root certificate for the utility in the Trusted Root Certificate Store of your computer and enable HTTPS decryption before collecting anything. You can also use the [native functionality built into browsers like Edge](#) to capture logs (replace the references in the article to the Azure portal with Teams).

If asked to provide diagnostic information to Microsoft support, the easiest method to collect diagnostic information is to click the Teams icon in the system tray and select the **Collect support files** option. This command captures information used by Microsoft (like the debug log) and stores the data in a folder in Downloads. The folder is called *MSTeams Diagnostics Log* together with the date and time of its creation.



# Migration to Teams

Microsoft has no migration tools to import existing content into Teams in such a way that the migrated information is usable within Teams. Instead, if you want to import content into Teams, you must do so manually. For example, you can email content to a channel from any application that supports SMTP, or you can move documents into the folders in the SharePoint document libraries belonging to Teams. Several ISV products are available from companies such as AvePoint, Quest, and ShareGate to import documents from SharePoint on-premises servers, file servers, and other sources, or you can use Microsoft's free [SharePoint Migration Tool](#).

If you want to move content from another chat platform like HipChat or Slack to Teams, you can export data from these platforms. The difficulty then arises in how to import that data into Teams in the form of conversations and associated documents. The availability of [a Graph API to import third-party messages into Teams](#) has encouraged ISVs to investigate the area of Teams migration, so it's worth doing an internet search to discover what migration products are currently available.

All trans-platform migrations involve some form of data manipulation to ensure that the data imported into the target platform is usable. The migration code must parse the information taken from the source platform and update it to fit the data requirements of Teams. For instance, when the items forming a conversation go into a channel, the items must be in the correct order. The migration must apply permissions to attachments, and so on. Usernames are probably different, so some process of fixing these is necessary as otherwise, permissions might not work.

# Chapter 14: Managing Teams Calling and Devices

**Ben Lee**

Previous chapters covered the structure and management of Teams. Here, we expand on that by looking at how to use Teams to make good quality calls and explore the different devices available, from phones to full-blown meeting room systems.

## Teams Calling Fundamentals

Traditionally, IT administrators might have viewed PSTN (Public Switched Telephone Network) calling from Teams as a black art. Microsoft has put great effort into making Teams calling as straightforward to use as they can, but things still need to be configured to ensure everything is working optimally.

### Teams Calls, Teams Meetings, and Teams Phone

When we talk about calling in Teams this is a “catch-all” term and could refer to a few different things. Essentially it means making a voice, or video, call from your device to another endpoint but let’s go a bit deeper by breaking those calls up into some specific definitions:

- **Teams Calls** is where users call from one Teams endpoint to another Teams endpoint. This could be between two internal users, or an internal user and an external user. The important thing is that Teams is used on either end of the call. Media for a Teams call can either pass directly between the Teams clients or be relayed through components of the Teams service.
- **Teams Meetings** are when more than two people are in a call together. All communications occur between the Teams endpoint and the Teams meeting service. Teams automatically converts any call with more than two participants into a meeting regardless of whether it was a scheduled or ad-hoc call.
- **Teams Phone** (PSTN – public switched telephone network) is when a call is between a Teams endpoint and a phone number. This could be a local number, an international number, a premium rate number, a mobile phone number, or even a call that terminates in another tenant. At some point in its journey, the call passes over the PSTN network. This is the global interconnection of phone carriers that provide telephony services to customers.

The first two types of Teams calls are similar in terms of their path over the network and the codecs used. They are also the types of calling included “out of the box” as part of the standard Teams product. To make calls with Teams Phone you need to have some extra licenses, which may be included in your licensing plan, and the traffic paths and codecs used can be very different from the other call types. The three main technologies that you can use to access the PSTN with Teams Phone are; Calling Plans, Direct Routing, and Operator Connect. We will cover all these technologies and their differences over the next few sections. But first, let’s look at how to prepare your network to best support calling in Teams.

## Networking Preparation

Almost everything you do inside Teams sends traffic backward and forwards from your endpoint to the Microsoft service. In non-calling workloads, we usually aren’t too concerned about what path that traffic takes

or how fast the connection is (unless things have gotten really bad). However, any of the calling workloads described above are very sensitive to bad network conditions as the voice and video are being processed in near real-time. If we end up on a network connection where, for example, there is very high latency, our calls have a lot of delay in them or garbled voice or video. Teams does an excellent job at optimizing itself for any network it is using, but we need to give it a helping hand to ensure that things are operating as smoothly as possible.

There are two basic design principles that we should apply when working with network optimization for Teams, summarized as follows:

1. Ensure you move traffic onto the Microsoft network as quickly as possible:
  - a. Use local internet breakout.
  - b. Use local or regional DNS resolution. Microsoft 365 and Teams use GEO DNS for name resolution, meaning that you need to resolve the Teams services using local DNS to get the IP addresses of services closest to you.
  - c. Avoid sending traffic over internal or multiple WAN links unless they are very low latency.
2. Interfere with the network traffic as little as possible:
  - a. Open all the correct IP and port ranges that Teams needs found in this [link](#). The most important ones for Teams Calling are ID 11 (required), 12, and 13 (recommended).
  - b. Disable any deep packet inspection on the traffic.
  - c. Disable any WAN optimization technologies.
  - d. Avoid using VPNs with Teams traffic.

Following these two basic principles helps you get the traffic from the client and over to Microsoft where it becomes their problem as quickly and cleanly as possible.

**Note:** Teams meetings and calls route to local media relay services and are independent of which region hosts your tenant. For example, If your tenant is in Europe and a team of four people in the US meet, the media flows through US media relay services to ensure the best possible quality for the call

Microsoft has documented [the range of IP addresses, ports, and FQDNs used by Teams](#) and other Microsoft 365 workloads accessed when consuming services. They break these up into three categories:

- **Optimize:** this traffic is extra sensitive to network performance, latency, and availability. The IP addresses listed are guaranteed to come from Microsoft-hosted IP ranges. This category covers the endpoints used to process Teams media.
- **Allow:** is not as sensitive to network performance and latency but still provides connectivity to Microsoft 365 services and features. This category includes the main teams.microsoft.com URL.
- **Default:** should be handled as regular internet traffic, no optimization required. Non-Microsoft data centers could host these IP address ranges. The traffic could be icons inside SharePoint or other supporting types of data.

We cannot stress how important optimizing these categories is for Teams call quality as audio, video, and desktop sharing are all consumed in real-time, so any issues are very noticeable. If traffic is not optimized you may get random connectivity issues, call drops, or poor audio, video, and desktop sharing quality. The Optimize and Allow ranges should be treated as trusted data center locations by your network and be allowed and bypass any advanced network security features. Teams encrypt all traffic that goes to these IP ranges.

Is ExpressRoute a requirement for optimizing Teams traffic? The short answer is no, you still connect to the same IP addresses and services regardless of using ExpressRoute or the internet. ExpressRoute is an option when you want the bandwidth to be reserved and guaranteed to the Microsoft 365 services, have an SLA on connectivity, and need to preserve DSCP tagging for QoS to the Microsoft Network. If you have constraints on existing internet connectivity, setting up an additional connection using ExpressRoute for Microsoft 365 traffic may be an option. However, keep in mind that the recommendation is to have local internet breakout per site,

which could mean multiple ExpressRoute connections would be needed. However, keep in mind the recommendation to have local internet breakout per site, which could mean multiple ExpressRoute connections would be needed. Read more about [ExpressRoute and Teams](#).

The IP ranges may change at any point in time. Microsoft has published a [web service](#) where you can always get up-to-date IP ranges and ports. To stay abreast of changes, you can set up a routine to fetch this feed and send it to a suitable team, and some network vendors support auto-updating of rules using this service.

## Documenting Internal Networks and Subnets

Documenting your network topology can be time-consuming. However, the networking data is usable in a range of Microsoft tools to help plan, review, and monitor the performance of your calling deployment. The tools where you can reuse this information are:

- Teams Network Planner.
- Network Topology in the Teams admin center.
- Microsoft 365 Network connectivity, part of the Productivity Score.
- [Microsoft 365 Network connectivity test](#).
- Call Quality Dashboard.
- Reporting Labels in the Teams admin center.

The sooner you start gathering information, the better. Not all the tools need the same information, but if you create a complete data set in, for example, an Excel spreadsheet, you can reuse the information for each tool by creating the necessary csv files. Here is the information you need for each subnet:

- Site name.
- Internal network subnet range.
- Network Mask.
- Public IP(s) the client will present themselves with.
- Physical location address.
- Physical location City.
- Physical location State.
- Physical location Region.
- Physical location Country.
- Building name(s).
- The approximate number of users.
- Expected number of Meeting Room Devices.
- Internet breakout with bandwidth if local.
- If ExpressRoute is being used for connectivity to Microsoft.
- If branch site, which site is it connected to for internet breakout.
- If branch site, WAN link bandwidth.
- If using Teams Phone, how will it be delivered (Calling Plans, Operator Connect, or Direct Routing).
- The FQDN of the local Session Border Controller (SBC) if you are going to use Direct Routing.
- The proxy SBC FQDN, if you plan to use this with Direct Routing.

After gathering the data, you are ready to populate the tools.

## The Network Planner

The Network Planner tool (available under **Planning** in the Teams admin center) helps to estimate the network capacity required to support Teams based on the number of users, links between sites, and internet breakouts. The inputs to the Network Planner are personas that model the type of users you have, and subnet sites that describe the network topology and how they connect to Teams. You can even add Microsoft Teams Rooms (MTR) as personas per site, which is important as they may require more bandwidth than a normal

user as they are usually used for video first meetings which can consume more bandwidth. The output is a report listing estimated constraints in your network. Use this tool to get an overview of your environment and a starting point for a network assessment for Teams. Read more about how to get started [here](#).

**Note:** To save some time, a [PowerShell script is available](#) to automate the population of the Network Planner.

When you have populated the Network Planner, you can use the tool to identify any expected network constraints. A good result is when the tool finds no expected constraints but remember this is only based on anticipated usage with the inputs you have supplied. You should still then refer to tools such as Call Quality Dashboard (CQD) to monitor actual usage and user experiences after deployment.

The planner measures the centralized internet breakout with the number of sites connected to it so that you can capture your total expected bandwidth usage. If you experience bandwidth issues for sites that Network Planner does not flag, you should look for other issues outside Teams to investigate why this is the case. If you find sites with network constraints, then you can either improve the sites for better connectivity or use the Teams Network Roaming policy to limit bandwidth usage for those sites.

## Teams Network Roaming Policy

A challenge for branch sites with network constraints is that calls and meetings can consume too much bandwidth for video and screen sharing. The Network Roaming policy is a per-site policy that ensures that branch offices with low bandwidth do not get overloaded when users are using Teams, regardless of their Meeting Policies. You can then avoid having to rely on users remembering that they shouldn't be using video in a particular location and just inform them that video does not work in this location because of network restrictions.

There are two settings we can control with the Teams Network Roaming policy: *MediaBitRateKb* and *AllowIPVideo*. You can also find these settings in the Teams admin center under **Meetings**, then **Meeting policy**, these values are used if there is no site-specific Teams Network Roaming policy applied. You can create multiple meeting policies in an organization, each with different values, and assign the policies to different sets of users as appropriate.

Before you can create a Network Roaming policy you need to define the following:

- The public IP address that a client will present to the Teams service, so that Microsoft knows that this client is internal.
- A region (can span multiple sites).
- A site (can span multiple subnets).
- One or more subnets assigned to the corresponding site.

These configuration items are used by the Network Roaming policy and for Direct Routing if using Local Media Optimization or Survivable Branch Appliances (SBA). These concepts are described later.

External trusted IPs are the Internet external IPs of the enterprise network used to determine if the user's endpoint is inside the corporate network. This must be properly defined since the Network Roaming policy will only apply if the client has been identified as "internal" in the first place.

Use the Teams PowerShell module with the following cmdlets to configure the required policies:

```
[PS] C:\> New-CsTenantTrustedIPAddress -IPAddress 172.16.240.110 -MaskBits 32 -Description "Site1 or Region1 Public IP"
```

After adding the public IP addresses, you can create a region. You don't need to connect a region to the public IP addresses:

```
[PS] C:\> New-CsTenantNetworkRegion -NetworkRegionID "Region1"
```

Now you can create a site and connect it to a region, a region can have multiple sites:

```
[PS] C:\> New-CsTenantNetworkSite -NetworkSiteID "Site1" -NetworkRegionID "Region1"
```

Before you create the Network Roaming policy, create a subnet for the site. You can assign multiple subnets to a site:

```
[PS] C:\> New-CsTenantNetworkSubnet -SubnetID 192.168.1.0 -MaskBits 24 -NetworkSiteID "Site1"
```

After documenting the network topology, you can create the policy and define if video is available and the maximum amount of bandwidth Teams can use in that location:

```
[PS] C:\> New-CsTeamsNetworkRoamingPolicy -Identity "Site1RoamingPolicy" -AllowIPVideo $false -MediaBitRateKb 2000 -Description "Site1 roaming policy"
```

This creates a roaming policy that prevents video and limits the total bandwidth used within that site to 2000 kbps. The default limit in a Teams Meeting policy is 50000 kbps per user.

To finalize the setup, modify the site you created representing the location and add the new Network Roaming policy:

```
[PS] C:\> Set-CsTenantNetworkSite -Identity "Site1" -NetworkRoamingPolicy "Site1RoamingPolicy"
```

You can validate the site configuration with this command:

```
[PS] C:\> Get-CsTenantNetworkSite
```

Users configured for Teams Phone who visit the location now cannot consume more bandwidth than you have assigned. To include non-telephony users, configure the *AllowNetworkConfigurationSettingsLookup* in your meeting policies. For example, this command updates the global Teams Meeting policy:

```
[PS] C:\> Set-CsTeamsMeetingPolicy -identity Global -AllowNetworkConfigurationSettingsLookup $True
```

You have now created the necessary regions and sites, added the subnets for the location, assigned a roaming policy to the site, and made sure it covers most users. T's.

## Network Topology in the Teams admin center

Network information is in the Teams admin center under **Locations > Network Topology** and **Network & Locations**. You will see a tab for trusted IPs related to your tenant, and you can update the addresses through the admin center. When setting up Call Analytics, we will populate the **Reporting Labels**, which you will also find under **Locations** in the Teams admin center.

## Microsoft 365 Network Connectivity

The Microsoft 365 Network Connectivity report is under Health in the Microsoft 365 admin center. It is an important part of the tenant's Productivity Score discussed in Chapter 4. This tool can validate site performance as part of the planning process and monitor performance as part of the operational process. Network performance for Exchange Online, SharePoint Online and Teams is highlighted by drilling into the Network Connectivity under Technology experiences in the Productivity Score. In the Points breakdown on the right side, you see your Productivity Score rating for Teams performance. The performance is based on the same network metrics (packet loss, jitter, and latency) highlighted by the Call Analytics and Call Quality Dashboard. Administrators use these same metrics to validate if a tenant is ready to deploy Teams, and the Locations tab also shows how your sites are performing.

At first, you may see some locations which do not make sense. The information is a mix of locations where you have offices, and it also includes home offices and places where users connect to Office 365 on the road. Here we can use the information you gathered to add your sites, subnets, and associated public IP addresses. By clicking Add location, you can add details of the name, physical office location, LAN subnets, and the

public IP addresses presented from that location. When you have added your locations, you can filter the view to show only locations you have added and get a more controlled view. It can take up to 72 hours before data is available for a location after creation.

## Microsoft 365 Network Connectivity Tests

As part of planning for Teams calling and meetings, you can manually capture and report network data using the Microsoft 365 Network Connectivity tool. Selecting Network connectivity tests in the top right corner of the Microsoft 365 Network Connectivity page leads you to <https://connectivity.office.com/>. From there, you can run local connectivity tests from your current location. When the test finishes, you can upload the results to the Microsoft 365 Network Connectivity report. You can use the overview to see the network readiness for real-time communications. Over time, this report is a great way to validate Microsoft 365 performance because it uses actual traffic and tells you if you are following the two principles for network connectivity.

When running the tool, you see how you connect to the Microsoft network, the public DNS used, and the connectivity performance for Exchange Online, SharePoint Online, and Teams. It will also perform a traceroute to some Microsoft services to measure the number of hops and show you if you are spending unnecessary time routing traffic over the internet before you hit the Microsoft network. The number of traceroute hops should be between 6-20; anything over 20 could result in increased latency.

You should run connectivity tests in two places: close to your network edge and where users are. By comparing the results, you can see if there are any differences in quality between your network edge and user locations that might indicate bottlenecks in your network. Remember that a single result may not be a good enough measurement to launch a network investigation. By running many tests at different times of the day over a sustained period, you will gather sufficient data to highlight trends within your network. By understanding the trends, you'll know if a more thorough investigation of internal network quality is necessary.

When running the tool, the admin center prompts you to download a file. After downloading, this executable file runs connectivity tests from your computer and sends the results back to the service. You will see that the test almost seems to stop towards the end. This is because the test makes some sample calls that are 17 seconds long to measure audio, video, and desktop sharing performance. The captured metrics then become the basis for measurement. By clicking details, you see data for the performance and see if any of the expected metrics are below expectations. Figure 14-1 shows a capture, and you also see the expected values for the metrics.

### Microsoft Teams

Test	Result
✔ Media connectivity (audio, video, and application sharing)	No errors
✔ Packet loss	0.00% (target < 1% during 15 s)
✔ Latency	23 ms (target < 100 ms)
✔ Jitter	3 ms (target < 30 ms)

Figure 14-1: Microsoft 365 Connectivity test results for Teams

If you are signed into Microsoft 365 when you run the test, the test automatically uploads the data to the Microsoft 365 Network Connectivity report. The data is available under the office location you have already

added, and an administrator can see the results and view trends for that location. The uploaded data should be available on the site within minutes, but delays might occur if the service is busy.

The [Teams Network Assessment Tool](#) is a manual method to gather network insights for a location. This tool can measure Teams performance but does not upload its results to the Microsoft 365 admin center. It is a command-line tool downloadable from the Microsoft Download Center. If you run the tool with no command-line switches, it will test whether the ports and protocols that Teams needs are open outbound from the client. If you run the tool with the `/qualitycheck` switch, it generates sample 17-second calls and measures the key quality metrics for those calls. This tool can run without having any Teams components installed on the machine, so it can be run from servers in datacenters to help map out your network paths. You may recognize that the results are the same metrics that Microsoft 365 Network Connectivity highlights for Teams. Here is an example:

```
PS C:\Program Files (x86)\Microsoft Teams Network Assessment Tool> .\NetworkAssessmentTool.exe
/qualitycheck
Microsoft Teams - Network Assessment Tool

Initializing media flow.

*****
Starting new call

Media flow will start after allocating with relay VIP FQDN: worldaz.tr.teams.microsoft.com
If user wants to allocate with a particular relay VIP IP address, please specify in
NetworkAssessment.exe.config.

Waiting for call to end after 300 seconds, displaying call quality metrics every ~5 seconds.
Change the 'MediaDuration' field in the NetworkAssessmentTool.exe.config file to change the media
flow duration.

TIMESTAMP is in UTC. LOSS RATE is in percentage, out of 100.
LATENCY and JITTER are in milliseconds, and are calculated as averages in ~5-second windows.
PROTOCOL displays whether UDP, TCP (PseudoTLS/FullTLS), or HTTPS protocol was used to allocate with
the relay server.
Note that for PROTOCOL, UDP protocol is attempted first to connect to the relay, by default.
LOCAL ADDRESS is the local client IP and port that media is flowing from.
REMOTE ADDRESS is the peer (relay server) destination IP and port that media is flowing to.
IS PROXIED PATH shows whether a proxy server is used to connect to the relay, only applies to
TCP/HTTPS connections
LAST KNOWN REFLEXIVE IP shows what your latest public (NAT translated) IP and port is that the relay
sees during media flow.

[If LOSS RATE is 100%, the output lines here will be in red]

Quality check source port range: 50000 - 50019

Call Quality Metrics:

2022-06-15 05:46:35          Loss Rate: 0          Latency: 25.63          Jitter: 13
Protocol: UDP
Local IP: 192.168.1.53:50016          Remote IP: 52.114.231.164:3478
Is Proxied Path: False          Last Known Reflexive IP: 212.159.102.162:14846
```

The tool runs for 300 seconds by default and saves the output to `c:\Users\<username>\AppData\Local\Microsoft Teams Network Assessment Tool\<date_timestamp>_quality_check_results.csv`. You can open the results file to calculate your average for that run. Read the Usage.docx file found at the tool install location for information on how to configure longer iterations.



You can also measure the number of hops to the same IP address as shown in the Network Assessment Tool to check that you have an optimal path over the internet to the Microsoft data centers. Using `tracert` in the example below gives us the following result:

```
C:\>tracert 13.107.8.20
```

```
Tracing route to 13.107.8.20 over a maximum of 30 hops
```

```
  1    8 ms    8 ms    8 ms  192.168.10.1
  2   16 ms   14 ms   19 ms  1.51-175-128.customer.lyse.net [51.175.128.1]
  3   21 ms   18 ms   18 ms  111.81-166-123.customer.lyse.net [81.166.123.111]
  4   13 ms   29 ms   18 ms  microsoft.dix.dk [192.38.7.76]
  5   24 ms   29 ms   18 ms  tor3.ber30.msedge.net [104.44.80.170]
  6    *      *      *      Request timed out.
  7   29 ms   18 ms   18 ms  13.107.8.20
```

We entered the Microsoft network at hop 5 and reached our destination after seven hops. You would normally see between seven and twelve hops before reaching the Microsoft network. More than twelve hops are not bad, but if you see more than twenty hops you should work with your ISP to get a more optimal path. More hops can introduce latency and jitter and may be suboptimal in terms of the media path. Depending on what the IP is that you are tracing against you may end up with a string of "Request timed out" messages. This is normal as not all devices inside the MS network will respond to pings.

**Note:** It is possible to use Resource Monitor on your PC to monitor Teams client network traffic. You get an overview of how much bandwidth Teams is using at any given time. You can use the same metrics to upload information to Azure Monitor. This could be useful for organizations that want to validate Teams performance in different locations by for instance using support personnel computers. This [blog post](#) goes into more detail.

## Adding Network Information to the Call Quality Dashboard (CQD)

The CQD portal shows you call quality trends in your environment. The same information you initially gathered about your network is useful with the CQD portal. The portal can be accessed from the Teams admin center. When opening the dashboard for the first time, you may feel that it does not give you much feedback. Only when you upload your subnets to the portal will the dashboard become useful because that is when you can differentiate on-call quality per location. You will also see the difference between outside traffic (typically unmanaged networks) and inside traffic (typically internal networks). When uploading subnets to the portal, you should also specify the IP range for VPN traffic by creating a region for that. The source of an unexpectedly high volume of VPN traffic should be investigated. VPN traffic is not optimal since it is better when the media flows directly from the external client to the conference bridge.

To import subnets to CQD you need to fill out a CSV file and upload it to the portal by navigating to <https://cqd.teams.microsoft.com/spd/#/TenantDataUpload> and using these steps:

1. The CSV file you are must only contain only the actual data, headers are not needed.
2. To upload the file, go to the cogwheel in the CQD portal and choose tenant data upload.
3. Specify your file and upload the data, unless you want a specific start and end date for your new upload, just leave the default options.
4. When imported successfully you will see "Successfully uploaded file."
5. It may take up to 24 hours before the reports reflect the new subnets.

The headers used are in the file are:

```
Network,NetworkName,NetworkRange,BuildingName,OwnershipType,BuildingType,BuildingOfficeType,City,ZipCode,Country,State,Region,InsideCorp,ExpressRoute
```

Here is an example showing how to populate the csv, using some of the information we gathered at the beginning of this chapter:

192.168.1.0,US/Seattle/SEATTLE-SEA-1,24,SEATTLE-SEA-1,Contoso,IT Termination,Engineering,Seattle,98001,US,WA,MSUS,1,0

## Reporting Labels for Call Analytics

In the Teams admin center under Analytics & reports, you will find the Reporting labels used by Call Analytics to make the reports easier to read by adding network names to internal sites. Here you can upload the same csv file that was used with the CQD portal. Call Analytics shows the same metrics information you found when running the Microsoft 365 Connectivity tests. You can use it to investigate individuals to see how their clients performed in calls and meetings, we will cover it more when we discuss troubleshooting.

Now that you understand the importance of documenting your network, make sure that you periodically update the information in Network Planner, Microsoft 365 Connectivity, CQD and Call Analytics to keep it current

## Licensing for Teams Calling

Teams products are available in three (possibly four or even five depending on your choices) licensing types. The core Teams license covers you for all the Teams collaboration workloads and Teams Calls, but there are more license types when you want to do something with the PSTN (this can be referred to as telephony). This could be the Audio Conference license to add PSTN capabilities for meetings or Teams Phone license to let you make and receive PSTN calls. There are also licenses aimed at some specific device use cases.

**Note:** When considering the overall benefits of migrating your telephony workload to Teams it is important to consider more than just the price point. In January 2020 Forrester Consulting conducted a [Total Economic Impact™ \(TEI\)](#) study for Microsoft that helps provide a framework for organizations to evaluate the potential financial impact of Microsoft 365 Cloud Voice.

### Audio Conferencing

The Audio Conferencing license lets users dial-in and dial-out of meetings (subject to the appropriate configuration) using a normal phone. This can be helpful when you do not have suitable internet connectivity to sustain a full Teams call, or it can be a quick and easy way to allow external users to join a meeting. Remember that a dial-in participant can only hear the audio portion of the meeting.

By default, when a user is enabled for Audio Conferencing they have access to a set of Microsoft-provided shared numbers in over 180 Countries. You can also add more numbers to the service if you have particular numbers that you want to bring with you from a legacy platform. It is also possible to use toll-free numbers with the dial-in service, however you need a license called Communication Credits where the cost of the incoming call costs can be taken.

The Audio Conferencing license is standard in the Office 365 E5 plan but needs to be added for users with an Office 365 E3. Alternatively, a pay-per-minute version is available for companies with a volume licensing agreement.

### Teams Phone

Teams Phone is the basic license required for an account to have a PSTN number. This license covers the ability to make and receive calls via your Teams account. You still need to provide the means of actually placing calls which could be through one of several different delivery methods which we will talk about soon.

The Teams Phone license is included as standard in Office 365 E5 but must be added to Office 365 E3 accounts.

## Calling Plans

Once you have a Teams Phone license you still need to provide numbers and dial-tone. Calling Plans is a way to use Microsoft as your carrier and have them provide numbers for you. The numbers come from a pool owned by Microsoft and this service does not integrate with existing telecommunications services but you can port numbers to and from Microsoft in supported regions.

Calling plans are subject to geographical restrictions where Microsoft has been licensed to provide telephony services which currently includes thirty-one countries. You can view a full list of the telephony services available [here](#).

As well as letting you take a number from Microsoft, a Calling Plan acts as a pool of credit that can be spent on calls. Several different types of Calling Plans are available to provide some flexibility depending on your calling requirements. Typically, plans include per-user domestic calling minutes (3,000 in the US and 1,200 in EU countries) and 600 tenant-wide international calling minutes.

The minutes in plans combine across all users with the same type of Calling Plan and in the same geography. For example, if you have two Calling Plans subscriptions for US users, you would have 6,000 domestic calling minutes each month and this would stay the same even if you added more Calling Plan users in say the UK. Unused calling minutes do not carry over to the next month.

Calling Plans have two categories of licenses:

- **Domestic Calling Plan** allows calls to the country/region where users are assigned. These are also available in either small, medium, or large variations.
- **Domestic and International Calling Plan** allows calls to a users home country/region and to international numbers in 196 countries/regions.

When you buy a Calling Plans license, you can order user numbers for that country. User Numbers are numbers you assign directly to user accounts and for use by individuals. To order User Numbers, you must have at least one Calling Plan subscription assigned to your tenant.

When you assign a Calling Plans license to a user account, the following happens:

- The user's Voice policy is set to *BusinessVoice*.
- You can assign a user number to the account using either PowerShell or Teams admin center.
- After assigning a Calling Plans license, it may take up to 48 hours until the dial-pad appears in the calling tab in the user's Teams client. Also, the user may need to sign out and back in to clear Teams client cache.
- The ability to assign a Calling Plan license to a user requires that the *UsageLocation* attribute is set to a country where Calling Plans are available.

**Note:** Calling plans are not your only option when adding numbers and dial-tone to your Teams service. You still need to have a Teams Phone license but you can bring your own numbers through Direct Routing (more on that later), or you use a partner/provider using either their hosted Direct Routing platform or something called Operator connect.

## Communications Credits

Communications Credits are used as a "pot" of money that is then used to pay for telephony-related functionality that isn't covered elsewhere. Funds added apply across the tenant, but only the users assigned the license can access it if needed. The Communication Credits are used to:

- Add toll-free numbers to Audio Conferencing meetings, auto attendants, and call queues. Toll-free calls are billed per minute.
- Cover costs for dial-out in a meeting when not covered by another license type.

- Dial-out from a Teams meeting to your mobile.
- Pay for any call-back features not covered by another license type.
- Cover international calls when a user is assigned a Domestic Calling Plans license.
- Covering telephony calls for a user where their pot of Calling Plan minutes has been used.

Communications Credits are available in the same countries where Audio Conferencing is available. You can buy and assign the license once you have either an Audio Conferencing license or a Phone System license in your tenant. The amount of funds you want to allocate varies greatly depending on expected usage so you should keep a close eye on it for the first few months. The funding can be auto-recharged to avoid any service disruption.

**Note:** Funds are used for Communications Credits at Microsoft published rates only when the services are used. Any funds not used within 12 months of the purchase date expire and are lost. You can monitor Communication Credits under **Billing > Your products** in the Microsoft 365 admin center.

It can take a bit of time to set up Communications Credits so you may want to add a small amount to the tenant so that you have it available at short notice if required.

## Number types

Numbers in Teams are split into two main categories depending on their intended use:

- **User Numbers** are allocated to users for use with Teams Phone and are used to route calls in and out. User Numbers only apply if you are consuming Teams Phone via Calling Plans or Operator Connect.
- **Service Numbers** are used in several different places, such as dial-in conferencing numbers or for VoiceApps (telephony-enabled workflows) with Attendants and call Queues.

**Note:** All number management in Teams is done using E.164 format where numbers start with a + symbol, followed by a country code. Teams converts user-entered numbers from their local country format into E.164 for you (you can modify this to some degree as we will cover later).

## User Numbers

To be able to order user numbers (from Microsoft) you must first buy and assign Calling Plans to all users you want numbers for. The amount of user numbers you can order is calculated as follows:

$\langle \text{Number of Calling Plan Licenses} \rangle \times 1.1 + 10$

In other words, if you have 50 Calling Plan licenses, you can order 65 numbers. Because they come from a pool of available numbers it is improbable that you will get consecutive numbers. However, these days when most calling is done via contact cards having consecutive number ranges is less important. If you do need specific ranges of numbers for your tenant, you can raise a ticket to Microsoft, and they will see what can be done.

Ordering user numbers from Microsoft is done through the Voice section of the Teams admin portal, where you specify the type (user) and quantity of numbers you need. If you have not added a location for your users, you must provide a location during the ordering process. The location describes where you are in case of an emergency and defines the area code you get for your numbers.

After finding a suitable block of numbers, you have 10 minutes to claim the numbers before the system releases them back to the available pool. After claiming the numbers, you can assign numbers to users through the Teams admin center (**Voice** and then **Phone numbers**) or using the `Set-CsPhoneNumberAssignment` cmdlet. For example:

```
[PS] C:\> Set-CsPhoneNumberAssignment -Identity Ben.Lee@office365itpros.com -PhoneNumber +44191375xxx -PhoneNumberType CallingPlan
```

**Note:** If you need to port existing numbers, you can read about the process [here](#).

## Service Numbers

Service Numbers are numbers available in a limited supply per tenant when necessary to handle a large call volume. As with user numbers, you can either take them from the Microsoft pool or port them from your existing provider (see note in previous section).

There are two types of Service Number: toll numbers and toll-free numbers. The most common type of service number is a toll number. All that you need to order toll numbers is a Phone System license in your tenant. Toll-free numbers can only be ordered if Communications Credits are configured (see above).

When assigning toll-free numbers to an auto attendant, the call is free for the customer dialing into the service and the company pays that cost as a charge against the Communications Credits prepaid amount. When using toll numbers, the user dialing into the service pays the regular fee. Service numbers are available in ~80 countries today, and the markets will grow over time.

To get new service numbers, go to the Teams admin center, navigate to **Voice, Phone numbers**, and click **Add**. From there you define the country, type, and region to find suitable numbers. The amount of service numbers a tenant can get is also based on the number of Phone System and Audio Conferencing. A sample of the table is shown here, for the complete list [see here](#):

- If there are 1-25 licenses, 5 service numbers can be requested.
- If there are 100-149 licenses, 30 service numbers can be requested.
- If there are 500-749 licenses, 90 service numbers can be requested.
- If there are 1000-1,249 licenses, 125 service numbers can be requested.
- If there are 7,000-9,999 licenses, 500 service numbers can be requested.
- If there are 50,000+ license, 1500 service numbers can be requested.

## Teams Administrator Roles

In larger organizations, you may have different people looking after communication platforms such as telephony and video conferencing and collaboration platforms such as SharePoint, Yammer, and the intranet. Office 365 supports a set of RBAC roles to allow people to have the necessary permissions to manage different services, including Teams.

Because the Teams Calling workloads are pretty specialized in their own right, Microsoft provides the following Admin roles specifically for managing them:

- **Teams Service Administrator:** can manage everything Teams-related without needing Global Admin rights, including all Teams Calling configuration.
- **Teams Communications Administrator:** can manage calling and meeting features within Teams.
- **Teams Communications Support Engineer:** can troubleshoot communications issues with access to advanced tools.
- **Teams Communications Support Specialist:** can troubleshoot communications issues within Teams by using basic tools, is not allowed to see user identifying information such as SIP addresses or phone numbers.
- **Teams Device Administrator:** can manage Teams certified devices from the Teams admin center.

The RBAC role best suited for working with calling and meetings is the Teams Communications Administrator. With this role you will be able to configure everything related to voice, even create resource accounts, and you will get full access to Call Analytics and CQD.

If you have AV partner deploying shared devices for you, they will usually only need the Teams Devices Administrator role which grants them the ability to manage the devices in the tenant. They do not have access to Call Analytics or CQD, but If that is a requirement, you can combine the role with the Teams Communications Support Engineer. Read more about Teams Administrator roles [here](#). You can also use Administrative units in Azure to define the scope of which devices they can manage, see [here](#) for more details.

## Teams Meeting Enhancements

Chapters 12 and 13 cover Teams meeting configuration and policies. Here we consider some improvements we can make to the calling capabilities discussed in the previous section.

In recent years, working from home for many became the new normal, and conducting meetings through Teams became common. As offices embrace hybrid work models, there is a new expectation that people can be first-class participants in a meeting no matter how they choose to join.

There have also been steps towards better interoperability between the key players in the UC space to allow some elements of cross-platform interoperability. The goal should be a universal invitation that works if you want to call in with a conference phone while driving home from work or even a video system from a vendor such as Cisco or Zoom. We should remove any technology barriers to joining meetings and make them as universal as possible.

How do we do that? First, we ensure users have a Teams Audio Conferencing license that enables dial-in scenarios from any phone. We can also add a Cloud Video Interop integration that supports joining from legacy conferencing systems.

Figure 14-2, shows a meeting invite broken down to show what is being provided by which license and also where you can customize it with your own information.

### Microsoft Teams meeting

Join on your computer or mobile app  
[Click here to join the meeting](#)

Join with a video conferencing device

[teams@vc.equinox.com](mailto:teams@vc.equinox.com)

Video Conference ID: 128 294 834 2

[Alternate VTC dialing instructions](#)

tel:+4721402221,,362187906#

Or call in (audio) [Click or tap to follow link.](#)

+47 21 40 22 21,,362187906# Norway, Oslo

Phone Conference ID: 362 187 906#

[Find a local number](#) | [Reset PIN](#)

<Add your logo or a descriptive picture here>

If you are using Skype for Business, you can join the meeting by calling [teams@contoso.com](mailto:teams@contoso.com) and use the VTC Conference ID

[Learn More](#) | [Meeting options](#)

Teams meetings can be joined from the clients on PC, Mac, Mobile and as guest users via Edge or Chrome

Cloud Video Interop, which is part the of the Phone System license and via the three providers enables users to join from any system into the teams meeting. This is the service that helps your Teams invite transcend all technology barriers

The Audio Conferencing license enabled dial-in number in the meeting. You can find a local number from more than 70 countries, gives the ability to join the meeting on the go, notice that conference ID is part of the dial in link!

Add a custom text, disclaimer and logo to the meeting, here we describe that you can call in to the meeting using Skype for Business client via the interoperability service

Figure 14-2: The anatomy of a Teams meeting invite

To change the custom text and disclaimer shown in Figure 14-2, go to Meetings in the Teams admin center and open Meeting Settings. You can add a logo or an explanatory picture, a link to a support site, and maybe a disclaimer that meetings can be recorded.

## Audio Conferencing for Teams Meetings

Assigning an Audio Conferencing license lets users create Teams meetings with a dial-in telephone number, but why might you want to do this Teams supports many different clients and devices? The main reason is that the telephony network is one of the few services that is (mostly) global and doesn't need an internet connection.

The dial-in option increases the usability of Teams meetings. It is a popular feature in organizations that use conference rooms where attendees want to use existing phones in the room to join meetings. Audio Conferencing also allows external users to dial into meetings from a mobile phone, which can be beneficial for joining from locations with poor internet coverage.

## Licensing Audio Conferencing

The Audio Conferencing license operates independently of Teams Phone, it does not require the user to have a Teams Phone license and covers any meeting they create.

You assign Audio Conferencing licenses to user accounts via the Microsoft 365 admin center or as part of your license lifecycle management via PowerShell or Azure AD Group-Based Licensing (see Chapter 5). Teams automatically assign users a primary Audio Conferencing number based on their configured location (as specified in the Microsoft 365 admin center). You can change the assigned Audio Conferencing number by editing the user in the Teams admin center.

It is recommended that you enable Audio Conferencing for all users so that all their generated meetings include dial-in details. If you do not have a suite such as Office 365 E5 where this license is already included, this could be an expensive outlay however, there are two ways to purchase it:

- Audio Conferencing service plan add-on to Office 365 E3 (included in Office 365 E5).
- A pay-per-minute license acquired through a volume license agreement.

The difference between the two is that the regular add-on has a cost per user, as well as enabling dial-in, it also includes 60 dial-out minutes per licensed user per month pooled on a tenant level to [zone A countries](#). When considering this option you should focus on the expected volume of minutes dialing into meetings as this is the main use case for Audio Conferencing.

The pay-per-minute license type is a free license you can assign to all users, but then you need to use Communication Credits to pay for both dial-in and dial-out capabilities. Therefore, you should calculate how many dial-in minutes you might expect to use to help determine which license type is a good fit. Most volume licensing customers find the pay-per-minute model works out best, sometimes with a 90% cost reduction if covering all users. To help calculate costs for Audio Conferencing minutes, you can find a form near the bottom of this [link](#) that lets you view the rates used.

After deploying Audio Conferencing, you can monitor how many minutes are available and the current usage from inside the Teams admin center. The report can be found under: **Analytics & Reports > Usage reports**, and run the **PSTN minute & SMS (preview) pools** report.

## Configuring Audio Conferencing

If you do go ahead and deploy Audio Conferencing, there are some things to consider about how it operates.

Firstly, the default setting is that all participants in the meeting hear a tone when someone joins a meeting using the Audio Conferencing bridge. Some companies consider this a must-have requirement so that no one can join a meeting without everyone being aware. You can even force dial-in users to record their name that is played into the meeting. However, suppose you plan to have large meetings with many dial-in participants. In that case, these announcements can soon distract from the audio in the meeting itself with a constant stream of join/disconnect notices, so configure accordingly.

The setting is in the Teams admin center under **Meetings, Conference Bridges, and Bridge Settings**.

Secondly, when you assign the Audio Conferencing license as part of an onboarding process Teams sends a “welcome” email to the configured user. This email gives the user some information about what Audio Conferencing is and their PIN (used to authenticate them if they were to dial-in to a meeting). However, this is not a commonly used scenario so it can cause more confusion than benefit.

If you want to disable the notification email, this can be done before configuring users with the `Set-CsOnlineDialInConferencingTenantSettings` cmdlet in the Teams PowerShell module:

```
[PS] C:\> Set-CsOnlineDialInConferencingTenantSettings -AutomaticallySendEmailsToUsers $false
```

Thirdly, if you have Direct Routing deployed you could use On-network Conferencing for Audio Conferencing that allows PBX users to dial-in to the meeting via the Direct Routing. This is designed to allow users who have not been migrated to Teams to dial-in to Teams meetings from their legacy PBX phone without incurring any billable calls. It is not intended to let you bring your own conferencing numbers to the Teams platform. You can read up further about On-network Conferencing [here](#).

And finally, you can control how the phone numbers of external participants are displayed in the Teams meeting. By default, internal users can see the full phone numbers of those who join Teams meetings via audio conferencing. For external users Teams applies a mask so that the numbers show up like +35\*\*\*\*\*85.

The `MaskPstnNumbersType` setting managed with the `Set-CsOnlineDialInConferencingTenantSettings` cmdlet controls how Teams deals with participant numbers. The setting can be:

- **MaskedForExternalUsers:** This is the default value and means that external users see masked numbers.
- **MaskedForAllUsers:** Teams masks the number for any user joining a meeting with audio conferencing.
- **NoMasking:** Teams doesn't apply masking to participant numbers.

As an example, here's how to enable masking for participant numbers:

```
[PS] C:\> Set-CsOnlineDialInConferencingTenantSettings -MaskPstnNumbersType "MaskedForAllUsers"
```

The cmdlet is only available to tenants with the Audio Conferencing license.

## Teams Meetings and Cloud Video Interoperability (CVI)

To complete our mission of building a universal meeting, the last barrier to overcome is the ability to join a Teams meeting from VTC (Video Teleconferencing) endpoints (sometimes referred to as legacy video conferencing devices). Even though Microsoft Teams Rooms and Surface Hubs offer a great meeting room experience, it may still be too expensive for an organization to replace fully functional VTC's in meeting rooms with Certified for Teams devices. CVI helps by allowing any legacy VTC device to connect to a Teams meeting. VTCs must always join the meeting with the details provided in the CVI enabled meeting invite, a Teams client cannot dial out to them through the CVI service.

Four CVI providers are available today and are all pre-configured by Microsoft, so you just need to select which ones you want to purchase and then enable it. You can access details of the providers via PowerShell:

```
[PS] C:\> Get-CsTeamsVideoInteropServicePolicy | Format-Table Identity, ProviderName
```

Each provider has its flavor to CVI:

- **Poly** is the only provider where the licenses integrate into the Microsoft 365 admin portal. It is a shared service and all VTCs join either via the lobby or directly into a meeting as Internal and external endpoints cannot be differentiated.
- **BlueJeans** is also a shared service but licenses are bought directly from BlueJeans.
- **Pexip** is the most flexible of the solutions. You can choose to host the service yourself, use dedicated hosting from a suitable service provider, or access a shared service delivered by a service provider. Licenses are purchased through a Pexip partner. Pexip is the only solution that can differentiate



between your own internal VTCs and external ones, allowing yours to join directly while making external VTCs enter the meeting via the lobby.

- **Cisco** is the latest addition to the CVI platform, and their offering is run as part of their Webex cloud.

The quickest to configure is Poly, as the management of their licenses happens in the Microsoft 365 admin center, which suits those wanting a simple solution to enable CVI. You need one license per user who needs to include the CVI data in their meeting invitations.

Pexip differentiate themselves with the ability to choose a flexible deployment method which suits the enterprise market where there is a higher demand for integration with existing video infrastructure, customization, and control over media traffic. Additionally, with Pexip you pay for the capacity you need based on concurrency rather than per enabled user, so you can assign the capability to all users without extra licensing costs.

After choosing your preferred provider, you must run the *New-CsVideoInteropServiceProvider* command to configure the provider and its settings. You can set the provider up with false info in advance just to see how it looks in your tenant. You can create the provider and grant the policy tenant wide or on a per-user basis:

```
[PS] C:\> New-CsVideoInteropServiceProvider -Name CVI -TenantKey "teams@yourdomain.com"
```

You can then grant the policy at a tenant wide level or per user. For example:

```
[PS] C:\> Grant-CsTeamsVideoInteropServicePolicy -PolicyName PexipServiceProviderEnabled -Global
[PS] C:\> Grant-CsTeamsVideoInteropServicePolicy -PolicyName PexipServiceProviderEnabled -Identity ben.lee@office365itpros.com
```

Granting the policy formats your meeting with info from the correct predefined provider.

**Note:** To integrate with a CVI partner, you need to give the partner access to meetings in your tenant. This may pose some security concerns. Pexip has a good writeup of what the consent gives access to [here](#).

## Content Delivery Networks for Live Events

Live events are a special meeting type designed to cover a broadcast-style scenario where a few presenters send content to many watching users. From a Teams point of view the same principles of networking optimization apply, however we also have the option to integrate third-party Enterprise Content Delivery Network (eCDN) solutions to minimize the streaming footprint on your internet connection.

eCDNs work as local caching repositories so that the required media is streamed once from the internet to a local destination. Clients nearby access that copy instead of downloading their own version. For Teams where all traffic is encrypted deploying your own eCDN solutions would not be possible. This is why Microsoft has partnered with some vendors to bring an integrated solution to the Teams service (similar to how we just discussed using third-party CVI solutions for meeting interop). Using eCDNs can help reduce your internet bandwidth by as much as 98% when viewing Live Events.

Hive Streaming and Kollektive are examples of eCDNs that are already integrated with the live event service. You still need to install and configure the third-party software on each client to set up a peer-to-peer streaming network inside your organization per the vendors instructions, but once installed you can configure the selected solution for your tenant. This configuration enables an eCDN provider and redirects media streams to use their service.

In the following PowerShell command, we enable Hive Streaming for the tenant by amending the broadcast configuration.

```
[PS] C:\> Set-CsTeamsMeetingBroadcastConfiguration -AllowSdnProviderForBroadcastMeeting $True -SdnProviderName hive -SdnLicenseId {license ID GUID provided by Hive} -SdnApiTemplateUrl "{API template URL provided by Hive}"
```

If you run large Teams meetings with more than 1000 attendees, then eCDN is used for all users in the overflow experience for user number 1001 and up joining the meeting. Here's where [to read more about eCDN](#).

# Teams Phone

Teams Phone is used for making and receiving calls over the telephony or PSTN network from inside Teams. Teams Phone is a very mature solution with many organizations worldwide relying on it as their only phone system. Many have undergone projects to remove costly legacy phone system solutions that require lots of support and upkeep.

## Teams Phone delivery methods

To make and receive PSTN calls in Teams, besides having a Phone System license, you need to have some way of bringing dial-tone to your Teams environment. There are three broad categories of how you can manage this:

- **Calling Plans** with numbers from Microsoft if they are available in your region. Call charges and numbers are provided by Microsoft
- **Operator Connect** lets certified voice carriers/partners bring their telephony services directly to your tenant. Call charges and numbers provided by the carrier
- **Direct Routing** provides a BYO approach where you (or a partner) provide SBCs and PSTN connectivity. Call charges and numbers provided by you or the chosen hosted DR provider

Calling Plans were covered in detail in the previous section as little other configuration is needed.

### Operator Connect

Operator Connect bridges the gap between the management simplicity of Calling Plans (where no tenant configuration is required) and the freedom of Direct Routing (selecting your providers, negotiating minutes outside of what Microsoft can provide). It is a certification program where voice carriers or Partners (operators) must meet minimum requirements for connectivity into Microsoft 365 and maintain certain levels of calling quality. Once a partner has been certified, they are available for selection via the Teams admin center.

Operator Connect settings are shown under **Voice** and **Operator Connect**. Here you can send a request to one of the operators offering services in your required countries. A notification goes to the operator, who will contact you for further steps to enable their service. Microsoft is not involved in the billing aspects of Operator Connect services. The certification program only covers the technical aspects of the carrier to Teams integration, all billing for PSTN services is handled directly between you and the provider.

The advantage of using Operator Connect is that the operators have large-scale direct SIP trunks connected to the Microsoft backend with end-to-end QoS. As an administrator, you must assign users the Phone System license and a number. No additional configuration is necessary, much like Calling Plans. Operator Connect is a great way to add local connectivity to the Teams offering, with next to no configuration, and might allow you to reuse existing telco agreements.

As part of Operator Connect, you can also get Operator Conferencing, letting the operator provide local dial-in conferencing numbers. It may be beneficial if you expect large numbers of employees to dial into meetings, and you want to cover that traffic with your existing telco agreement.

The testing process behind Operator Connect is very rigorous and lengthy. To help rapidly onboard more carriers, Microsoft introduced a program called the Operator Connect Accelerator (OCA). OCA allows smaller partners to bring their PSTN offerings to Teams through several pre-certified providers. This shows that

Microsoft has the ambition to quickly onboard a broad set of operators. As a result, you may find your favorite operator on the list soon if they are not already there.

**Operator Connect Mobile:** Microsoft announced Operator Connect Mobile in March 2022. This enables the assignment of mobile numbers to Teams users. Inbound calls ring on Teams or a mobile device as a cellular call. Operator Connect Mobile requires a special eSIM solution from carriers supporting the program in your mobile device. The intended use case is to help Frontline Workers stay connected with a single number. OCM is not a simple call forking scenario but deeper integration where the same number is accessed across a user's mobile phone and their Teams client. Read more about Operator Connect Mobile [here](#).

## Direct Routing Architectures

Direct Routing is the fully flexible way to bring existing PSTN connections into Teams. Teams does not care what the actual PSTN connections are, they could be legacy Integrated Services Digital Network (ISDN), analog connections or SIP trunks to either another PBX or your carrier. What Teams does care about is that you have:

- Have a [compatible SBC](#).
- Provide internet access to the SBC.
- Have an internet resolvable DNS name for the SBC.
- Provide a public certificate for the SBC.
- Configure your firewall to allow Teams traffic to the SBC.

This SBC can be hosted in many different ways, and you may not even need your own SBC but can simply pay to use a service provider/partner's hosted SBC with whatever PSTN connectivity they provide.

Three main scenarios exist for deploying Direct Routing:

- Installed in your data center.
- Installed in Azure.
- Hosted by a Service Provider.

**Note:** Some SBCs are also certified for connecting analog devices and make them reachable from Teams read more about it in the [Microsoft documentation](#). Read more about it in the [Microsoft documentation](#).

### Installed in Your Datacenter or Offices

You may install your own SBCs in a datacenter or local office depending on what needs to connect into the SBC. This option is for companies that need to manage and maintain their own SBCs and have full control of call paths. This could be for a number of reasons such as; providing coverage in countries where neither carriers nor Microsoft can deliver Teams numbers (commonly countries such as Russia, China or Brazil); or wanting media to stay on their network as much as possible.

### Installed in Azure

You could install virtual SBCs in Azure to still have full control over them and include them as part of your managed infrastructure but do not have to maintain any hardware. If you have a strategy of removing on-premises equipment and are reducing your datacenter infrastructure this is a good choice. The leading SBC vendors (Ribbon and AudioCodes) have software-only versions of their SBCs and AudioCodes have a [whitepaper on how to set up their virtual SBC in Azure](#).

The requirements from Microsoft's side are still the same regardless of where the SBC is running and if it is a physical or virtual device.

**Note:** When you have set up your Direct Routing solution and want to validate and test it, you can use the SIP Tester client script found [here](#). This tool requires that you authenticate with licensed users who have numbers assigned. The script does not support modern authentication, so the account must bypass any conditional access policies if implemented.

### Hosted by a Service Provider

Since the integration with Teams using Direct Routing is just a SIP trunk connection, service providers can easily connect to multiple tenants from a centralized SBC deployment. Microsoft supports a specific variation of Direct Routing designed for just such a scenario. The service provider simply registers a subdomain per customer from their infrastructure and uses this to connect to the customer tenant.

You might wonder how this can operate when the domain used by the SBC is different from the tenant domain? Remember that from Teams point of view the Direct Routing connection is just a SIP trunk; all it needs to see is an FQDN registered in Microsoft 365 for connectivity. You can read up further on supported multi-tenant Direct Routing scenarios [here](#).

Consuming Teams Phone from a service provider is a good choice for organizations that do not want to manage any SBCs with their upstream connections and do not need to keep close control of the calling traffic.

### [When to Use Calling Plans, Operator Connect or Direct Routing](#)

From an end-user perspective, all of these delivery methods look the same. The differences are in what is happening behind the scenes to make calls work. The good news is that you can mix and match at will, depending on your requirements. For example, Calling Plans can be a quick and easy way to test out Teams Phone, but for a wider deployment, you might find Operator Connect a better deal, or you may want the flexibility to manage your own SBCs and migrate your current numbers and contracts.

**Note:** Emergency calling is available for Calling Plans, Operator Connect and Direct Routing. For example, you can configure policy settings to automatically notify your organization's security desk and have them included in the call, you can also define an in-client notification message that your users must acknowledge. Read more on emergency calling configuration [here](#).

## Direct Routing Configurations

When dialing from a Teams client, the following are all configuration elements used to determine where Teams will send the call:

- **Online PSTN Gateway:** contains the DNS name for an SBC, and any configuration to be used when sending calls to it, such as forward P-Asserted-Identity (PAI) or Preferred Codecs.
- **Voice Routes:** matches the pattern of the number dialed against Online PSTN Gateway entries.
- **PSTN Usages:** hold one or more Voice Routes. They can be shared across different Voice Routing Policies
- **Voice Routing Policy:** contain one or more PSTN Usages. These are assigned against users.

The configuration can be found in Teams admin center under **Voice** and is split between the **Voice Routing Policy** and **Direct Routing** sections.

This call routing can be confusing to set up initially and can become very complex when you have multiple SBC gateways and routes. This [docs article](#) provides a good reference for how to handle failover and other scenarios.

**Note:** You can find an in-depth overview of how routing works in Teams in [this Microsoft diagram](#).

By default, the traffic path taken by media is from your Teams client, to the Media Processes (in the Teams service) and then to the SBC before being placed onto the downstream PSTN connection. However, this may

not be ideal in scenarios where the user is already inside the firewall with the SBC, or if we want to reduce the volume of internet traffic generated.

With Direct Routing, three techniques can help with optimizing the media paths:

- Media Bypass for Direct Routing.
- Local Media Optimization for Direct Routing.
- Survivable Branch Appliance (SBA) for Direct Routing.

These scenarios do not apply to Operator Connect and are unlikely to be helpful when using Direct Routing hosted in Azure or Direct Routing hosted by a service provider as the SBCs are outside your core company networks.

### Media Bypass for Direct Routing

Media Bypass shortens the path of media traffic and reduces the number of hops in transit for better performance. This is a great option if you host an SBC in your data center or in Azure using Express Route. With media bypass, media is kept between the Session Border Controller (SBC) and the client instead of sending it via the Microsoft Phone System relay service. By enabling Media bypass, media traffic can negotiate to travel between the Teams clients and the external public leg of the SBC.

You must ensure that the SBC and firewall are configured to allow the correct ports through (even if it is just from inside your network). Internal clients also need to resolve the public leg of the SBC and send media traffic to that IP address on the port range of 50 000 – 59 999. On some firewalls, this can cause hairpinning issues where internal traffic goes to an IP on the outside edge of the firewall. It is only the media that goes to the SBC, signaling still goes through the Teams services.

When Media Bypass is enabled, you can restrict access to the SBC from non-Microsoft [IP ranges](#), forcing internet clients to still relay via the Teams services and reducing your external attack surface area.

To support Media Bypass the SBC must have ICE light negotiation configured, the SBC vendors have instructions specific to their devices with how to configure this. You then must enable it in Teams against the SBC object, with the following PowerShell cmdlet:

```
[PS] C:\> Set-CSOnlinePSTNGateway -Identity "SBCSite1.office365itpros.com" -MediaBypass $true
```

### Local Media Optimization for Direct Routing

Local Media Optimization is very similar to Media Bypass but takes things a step further by allowing the internal clients to send media directly to an internal IP address on the SBC (rather than relaying to the external interface). The advantage is that you can avoid firewall hair pinning and apply end-to-end QoS on the media traffic when making and receiving PSTN calls. It can also be used to have multiple internal SBCs (for example at branch sites) that can be accessible to Teams via one master SBC (in the datacenter) so that you do not have to publish multiple gateways to the internet.

Earlier, we talked about Teams Network Roaming policy and defined Public External IP addresses, regions, sites, and subnets. Local Media Optimization uses the same configuration data, and the clients needs to be identified as one of the defined local networks. If, for example, the public IP range does not match, then the client is deemed external and Local Media Optimization is not used. When using Local Media Optimization, you have two configuration choices:

- **Always** will send media to the internal leg of SBC regardless of where the user is located if they are classified as inside. This would require the SBC internal leg to be available from all internal networks.
- **Only for local users** means that the user subnet needs to match one of the subnets added to the *GatewaySiteID*. This applies if you have Direct Routing deployments with multiple SBCs in different countries. If that is the case, you also need to create a region, a site and assign the internal subnets for the site where the SBC should be accessed. You can re-use the configuration described for *Teams*

*Network Roaming Policy*, the *GatewaySiteID* is the same as the *SiteID* we created as part of that process.

To enable Local Media Optimization, connect to the Teams PowerShell module and use the following cmdlet:

```
[PS] C:\> Set-CSOnlinePSTNGateway -Identity "SBCSite1.office365itpros.com" -GatewaySiteID "Site1" -MediaBypass $true -BypassMode "OnlyForLocalUsers" -ProxySBC $Null
```

In the above cmdlet you also see that the *ProxySBC* attribute is set to *\$Null* as this setting is required. This attribute is used when you have a branch site SBC without internet connectivity. The branch SBC routes its calls to a centralized SBC and then to Teams, using the central SBC as a proxy. You can still use Local Media Optimization in this scenario, but clients need to be able to send media traffic to the internal leg of the branch SBC. To enable the *ProxySBC* functionality, the SBC is registered in Teams and a relay SBC configured. Here is an example PowerShell command to run to create a branch SBC using our existing SBC as the proxy.

```
[PS] C:\> New-CSOnlinePSTNGateway -Identity "SBCBranchSite2.office365itpros.com" -GatewaySiteID "BranchSite2" -MediaBypass $true -BypassMode "OnlyForLocalUsers" -ProxySBC "SBCSite1.office365itpros.com"
```

To complete the configuration, you would need to have a site called Branchsite2 with the correct local network subnets defined. Read more about configuring Local Media Optimization for Direct Routing [here](#).

### Survivable Branch Appliance for Direct Routing

Teams Phone relies heavily on an internet connection to make calls so you are stuck if a site's internet link goes down. Survivable Branch Appliances (SBA) are designed to help solve this issue.

An SBA is an SBC running a Windows Server based appliance which enables limited call routing to be maintained even if there is an internet outage. More precisely, it handles signaling for setting up the call, and directs the Teams client to route media to the local SBC. For this to work, Local Media Optimization needs to be set up and validated before the SBA is configured for that branch site. Setting up and configuring the SBA is well documented by the SBC vendors and usually requires an extra license from the SBC vendor.

The configuration in Teams is simple, and clients need to have signed in once for the configuration to be picked up before the internet link is unavailable. Run the *New-CsTeamsSurvivableBranchAppliance* cmdlet to create the SBA in Teams.

```
[PS] C:\> New-CsTeamsSurvivableBranchAppliance -FQDN "SBABranchSite2.office365itpros.com" -Description "SBA Branch Site 2"
```

After creating the SBA, you need to create an SBA policy and assign it to users using the *Grant-CsTeamsSurvivableBranchAppliancePolicy* cmdlet or you can use the *New-CsBatchPolicyAssignmentOperation* cmdlet for multiple users.

```
[PS] C:\> New-CsTeamsSurvivableBranchAppliancePolicy -Identity "Branch Site 2" -BranchApplianceFqdns "SBABranchSite2.office365itpros.com"
[PS] C:\> Grant-CsTeamsSurvivableBranchAppliancePolicy -PolicyName "Branch Site 2" -Identity "ben.lee@office365itpros.com"
```

Read more about configuring Survivable Branch Appliance [here](#).

**Note:** An SBA seems like a great solution to maintaining calling during an internet outage. In reality, most customers choose to improve resiliency in their internet connections (such as providing backup lines). This saves spending on hardware appliances that need to be maintained, and improving WAN connectivity helps sustain other services, such as access to documents and email.

## Calling Configuration

After you have decided on how to consume telephony in Teams with Calling Plans, Operator Connect or Direct Routing, it is time to assign the Phone System license to users to enable calling capabilities. One account setting that is easy to overlook is the *UsageLocation* attribute. The location is assigned at account creation and influences the services available to accounts, especially phone services.

When Teams makes a call the number format needs to be in the [E.164 format](#), a universal format that includes '+', country code, and subscriber number. Before the call is routed the number is normalized into the E.164 format by dial plans. All users with a Phone System license are assigned a dial plan by default. This is a basic dial plan provided by Microsoft that covers simple country based number conversions.

The *UsageLocation* attribute of a Microsoft 365 account controls the default dial plan and Audio Conferencing number. For Phone System, it defines how users normalize numbers and if they configured with a location where Calling Plans are not available you cannot assign that license. To check *UsageLocation* and the location based *DialPlan* assigned, use the *Get-CsOnlineUser* cmdlet:

```
[PS] C:\> Get-CsOnlineUser ben.lee@office365itpros.com | Select-Object DisplayName, UsageLocation, DialPlan, TenantDialPlan
```

To change the *UsageLocation*, use the *Update-MgUser* cmdlet:

```
[PS] C:\> Update-MgUser -ObjectId ben.lee@office365itpros.com -UsageLocation US
```

**Note:** Typically, adding the country code and a plus is not enough for localized country specific normalization of local or special numbers. <https://UCDialplans.com> has a comprehensive set of scripts to implement proper normalization rules. The site generates the necessary code to create *TenantDialPlans* that you can assign on a user level for the correct normalization of numbers. *TenantDialPlans* works together with the default dial plan. When dialing out, the normalization process looks at *TenantDialPlans* first and then at the default *DialPlan*.

## Calling Policies

Calling Policies control the calling features available to users, these do not apply to Meetings. A default global calling policy is available within a tenant found under **Voice** then **Calling policies** in the Teams admin center, but admins can also create and assign custom calling policies. The settings controlled in a policy are:

- **Make private calls:** Acts as a control switch for all calling functionality in Teams.
- **Cloud recording:** Allows calls to be recorded
- **Transcription:** Allows calls to be automatically transcribed by Azure voice services
- **Call forwarding and simultaneous ringing settings:** Controls if incoming calls can be passed on to either other internal users, or external PSTN numbers.
- **Voicemail is available for routing inbound calls:** If inbound calls can be sent to voicemail. It can be either enabled, disabled or left to the end-user to decide
- **Inbound calls can be routed to call groups:** Controls if call groups can be configured to receive incoming calls
- **Prevent toll bypass and send calls through the PSTN:** Helps meet any legal requirements if international calling cannot pass over an IP network and must use the PSTN.
- **Music on hold:** If system music should be played when a call is on hold.
- **Busy on busy when in a call:** What should happen if a user is already on a call and a second one comes in. Setting this to unanswered lets users decide in their client.
- **Web PSTN calling:** Allows PSTN calls to be made from the Teams web client.
- **Real-time captions in Teams calls:** Allows in-call captioning to be enabled

- **Automatically answer incoming meeting invites:** Useful to configure against a meeting room account where you want the room to join if invited automatically.
- **Spam filtering:** Options to include SPAM notifications as part of the incoming call toast if Microsoft identifies the number as suspect.
- **SIP devices can be used for calls:** Ability to place calls from SIP registered devices.

Some of these settings can have an impact on the traffic path, for example, if a user enables services that are cloud-powered like recording, transcription or captioning then the media must go through the Teams service for processing to enable the feature. Other features you may want to control at a more granular, country or regional level (with corresponding policies), busy on busy being a good example where different countries naturally expect this to behave a certain way.

## Cloud Voicemail

Another setting controlled by the **Calling policies** is Cloud Voicemail, which is available by default in the Teams client, even without the Phone System license. The voicemail capability replaces the old voicemail hosted by Exchange Online and Exchange on-premises servers. Microsoft hosts the voicemail service in Azure, and you control the availability of voicemail for users through the Teams calling policy assigned to their accounts. You can decide if users can enable voicemail manually, make it mandatory or disabled. If it is set to user-controlled, users can decide if unanswered calls go to voicemail by configuring it in the Teams client under settings. Sending a call to voicemail in Teams does not mean someone has to leave you a message, you can customize the behavior from the Teams client by clicking **Configure voicemail** in the **Settings** menu.

Integrating cloud Voicemail with a hybrid Exchange solution is possible, but the user's mailbox must be in Exchange Online for voicemail to work with Teams. Voicemail supports depositing voicemail messages only in an Exchange Online mailbox and doesn't support any third-party email systems. As a fallback mechanism, Cloud Voicemail can resend messages using SMTP, which means third-party email system could receive voicemail messages, but this has no guaranteed availability or support for features, such as changing their greeting so is not recommended.

Transcription of voicemails is enabled by default for all users enabled for Cloud Voicemail and can be turned off using the `Set-CsOnlineVoicemailPolicy` command. Twenty-seven languages are available for greetings, and transcription is available for ten of those languages. Read more about the set of supported languages [here](#).

**Note:** Some countries have language laws requiring that callers must have the option to choose which of the country's official languages they wish to be served in. Voicemail has Dual Language support to meet this need, read more about the capability [here](#).

## Toll Bypass Restrictions

Over 50 countries including India, China, Brazil, Algeria, Bahrain, and UAE have some form of toll bypass restrictions in their local law or telephony regulations. Essentially, they mandate that all international PSTN calls should go through a local endpoint so that the (usually state owned) telecommunications provider handles the call. For instance, if an organization has a support center in India, calls to international numbers must dial out from India. Modern routing methods such as least cost routing are therefore unhelpful in these scenarios as essentially for users in those countries we need to apply most cost routing.

Microsoft has introduced Location-Based Routing (LBR) to support this in Direct Routing based scenarios, LBR does not apply to Calling Plans as Microsoft does not offer services in these locations, and with Operator Connect your provider needs to manage it via their network. LBR is assigned on a site level and ensures that all international calls break out from the local SBC, enforcing international toll on international calls. Here is an example:



If you have users in India dialing a UK number and have SBCs in both India and the UK, ideally, you would route that call to the UK SBC and only pay the cost of a UK local call. However, as India has toll bypass restrictions this is not legally allowed. LBR would be configured to route these calls to the India SBC where it would be at full international cost. If one of those Indian users travels to the UK, LBR will match the new user and route calls to the UK SBC for the duration of their trip. If a UK user travels to India for a short period, they can still have their calls routed via the UK gateway.

LBR re-uses the same External IP, region, site, and subnets described at the beginning of this chapter. To use LBR, we need to enable the capability against the site we configured and assign a pre-defined Calling policy using the Teams PowerShell module.

```
[PS] C:\> Set-CsTenantNetworkSite -Identity "Site1" -EnableLocationBasedRouting $true
[PS] C:\> Set-CsOnlinePSTNGateway -Identity "SBCSite1.office365itpros.com" -GatewaySiteLbrEnabled $true -GatewaySiteID "Site1"
[PS] C:\> Grant-CsTeamsCallingPolicy -PolicyName "AllowCallingPreventTollBypass" -Identity "ben.lee@office365itpros.com"
```

The last command grants the built-in Calling policy called *AllowCallingPreventTollBypass* (where *PreventTollBypass* is set to *\$true*) to a user. If you have a list of all users in a site, you can use the *New-CsBatchPolicyAssignmentOperation* to assign it to multiple users at once. Read more about planning and configuring LBR [here](#).

## Number Management

There are three ways to consume phone numbers in Teams:

- Ordering them via Microsoft for Calling Plans
- Getting them from carriers via Operator Connect
- Acquiring them from carriers via sip trunks for Direct Routing

Service numbers, Calling Plan numbers and Operator Connect numbers can be found in the Teams admin center by going to **Voice** and **Phone numbers**. However, Direct Routing numbers are not shown in the Teams admin center as Teams does not hold a record of what your ranges are. You will see the location they belong to, if they are available and details of how they can be assigned (available usage).

To manually assign a number to a user, use the filter in the top right to search by Unassigned status and match the location to your user. After finding a number you want to use, select it and use edit to assign it. Using these filters is a good way to investigate how many available numbers you have at any given time for a location. This page also gives you access to your Microsoft number order history and you can get support using the link under the Actions menu for number-related issues, such as porting queries or changing number assignment types.

## Automating Phone Numbers

The Phone number portal in the Teams admin center is a great start, but the information found there is not available via Graph API yet, so you cannot use the same information in your automation routine. If you have numbers via Direct Routing or want to classify numbers by desirability, such as gold or silver, then you need to create your own routine. If you open PowerShell and connect to the Teams module, you can run one simple command to get all users with numbers in your deployment:

```
[PS] C:\> Get-CsOnlineUser -Filter {LineURI -ne $Null} | Format-Table DisplayName,LineURI
```

The command lists all numbers used in your Teams environment, including numbers assigned to CAP, MTR, auto attendant and call queue users. With this, if you also know your full number ranges, you could easily and programmatically find the next available number. There is a script available to get you started with this routine, and it can output all available numbers in your company as a CSV, grid view, and variable. You could integrate the script into an existing identity management procedure or use it as part of a standalone number

management process. For a larger environment, you might want to store this information in a database to keep additional information with the numbers. For example, if you need to mark the numbers of recent leavers so that Teams does not automatically assign the numbers to new users. You can add your own numbers for retention, and it will not offer numbers classified as gold or silver. Read more about this script [here](#).

## Unassigned Numbers

Organizations often have vanity phone numbers, which have a specific meaning to your organization, such as a previous sales queue number or an old main number. These are numbers not assigned to any auto attendant, call queue, or user, but you need them to be answered. Unassigned numbers in Teams can catch numbers that are not associated with a specific user or account and then redirect them to another user, resource account or play an announcement. An additional scenario is when people leave your organization, and you may want to route the calls to reception for a few months. You can configure unassigned number ranges that cover large portions of your full range, numbers assigned to users, resource accounts, or conference bridges are ignored and are routed to their correct destination.

Today, administrators can create unassigned numbers only using PowerShell with the *New-CsTeamsUnassignedNumberTreatment* cmdlet, and you can only assign a number or number range using regex. You can route the call to a number on Operator Connect or Direct Routing but be aware that if you send it to a Calling Plan user, Communications Credits must be available in the tenant to pay for it.

There are some good examples provided in the [Microsoft docs](#).

**Note:** [This](#) site builds a nice visual from your regex code so you can see at a glance if it looks correct.

## Blocking Numbers

Like the regular phone system, your users might receive calls from spammers or others offering dubious services. Users can block callers in the calling app in the Teams client, but it can be more effective to put an organization-wide block in place. This is configured using PowerShell by creating a block rule with the *New-CsInboundBlockedNumberPattern* cmdlet. As the name suggests, the block works by defining a number pattern for Teams to recognize inbound calls to block. The pattern can block a single number, a sequential range, or multiple separate numbers. For example, here's how to block a single number:

```
[PS] C:\> New-CsInboundBlockedNumberPattern -Name "Spam Block" -Enabled $True -Description "Blocks spammer from Tunisia offering Microsoft support services" -Pattern "\+21690373633"
```

After setting up a rule, you can test it using the *Test-CsInboundBlockedNumberPattern* cmdlet.

## Managing Users Calls

One of the primary benefits of deploying Teams Phone for users is that it gives them a lot of flexibility to choose how to manage their phone calls. The availability of these features is controlled through Teams calling policies covered in earlier.

Users can choose to forward their calls to other users or numbers in several different ways:

- **Call forwarding** sends the call straight to the destination number.
- **Simultaneous ring** forks the call and rings both Teams and the destination number, whoever answers first takes the call.
- **Call Groups** allow users to create a mini group of users they are working closely with who will receive the call on their behalf.
- **Delegation** for calling lets another user make or receive (subject to configuration) calls on the original users behalf.

**Note:** When a call is forwarded to another Teams user the incoming call notification will include a visual indicator to tell the user that it has been forwarded from another source. Teams also will not apply any call forwarding settings from the destination user to avoid a calls looping and forwarding forever.

## Configuring Call Forwarding for Users

These settings can be updated through Teams admin center under **Manage Users**, select the user then in the **Voice** tab, or via the `Get-CsUserCallingSettings` cmdlet. As well as seeing the current configuration, you will be able to update the timeouts before actions happen, for example allowing calls to route to the users Call Group after 20 seconds of trying the original user first. The documentation for the cmdlet is [available online](#).

## Call Delegation

Call delegation is when a user delegates the ability to make and receive calls on their behalf to someone else. This setting can be controlled through PowerShell, the Teams admin center or by end-users directly. It is best practice to try and get users to take responsibility for defining their call behavior and delegates are a part of that. When configuring a delegate, users can decide what a delegate can do. Options include; make calls, receive calls, or change call and delegate settings. If a delegate can change the call and delegate settings, they can add and remove other delegates on behalf of the original user.

A typical delegation scenario is when executives want personal assistants to be able to handle incoming and outgoing calls on their behalf. The executive can configure their first delegate through their Teams client, if allowed in their calling policy, by going to **Settings, General, and Delegation**. Once the first delegate has been assigned, they can configure any other delegates that are required. Alternatively, as executives do not usually want to handle setup tasks, you could use the Teams admin center to find the executive under **Users, Manage users** then configure the first delegate in the **voice tab**.

You can add up to 25 delegates for a user, and delegates can have up to 25 managers. There is no limit to the number of delegation relationships created in a tenant. Read more on how to configure it as an end-user [here](#) or as an administrator using PowerShell [here](#).

## Group Call Pickup

Group call pickup is configured in the Teams desktop client under **Settings, Calls** and then choosing a forwarding scenario that uses it. Here you can add users to your **Call group** as well as choosing the order to call them. When you have defined your call group, you can choose call groups under call answering rules and it rings to your group of users if you choose the call group if your call is unanswered.

As an administrator, you can manage users call groups in the Teams admin center under the voice tab in user administration, as for delegates. There you can define members of the personal call group and define the notifications the members get such as ring, muted, or in client banner. A tenant can contain a maximum of 32,768 call groups. [This documentation explains](#) how users can configure group call pickup.

## Call Park

Call Park (CP) is not strictly a user-calling action as you cannot park calls automatically. However, it is important to know about it as it is used in some niche, but important scenarios. It is used by companies where employees are hard to contact when they are working. Users (such as reception) can assist the caller and help connect them to the at-large employee. When the call comes in, reception (or any user) can park the call by using the "more actions" icon after answering it. A unique code will be given to the person who parks the call. An intercom, pager, or even texting can be used to send the code to the target employee. The target employee can then retrieve the call from a Teams client using the code provided. The code is entered in the Teams client on the calling page by clicking on **Unpark** button and entering the code. Call Park is available to users with a Teams Phone license.

Call Park is controlled via the call park policy assigned to user accounts. New call park policies are created using the *New-CsTeamsCallParkPolicy* cmdlet, or in the Teams admin center under **Voice, Call Park Policies**. Administrators assign policies to specific users using *Grant-CsTeamsCallParkPolicy*. If you want to remove the policy from a user, you still use the grant command but set the policy name to *\$Null*. Here is how you create the policy and grant it to a user:

```
[PS] C:\> New-CsTeamsCallParkPolicy -Identity "SalesPolicy" -AllowCallPark $false
[PS] C:\> Grant-CsTeamsCallParkPolicy -PolicyName SalesPolicy -Identity
"ben.lee@office365itpros.com"
[PS] C:\> Grant-CsTeamsCallParkPolicy -PolicyName $Null -Identity "ben.lee@office365itpros.com"
```

## Voice Apps

So far, we have covered scenarios for “individual calling” where calls are to/from specific users. However, it is very common for businesses to need to direct calls to groups of users, or have workflows associated with PSTN calls. For example, to route calls among a support team or to handle incoming calls to a main office number where you might also need to take a message after business hours. Fortunately, Teams allows for these scenarios with two types of Voice Apps:

- **Auto attendants** let you define welcome messages, interactive menu choices, and program different behavior based on opening hours.
- **Call queues** determine how calls are sent to different users (agents). You can define the group of agents, type of call distribution, and overflow options for calls in the queue. call queues can be mandatory, where users are always in a queue when online, or you can allow opt in and out so users can control their availability for receiving the incoming calls.

All the features discussed under group-based calling require that users and agents have a Phone System license and have a number assigned, either through Calling Plans, Operator Connect or Direct Routing.

**Note:** APIs to build customized switchboards and call centers are available inside the Teams service, allowing integrated switchboards that can deliver better reporting, wallboards, and more advanced IVRs than is available with the standard auto attendant and call queue functionality in Teams. The APIs are [Cloud Communications APIs](#) and [Presence API](#). Now that the *UpdateRecordingStatus* API is available, third parties can incorporate recordings of audio, video, screen share, and chat. Read more about [supported contact centers](#) and the recording API [supported vendors](#). To use [Policy Based Recording](#) Graph API Access and Contact Center solution, Graph API Access users need Microsoft 365 E3, E5, A3 or A5 licenses. Microsoft has announced a voice and Teams integration for their Dynamics 365 Customer Service which integrates with Power Virtual Agent. This integration can be used as an interactive voice response (IVR) for the voice channel, chatbot for SMS, live chat, and social messaging channels, all landing in Dynamics 365 and Teams. Read more in this [announcement](#).

## Resource Accounts

Behind each voice app lies a resource account. The resource account is a disabled user account configured with the number used so Teams knows where to route calls to. The settings for resource accounts are located under **Voice** in the Teams admin center. Resource accounts must have a special kind of license called the Teams Phone Resource Account. This is a free license, and you can get 25 licenses with the first Phone System license in your tenant, together with one new virtual license per ten Phone System licenses acquired. You can get the licenses from the Microsoft 365 admin center, from your CSP provider or as part of your Enterprise license Agreement.

After creating the resource account in the Microsoft 365 admin center, you can license the resource account and assign numbers. You can assign service numbers directly through the Teams admin portal, but you need to assign Direct Routing numbers via PowerShell.

Here is a summary of the process to create a resource account:

1. Obtain two or more service numbers as described earlier.
2. Obtain some free *Teams Phone Resource Account* license by going to the add-on licenses section in the Microsoft 365 Admin portal, under **Billing**, then **Purchase services**. You can have as many as you want but must provide valid billing details.
3. Create at least two Resource accounts under **Voice** in the Teams admin center and specify one as call queue and the other as auto attendant.
4. Assign the Virtual user license to the resource accounts, preferably using Group Based licensing in the Azure portal.
5. Assign a service phone number to the Resource accounts by going to the Teams admin center > **Voice** > **Resource account**. Select the resource account you are working on and click **Assign/unassign**. Next, select the number you want to assign and click **Save**.
6. Create a call queue, assign a Resource account, assign agents, and choose the routing method.
7. Create an auto attendant with a welcome message and open hours. Assign a Resource account and route the auto attendant calls to the call queue that contains the agents.

The example below shows how to create a resource account using the Teams PowerShell module. Give it a name and define the type of service by setting the *ApplicationID* which defines which service (AA or CQ) the resource account is to be used for. These IDs are the same across tenants as they reference the global M365 IDs:

- Auto attendant: ce933385-9390-45d1-9512-c8d228074e07
- Call queue: 11cd3e2e-fccb-42ad-ad00-878b93575e07

```
[PS] C:\> New-CsOnlineApplicationInstance -UserPrincipalName
"_service_TeamsRouting_GSDDR@office365itpros.com" -DisplayName "Service Desk" -ApplicationId
"11cd3e2e-fccb-42ad-ad00-878b93575e07"
```

After assigning the account a Phone System Virtual User license, we can assign it a Service Number by using the *Set-CsPhoneNumberAssignment*. The *PhoneNumberType* defines the source as *CallingPlan*, *OperatorConnect* or *DirectRouting*:

```
[PS] C:\> Set-CsPhoneNumberAssignment -Identity "_service_TeamsRouting_GSDDR@office365itpros.com" -
PhoneNumber +4714073200 -PhoneNumberType CallingPlan
```

Remember that you can mix and match calling delivery so even if your users numbers come via Direct Routing, you can still take Service Numbers directly from Microsoft. This can be a good idea for some queues, like an internal support queue where you would want that to be operational even if you have a fault with your infrastructure. After assigning a number to the resource account, you can then associate it to an auto attendant or call queue in the Teams admin center.

## Call Queues

Call queues are simple hunt groups configured in Teams admin center under **Voice** and **Call queues**. You would typically have an auto attendant in front of a call queue to enable opening hours and interactive routing. Call queues are used to distribute many calls (up 200 calls per instance) and use the first-in, first-out principle by default, there can be up to 50 agents per call queue. Users can opt in and out of the queue from the Teams client under **Settings** and **Calls** if allowed by the queue.

Microsoft introduced conference mode to speed up the time it takes to answer a call as an agent. The inbound call is held in a conference with audio established, so when an agent joins, they only need to add

their audio instead of negotiating a media path end-to-end. Without conference mode it can take longer before the agent can hear the audio of the inbound call.

The following routing methods are available for distributing incoming calls:

- **Attendant** – sends the call to all agents at the same time.
- **Serial** – sends the calls in the published order each time.
- **Round-robin** – distributes the calls at random between the agents.
- **Longest idle** – sends the call to the agent who has had no call for the longest.

If you turn on presence-based routing, calls will only be presented to agents who have Available as their presence status and is only applicable for users in Teams only mode.

Microsoft supports the mobile clients for the auto attendant and Call queues features so users on the move can still take calls. However, you should ensure they are fully aware of how to update their presence to avoid calls coming in at inappropriate times.

Another big customer request is the ability to redirect calls to a voicemail or external phone number as part of a call workflow. While an external number cannot be part of the group of agents called first, you can achieve this using queue timeouts and overflow configuration.

**Note:** Microsoft has released Power BI reports to report on auto attendants and call queues. These are not real-time reports as the data is delayed by at least 30 minutes. Even so, the data lets you review the past performance of queues. Read more [here](#).

Auto attendants and call queues are covered by the Phone System license and do not need additional licenses. Both services need service numbers to be assigned to your tenant for setup. You can assign multiple numbers to a single auto attendant or call queue using resource accounts.

Three ways exist to add agents to call queues. You can:

- Add users one by one.
- Add sets of users using a distribution list or Microsoft 365 Group.
- Enable collaborative calling by selecting a channel within a team.

Agents must have a Phone System license and have a phone number assigned to become call queue agents. Collaborative Calling adds a visual experience of the call queue as a tab in the channel named Calls. Users can see an overview of all calls and see active agents for the queue. They can see the call history, who answered calls, and what calls have been sent to voicemail. Via the call history, agents can listen to voicemails and choose to call back. Keep in mind that it may take 24 hours for the Calls tab to show in your chosen channel.

When agents call back or use the built-in dial pad, agents can display their personal number or choose to display a number assigned to any resource account. Numbers available are specified in the call queue configuration in the Teams admin center under **Assign calling ID**. This is separate from Caller ID policies and is configured per call queue. If you have a setup with Auto Attendants in front of call queues, you may wish for the agents to show the auto attendant number when dialing out. This can be done by adding the resource account for the auto attendant to the Assign calling ID configuration on the agent's call queue.

## Auto Attendants

The auto attendant is an automated interactive system for incoming calls that can generate welcome messages, give opening hours, provide interactive menu choices, and route calls to other auto attendants, call queues, and users directly. It also includes a voice-driven Directory Search to reach a person by name. Directory Search is a niche feature, but if you need it, be aware that the default configuration allows searching across the entire directory. To configure an auto attendant, go to the Teams admin center under **Voice** and

find **Auto attendant**. To give the auto attendant a number, you need to configure a Resource account as detailed earlier.

It is common practice for organizations to combine auto attendants (to allow time based routing and providing users a menu of call options) with call queues into a single voice workflow. This is a fully supported scenario.

A common requirement is to route calls to voicemail after hours, during holidays, or whenever the business is closed. The voicemail receiver is a Microsoft 365 Group; all members of that group will be able to see and play the message in any Outlook client. Auto attendants can also route calls to an external phone number, if you want the call to go to on-duty person for example.

Holidays are managed in the Teams admin center under **Voice** settings and **Holidays**. Here you can define a set of holidays and assign them to multiple auto attendants. It is possible to add holidays for a country programmatically as described in this [blog](#).

## Teams Devices

Teams supports a rich ecosystem of third-party devices where manufacturers have produced devices that are compatible with and certified for use with Teams. We spent a lot of time earlier this chapter talking about how to prepare your network and making sure you have a suitable connection to handle your traffic going to Teams, but a significant, but often underestimated contributing factor to call quality is the device you are using. If you put bad audio or video into Teams, you will only ever get poor quality out the other side. Audio quality can significantly impact the overall perception of the quality of experience.

It is not uncommon to get support cases where users complain about poor call quality in a meeting, only to find that their call streams were all good, but it was someone else broadcasting bad audio into the call (or picking up the conversations of everyone else nearby!).

The certification program not only guarantees a minimum level of quality for the key components of the device, but it also means that all the expected functionality works. This includes proper mute/unmute synchronisation between the device and your Teams client or answering a call directly from the device by pressing a button if your PC is locked.

## Device Types

Traditionally devices were split into two broad categories:

- **Personal Devices** are devices designed to be used by one user and typically refer to headsets, webcams or individual phones.
- **Shared Devices** on the other hand refers to equipment that more than one person uses, such as hardware in a meeting room or a phone in a communal workspace.

In the last few years Microsoft has been introducing several newer categories of devices such as Teams Displays and Teams Panels, as well as broadening their cross-device compatibility with the introduction of services like the SIP gateway. For this reason, I think that in this chapter it is helpful to move Phones and their specific nuances into their own category of Phone Devices as you. You can choose to either jump straight to, or skip over, that section depending on where you are in your Teams Journey.

You can find a full list of the current certified device & device categories in the [Teams Enabled Devices catalogue](#).

## Personal Devices

Personal devices typically refer to devices owned by, or signed in as, one user. It is an important device category as the chances are these devices will be the ones used most frequently by your users, and while there is a temptation to allow users to bring their own headsets, such as AirPods, while these headsets work, they might not provide the best experience for either the user in question, or the others in the call.

## Headsets and Webcams

The Teams client will accept the video feed from any USB camera that the underlying Operating System can access. The main consideration with webcams should be whether you want to give laptop users an external camera. If you want to drive the use of “video first” meetings, you may find that relying on the webcam in laptop lids doesn’t provide the best quality video (or the most flattering up-nose shots).

For headsets, it is a different story, because there are some specific benefits to having a certified device:

- Dedicated Teams button and LED indicator light. The Teams button provides context-sensitive actions; for example, press it when a call is ringing and it will answer, when a meeting starts and it will join the meeting, or long hold and it will raise your virtual hand in the meeting.
- User notification of events and alerts from the Teams client (e.g., meeting starting, voicemail or missed call).
- Ability to activate the Teams client or respond to notification with a button press (e.g., to join a meeting).
- Certification designed for extensibility (new device/client features to be delivered via firmware update).
- Automatic selection as the default audio device in Teams because they are USB based.
- Basic call control means that you are muted in Teams when you click mute on your headset.
- Audio quality such as no echo or excessive glitches, echo cancellation across devices, wideband audio support, comfort noise support.

Microsoft lists some benefits of using certified Teams devices [here](#), and you can also find the full specification for certification for all Teams devices [here](#).

Recently the Teams desktop client has been gaining many capabilities around audio processing where it can apply noise suppression to outbound audio generated during online meetings. Does this replace the need to deploy certified headsets? Not at all, it will simply complement the already good capabilities of your headset and further improve things.

**Note:** If you use headsets for both Teams calls and to listen to music, keep in mind that you have a secondary ringer option. You can define the device where you want to hear incoming calls. This setting is found in the Teams desktop client under **Settings**, and then **Devices**.

Most vendors have management tools you can use to centrally manage firmware upgrades and push configuration to their devices, [Jabra Direct](#) and [EPOS Manager](#) are examples such tools.

To know which headsets are used in your organization, you can create a custom Power BI report. The Call Quality Dashboard (CQD) database has this information, and you can create a view by looking at two fields, Audio Call Count and Second Capture Dev. To connect to the CQD and read this data, you must have at least the Teams Support Engineer admin role. At the end of this chapter, you will find all you need to know on how to implement the Power BI report and show it in Teams through a channel tab.

If you had previously invested in certified headsets for Skype for Business, those devices are still acceptable to use with Teams today. Anything Optimized for Lync 2013 or above works with Teams but may not support all of the correct button press actions.



For headsets certified for Teams with Bluetooth capabilities, you will find that you get a Bluetooth dongle with the headset. The built-in Bluetooth in laptops today is meant for keyboard and mice peripherals, not for Teams codec support or advanced HID functionality which enables the headset to control the Teams client. Bluetooth support without a dongle was recently improved allowing button press support but does not yet support the SILK codec.

## Teams Displays

Teams displays are a new category of device designed to act as a sort of “companion” alongside your primary computer. They are Android-based devices with a touch screen and speakerphone built-in. The offering from Lenovo is based on the same hardware platform as their Google smart assistant device.

They are personal devices intended to offload your computer and be an always-on, always-available Teams device, allowing users to join meetings without needing to disturb whatever’s happening on their workstation. The functionality available in Teams displays includes chat, channel conversations, switching tenant, meetings, and making and receiving calls. You can connect a Teams display to your computer, making the device part of the joining experience when you connect to a meeting. This means that the Teams display acts as an extension of a desktop client to deliver video and audio while shared content (as necessary) comes from your computer. These devices can also be used with multiple account logins, so they support hot desking scenarios.

While these devices may not seem like they have an obvious use case, they can prove popular for small offices or spaces where an executive spends a lot of time in Teams meetings. However, they are unlikely to go down well in open office type scenarios as they are intended for hands free use and can generate a lot of noise when used as a speakerphone.

From a management perspective, they benefit from the same management capabilities as the other dedicated Teams hardware devices we will cover next. In Teams admin center you can find them listed under the **Teams displays** tab under **Devices**. You can create and assign the same type of configuration profiles as you do with IP phones and control settings such as time zone and language. A list of all certified Teams displays is available [here](#) with a list of latest features available [here](#).

## Meeting Room Devices

Meeting room devices are the collective term given to hardware found in a shared space where people need to work together. These can vary in size and type from very large board room type spaces to smaller fluid huddle rooms or pods. When looking at what hardware to put in your meeting room you should consider its primary purpose, if it is mainly used for voice & video or if there is a requirement to also have some form of collaboration via whiteboards or some other medium.

The meeting room is one of the areas that is seeing the fastest evolution and growth with new features being added rapidly and a strong ecosystem of hardware partners producing devices.

### Phones

We will cover phones in more detail in the next section, but it is worth mentioning here that the most basic type of meeting room device is a phone or speaker phone. Teams phones come in many shapes and sizes, even one that matches the traditional “spider phone” device that you may have seen in many standard meeting room layouts. Teams phones have a special sign in mode for phones better suited for use in a meeting room. More on this and their modes later.

### Teams Panels

Teams Panels are Microsoft’s solution to the growing requirement to have status displays outside of meeting rooms. They display the current status of a room and show future meetings and availability. In addition, the panels allow users to search for an available meeting room and make an ad-hoc booking via the panel.

The panels support advanced scenarios such as meeting room check-in where a user taps on the panel as they enter the room to acknowledge that they are using the room for the booked meeting. If the room is not used after a predetermined length, the booking can be removed from the room and made available for someone else to use without worrying about clashes. Also, if a room is currently free, you can use the panel to book an ad-hoc meeting to claim the space as your own.

Teams panels are Android-based devices that can be managed in the Teams admin center under **Devices**. No extra user account or license is necessary for Teams panels as they are normally logged in with the same user as the Teams Room or Surface Hub devices already in the room. A list of all certified Teams panels is available [here](#) and more info on how to prepare for Teams panels can be found [here](#).

### Teams Rooms on Windows

Microsoft Teams Rooms on Windows (MTRoW – yes, it’s a mouthful), are designed to be a premium in-room solution for joining Teams meetings. Several vendors produce these devices. Sitting at the heart of each room setup is a Windows 10 PC connected to the other devices in the room (cameras, screens, microphones, etc.). This runs a special version of the Teams client designed to be operated full screen and via touch.

There are many kinds of MTRoW devices from vendors and many different possible configurations of cameras and microphones to suit different scenarios. Some vendors, such as Crestron, also bring their own customizations or integrations. For Crestron this allows managing the in-room devices (like raising or lowering projector screens, or diming lighting) but they all have to have a touch panel in the room to control the Teams application.

Teams Rooms can also join meetings hosted natively in Cisco Webex or Zoom using WebRTC. This can be a very useful feature to help maximize the use of your meeting rooms, because your users might be invited to meetings at other companies who use a different technology stack.

### Teams Rooms on Android

Microsoft Teams Rooms on Android (MTRoA) started out life as Collaboration Bars. They were intended to be lower specification meeting room hardware for smaller spaces where an investment in MTRoW wasn’t worthwhile. As time has passed, the hardware and software capabilities of the MTRoA devices have caught up somewhat with their larger MTRoW siblings. For some companies, they are becoming the preferred hardware choice as Android-based devices can be easier to manage and have (in theory) a smaller attack surface area. For example, if an MTRoA device develops a fault and needs to be factory reset these appliance-like devices are easier to restore than something running a full Windows 10 OS. There are some features that MTRoA do not have, however if you look at the latest [comparison list](#), you will see that there is not many. The main things missing from MTRoA are currently:

- Lack of support for content cameras. These are cameras pointed at traditional whiteboards to support better collaboration scenarios in a meeting
- Coordinated meeting join, where separate meeting room devices can be combined into the same meeting, usually with a Surface Hub
- Peripheral health management reporting

For most companies looking to deploy meeting room devices, the MTRoA appliance approach outweighs the slightly feature-rich MTRoW experience.

### Surface Hub

The Surface Hub is Microsoft’s own variant of meeting room hardware. It is a 55” or 84” touch and pen-enabled screen that can be used as a giant digital whiteboard. The Surface Hub can also run other Windows “modern applications” such as maps or a browser, so it operates almost like a large tablet at the front of the room. Some models of the Surface Hub can be mounted on a wheeled stand with a battery allowing them to be moved about freely, although your mileage may vary depending on your office layout.

The Surface Hub used to have a different meeting interface to the Teams Room devices as it allows more usage outside of just Teams, however with recent [updates](#) a new Teams application was installed which makes the room interactions align more closely with the Teams Rooms experience.

You can also configure a Surface Hub and a MTRoW to work in tandem with each other. For example, the Surface Hub provides digital whiteboarding for meetings and an MTRoW device provides the voice and video streams.

### Teams Meeting Room Licenses

The MTRoW/MTRoA/Surface Hub devices have their own license SKU for Teams which comes in a standard or Premium version. This license includes:

- Teams license
- Phone System
- Audio Conferencing
- Intune (for management)
- Azure AD P1

The premium license includes a managed service license from Microsoft where the devices are monitored and proactively kept up to date by Microsoft. Of course, you need onsite support available, but Microsoft can provide 24x7 escalation for issues relating to the meeting rooms. You can read more about the licensing for Teams Rooms [here](#), or about this unique managed service offering in the premium tier [here](#).

### Phone Devices

Teams supports several different types of phones but only one category offers full native support, the others operate through compatibility gateways. If you have legacy devices that are supported by one of the compatibility services, then analyse what functionality you need from your phones and test that they work how you need. If you have no existing phones that are compatible with Teams, then you should only consider deploying native Phones for Teams as they provide the fullest user experience.

### Phones for Teams

Microsoft has had hardware phones that are compatible with their platform for many years. How those devices have been built and operated has changed over the years. For Teams, Microsoft reset the phone delivery model and now allows manufacturers to produce their own hardware, but they must run a version of the Android OS. Microsoft then provides a version of the Teams client to run on them. This way hardware manufacturers can keep up with current specifications, but Microsoft has control over the application and the UI, providing users with a clear and consistent interface across different device types.

Features found on Teams IP Phones include:

- Native Teams experience with hardware button integration and LED notifications
- Calendar integration and one-touch meeting join
- Support for touchscreens

As mentioned in the previous section, these devices can run in three modes depending on your requirements:

- User mode, with a focus on contacts, calendar, and voicemail
- Meeting mode, with a focus on the calendar
- Common Area phone mode provides quick access to just a dial pad

It is easy to get started as an end-user with IP Phones, just enter your email address, type your password, and answer the two-factor authentication challenge, and you are in! The IP Phone inherits the same capabilities you have as a user, so if you have the Phone System license, you get the same capability when you are signed in with the device.

## Skype for Business phones (3PIP)

For businesses that have previously invested heavily in Skype for Business 3PIP (3rd Party IP Phone) certified phones, there is a special interoperability gateway that Microsoft runs as part of the Microsoft 365 service. This gateway allows these devices to be registered against a Skype for Business service (transparent to the user and administrator), allowing them to make/receive calls in Teams.

This gateway service will be retired in July 2023, so it is not something you should be considering using unless you already have a significant estate of these type of devices. You can read more about the gateway and the functionality available [here](#). While this gateway does provide a suitable method for connecting older devices, depending on the models, you might find that using the SIP gateway (see below) is a better solution with improved management capabilities.

## SIP Gateway

The SIP gateway is another service provided by Microsoft out of the Microsoft 365 cloud where certain compatible devices can be SIP registered against the Teams service. These phones will download a specific firmware version and then log on using the SIP protocol against the gateway and will have capabilities such as make and receive calls, hold, resume, mute, unmute and call transfers. A presence indicator is not available on the device itself, but the presence will update for other users in the Teams client.

You can see a list of compatible devices on this [page](#) from Microsoft.

## Common Area Phone License

You will need to provide an account with basic licensing when you need to use phones in shared areas (but not meeting rooms) such as canteens, reception waiting areas or factory floors. Microsoft has a specific license for this shared usage scenario.

Instead of giving these shared devices a full Office 365 E5 or Office 365 E3 license with a Phone System add-on, the Common Area Phone license provides just a Teams and Phone System license, so it is cheaper to purchase.

The Common Area Phone license, however, does not include an Intune license, which means that these devices cannot enroll into Intune. Therefore, depending on the volume of devices you have and your approach to device management, you may find it preferable to use a Teams Meeting Room Standard license.

# Device Management and Security

Teams device management is evolving quickly, more options are becoming available directly in the Teams admin center, making common tasks much simpler. For example, you can now remotely provision Common Area IP Phones, check the health status of MTRs and receive alerts about status issues. However, a downside to these improved admin capabilities is that some features are only available via the Teams admin center and not via PowerShell or Graph API where they could be automated.

Most of the devices we have talked about run either Windows or a variant of Android and are capable of being included in any Intune management of your estate (subject to the right licensing). However, you would not normally want to apply the same policies as you do for standard user devices. The challenge is that they are Android based but not a phone, or Windows based but not a PC. Policies in Intune or conditional access can block access, rendering the devices useless. Here are the steps you should take to cater for these devices.

## Managing Meeting Room Devices

Meeting room devices (or other shared hardware) such as Conference Phones, MTRoA, MTRoW and Surface Hubs are typically used in open shared areas, so we need to ensure they are properly configured.

For example, you might consider not allowing these devices to make calls to international numbers. If using Calling Plans, this is easy, just assign them a Domestic only Calling Plan and do not assign communication

credits. However, if you are using Direct Routing, you must create an *OnlineVoiceRouting* policy per country for CAPs and MTRs. Below is an example from Netherlands which enables dialing only to numbers starting with the correct country code.

```
[PS] C:\> New-CsOnlineVoiceRoute -Identity "NL-Amsterdam-NationalOnly" -NumberPattern
"^+\d{3,14}$" -OnlinePstnGatewayList @{{Add="SBCSite1.office365itpros.com"}} -Priority 1
-OnlinePstnUsages @{{Add="NL-Amsterdam-NationalOnly"}} -Description "For CAP and MTR will be able to
call Netherlands numbers except premium numbers and international numbers"
```

You might also consider blocking premium numbers for these shared devices. <https://www.ucdialplans.com/> has information on premium numbers for more than 200 countries. You can use this information to create PSTN Usages and Voice Routes which restricts the use of premium numbers for shared devices.

### Configuring Resource Accounts MTRs

Configuring accounts for devices can be complicated, as there are many different elements to check, especially if you are running hybrid identity or hybrid Exchange. Microsoft has a good starter article [here](#). Typically, when introducing a Teams component to a meeting room, you want to enable the existing Exchange meeting room resource account for Teams. Here is a checklist when configuring meeting room accounts:

- Make sure the network is ready with proxy exceptions and firewall configuration. These are the same as the Teams client on a PC.
- Create Intune enrollment token, so that all MTRs are enrolled in Intune as discussed earlier.
- Create conditional access rules for MTR accounts to avoid requiring MFA as discussed earlier.
- Assign Teams Rooms license to the synced user.
- Change UPN to have a common prefix, such as MTR.
- Enable the Exchange meeting room account for sign in and give it a unique password that never expires. Do this from Active Directory in a hybrid deployment or through Azure Active Directory if it is a cloud-only account. Migrate user to exchange online if part of a hybrid environment.
- It is usually enough to assign the Teams license to enable the room for Teams. If you have a hybrid SfB deployment, then you must migrate the account to Teams. If the resource account was not originally enabled for SfB and hosted on-premises in Exchange, consider enabling it as an online meeting room in the SfB deployment so that all attributes are configured correctly since this is a synced account. Run the below cmdlet from the SfB PowerShell module from one of the SfB Front Ends:

```
[PS] C:\> Enable-CsMeetingRoom -Identity "username@domain.com" -SipAddressType "EmailAddress"
-HostingProviderProxyFqdn "sipfed.online.lync.com"
```

- After the MTR is set up and Intune enrolled, evaluate changing the admin account password after deployment, as this password is the same for all MTRs by default. This [YouTube video](#) shows how to change the password.

When reusing an existing Exchange room resource account, you need to ensure calendar processing is configured correctly to work for MTR meeting invitations. It is important that the settings *DeleteComments* and *DeleteSubject* are set to *\$False*, otherwise the join button for the meeting may not show up on the MTR. By using Exchange Online PowerShell, you can check this with the following cmdlet:

```
[PS] C:\> Get-CalendarProcessing -Identity mtr-room1@office365itpros.com | Format-List
AutomateProcessing, AddOrganizerToSubject, AllowConflicts, RemovePrivateProperty, DeleteComments,
DeleteSubject, AddAdditionalResponse, AdditionalResponse
```

```
AutomateProcessing      : AutoAccept
AddOrganizerToSubject   : False
AllowConflicts          : True
RemovePrivateProperty   : False
DeleteComments         : False
```

```

DeleteSubject           : False
AddAdditionalResponse   : True
AdditionalResponse      : This room has Teams enabled room equipment for up to 10 people

```

You can use a string stored in the *AddAdditionalResponse* property to tell users that this meeting room has Teams Room equipment installed. A common misconception is that it should be possible to forward external Teams meetings to the Teams Room account. This does not work by default and the reason is that *ProcessExternalMeetingMessages* is set to false. To fix this, *ProcessExternalMeetingMessages* needs to be set to true, which enables anyone to send meeting invitations to the room. To enable the feature, run the following cmdlet:

```

[PS] C:\> Set-CalendarProcessing -Identity mtr-room1@office365itpros.com
-ProcessExternalMeetingMessages $True

```

In addition, you may see that the forwarded meeting invitation arrives at the MTR, but there is no join button. This could be because of the way that Microsoft Defender for Office 365 safe link processing works. You need to create an exception for your mtr-\* devices and not rewrite [https://teams.microsoft.com/\\*](https://teams.microsoft.com/*) links. When this exception is in place, you should see a forwarded meeting arrive at the MTR with a meeting join link. If you are using guest join with the MTRoW, you should also add Zoom and Webex meeting links to the allow list.

**Note:** A Microsoft article is available about [MTR security](#), covering hardware, software, account, and network security. The text is updated regularly.

## Intune and Conditional Access for MTRoW and Surface Hubs

The recommendation for MTRoW and Surface Hub devices is that they are not enrolled in Active Directory, but as cloud native devices. Teams is a cloud-only service, so it does not make sense to administer them with on-premises legacy techniques such as Group Policy. Here are the planning considerations for MTRoW:

- **The account:** should have a UPN prefix such as MTR-xxxx to let you create an Azure AD dynamic group for all room accounts. It is acceptable to rename the UPN of existing meeting rooms since users search for the room using display name, and email clients like Outlook use email proxy addresses to find accounts, which will not change.
- **Conditional access policy:** Exclude the group you made from the need to use MFA but require that it logs on from a trusted public IP address.
- **Compliance policy:** Create a Windows based compliance policy for Teams Rooms that includes BitLocker, secure boot and MS Defender.
- **Intune Bulk enrollment for Windows devices:** It is recommended to use bulk enrollment for Windows for both MTRoW and Surface Hub. The advantage with this is that the person setting up the device does not require Intune enrollment rights in your environment. This can happen after the first time setup for MTRoW but must be done as part of the initial bootup for Surface Hubs. Read more about bulk enrollment [here](#).
- **Bulk deployment of MTRoW:** Autopilot will soon be supported for MTRoW which will be the preferred way over using Endpoint Manager Configuration Manager. Using the autopilot approach may be able to automatically enroll the device to Intune as well.

## Managing Teams Rooms

MTRoW shows up under **Teams devices, Teams Rooms on Windows** in the Teams admin center, while MTRoA show up under **Teams Rooms on Android**. From there, you will be able to see the health of the devices, peripherals connected, call activity, and management history of your MTRs. A typical scenario is to discover when a peripheral is disconnected or what software version the device uses.

You can interact with the device by gathering device logs and restarting the MTR, be careful though as it will restart even if in a call. You can also control most of the settings you find in *SkypeSettings.xml* (a file used during manual deployment of the devices) such as **Account, Meetings, Device, Peripherals, and Theming**.

For instance, you can enable modern authentication under **Account**, turn on proximity join via Bluetooth beaconing under **Device** (beaconing is also required to let screen casting to the device work), and set the current theme under **Theming**. The custom theme still requires an upload of the image file to the device as a separate process.

The Teams admin center shows the MTRoW status under **Teams Devices, Teams Rooms on Windows** where administrators can view the health of their devices. If you click on a device, as shown in Figure 14-3, you can see the status of individual components that make up the MTRoW. You can use the [...] menu for each device and mark it as non-urgent, no-impact, critical or reset to default. In this example, you see that the Microphone and Conference speaker have been set to “No impact.” The room is configured to perform coordinated meeting joins with a Surface Hub, so they are not needed. Read more about the management options [here](#).

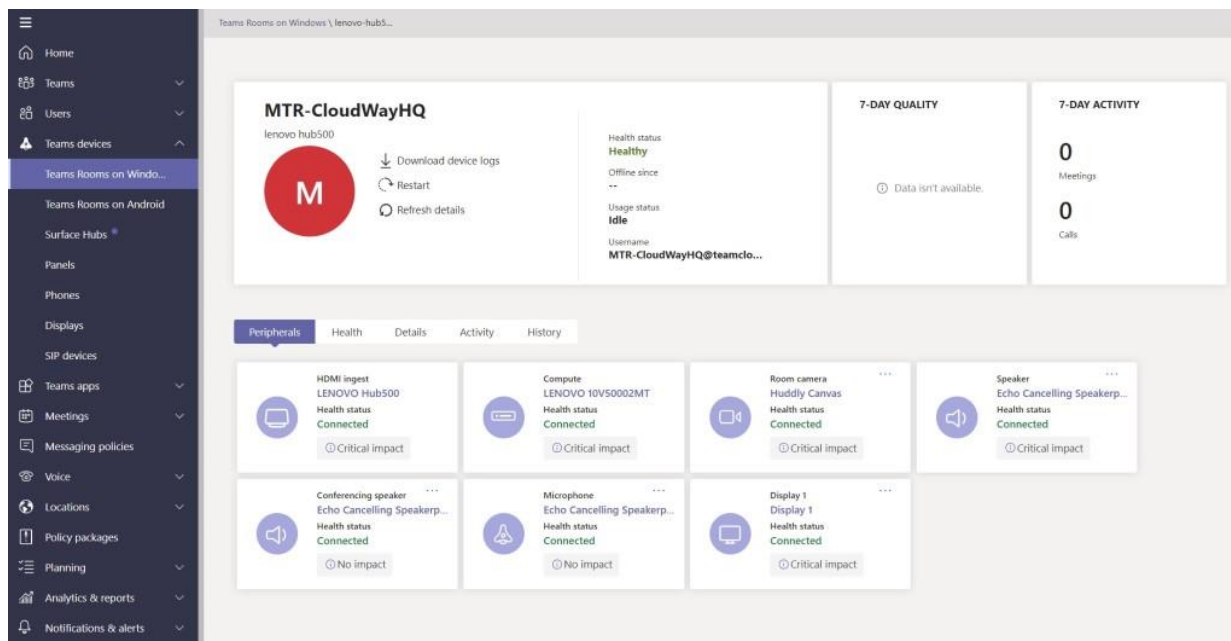


Figure 14-3: MTR management in the Teams admin center

**Note:** MTRoW devices are rebooted every night at 2am local time. During that reboot the device checks for updates, installs any if available and does necessary maintenance. This is a feature to make sure the device functions at peak efficiency. Microsoft wrote a [detailed article](#) about the difference between the quarterly updated Teams Room UWP application, which gets pushed via Windows store, and the Teams web app component which updates at a higher interval with new features.

## Managing Surface Hubs

Surface Hubs are discoverable and available the Teams admin center, but this is a preview feature so subject to change. By going to **Teams devices > Surface Hubs**, you can see the inventory of hubs that have signed into your tenant. You can see the software version and network connectivity and if the devices run the latest software version. If a device is connected correctly, it is marked as healthy. Unfortunately, you cannot publish any settings to the devices, but you can restart them and collect device logs for troubleshooting. Hopefully, more capabilities will be available in the future.

## Remote Provisioning for Teams Android Devices

Remote provisioning lets you safely sign in devices without having to give anyone access to account credentials. To remotely provision a Teams Phone, Teams Panel or MTRoA as a shared device, the MAC address must be added to the Teams admin center. Do this by selecting **Teams devices**, choose the sub-section for the device type you want to add, then in the top right corner click **Actions** and select **Provision devices**. Either add MAC addresses manually or upload a .csv file containing MAC address and location. The menu here contains a CSV template you can download. After adding the MAC address(es), select the phone or

phones and click on **generate verification code**. This will generate a verification code that is used at the initial stage of signing in. Distribute the list of MAC addresses and verification codes to the on-site technicians for deployment.

When the onsite personnel signs into the phone, click on the cog icon in the top right of the device screen. From there they select provision phone and enter the provided verification code. If the operation was successful a welcome screen will then appear. When a Teams phone is successfully provisioned it can be remotely signed in to using Teams admin center.

Go to the same section you just used to add the MAC address, except this time move to the **Waiting for sign in** tab. Select a Teams Android device to which you want to sign in and click Sign in a user. Follow the instructions on the pop-up screen, navigate to <https://microsoft.com/devicelogin> and paste the code. Select which account should sign into this specific phone and please note that it is only possible to sign in using a shared account. This is a great approach as it ensures that no need exists to share the password for the shared account with the onsite personnel. The list of supported devices [is available online](#).

### MTR Notification and Alerts

A default alert rule is available to get notification on changes to MTR health status which can be found in the Teams admin center > **Notifications & alerts** > **Rules**. These rules are near real-time, but you may have up to a 30-minute delay before you get a notification of an MTR being offline or unhealthy. It isn't ideal but does allow you to respond more or less proactively to an incident.

You need to manually add the MTR user accounts to the rule, and there is no PowerShell or Graph API approach to automate this. The monitoring rule looks at all devices the account is signed in to. You can receive the alerts in a Teams channel and/or via a webhook. To configure the JSON for the webhook, check out the docs article [here](#).

### Managing Phone Devices

Traditionally phone devices used to be placed in segregated sections of the network (who remembers the Voice VLAN) where they had direct access to the PBX onsite. While you can still use this technique to separate your phones from user traffic, it can be less helpful as the devices connect over the internet to the Teams service, and generate more than just voice traffic. Therefore, you need to make sure any QoS rules you have for user networks also apply for these devices. Removing this type of VLAN separation can help simplify your overall network topology.

### Networking Considerations

If you use 802.x authentication on your network, the recommendation is to pre-register the MAC addresses for Teams devices so that they can access the network. Most of the devices can be configured with Wi-Fi enabled, but keep in mind that while this might work, it is not technically supported, so you should plan to have ethernet connectivity available to ensure good quality calls. Most of these devices also support Power over Ethernet which helps to reduce cabling requirements.

### Managing IP Phones and Teams Displays

Management of IP Phones is performed in the Teams admin center, under **Teams devices** then **Phones**. There you can differentiate the view on All phones, User phones, Common area phones, and Conference phones. From here you can download device logs, update software and firmware and restart the devices. You can get the details on each device, including MAC address and call history. If you have people who specialize in device management and setup, then the Teams Devices Administrator role gives them access to just this part of the Teams admin center. Be aware that holders of this role do not get access to call history, so combining this role with the Teams Communications Support Engineer could be beneficial.



You can create configuration profiles to make sure all phones have the correct time zone, language, and time format. You can configure device settings such as display screensaver timeout, backlight timeout, and office hours. The recommended network setup on these devices is DHCP, but you can also manually configure network settings.

To assign a configuration policy for multiple devices mark all the devices you want and then assign a configuration policy to them. Sorting by IP address, for example, can help show all devices in one location. Make sure to name the configuration policy to reflect the location you want the devices assigned to, as you need to search for them when assigning. Unfortunately, there is no way to automate this procedure, so you should make this process part of a device onboarding routine.

### Intune and Conditional Access for Teams Phones and MTRoA

As stated earlier, MTRoA devices run Android but are not phones. This means that we need to make exceptions for these devices if Intune and Conditional Access are in use. There are two perspectives to look at when making sure these devices can sign in:

- For IP Phones that are for personal use, administrators should create an Intune enrollment restriction policy to require the preregistration of the serial number in Intune before the devices can enroll. This is how IT can stay in control of which devices can access corporate resources.
- For IP Phones or MTRoA that are for shared use such as Common Area Phones, we need to look at the account signing in to the device in addition to the above enrollment restriction:
  - **Account prefix:** Add a prefix, for example CAP-xxxx so that we can identify them with a dynamic group. We discuss account setup in the configuration part later.
  - **Compliance policy:** Exclude the dynamic group from other policies and require that the device is not rooted.
  - **Enrollment restriction policy:** Target the dynamic group and allow for “legacy” Device Admin management as IP Phones do not yet support Enterprise Admin Device management. It is also recommended that you use a corporate device identifier by pre-registering the serial numbers of the devices to control access to the environment.
  - **Conditional access:** Exclude the dynamic group from all other Conditional Access policies and create a dedicated policy that blocks all platforms except Android and requires a compliant device to bypass MFA. In addition, you can also require that the device logs on from a trusted public IP address so that users cannot take a common area phone home and continue using it.

### Intune and Conditional Access for SIP Gateway

Devices connecting via the SIP Gateway need to bypass conditional access policies because the actual sign in does not happen from the phone but rather from a Microsoft trusted IP addresses (that of the SIP Gateway). You should create a separate conditional access policy and add the accounts you expect to sign in on these devices, be it regular users or common area phone accounts. In this rule you should add the six IP addresses used to sign in the accounts. There are two IP addresses for EMEA, two for North America, and two for APAC. Your organization might only need to sign in from one or two regions, which means you can limit the authorized set. Review the current IP addresses [here](#).

### SIP Gateway Registration

You can get a full list of all SIP devices logged in to your environment through the Teams admin center by going to **Teams devices** and **SIP devices**. From here you can provision new devices by adding the MAC addresses to make sure that only devices approved by administrators can sign in. Limited capabilities are available to manage these devices: you can restart the device and see if the device is signed in. No configuration policies are available. You cannot define the device’s language as this is done by appending a

code string at the end of the registrar URL (read about supported languages and find the code [here](#)). To sign into a device the following conditions must be met:

- The user must have a Teams and Phone System license (or Common Area Phone license).
- The user must be assigned a phone number through Direct Routing, Operator Connect or Calling Plan.
- SIP devices used for calls must be enabled in the Teams Calling policy assigned to the users.
- MAC address must be pre-registered in the Teams admin center under **SIP devices**.
- URLs and IP ranges must be opened in the network where the phone is located (documented [here](#)).

The PowerShell code below is an example of creating an appropriate Calling policy. Note that these phones depend on call redirect, which must be set to enabled.

```
[PS] C:\> New-CsTeamsCallingPolicy -Identity "SIPPhones" -AllowSIPDevicesCalling $true
-AllowCallRedirect Enabled
```

After registering the MAC address of the devices, assigning the correct licenses and calling policy to user accounts they are ready to sign in. The device must be set up with the correct registrar address, either through DHCP or manually. The registrar addresses for each region are documented [here](#). When the user signs in, a pairing code appears and the user must navigate to the URL shown, type in the pairing code, and sign in using their credentials. For Common Area Phones, the process is the same, except that an administrator can perform the final sign in steps from the Teams admin center as these are special accounts.

### Configuring Common Area Phones

The Teams IP Phone policy, *CsTeamsIPPhonePolicy*, applies configuration settings specific to Teams common area phones. Available policy settings are:

- **AllowHomeScreen** – A Teams phone home screen displays a summary of contacts and future meetings.
- **AllowBetterTogether** – When connected to a computer it will sync the users' settings and experience.
- **AllowHotDesking** - Allows for any user in the organization to temporarily sign into the phone enabled with this feature.
- **HotDeskingIdleTimeoutInMinutes** – How long before a temporary logged in user will be logged out.
- **SignInMode** - Determines the sign in mode/experience for the phone: User, CAP, or Meeting room.
- **SearchOnCommonAreaPhoneMode** - Determines whether it is possible to search the Global Address List when CAP sign in mode is set.

You may not want to enable hot desking on CAP devices since there is a long timeout. By default, the timeout is 120 minutes. If the user goes to lunch and the phone remains signed in as that user, then anyone can answer calls or make calls on behalf of this user. Therefore, you might want to consider any risks of enabling hotdesking for CAPs.

To disable it, use this command:

```
PS C:\> New-CsTeamsIPPhonePolicy -Identity "CAPNoHotDesking" -SignInMode CommonAreaPhoneSignIn
-AllowHotDesking $false -SearchOnCommonAreaPhoneMode Disabled -AllowHomeScreen Disabled
-AllowBetterTogether Disabled
```

As you see, we disable *SearchOnCommonAreaPhoneMode* so that those using the phone do not have access to search in the Teams global address list. To apply an IP Phone policy to a CAP account, use the *Grant-CsTeamsIPPhonePolicy*:

```
PS C:\> Grant-CsTeamsIPPhonePolicy -Identity CAP-MainOffice@office365itpros.com -PolicyName
CAPNoHotDesking
```

# Troubleshooting and Monitoring Calls

So far, we have covered a lot of technical details about calling configuration, devices types and setup, as well as how to prepare your environment to support calling. Lastly, we will talk about how to keep on top of all these things, looking at what tools are available to help pro-actively monitor and troubleshoot any issues as and when they occur.

## Validating functionality for Teams Phone

When implementing Teams Phone it is important to test all features you are expecting to use before making it widely available. The goal is to find niche scenarios that do not work as intended as it can be difficult to troubleshoot when the system is fully live.

You should develop a structured testing process to document what is tested, how it was tested, what the result was. Then you will have a good overview of what was working when telephony was implemented. If it should stop working at some point, you know how it was working initially. Typical scenarios you want to verify when testing include:

- Normalization of numbers per country and region when dialing out.
- Forwarding and transferring of calls to mobile phones or Teams users.
- That your PSTN usages are working as expected, for instance, restricting calls to premium numbers.
- Verify that you can have calls longer than 30 minutes when using Direct Routing, sometimes, firewall filtering can interfere with the call and disconnect long calls.

Having a structured test to verify call handling may save you a lot of time further down the road. For example, when you know it was working initially but has stopped working, you know something may have changed to break the feature. You can find a very basic sample test schema [here](#), which focuses mainly on Direct Routing testing, but can be adapted to Calling Plans.

## Troubleshooting Teams Phone for Users

There are a lot of settings that need to be in place for telephony to function. Users must have the correct policies, the right set of licenses, and the correct settings in Azure Active Directory. This PowerShell one-liner will help you gather the key information you need for an account so that you can evaluate a user's settings all in one go:

```
[PS] C:\> $user="stale.hansen@office365itpros.com"

[PS] C:\> Get-CsOnlineUser $user | Format-List UserPrincipalName, DisplayName, SipAddress,
OnlineVoiceRoutingPolicy, TenantDialPlan, DialPlan, TeamsVideoInteropServicePolicy,
TeamsUpgradeEffectiveMode, EnterpriseVoiceEnabled, AccountEnabled, LineURI, OnPremLineURI,
FeatureTypes, TeamsCallingPolicy, UsageLocation, City, HostingProvider, InterpretedUserType

UserPrincipalName           : stale.hansen@office365itpros.com
DisplayName                 : Stale Hansen
SipAddress                  : sip:stale.hansen@office365itpros.com
OnlineVoiceRoutingPolicy    : AzureSBC1
TenantDialPlan              : US-NY-NewYorkCityZone01
DialPlan                    : US
TeamsVideoInteropServicePolicy : PexipServiceProviderEnabled
TeamsUpgradeEffectiveMode   : TeamsOnly
EnterpriseVoiceEnabled       : True
AccountEnabled              : True
LineURI                     : tel:+19175428xxx
OnPremLineURI               : tel:+19175428xxx
FeatureTypes                 : {CallingPlan, AudioConferencing, Teams, PhoneSystem}
TeamsCallingPolicy          :
UsageLocation                : US
```

```
City : New York
HostingProvider : sipfed.online.lync.com
InterpretedUserType : PureOnlineTeamsOnlyUser
```

Let's go through each attribute and discuss how to validate them:

- **UserPrincipalName:** This is the UPN of the user as found in Azure AD.
- **DisplayName:** Not much can go wrong here, it is based on *DisplayName* as set in the Azure AD account.
- **SipAddress:** The SIP address is based on UPN. If you see that the SIP address is not populated, then it may be that another user is configured with the same SIP address. Use the following cmdlet to locate the user:

```
[PS] C:\> Get-CsOnlineUser | Where-Object {$_.SipAddress -match "stale.hansen@office365itpros.com"}
| Format-List UserPrincipalName, DisplayName, SipAddress
```

You can manually overwrite the SIP address by populating the *SIP:ProxyAddresses* value in Exchange or the *msRTCSIP-PrimaryUserAddress* in AD.

- **OnlineVoiceRoutingPolicy:** To be able to dial out via Direct Routing, this attribute must be populated with the correct policy. To assign the correct value, use the following cmdlet:

```
[PS] C:\> Grant-CsOnlineVoiceRoutingPolicy -Identity $user -PolicyName "AzureSBC2" -Verbose
```

- **TenantDialPlan:** Make sure the correct *TenantDialPlan* is set on the user. This setting is configured manually. A good attribute to match with is the *UsageLocation* or *City*, depending on if you have multiple locations in the same country:

```
[PS] C:\> Grant-CsTenantDialPlan -PolicyName "US-NY-NewYorkCityZone01" -Identity $user -Verbose
```

- **DialPlan:** This is the default dial plan assigned to the user based on the *UsageLocation* attribute. If it is blank and *UsageLocation* is set correctly, make sure the user has a Phone System license assigned.
- **TeamsVideoInteropServicePolicy:** To make sure users get the correct CVI settings in their meeting invite if it is configured in the tenant. Set one of the three correct policies if you need to:

```
[PS] C:\> Grant-CsTeamsVideoInteropServicePolicy -PolicyName PexipServiceProviderEnabled -Identity $user -Verbose
```

- **TeamsUpgradeEffectiveMode:** This shows the mode the user is in, all the modes are documented [here](#), and to use Teams Phone should be set to TeamsOnly. To set the correct mode, run the following command:

```
[PS] C:\> Grant-CsTeamsUpgradePolicy -PolicyName UpgradeToTeams -Identity $user
```

- **EnterpriseVoiceEnabled:** If it is set to false, ensure the user has a Phone System license. If you migrate from SfBS and the user was disabled for enterprise voice, but should be enabled online, this attribute is still set to false. To enable the user, run the following command:

```
[PS] C:\> Set-CsPhoneNumberAssignment $user -EnterpriseVoiceEnabled $true
```

- **AccountEnabled:** indicates if the account is enabled for login in Azure AD. If the account is disabled, you can enable the user with the following command (after connecting with *ReadWrite* Graph API permissions):

```
[PS] C:\> Update-MgUser -UserID $user -AccountEnabled:$true
```

- **LineURI:** Defines the actual telephone number a user has and shows when dialing out from Teams. If you use Direct Routing, this attribute is synced with *OnPremLineURI*. If the *LineURI* attribute does not

sync, validate that the user is not also assigned a Calling Plan in addition to using Direct Routing. If Calling Plan is used, use this cmdlet to set the correct phone number:

```
[PS] C:\> Set-CsPhoneNumberAssignment -Identity $user -PhoneNumber +15555428572 -PhoneNumberType CallingPlan
```

- **OnPremLineURI:** Is the phone number provided to the user through Direct Routing. If you have migrated from SfBS to Teams, you need to clear out the *msRTCSIP-Line* attribute in Active Directory.

```
[PS] C:\> Get-ADUser -Filter "UserPrincipalName -eq '$user'" -property * | Set-ADUser -clear 'msRTCSIP-Line'
```

Wait until Azure AD Connect has synced the setting to Azure AD and then set the phone number using *Set-CsPhoneNumberAssignment*, which almost instantly populates the number. After the attribute is cleared, you can use the *Set-CsPhoneNumberAssignment* cmdlet to manage phone numbers. The *PhoneNumberType* option defines the source of the assigned number, this can be either *CallingPlan*, *OperatorConnect* or *DirectRouting*.

```
[PS] C:\> Set-CsPhoneNumberAssignment -Identity $user -PhoneNumber +15555428572 -PhoneNumberType DirectRouting
```

- **FeatureTypes:** Shows an array of values depending on what features are enabled. Values you may expect to see here include; Teams, AudioConferencing, PhoneSystem and CallingPlan. This tells you what features Teams thinks a user has activated.
- **TeamsCallingPolicy:** Defines the calling policy assigned to the user which defines calling features such as voicemail.
- **UsageLocation:** Based on *UsageLocation* in Azure AD. This attribute affects Calling Plan availability and the default dial plan that is assigned to the user. Use the *Update-MgUser* cmdlet to update the *UsageLocation* attribute:

```
[PS] C:\> Update-MgUser -UserID $user -UsageLocation GB
```

- **City:** A typical location defining attribute. A great way to determine which *UsageLocation* to assign to the user account.
- **HostingProvider:** If a different value than sipfed.online.lync.com is found here, the user is considered an on-premises user. If you have migrated all users online and decommissioned your on-premises environment, use this Active Directory PowerShell cmdlet to clean up the attributes for SfBS. First, to check the current value of the attributes:

```
[PS] C:\> Get-ADUser -Filter "UserPrincipalName -eq '$user'" -property * | Format-Table 'msRTCSIP*'
```

Clean up at least the *DeploymentLocator* attribute. The *PrimarySIPAddress* should remain as it helps the SfB client with SSO the first time a PC logs on. *LineURI* should remain as it is the user phone number when using Direct Routing:

```
[PS] C:\> Get-ADUser -Filter "UserPrincipalName -eq '$user'" -property * | Set-ADUser -clear 'msRTCSIP-DeploymentLocator'
```

- **InterpretedUserType:** This is a great source of information on the state of your user in a hybrid environment and pure online environment. Microsoft has added more states to this attribute, and it is possible to show its value in the Teams admin center in the user view. This [GitHub article](#) has descriptions of the different states for this attribute. For instance, it may show that the user is a disabled user in the local Active Directory, configured as an on-premises user when it should be an online user. It will also show if you are missing the correct licenses. If the user shows as *HybridOnpremSfBUser* when it should be a *DirSyncSfBUser* user, it may be an indicator of residual SfBS

attributes. To see the SfBS attributes for a user, run this command on a server that has the Active Directory PowerShell module installed:

```
[PS] C:\> Get-ADUser -Filter "UserPrincipalName -eq '$user'" -property * | ft 'msRTCSIP*'
```

It is a best practice that all attributes, except *msRTCSIP-PrimaryUserAddress*, should be cleared out. *msRTCSIP-PrimaryUserAddress* is used by the SfB client to discover the user's SIP address and populates the Sign-In address for newly deployed PC's. If *Disable-CsUser* was not run as part of decommissioning SfBS when moving online, these attributes are still populated. Here is how to clear out the common msRTCSIP attributes, make sure you capture all *msRTCSIP-Line* values if they are configured, so that you can configure users online with the same number:

```
#Get all msRTCSIP properties for a user that has a value
[PS] C:\> $Properties = Get-ADUser -Filter {UserPrincipalName -eq
"stale.hansen@office365itpros.com"} -Properties * | Select-Object -Property 'msRTCSIP*'

#Clear all properties for a user
[PS] C:\> Get-ADUser -Filter {UserPrincipalName -eq "stale.hansen@office365itpros.com"} -Properties
* | Set-ADUser -clear ($Properties | Get-Member -MemberType "NoteProperty" | % { $_.Name })
```

**Note:** Sometimes, users might complain about not seeing a dial pad for Teams Calling in their client. There could be multiple reasons for the problem. After enabling calling for a user, it can take up to 72 hours before all the settings synchronize across Microsoft 365 and another 24 hours before the desktop clients receive the new policies. If your users are in a hurry, ask them to check if a dial pad appears in the web client. If it is not there, refer to [the documentation](#) to make sure all necessary settings are present. There is even an administrator test button available to run a self-diagnostic test in the Microsoft 365 admin center.

## Call Quality

To succeed with a Teams voice deployment, you need to make sure calls are properly connected and that call quality is good. What constitutes good call quality? This is highly subjective, but packet loss, jitter, and latency are key network metrics you can measure. At the start of this chapter, we covered basic best practices for approaching your networking setup, but here we are going to go deeper into how Teams actually uses the network. This is where Teams gets challenging to understand and master, but luckily there are several tools available that we can work with, such as the Network Assessment tool, Call Quality Dashboard, and Call Analytics, that can help. These are all covered in the final section of this chapter.

### Understand Signaling and Media

Every call has two parts, signaling and media. Signaling is the part that manages your call such as establishing, maintaining, and terminating it. The concept used in Teams for signaling is the same as SfB, even if the protocol differs. SfB uses the Session Initiated Protocol (SIP) while Teams uses HTTPS REST signaling. Media communication is still based on UDP and travels as directly as possible between endpoints in all scenarios.

In 1-to-1 calls, media will try to go directly between the clients after signaling is established. For clients on the same corporate network, or subnet, this is likely to succeed as shown in Figure 14-4. Voice and video here will have a higher chance of being good quality with fewer packet drops than when passing through the internet.




Figure 14-4: Signaling goes via the Call Control and media goes directly between the clients

In 1-to-1 calls where media cannot go directly, because a firewall is between subnets or one user is external, the media will be relayed as shown in Figure 14-5.




Figure 14-5: Signaling goes via the Call Control and media goes via the relay service in Azure

In an ad-hoc multiparty call with more than two participants or a scheduled meeting, the media will go to the conferencing service as shown in Figure 14-6. Regardless of who is invited to the scheduled meeting, the first participant to join the call dictates where the conference is hosted. For instance, if you have three participants in the US and 10 participants in the EU and a US participant joins first, the conference will be hosted in a US data center, and all 10 EU participants must join the call hosted in the US.

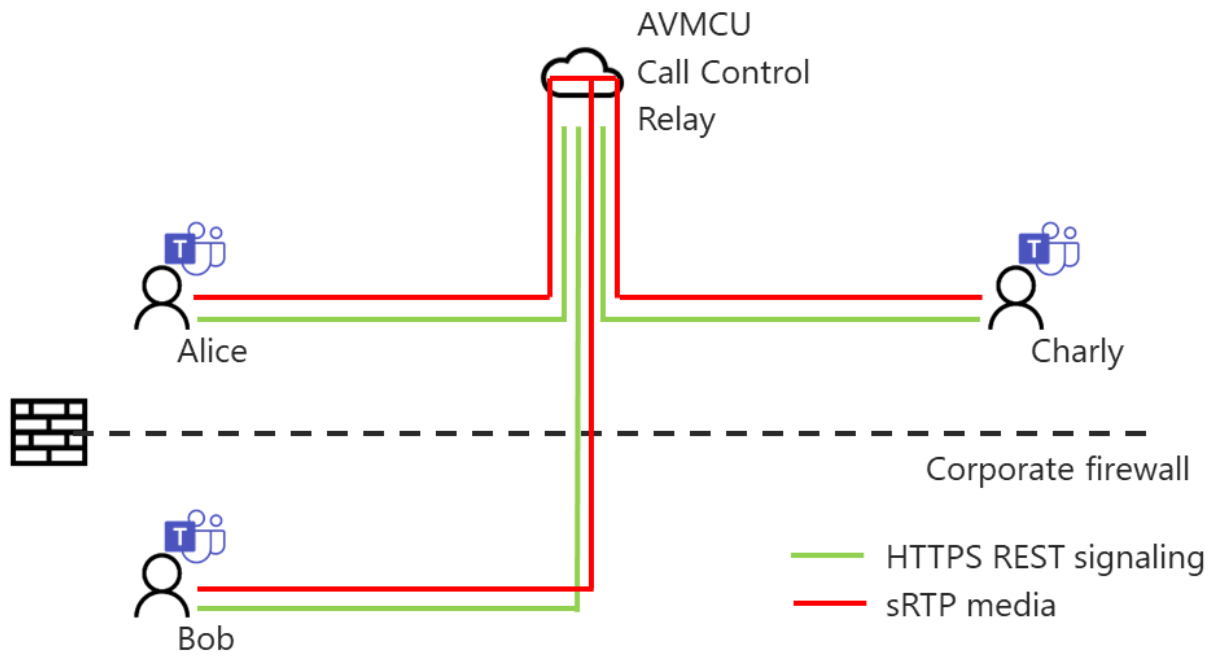


Figure 14-6: Signaling goes via the Call Control and media goes to the AVMCU where the first joiner is homed

If you have a VPN configured, you could see what effect this would have on the traffic. For example, in Figure 14-6, Alice's media traffic must first come inside the firewall before going to the AVMCU. The best practice for Teams is to separate the Teams media traffic to go outside the VPN tunnel wherever possible. The traffic that Teams puts out on the network is always encrypted and so does not need the additional protection of a VPN tunnel as this just adds overhead to the packets (for processing and hops).

If you must use a VPN try to at least allow the Teams Optimize media ranges to go outside the VPN tunnel and keep in mind, those are just IP ranges and do not have URLs associated.

### Codecs and their impact

There are four codecs used for Teams depending on the scenario and devices used on either end of the connection:

- **Satin** is the primary codec for 1-to-1 calls and soon for meetings. It uses AI to optimize for high quality under high packet loss and starts at a bit rate for wideband voice at 6 kbps. At 17 kbps it can produce full-band stereo music. Check this [article](#) for some audio samples.
- **SILK** is the primary codec in meetings and is both wideband and narrowband across clients except for Edge and Chrome.
- **G.722** is the secondary codec that is available in all scenarios across clients.
- **G.711** is the secondary codec in PSTN calls.

Table 14-1 shows the typical bandwidth usage of Teams codecs, however Microsoft has not released data statistics for Satin yet.

**Note:** Forward Error Correction is when you start having packet loss in the network and Teams compensates by sending overlapping information per packet over the network so that we lose less information per packet lost.



<b>Audio codec</b>	<b>Scenarios</b>	<b>Audio payload bitrate</b>	<b>Bandwidth Payload, IP header UDP, RTP</b>	<b>Bandwidth payload, IP header, UDP, RTP, SRTP</b>	<b>Bandwidth Payload, IP header, UDP, RTP, SRTP, Forward Error Correction</b>
SILK Wideband	Meetings	36.0	52.0	64.0	100.0
SILK Narrowband	Meetings	13.0	29.0	41.0	54.0
G.722	Secondary in all	64.0	80.0	95.6	159.6
G.711	PSTN	64.0	80.0	92.0	156.0

Table 14-1: Teams and codec bandwidth in kbps for audio

Video is based on X.264 and hasn't changed much in recent years, Teams can scale video by resolution and framerate. Resolution may vary based on the screen the participants use in the meeting and the resolution they request. Each client sending video will send a stream to the Teams service at the best quality that can be supported (or is needed). The participating clients in the meeting then request different quality streams from the server depending on what layout is being shown to the user, for example there is no point in the client requesting a full 30FPS 1080p video feed from all users when they are only being shown as thumbnails alongside a screen share. Usually, the more participants there are in the meeting, the smaller the requested video gets as the more thumbnails are being displayed.

You only see an option for Large Gallery view when more than 10 people are in a meeting. In gallery view the meeting service stitches together the video streams from participants and then sending one video stream of this combined image to the users client.

Screensharing is done by Teams using Video Based Screen Sharing (VbSS), which means that the share is another video stream in the meeting that is handled differently when displayed by the client.

PowerPoint decks can still be presented through the Office Online Server and are not video based, so PowerPoint content is not recorded when using Cloud Recording.

Table 14-2 shows some video scenarios and potential bandwidth when using a 1080p resolution. Be aware of the equipment your users are using and where most of them may run video meetings and make sure you scale available bandwidth to anticipate the estimated load. This is where the network planner will help.

<b>Participants/Activity</b>	<b>Max resolution</b>	<b>Total max download bit rate</b>	<b>Total max upload bit rate</b>
2 Participants	1 * 1920x1080	4	4
3 Participants	2 * 1920x1080 (Full Bleed) 2 * 1280x720	8 5	6.5
4 Participants	1 * 1280x720 + 2 * 960x540	5.5	4
5+ Participants	4 * 960x540	6	1.5
Video Based Screen Sharing (Only)	1 * 1920x1080	4	4
N Participant + VBSS [N=0-4]	1 * 1920x1080 + N * 424x240	4 + (N*350 Kbps)	~4.34

Table 14-2: Teams and bandwidth for video (bit rate in Mbps)

Keep in mind that when Teams uses the large gallery and together mode in meetings, it is one video stream.

## Network Factors Affecting Voice Quality

Network metrics such as packet loss, jitter, and latency affect voice quality. These network metrics affect real-time media and are experienced in different ways.

Users experience packet loss in two ways. If you hear metallic and variable sound quality, it can be caused by random packets being dropped. Sometimes you may hear the person talking go silent for several seconds, which is caused by burst loss of packets, where contiguous packets are lost.

**Packet loss** is typically caused by transmission errors and router congestion. If contiguous packet loss exceeds 10 seconds, the call may be disconnected. This is typically a problem over Wi-Fi that is designed for access and not throughput. Especially if you walk around with an active call and need to switch access point, the handover time may be too long, and the call gets disconnected.

**Jitter** is caused by packets arriving in a different order than the order they were put on the network and with longer intervals. It is typically caused by packet taking different routes due to load balancers or re-direction due to router congestion. You experience jitter when you hear a short silence and then hearing the person talking faster than normal. That is the Teams client buffering packets and playing them back when enough packets have arrived. If there are more than 20 ms between packets the audio healer will start dropping packets and the experience will be the same as packet loss.

**Latency** is the time it takes for a packet to travel from the sender to the receiver. When measuring latency, it is important to put the result in context. Latency higher than 100 ms within the same country is not good. A latency of 150 ms across continents is very good. Latency is typically caused by distance, queuing, and buffer overflow on the network. You experience latency when people you talk to on a call seem to take a long time to answer. A call with a lot of latency can be quite difficult to manage and those on the call may end up talking at the same time.

Table 14-3 shows the target network metrics in an unmanaged network, which include all the locations the network administrator has no control over such as the internet, home office, or favorite coffee shop. Table 14-4 shows the target network metrics in a managed network, typically inside corporate offices. Meeting these network metrics is important to meet user expectations of good voice quality and dependent stability of calls. We look at managed networks in two scenarios, where the users are and the corporate network edge. The goal is to know when it is an internet problem or where it is a local network problem.

<b>Unmanaged Network Voice</b>	<b>Optimal</b>	<b>Acceptable</b>	<b>Poor</b>
Inter arrival packet jitter (average)	≤ 5ms	≤ 10ms	> 10ms
Inter arrival packet jitter (maximum)	≤ 40ms	≤ 80ms	> 80ms
Packet loss rate (average)	≤ 1.0%	≤ 5.0%	> 5.0%
Network latency one-way	< 100ms	≤ 100ms	> 100ms

Table 14-3: Unmanaged network targets

<b>Managed Network Voice</b>	<b>Teams client to Microsoft network edge</b>	<b>Corporate edge to Microsoft network edge</b>
Latency (one-way)	<50 ms	<30 ms
Latency (round-trip)	<100 ms	<60 ms
Burst packet loss	<10% during any 200 ms interval	<1%
Packet loss	<1% during any 15s interval	<0.1% during any 15s interval
Packet inter-arrival jitter	<30ms during any 15s interval	<15ms during any 15s interval
Packet reorder	<0.05% out-of-order packets	<0.01% out-of-order packets

Table 14-4: Networking requirements for Teams media in managed networks (source: Microsoft)

## Teams and QoS Tagging

Two QoS port ranges are used for Teams, 1-to-1 port ranges and meeting port ranges. 1-to-1 calls use the port range of:

- Audio using TCP/UDP: 50,000-50,019
- Video using TCP/UDP: 50,020-50,039
- Screen Share using TCP/UDP: 50,040–50,059

As well as using network based tagging, for your windows clients you can create a GPO to make Windows apply DSCP tags against the Teams.exe client. You can add different markings for audio, video, and screen sharing. Audio is typically DSCP value 46 which is the highest priority.

If you use Direct Routing with Media Optimization, the audio port range is UDP 50,000 to 59,999 and the traffic goes directly to the SBC. In the future, media bypass ports may change to 3478 and 3479.

Conference media traffic goes via the Teams media relay services in Office 365. Some of that transport is over the internet, but you are still able to add QoS tagging and prioritize the traffic while it is in the internal network. All media that goes via the Teams media relay use port UDP 3478 by default. To separate audio, video, and screen sharing traffic you need to go to the Teams admin center, **Meetings**, and **Meeting settings**, and turn on **Insert Quality of Service (QoS) markers for real-time media traffic**.

When this is done, port 3479 will be used for Audio, 3480 for video, and 3481 for screen sharing. The QoS markings will be persistent if used with ExpressRoute and traffic from Android, iOS, and Mac clients will have DSCP tagging on these ports. Read more about QoS for Teams [here](#).

**Note:** The [Media in Teams - Media Flow YouTube video](#) does a great job of explaining Teams and network impact and routing in detail.

## Monitoring and Validating Call Quality

Teams includes two ways to monitor and troubleshoot call quality problems: Call Analytics and Call Quality Dashboard (CQD). Call Analytics is meant for reviewing call quality for individual calls and you can see and explain why a call was experienced as poor. You get information about what equipment was used, if it was a wired or wireless call, and information on network metrics. CQD is meant to look at trending quality within your environment. CQD has a separate portal, found in the Teams admin center, but the CQD reports are best consumed via Power BI reports as described at the end of this chapter.

### Reporting Teams Phone Activity

When an organization uses Teams Phone, you will want to report on how people are calling, by understanding your usage you can make sure you are not paying too much for your PSTN services. A usage report of PSTN calls for Calling Plans and Direct Routing is available in the Teams admin center. The report is good, but you can also access the data programmatically to create custom reports. You can do this using the *callRecord: getPstnCalls* API (documented [here](#)). The information retrieved includes who called, duration, dial plan used, and destination, amongst other attributes. You might be unfamiliar with the Graph API, but the Teams community is here to help you out. You can find a PowerShell script to get this data [here](#) and a PowerShell module to simplify access to the data [here](#).

### Call Analytics

Call analytics, which can be accessed from the Teams admin center by going into **Users, Manage users** and then in a users details picking **Meetings & calls**, shows the users most recent calls and meetings. Call analytics focuses on single calls and meetings, providing a detailed analysis of the selected call or meeting for issue resolution by the helpdesk and requires one of the Teams Administrator, Teams Communications Support Specialist, or a Teams Communications Support Engineer roles.

You can analyze four core areas:

- **Device** shows you specific information about the capture and renders device the user was using during this call, and you can determine if it was a Teams certified device.
- **System** shows you system statistics such as OS and client patch level.
- **Connectivity** tells you how the system is connected, for example, via Wi-Fi.
- **Network** shows you networking statistics such as packet loss, jitter, and delay.

Figure 14-7 shows an example of the report where we see a good call where network performance is acceptable. The report shows you real actionable statistics in each category described above. You can use what you have learned here to analyze the metrics collected.

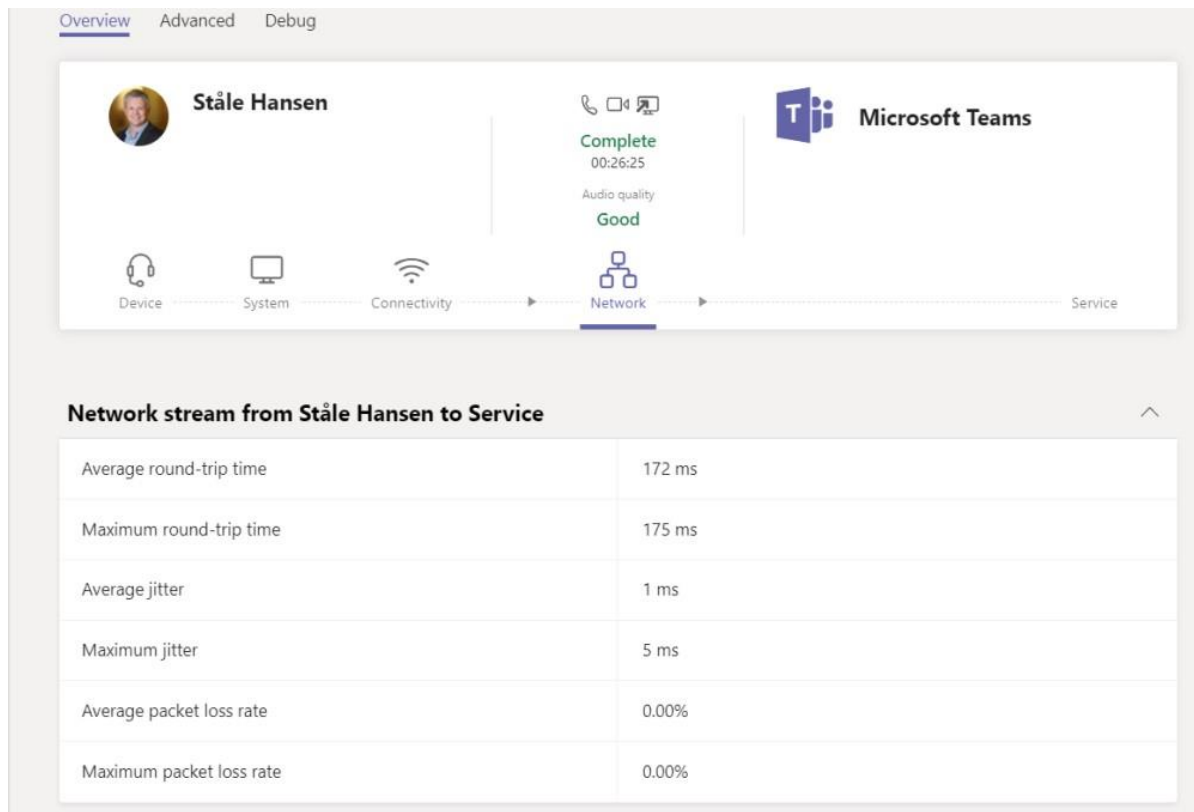


Figure 14-7: Call Analytics example of a conference call

Analyzing the metrics, you can tell the user why the call failed, whether it was because of packet loss on a Wi-Fi connection or if they used an unsupported audio device. You can also use the data to investigate further trending on the network location or the user's driver version. The report highlights poor calls for individuals and is only good for investigating individual calls. If you want to review further trends in your subnets, you should use the Call Quality Dashboard.

You can also see near real-time statistics for meetings when a user is in them. From the same part of Teams admin center with a user selected under **Meetings & calls**, any active meeting is highlighted as **In progress** and from there you can see user network metrics using the **Real-Time Analytics** capability. Having access to real-time information about calls allows the IT/Helpdesk to deliver live support during high stakes meetings. The data you will see is jitter, latency, packet loss, and frames per second for audio, video, and screen sharing. Based on this insight it is possible to take immediate action during meetings or use it as an active troubleshooting tool. There are some known real-time limitations which you can read more about [here](#).

You can also find a **Call Health** feature in the Teams client when in a call. This is accessed via the **More** menu. Here you can verify what codec is used, compare the network metrics to the guidelines in Table 14-4 and dive into details broken into four categories: Network, Audio, Video and Screen Share. Be aware that this

information is only what your client sends and receives and is not a total overview of all participants in the call.

**Note:** Users should be encouraged to use the **Make a test call** feature in the Teams client under **Devices** in the client **Settings** menu before meetings or important calls to verify that your sound is working properly.

## Call Quality Dashboard

Call Quality Dashboard (CQD) is a tool that helps you identify and troubleshoot trends in your organizations calling usage. It is there to capture symptoms and help you get a feel of the solution and where to focus your improvement efforts. As such, you should try and build a habit of exploring CQD at least once a month as part of your maintenance processes. If you have only recently deployed Teams or added a new workload such as telephony then you should be doing more regular weekly reviews to ensure things are working smoothly.

CQD is accessed via its own [URL](#) and is also found under **Analytics & reports** inside Teams admin center. All Teams admin roles can access CQD, which updates on average every 30 minutes. Due to compliance reasons, Personally Identifiable Information (PII) data is kept for only 30 days, and all other CQD data is available for up to 90 days.

CQD is all about streams. A stream is a one-way direction of a media modality in a call. You can have multiple streams within a call, so the number of streams in a location is much higher than the number of calls. It is dependent on how many users were in the call and how many streams there were per user. The streams are classified into three categories which are Good, Poor, and Unclassified. The percentage of poor streams indicates how many of the total streams were graded as poor. However, make sure the total stream count is significant, or this can throw out the average. Unclassified streams are too short to collect data or calls with federated parties. Data from federated parties is not visible inside your tenant. Figure 14-8 shows a summary report from the CQD portal, showing the overall number of streams for any given month.




Figure 14-8: All Audio Streams default report in CQD

A good place to start is looking at the Quality of Experience Reports. Here you can follow a breakdown of calls split by those with quality issues or reliability problems. Each section lets you dig into calls by type, such as conference, 2-party or PSTN calling. From the previous section on Call Quality, you should now be able to help identify where issues may be on the network. One sub-report here that can be particularly useful is TCP Usage. If you see a high volume of TCP traffic in your media calls, this indicates something is blocking Teams traffic. Teams always tries to use UDP for media first.

While you should be aiming to drive down the overall percentage of poor calls, it is unlikely that you will get it close to zero. Quite simply, too many factors can influence call quality and too many scenarios where users

are making calls. An administrator's job is to optimize as many of these factors as possible. Having a poor percentage figure below 4% means you are doing a good job.

**Note:** When working with CQD you will see two sets of most metrics. Anything generated by the server (so Teams) is always referenced as "first" and the client as "second" so when working with filters make sure you use the "Second" value, for example, *Second Building Name*.

### What classifies as a poor stream?

Five factors are evaluated to determine if a stream is considered poor, these are slightly different to the metrics we used previously for call troubleshooting. The stream is classified poor if one or all metrics are above the threshold shown in Table 14-5. These metrics are higher than what is considered bad, this way you know that a poor stream was terrible. Degradation average is when the network gets worse during a call, if the degradation average is above 1 you have an unstable or constrained network, look for routers or switches that are underperforming or overloaded Wi-Fi. Ratio Concealed Average is a technique used to compensate for dropped network packets, and the Ratio (%) is the percentage of packets concealed in a call.

<b>Network Metrics</b>	<b>Poor Stream</b>
Packet Loss Rate	> 10%
Round Trip Time	> 500 ms
Jitter	> 30 ms
Degradation Average	> 1
Ratio Concealed Average	> 0,07 (7%)

Table 14-5: Network metrics and poor streams

**Note:** If you want to learn more about CQD customization, the SOF CQD training videos found on the [Skype for Business YouTube channel](#) are highly recommended and are still relevant in a Teams setting.

### Implement Power BI reports to view CQD Data

Microsoft has released a set of Power BI reports to view CQD data. These reports are better and more usable than the native CQD reports to look at trending data and give the helpdesk a better tool to look at call quality. The Power BI reports connect directly to the CQD dataset and may be a bit slow because of that, but they give valuable information, so it is worth the wait. To view the CQD data in Teams via the Power BI reports, you need to have one of the below admin roles assigned.

- Global Administrator
- Global Reader
- Teams Service Administrator
- Teams Communications Administrator
- Teams Communications Support Engineer
- Teams Communications Support Specialist
- Reports Reader

Note that the Teams Support Specialist and Reports Reader roles do not have access to see PII data such as SIP address and telephone numbers. The available reports are:

- CQD Helpdesk Report
- CQD Location Enhanced Report
- CQD Mobile Device Report
- CQD PSTN Direct Routing Report
- CQD Summary Report
- CQD Teams Utilization Report

- CQD User Feedback (Rate My Call) Report

The helpdesk report as shown in Figure 14-9 is especially useful for day-to-day support. To get started with the reports you can download them [here](#). There is a detailed instructions in the zip file which you can follow, you can also follow the instructions in this [blog post](#).

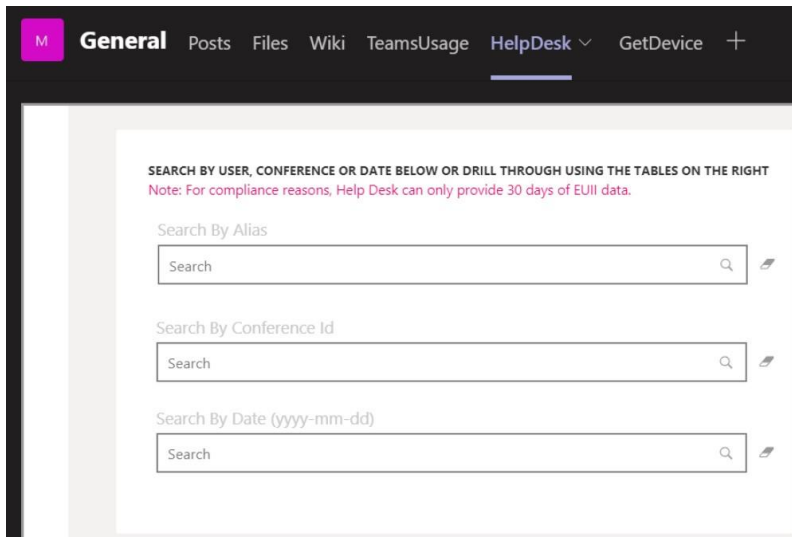


Figure 14-9: The CQD Helpdesk Power BI report in Teams

One important report is missing. This is an overview of all audio devices used in your environment as shown in Figure 14-10. As we mentioned earlier, Teams certified headsets should not be considered optional! We have created a report that shows all devices used in Teams and you can even drill down on specific users and see what devices they are using. We are looking at Audio Call Count and Second Capture Dev to accomplish this. You can download the CQD Get Device report [here](#). Follow the same procedure as for the CQD reports to publish this report to Teams. The report is provided as-is and you may see duplicate device entries in the report. When you have an overview of which devices are being used, you can see if it worth creating more user training to increase awareness around using certified for Teams devices.

Audio Call Count	Second Capture Dev
506	Sennheiser BTD 800 USB for Lync
471	Sennheiser SDW 5 BS - EU
365	Sennheiser SDW 3 BS - EU
297	Headset Microphone (Sennheiser BTD 800 USB for Lync)
290	Default input device
244	Headset Microphone (Sennheiser SDW 5 BS - EU)
220	Jabra Link 370
165	Headset Microphone (Plantronics BT600)
160	Headset Microphone (2- Plantronics BT600)
157	Audio
155	Realtek High Definition Audio(SST)
121	Headset Microphone (Sennheiser SDW 3 BS - EU)
95	Headset Microphone (Jabra Link 370)
91	Headset Microphone (Sennheiser MB 660 MS)

Figure 14-10: Reporting the devices used in your environment

**Note:** No centralized way exists to gather Teams Live Meeting reports. There is a manual way to import Live Event data into Power BI based on the post-meeting Q&A and attendee reports. Find out more [here](#).



# Chapter 15: Managing Clients

**Paul Robichaux**

## Many Clients, One Service

In the same way that users say things like “My Outlook is down” when they mean that their Exchange Server mailbox is broken, many of us think of cloud systems through the lens of the clients we use. Microsoft has a sometimes-confusing mix of clients for its cloud services, including desktop clients for macOS and Windows; mobile applications for iOS and Android; web applications; and embedded clients built into desktop phones, Surface Hub devices, and so on.

It makes sense to consider these applications as individuals for two reasons. First, the same application (say, Outlook) may have different features on different platforms. As an example, Outlook for Windows has several features, including voting buttons, that still haven’t made it to Outlook for macOS. Second, different applications on the *same* platform (say, Windows) may have significant differences too—you may remember how long it took for all of Microsoft’s desktop applications to fully support modern authentication.

With that in mind, we can put Microsoft 365 clients into three basic categories:

- *Browser-based clients* run in a compatible web browser and communicate directly with the service. Their supported features depend much more on the browser you choose than on the underlying host OS. Every workload in Office 365 has some level of browser-based client support.
- *Rich applications* run on macOS or Windows and provide, usually, more functionality than their browser-based counterparts. For example, the macOS and Windows versions of Outlook both support accessing multiple Exchange Online accounts in the same client, whereas Outlook on the web does not.
- *Mobile applications* run on Android or iOS and provide on-the-go access from tablets and phones to some, but not necessarily all, of the service workloads and their features.

Within these categories, you’ll notice some interesting differences. For example, Exchange Online, Teams, and To Do all have clients in all three categories, while Planner and Stream don’t have desktop clients. It’s important to keep in mind that Microsoft thinks of these clients as ways for users to interact with the data maintained in the “intelligent substrate,” the collection of all the data and services that Microsoft 365 delivers. If you think of Microsoft 365 as a single unified service powered by the substrate instead of a collection of standalone services, this blurring of client boundaries makes sense. As you’ll see in Chapter 9, a task is a task, no matter whether you access it through To Do, Planner, or Exchange Online, and no matter where it’s stored. As Microsoft adds new services to their cloud platform, they introduce clients for those services but the feature set and supported platforms for each service’s clients may differ.

## The Browser: The Base Microsoft 365 Client Platform

All Microsoft 365 workloads support web browsers. Most new services initially appear as web-only applications; some gain desktop applications later, but some do not. Some services launch on multiple platforms simultaneously. Teams is the best example as it launched into preview with web, desktop, and mobile clients from day 1.

For some time, Microsoft has stuck to a simple broad theme: the [Microsoft 365 system requirements page](#) recommends that you use the *latest* versions of the supported browsers. Right now, the “supported” list includes Edge, Apple Safari, and Google Chrome. Mozilla Firefox is labeled as supporting “most Microsoft 365 apps.” Other browsers, such as Opera or Brave, or older versions of supported browsers, may work, but Microsoft doesn’t guarantee that every app feature will work.

Notice what’s *not* on this list: Internet Explorer. Support for IE 11 formally ended June 15, 2022. If you’re still using IE11, shame on you. To help customers move away from it, Microsoft provides a feature in the **Setup** pivot in the Microsoft 365 admin center that allows you to customize a template to tell your users “hey, start using Edge,” including giving them instructions on how to reload their legacy web apps in Edge’s IE11 compatibility mode.

There are still a few edge cases to the broad outline of browser support. For example, [Teams audio/video calls won’t work](#) in Firefox just yet (see MC352622), and you may occasionally notice odd UI behavior or display problems in browsers not made by Microsoft. This situation is complicated by the fact that sometimes browser developers make changes that break things that worked properly before their intervention.

One more key thing you should know about browser-based clients: there are no special licensing requirements to use them (as there are for some of the desktop and mobile clients). There is also a specific license type, the kiosk or frontline worker license, which gives access mainly via web browser and not via desktop applications. Kiosk users can also access mailboxes using ActiveSync clients, as well as POP3 and IMAP4.

## Rich Office Clients

Microsoft 365 supports any version of Microsoft Office desktop software that is still in [mainstream support](#). The Microsoft 365 Business Basic and Standard plans, and Enterprise plans from E3 upward, include the right to install and run the Microsoft 365 Apps suite of enterprise applications (Word, Excel, PowerPoint, Outlook, and so on). Companies can also license the Office desktop applications separately, without any of the other cloud services such as Exchange Online or SharePoint Online included in the plan, and gain the benefits discussed below. But because you are reading this book, we’ll assume that you are doing more with Microsoft 365 than just using the desktop applications.

The desktop Office applications have historically been available in two forms. First is the familiar “perpetual license” version: you buy it once and keep it forever. Perpetual versions of Office for Windows have long been built on the Microsoft installer (MSI) format. The second, newer, form delivers Office as a subscription-based service. When you license Office this way, you only have the right to use it while you keep paying. In exchange for this ongoing fee, you get more rapid product updates. You will sometimes hear people refer to subscription-based Office as “click-to-run” or “C2R” because the Office apps themselves are built on an application virtualization technology (known as App-V) that allows application updates to be streamed on-demand. It is challenging to mix C2R and MSI versions of Office on the same machine, and Microsoft recommends that you not do it.

**Real World:** The Microsoft Teams desktop client doesn’t use either the MSI or C2R mechanisms even though it is officially now part of the Microsoft 365 Apps suite. You can download a platform-specific executable installer that streams the Teams installation bits to the machine using App-V, but after that initial installation, the Teams client will automatically update itself. Bowing to user requests, Microsoft also offers an [MSI-based installer for bootstrapping Windows Teams installations](#) through Microsoft Endpoint Configuration Manager or Group Policy, but once the product is installed that’s the end of your control over its updates.

In February 2021, Microsoft introduced a new perpetual version of Office, the [Long Term Servicing Channel](#) (LTSC) version, which was released in September 2021. As with the [LTSC branch of Windows 10](#), Microsoft targets Office LTSC at devices that cannot or should not be updated frequently. Microsoft cites examples such as process control systems on a factory floor or devices used in regulated industries such as healthcare or financial services. Office LTSC is intended for limited use on a subset of any organization's devices, not as the primary Office version in use. Instead of including the Skype for Business client, Office LTSC will include the Teams client, which raises some interesting questions about how Microsoft will balance the high frequency of Teams client and service updates with the concept of a stable, slow-changing LTSC version. In exchange for getting less functionality and less frequent updates, Microsoft is also raising the price of the LTSC version of Office Professional and Office Standard by "up to 10%." There's also a macOS version, known as "Office 2021 for Mac." Alongside these LTSC versions, there's also a perpetually-licensed consumer version known as "Office 2021."

Why would anyone ever be willing to pay an ongoing license fee for Office? That's a fair question, but this approach offers some concrete benefits. First is that Microsoft updates subscription-based applications regularly, and much more frequently than the perpetual versions. Secondly, and probably more importantly, getting Office apps as part of a Microsoft 365 subscription can be surprisingly cost-effective. Buying a subscription to Office apps offers a fixed, predictable monthly cost tied to the number of users you have. This has historically been true even though Microsoft in August 2021 announced a price increase for many Office 365 and Microsoft 365 plans—the additional US\$36/year for Office 365 E3 represents a 15% increase, but given that there have been no similar price increases in the past, it's not unwarranted.

Some people argue that the cost of a subscription is higher over the long term than the traditional model of buying licenses at a one-time cost. There are many variables to consider in that equation, including the upfront cost, the cost of buying new versions to remain in mainstream support, and the cost of deploying and maintaining the software with updates. It's important to make a fair comparison, and not try to compare prices for Home, Family, or Personal editions of Office with the full Microsoft 365 Apps suite, as the non-business versions are not licensed to connect to cloud services like Exchange Online. The Microsoft 365 Apps versions allow customers to install on up to 5 machines, which may further reduce your overall cost of deployment. For customers using Microsoft 365 cloud services, the additional requirement to stay fully within mainstream support will come into effect in the year 2023, requiring standalone Office users to upgrade to newer versions, whereas right now it's possible to stretch out the lifespan of an Office install and connect to Microsoft 365 services with versions of Office that are in extended support.

The third reason can be a bitter pill to swallow. Microsoft mostly releases new features, such as Focused Inbox in Exchange Online and integrated support for streamed language translations in PowerPoint, only in the subscription-based applications. This irritates many people who have purchased the perpetually-licensed versions, but there is no reason to expect Microsoft to add new features to a previously shipped version of a product for free. It's quite clear that the bulk of Microsoft's attention and focus going forward will be on the cloud, and that on-premises software, while still important, is no longer the most important thing to them. When you consider the likely future for Microsoft's perpetually-licensed desktop and server software, moving to cloud-based licensing seems much more attractive.

As a rule, the subscription-based client applications deliver the most up-to-date, feature-complete, and compatible version of Office for your users, and that will often outweigh any comparison based on price alone. The introduction of the LTSC version means that organizations that need a stable, slowly-changing version of Office can still get it, but Microsoft's positioning of LTSC as the solution for small, select deployments reinforces their argument.

**Real World:** Many client applications, such as browsers, Teams, and the C2R version of the Office desktop applications automatically update themselves when new versions are released. Some enterprises prefer to disable automatic updates and control the deployment of updates. That is fine... as long as you don't prevent the updates from *ever* being deployed. Your update strategy should be either automatic or manual. Not updating at all isn't a valid strategy for Microsoft 365. For smaller organizations that are content to allow automatic updates, the IT administrators should at least maintain awareness of the vendor releases by subscribing to RSS or email notifications from those vendors.

## Mobile Clients

Apple, Google, and Microsoft desktop and mobile operating systems all include basic applications for email and calendar access; they may also include apps for simple office productivity tasks, chat, and so on. Interestingly, Apple and Google also compete directly with Microsoft for productivity and cloud services dollars. That limits their willingness to put first-class support for Microsoft 365 into their native applications. That's perfectly OK, though, because Microsoft is doing that on their own.

Microsoft produces a wide array of Microsoft 365 client apps for iOS and Android, including:

- Authenticator (used for passwordless authentication, multi-factor authentication, and more).
- LinkedIn.
- Microsoft 365 Admin.
- Microsoft Lists.
- Microsoft Teams.
- Microsoft To Do.
- Microsoft Lens (formerly known as Office Lens).
- Office (collection of different capabilities including Word, Excel, and PowerPoint).
- OneDrive (the same app can connect to both OneDrive for Business and OneDrive consumer).
- OneNote.
- Outlook (described in more detail later in the chapter).
- Planner.
- SharePoint Online.
- Sway.
- Whiteboard.
- Word, Excel, and PowerPoint, although these apps have largely been superseded by the single Office app mentioned earlier.
- Yammer.

Companion mobile apps typically appear soon after an application reaches General Availability. For example, Microsoft Lists became generally available starting in July 2020, and later in the year Microsoft announced released the Lists mobile client for iOS, with the Android version following in late 2021. Some differences might exist between the iOS and Android versions of mobile apps due to the different capabilities available to developers on the two platforms. More details about how to manage mobile devices and application access are in Chapter 16.

## Managing the Microsoft 365 Apps Suite

The pace of innovation and development in Microsoft 365 makes it impossible to maintain support for older versions of client software without incurring high development and support costs that Microsoft would then need to pass on to customers. When you consider moving to the cloud, one factor in your decision must be

an implicit commitment to keep your clients up to date. If you allow your client applications to become outdated, you can expect the quality of the user experience to degrade over time. Eventually, users might be unable to use some advanced features because their client software is unsupported or obsolete. For that reason, in this book, we've chosen to focus on Microsoft 365 Apps, the "service" version of the Office desktop applications. There are other versions; the full list of supported versions of Office includes:

- Microsoft 365 Apps: The Microsoft 365 Apps for enterprise suite of desktop Office applications is available both for Windows and macOS.
- The Office LTSC version, described earlier in the chapter.
- Office 365 Professional Plus ("Pro Plus"): desktop applications before version 2004 retain the Pro Plus branding
- Office Professional 2019: Office 2019 is essentially a perpetually licensed snapshot of features previously introduced in Office 365 Pro Plus. It's also available for both macOS and Windows, but we won't discuss it further.
- Office Professional Plus 2016: the 2016 standalone version of the Office client apps.
- Office Professional Plus 2013: the 2013 standalone version of Office, and predecessor to Office 2016; this version is supported for use with Microsoft 365 until October 2023. (Note that this is not the same as "the 2013 versions of Office 365 Pro Plus," support for which ended in 2017).

These are just the product names. In terms of specific versions, Microsoft doesn't usually emphasize specific version numbers, and instead recommends that you apply the latest updates. There's no minimum update level specified at any given time, and if you are experiencing client problems with Microsoft 365 services the first step should always be to check for available updates and, if there are new updates available, apply them to your client installations to see if your problem is resolved.

Once a version of an Office suite or application enters extended support, it is no longer supported for use with the service. Microsoft says that it won't deliberately seek to block or prevent you from connecting with unsupported versions of Office applications, but no bug fixes or other updates will ever be applied to make them compatible with the service, and they always have the option to block specific older client versions from communicating. They also can, and have, deprecated protocols required for older clients, rendering them unable to connect.

## Understanding Office Update Channels

Microsoft 365 Apps updates are distributed in what Microsoft now calls *channels*. The channel model has evolved, with the most recent changes happening in May 2020 as part of the Microsoft 365 Apps rebranding. The [Microsoft website](#) now describes three channels, as shown in Table 15-1, and the support lifecycle for feature updates now varies between channels. To summarize, the new channels are:

- **Current Channel** releases features and bug fixes as soon as they're ready (but at least once per month), with security fixes released on the second Tuesday of each month ("Patch Tuesday") to conform with Microsoft's existing security release schedule.
- **Monthly Enterprise Channel** releases monthly updates on Patch Tuesday, along with bug and security fixes. Microsoft expects most enterprises to use this as the default.
- **Semi-Annual Enterprise Channel** releases updates twice a year (in January and July). Bug and security fixes are released monthly on Patch Tuesday.

This is a pretty dramatic change from the old model. In the new model, each channel releases fixes and security updates on the same day of the month, each month. The primary difference in the channels is how often feature updates appear.

<b>Channel</b>	<b>Feature Updates</b>	<b>Security Updates</b>	<b>Non-Security Updates and Fixes</b>	<b>Products that default to this channel</b>
Current Channel	When they're ready	Monthly (Patch Tuesday)	Monthly	Visio Pro Project Microsoft 365 Apps
Monthly Enterprise Channel	Monthly, on Patch Tuesday	Monthly (Patch Tuesday)	Monthly	Microsoft 365 Apps
Semi-Annual Enterprise Channel	Every six months, on Patch Tuesday in January and July	Monthly	Monthly	None

Table 15-1: Microsoft 365 Apps Branches and Channels

Whether you use the Office Deployment Tool, the Office Customization Tool, or install from the Microsoft 365 admin center, when you install the apps, they will default to the Current Channel unless you have chosen another default as described below. The Current Channel is suitable for businesses that are willing to have their clients updated to the newest features rapidly (but at least once per month). Typically, these are customers who do not have special macros or Office add-ins that are critical to their business processes. Historically, these automatic updates have not caused major, widespread issues, and can be considered safe to use if you don't have the type of customizations or integrations in your environment that might break after an update.

There are also preview versions of the Current and Semi-Annual Enterprise channels. The intent of these preview channels is to let you test upcoming releases before they hit the corresponding channels and go into wider distribution. For example, a new feature is first released to Current Channel (Preview), then to Current Channel. Once it meets Microsoft's criteria for a wider release, it will be pushed to the Monthly Enterprise Channel. At some point before the next Semi-Annual Enterprise Channel release, the feature may appear in Semi-Annual Enterprise Preview, so that organizations using the semi-annual channel can test it against their release processes. However, the progress of a feature through these channels may be slowed or interrupted if it fails to meet Microsoft's criteria for usability, functionality, or quality—so just because a feature appears in (say) the Monthly Enterprise Channel in May, there's no guarantee that it will appear in the July Semi-annual Enterprise Channel.

If you have devices configured to use the Monthly Enterprise channel, Microsoft allows you to roll them back to a previous build in the channel. You might want to do this if you find that a Monthly Enterprise change breaks some part of a line of business application, for example. You can also skip an upcoming release. You access both these options from the **Office installation options** page in the Microsoft 365 admin center.

To control which channel a Microsoft 365 Apps installation uses, you can set the channel during installation, or by using Group Policy. The Group Policy administrative templates for Office include an Update Channel setting in Computer Configuration. Setting the update channel using Group Policy has the advantage of allowing you to change clients to a different update channel than was used during the initial setup or allowing you to target different channels to different parts of your user population without designing a different installation configuration file for each of them.

**Note:** The update channel Group Policy setting is not available for Office 2013 clients. It also does not apply to Office 2016 or Office 2019 clients that were installed using the MSI package, only those installed using C2R.

Each update channel has a different period of support. For the Current Channel, support for a build (or release) is only applicable until the next build is released. If a security bug is found in the latest build of the

Current Channel, a patch will be released for the latest build, but no patches will be released for any previous builds in the Current Channel. This means that the Current Channel builds are normally supported for one month.

The Semi-Annual Enterprise Channel versions are released every six months, and each build is supported for 14 months. Semi-Annual Enterprise releases that ship in January are supported until March of the next year; the July releases are supported until September of the next year. This [list shows all the build numbers and release dates](#) for Microsoft 365 Apps going back to 2015.

Because Microsoft supports both the current and previous Semi-Annual Enterprise Channel versions, there is no immediate pressure to update when a new version is released to that channel as there is with the other channels. Microsoft recommends this channel only for “those select devices in your organization where extensive testing is needed before rolling out new Office features.” If a security bug is found in today's version of the Semi-Annual Enterprise Channel, a patch will be released for the latest version and the previous version, but no older builds.

In mid-April 2022, Microsoft reinforced this belief by announcing (in a little-noticed Message Center post, MC362760) that they were going to move clients from Semi-Annual Enterprise to monthly. If you didn't see this post, it may be because you didn't have any devices currently enrolled in that channel.

This update cadence and support period for Microsoft 365 Apps are important to keep in mind when you're planning your update strategy. If you decide to manually control updates, you need to ensure that you deploy new versions promptly and do not let your clients fall into an unsupported state that could put them at risk of a security vulnerability.

With multiple update channels to choose from, you also need to strike a balance between providing a stable version of the desktop applications for the bulk of your user population, while also ensuring that new builds are being tested on a sample of your user population before all users receive the updates. Microsoft recommends splitting up devices into two broad categories: one (which some Microsoft folks call “general purpose”) that doesn't normally have or use any applications, macros, add-ins, or other business-critical tools that cannot be allowed to break, and another (“business essential”) set of devices that run line of business applications, belong to key employees, or otherwise need to be protected. You should assign the general-purpose devices to update directly from Microsoft or Endpoint Configuration Manager using the Monthly Enterprise Channel, then keep the “business essential” devices on the Semi-Annual Enterprise Channel to reduce the number and scope of changes you need to adjust to at any given point in time.

**Other update channels you should know about:** Other teams at Microsoft, including the Windows, Exchange, and Teams product groups, have their release strategies which are broadly similar to how Office does releases... but there are enough differences that you should carefully examine them to ensure that you're getting the desired mix of update frequency and stability. In particular, understanding how the Windows Long-Term Servicing Branch (LTSB) works is critical for organizations that have strict change control requirements.

## Operating System and CPU Support

For the most part, Microsoft seems to believe that customers update their computers to the latest operating systems whenever a new version is released, and they act accordingly. In practice, there are [specific requirements](#) for the underlying OS for the full Office clients.

Windows 11, Windows 10, Windows 8.1, Windows Server 2022, Windows Server 2019, and Windows Server 2016 are all officially supported. There is no blanket requirement for a particular version of Windows 11 or

Windows 10, although at any time Microsoft might add dependencies that require a certain OS level for specific features.

For macOS, the three most recent major versions are supported. When a new major version of macOS is released, that major version of macOS and its two immediate predecessors remain supported. For example, now that Apple has released macOS 12.0 (“Monterey”), version 10.14 (“Mojave”) is no longer supported. Apple does not have the same predictable release schedule for its OS families that Microsoft does, so the gaps between major releases may be shorter or longer than you expect. Microsoft provides plenty of advance notice of their support requirements, though.

As a further complication, occasionally Microsoft will draw a new support boundary between a version of the host OS and the applications. For example, in January 2021, Microsoft announced that version 16.43 of the macOS apps would be the last supported version on macOS 10.13 and earlier. That is if you have an older Mac that can’t run 10.14 or later, you won’t be able to update the Office apps themselves past version 16.43. In the same vein, in August 2021 they announced the forthcoming end of support for iOS 13 and earlier in the Teams Mobile client. Apple in general does a superb job of supporting older hardware (for example, my 2011 MacBook Pro shipped with macOS 10.6 and was supported until the release of macOS 10.14) but all good things must eventually end.

Microsoft has released versions of Word, Outlook, PowerPoint, OneDrive, and Excel for macOS that are recompiled for the [Apple Silicon](#) CPU architecture. Although no timeline has been announced for Apple Silicon-native versions of most other macOS applications (including Teams and To-Do), reports are that these applications run quite well using the M1 Rosetta translator included in macOS 11.0 and later. Microsoft hasn’t said anything publicly about native versions of other applications yet.

## Faster Updates Through the Office Insider Program

Microsoft has three programs for those who like to add excitement to their lives by running early, and possibly unstable, versions of software. You’ve probably run beta versions of software for test purposes in the past; the Windows Insider and [Office Insider](#) programs offer a very similar experience. When you join one of these programs, you’ll have access to builds of Windows or Microsoft 365 Apps before they are generally released. When you enroll in the Office Insider program, you’ll get builds in one of two additional channels: Beta (formerly known as “Insider Fast”) and Current Channel (Preview) (formerly “Monthly Channel (Targeted)”). These channels [generate releases](#) roughly weekly, which means you’ll get earlier access to new features but also that you may run into problems with features that don’t quite work properly. In exchange for granting early access to Insider members, Microsoft wants to collect user feedback and bug reports, and they maintain a set of Insider forums for that purpose.

Interestingly, Office Insider exists both for macOS and Windows desktops; there isn’t an equivalent for the mobile or web-based Office applications, although Microsoft does occasionally conduct open beta testing for the Office, Word, Excel, PowerPoint, Outlook mobile, Whiteboard, and Microsoft To Do client on iOS and Android.

What about the third program? Microsoft runs a [separate preview program and channel set for Teams](#). As with Office Insider, the Teams Public Preview program features three channels: Beta, Private Preview, and Public Preview. The details of these channels, how to join them, and how to manage your Teams Insider membership are covered later in this chapter.



## Installing Microsoft 365 Apps

Microsoft 365 Apps is packaged as a “Click-to-Run” application, using application streaming and virtualization technologies to reduce the amount of time between beginning the installation of the software and being able to start using the applications. Microsoft 365 Apps is included in the Microsoft 365 Business Standard, Office 365 Enterprise E3 and E5 licenses, as well as a standalone Microsoft 365 Apps license (this used to be called Office ProPlus).

### Users May Self-Install the Apps

Microsoft 365 users who have a license for Microsoft 365 Apps can install the desktop applications by logging into the Microsoft 365 portal, navigating to their user profile, and using the **Manage** link in the **Office apps** section.

For macOS users, Microsoft 365 Apps are also available from Apple’s Mac App Store. This might seem like a pointless offering, given that licensed users can download the apps from the portal, but it allows organizations using [Apple Business Manager](#) to centralize and manage Office deployments. Just downloading the apps from the Mac App Store gives you read-only access to documents and email; if you want to create or modify documents, or send or receive emails with Outlook, you’ll have to activate the apps by signing into the service using an account that has an appropriate license.

On both Windows and macOS, installing or updating the Microsoft 365 Apps package will get you the Teams client too. Microsoft considers Teams to be a full-fledged peer of the other desktop applications and treats it accordingly. If you don’t want the Teams client installed, your options are limited. You can block that installation by building a custom configuration using the Office Deployment Tool (as [described here](#)), or you can use Group Policy, or you can allow the installation but use Group Policy to [stop the client from starting](#) when the user signs in. There’s no obvious way to stop this behavior in the user interface, but you can control it via a registry or [Group Policy setting](#).

### Controlling User Software Installs

You can disable software downloads from the portal for users, which is often preferred by organizations who are using existing software licensing for Office clients, or who want to fully manage the installation process and not allow users to install the software at all.

1. Log into the [Microsoft 365 admin center](#) with your administrator account.
2. Navigate to **Settings**, then **Org settings**, then click the **Services** pivot.
3. Select **Office installation options**.
4. On the **Installation** pivot, choose the user software options you want.
5. Click **Save** to apply the changes.

When users log into the Office portal with their normal user accounts, they’ll see installation options based on your selections above. Users must hold local administrator rights on the computer where they install the applications, which is not likely to be an issue in a BYOD environment but may present a challenge for organizations that do not grant local administrator rights to end-users.

**Real World:** Although your corporate-owned computers may prevent users from installing Microsoft 365 Apps themselves, each license entitles the user to install on up to 5 computers, so if they are allowed to download the software at all, they will be able to log into the admin center from home and install the Microsoft 365 Apps software there.

## Managing Updates for Click-to-Run Builds of Microsoft 365 Apps

Updates to Microsoft 365 Apps can be managed in different ways to suit the needs of the organization. However, Microsoft recommends that you let them do the management by allowing all clients to individually download updates from the [Office 365 content distribution network](#) (CDN). Microsoft maintains multiple CDNs (including one that's used purely for caching web libraries used in Microsoft web applications), but the Office 365 CDN is unique in that it is used only to push application updates for the Microsoft 365 Apps suite and other Office 365 components. Using the Office 365 CDN allows the network and client to intelligently negotiate exactly which updates are required and transmit them as efficiently as possible, as described in [this Microsoft article](#).

Users who install Microsoft 365 Apps from the Office portal will find that their clients' updates are automatically downloaded from the internet and installed as they are released by Microsoft, and no other action is required by the end-user other than restarting applications when prompted that an update is ready.

If you've deployed the applications using ODT, then the XML configuration file you used will determine how updates are applied by enabling updates and configuring an update path. If you want your clients to automatically update themselves using the Microsoft 365 Apps update behavior, you'll have to re-run the Office Deployment Tool for Click-to-Run to download the latest build of Microsoft 365 Apps to the appropriate location on the network. Once the updated build is available in the update path your clients were originally configured with, the computers will automatically apply the updates.

Many enterprises prefer to use Microsoft Endpoint Configuration Manager, System Center Configuration Manager, [Quest KACE](#), or other similar tools to provide them better control over the operating system and application updates. When Microsoft releases updates to the Microsoft 365 Apps suite, they also release installer packages that can be deployed using these types of tools. Keep in mind that if you're using the Current Channel, Microsoft will be releasing updates at least once per month... and they only support the current version of each application against the service, meaning that if you're managing updates yourself you may wish to force your clients to use the Monthly or Semi-Annual Enterprise channels.

If you need to disable automatic updates completely, the ODT XML configuration file you use must specify the exact build number of Microsoft 365 Apps that you want to install. If automatic updates are disabled in your XML file, any new builds that you download using the Office Deployment Tool for Click-to-Run will need to be manually deployed to end-user computers.

### Using Delivery Optimization with Click-To-Run

Peer-to-peer update delivery can reduce the amount of time required to deploy updates throughout a large organization. Apple and Microsoft have used this technology for some time for operating system updates. The basic idea is that one client on a network downloads the updates, then caches and redistributes them to other nearby computers on the same network. Compared to having every computer download updates from a CDN over the internet, peer-to-peer delivery (which Microsoft calls *Delivery Optimization*) promises to speed deployment significantly when implemented properly.

The good news is that, by default, Windows 10 computers automatically use Delivery Optimization for OS updates, and when you install Microsoft 365 Apps, so will those applications *if* the prerequisite requirements are met:

- The client device must be running build 1709 or later of Windows 10 Enterprise or Windows 10 Education to enable Delivery Optimization updates for Microsoft 365 Apps when they're pushed by Microsoft.

- If you want users to be able to request updates (by going to the application backstage and using **Account > Update Options > Update Now**), the client device must be running build 1908 or later of Windows 10.
- You must have version 1912 or later of Microsoft 365 Apps. That version was released in January 2020, so you should certainly already have it.

Keep in mind that Delivery Optimization only works with computers that are configured to download their application updates. If you're using Microsoft Endpoint Configuration Manager or another distribution-management tool, the client applications won't use it.

If you want to check whether an individual Windows 10 device is using Delivery Optimization, you can use the *Get-DeliveryOptimizationStatus* and *Get-DeliveryOptimizationLog* PowerShell cmdlets. Microsoft has lots more [documentation of the Delivery Optimization feature](#) available if you're interested but, in general, most Microsoft 365 administrators can ignore the feature and let it work silently in the background.

## Managing Updates for MSI Builds of Microsoft 365 Apps

The update mechanisms and Channels described so far apply only to the C2R versions of Office. Licensed Microsoft 365 customers also get access to the traditional MSI packages for Office client deployment. If you use these MSI packages, you'll need to manage updates yourself using tools such as Microsoft Update, Windows Server Update Service (WSUS), System Center Configuration Manager (SCCM), Quest KACE, or the Microsoft Endpoint Configuration Manager.

For C2R clients, all update types are released at the same time on the second Tuesday of each month, also known as "Patch Tuesday" among IT pros. The MSI clients are handled differently. Non-security updates are released on the first Tuesday of each month, while security updates remain on the second Tuesday of each month.

MSI deployments are still widely used by customers for a variety of reasons, but update management is more complex and time-consuming because you must do it all yourself with whatever deployment tool you're using. As with so many other aspects of cloud services, if you want a higher degree of control, you can get it, but at the cost of increased overhead and inconvenience. For your client deployments, consider the C2R service as your first choice, and only revert to the MSI package if you have a good reason to. It will save you time and effort overall that can be better spent elsewhere.

## Managing Updates for macOS Versions of Microsoft 365 Apps

You can get the macOS versions of the Microsoft 365 Apps suite by downloading them from the Office 365 portal, or Apple's Mac App Store. In both cases, you'll find that Microsoft has an app update mechanism. The Microsoft AutoUpdate (MAU) app is automatically installed when you install any of the Microsoft 365 Apps family; it runs periodically to download and install updates to the component applications, or you can trigger it manually by choosing **Help > Check for Updates** from the apps. The MAU app uses the same Office 365 CDN that Windows does, just with a different set of update bits. Interestingly, Teams client updates are now delivered through MAU.

## Edge WebView2 and Office Apps

Many parts of the desktop Office applications create or display web-compatible content. Some of the features you use in Outlook, for example, aren't built into the Outlook desktop client, but are instead loaded from a different Microsoft service and rendered inside Outlook. Examples include the Meeting Insights view and the

room finder view. These components are called OWA experiences (OCX) and use an Edge component called WebView2. The One Outlook “Monarch” client also uses WebView2, as will the Teams 2.0 client.

Microsoft loads the WebView2 component onto PCs through:

- The Edge browser.
- The Microsoft 365 enterprise apps.
- Windows 11 updates.
- Windows 10 updates (starting June 2022).

After deployment, a program called Microsoft Edge WebView2 Runtime (`msedgeview2runtime.exe`) is available on a PC. You can manually install the runtime in advance if you want to, or you can prevent it from being installed by changing a setting in the Apps admin center (both operations are [described here](#)).

## Troubleshooting Office Clients with SaRA

Although Microsoft 365 is a very reliable service, there will naturally be problems that occur from time to time. Part of that is due to the scale and complexity of the service; it is impossible to keep that much software and hardware in 100% healthy condition all the time. Part of it is also due to factors outside of Microsoft’s direct control. A large proportion of the support calls that Microsoft receives are client issues that are due to the software on the computer, network issues, DNS issues, or other factors that can be difficult to identify quickly. The more time Microsoft spends supporting these types of issues, the costlier it is to run the service.

That’s why Microsoft has released the [Support and Recovery Assistant](#) (SaRA). This downloadable troubleshooting tool has a wizard-driven interface to allow users to identify the issue they are experiencing, and then allow SaRA to perform diagnostic tests and suggest solutions to the most common client-side issues. SaRA can check for issues around mail flow, access to shared mailboxes and calendars, Outlook freezes, and repeated authentication prompts. Although it doesn’t exhaustively test for every possible cause, it’s still a useful tool. In May 2021, Microsoft added a [new command-line version of SaRA](#), which is useful for enterprises that want to allow technicians to remotely diagnose Office problems.

SaRA requires administrative rights on the computer to be able to install the software and run the tests. For BYOD users or environments where local admin rights are given to end-users, this won’t be a problem. However, you can also imagine that the existence of SaRA may not be known to end-users who do not have an IT focus. They’re likely to call Microsoft for support anyway, but if they were to try to raise a support ticket using the admin center SaRA can be suggested to them automatically as part of that process. For customers with IT support staff, it is a simple matter of the IT personnel installing the tool on client computers and running the diagnostic tests.

It’s easy to think that Microsoft should just provide working support tools like this to customers. And indeed, with all the data they have from telemetry and from analyzing support trends, Microsoft can make reasonable guesses about the most useful diagnostics to develop for SaRA. That said, it’s important that customers and IT professionals provide feedback to drive the development of tools like SaRA, by letting Microsoft know when the tools don’t have any information about our problems, or when they fail to detect and suggest fixes for issues. To that end, SaRA has multiple places where feedback can be submitted to Microsoft.

## Controlling “What’s New” Content Displayed to Users

The Microsoft 365 desktop apps can display pop-up notifications to users to tell them what’s new in the product. To see what this content looks like, go to the Help item on the Office ribbon and choose the “What’s New” icon; you’ll see the content appear as a pop-up navigation bar on the right edge of the window. There are two ways to look at this. One is that it’s a useful and user-friendly way to notify users of valuable new

features. Another is that it's an unwanted intrusion. Since individual end-users could turn this feature off, which view you took was a matter of individual preference.

However, since mid-2019, Microsoft has been much more aggressive about connecting directly with enterprise users—not the admins and other IT people involved in managing and supporting tenants, but the end-users who are using the service to get their work done. While the impulse behind it may be commendable, Microsoft's continual insistence that they should be able to talk to (or, more properly, *talk at*) these end-users has ruffled a lot of feathers among their proper customers. However, they are trying to strike the right balance between making sure users are aware of new features and respecting the rights of companies and their employees. One of the latest changes in this balancing act is a new set of controls that allow administrators to control what users see in the "What's new" sections of the desktop clients. In early 2020, Microsoft introduced a feature in the Microsoft 365 admin center that allows you to hide specific cards for new features. To control which items users see, go to the admin center and select **Settings** -> **Org settings** tab, then select **What's new in Office** and then you can hide or show specific feature cards for individual features in individual releases. Each Office release has its section; most sections have only two or three new feature cards.

## Using the Microsoft Apps Admin Center

Traditionally, there have been several ways of enforcing Office client configurations: you could build a custom installation of the MSI version, you could use the Office Configuration Tool to specify a configuration applied with the installer, or you could specify a configuration template and push it to clients using Group Policy or another similar configuration management tool. This latter method is probably the most common approach, but it has a significant limitation: the target machines must be able to apply Group Policy settings. For devices that are only joined to Azure Active Directory, or for BYOD-style devices that aren't joined to any directory, there was no good way to enforce or change client policy settings. As an alternative, Microsoft offers a suite of cloud-based tools for managing Microsoft 365 Apps clients, the Microsoft Apps Admin Center (<https://config.office.com>). You can use this admin center to manage several aspects of the Microsoft 365 Apps deployed to your end-users without MEM, Office 365 MDM, or Group Policy objects. Currently the Apps Admin Center is available worldwide, except for Office 365 GCC, GCC High, Office 365 DoD, and Office 365 operated by 21Vianet.

The combination of desktop applications delivered by a cloud-based content delivery network and then managed by a cloud-based service is interesting and offers several potential benefits to customers. For now, the Apps Admin Center contains six items in its left navigation bar (note that some of these services are marked as being in preview):

- **Servicing** allows you to create and manage servicing profiles. Earlier in the chapter, I mentioned Microsoft's recommendation for creating "business essential" and "general purpose" device sets—you can follow this recommendation by creating servicing profiles, assigning devices to them, and then assigning update channels to them. You can define profiles that exclude devices based on the amount of free disk space they have, whether Office application macros have been used in the preceding 28 days, what channel they're currently in, and whether they have Office add-ins installed. Each profile can have exclusion dates, during which devices in the profile won't receive any updates.
- **Customization** contains tools that allow you to define and deploy custom configurations and policies for the Office apps.
- **Health** and **Inventory** allow you to monitor which applications are installed, whether they are healthy, and whether they are receiving updates as you expect—read on to learn more.

- **Learn More** is a collection of links to documentation, videos, and so on that may be useful as you learn to manage your Microsoft 365 apps using the Apps Admin Center.
- **Settings** has very little in it; it contains a control for generating what Microsoft calls a tenant association key (a base-64-encoded key that lets you tie devices in a specific tenant to your portal) and a slider that allows you to adjust how long the tool will remember devices that haven't updated their inventory data (the default is 30 days).

## Customizing Microsoft 365 App Configurations

Some organizations are happy to allow users to install local copies of the Microsoft 365 Apps packages. Others want or need centralized control or have constraints that make it impossible or impractical to install using the normal process of launching the small bootstrap installer and letting it download the required software over the Internet. Fortunately, the setup files can be downloaded in advance and then distributed through a network share or other means to avoid a dependency on the Internet connection. In addition, you can customize exactly what parts of the Microsoft 365 Apps suite are downloaded, where they are installed, how they're configured, and several other settings. To do this, you'll use two tools:

- The [Office Deployment Tool](#) (ODT) is a command-line tool that you run to apply a configuration file (formerly known as a transform) to the installation files. The result: the settings specified in the *configuration.xml* file drive the installation.
- The [Office Customization Tool](#) (OCT) is now part of the Apps Admin Center, under the **Customization** section in the left nav bar. The OCT provides a web-based tool for creating and modifying configuration files to be used with the ODT. You can also create a configuration file and then apply it to existing installations (as you might wish to do if you don't want the Teams client to auto-start).

You can find a full reference for how to customize the XML configuration file on the [Microsoft Support website](#). However, Microsoft recommends that you use the OCT to avoid making mistakes in XML formatting that will cause your installation to fail. In addition, OCT allows you to upload and store configuration files in the cloud, making it easier to keep track of them.

Normally you'll use these tools by following a process like this:

1. Install the ODT. That will give you two files: *setup.exe* and a sample *configuration.xml* file.
2. Navigate to <https://config.office.com> and either create a new device configuration from scratch or modify an existing configuration. It's a good idea to create a new configuration every so often to verify that you are taking advantage of all available settings in the OCT; Microsoft occasionally adds new features that may not be included in your older configuration files.
3. Test the configuration file to verify that it does what you expect, and that the resulting installation is set up the way you want it.
4. Use OCT to adjust the configuration file as required.
5. Once you're happy with the results from your configuration file, use the ODT to complete the installation.

You can run the ODT to complete the installation using any tool that can execute a command-line program. This includes almost every endpoint-management tool out there, plus Group Policy, plus even having users run the command themselves from a desktop shortcut you've previously deployed.

**Real-world:** If your environment doesn't provide good Internet connectivity, you might want to host the installation files on your internal network. You can do so using a conventional network share, a DFS share, or even removable media. To install in this manner, you'll need to run the ODT with the */download* switch to download the files. To do this, after customizing the XML configuration file, copy it to the same folder

that you placed the Setup.exe file in. Open a cmd.exe prompt on a computer on the network and run *Setup.exe* with the */download* switch to download the Microsoft 365 Apps source files.

```
C:\> start /wait \\obcdc1\installs\Office365ProPlus\setup.exe /download \\obcdc1\installs\Office365ProPlus\configuration.xml
```

The source files will be downloaded and placed in a folder structure automatically. Do not change or rename the folder structure created by Setup. After the download is completed you can deploy the apps from the install point by running *Setup.exe* with the */configure* switch. You can also run this command line using software deployment tools such as Group Policy or Microsoft Endpoint Configuration Manager.

Pay close attention to the “Product ID” that you configure in your XML file. The correct product ID must be used for the Microsoft 365 license assigned to the user. If the wrong product ID is installed on the computer, the user will not be able to activate it with their Microsoft 365 credentials and will need to reinstall the correct version. You will normally use a value of O365ProPlusRetail for most product SKUs (except Visio and Project, which have unique product IDs). Microsoft publishes a complete list [of product IDs matched to license types](#).

If you do not specify an update channel in your configuration file, and none is specified via Group Policy, the default branch or channel will be used (refer to the earlier section about Update Channels).

[A tutorial and sample PowerShell script](#) are available to help in deploying Microsoft 365 Apps using GPOs.

**Real World:** If you’ve already purchased licenses and provisioned Microsoft 365 accounts, you can begin your client deployments before you start moving mailboxes and other data to the cloud. The client applications can connect to on-premises servers such as Exchange, within the limits of the normal client system requirements for those on-premises products. Outlook is generally smart enough to automatically deal with mailboxes that have been moved from on-premises servers to the cloud without user intervention.

## Creating and Managing Application Policies

The Apps Admin Center allows you to create policies that are applied from the cloud to desktop Office applications (**Customization > Policy Management**). In its current release, it doesn’t support all the policy settings available through Group Policy templates, although it does have a *lot* of settings (as of July 2022, for example, there are 287 individual policy settings just for Microsoft Word and 2,182 policy settings in total)! Perhaps more importantly, the Apps Admin Center only allows you to set up user-based policies, whereas with GPOs you can apply Office policies to machines.

**Note:** The only Teams-related policies currently implemented in the Apps Admin Center allow you to prevent the Teams client from starting automatically on Windows machines or to restrict which domains Teams may sign into. If you want to control other aspects of Teams client behavior, you’ll have to use the Teams policies described in Chapter 13.

Policies are applied to the applications in a way very similar to how Group Policy settings are applied. When a user signs into Office 365 on a device, the application she signs into immediately checks the policy service for any applicable policies. If no applicable policies are found, the application checks again every 24 hours; any time a policy is found and applied, the application will recheck for policy changes every 90 minutes while the app is open. Each time a Microsoft 365 App is launched, it will check for policies (this includes Visio and Project, but not Power BI or To Do). It is important to understand that policy changes are not applied until the app is *next* restarted though—so if a user launches Outlook at 8 am, and you update the policy at noon, the new policy will be downloaded but not applied to that user until the next time she launches an Office application.

As one concrete example of a policy that you can manage with the Apps Admin Center, consider the user feedback mechanism described in Chapter 5. There are “feedback” and “survey” policy categories ([described here](#)) that let you control whether users can submit feedback, including free-text comments and surveys, to Microsoft, and (separately) whether Microsoft is allowed to follow up with individual users on their feedback.

The Policy configurations page in the Apps Admin Center shows which policies you have defined, their relative priority, and their health. Microsoft is deprecating the policy health status feature starting in mid-March 2022; per MC335282, they will eventually replace it with “advanced health reporting and compliance monitoring features” but there is no date associated with this.

To use the service, you need version 1808 or later of the Microsoft 365 Apps suite, and the users you want to target must all have accounts that are either homed in or synchronized with Azure AD. You assign policies to users based on their group memberships, so you’ll need to create groups for whichever sets of users you want to assign policies to. You can also create a policy that applies to anonymous users who read or edit documents using sharing links sent by users in your tenant.

## Monitoring Application Health

The **Health** section of the Apps Admin Center contains four sections:

- The **Apps Health** section is meant to show you data about individual machines’ applications and how healthy they are. You’ll only see data for users who have an E3, E5, or equivalent license, and only if they are running version 2008 or later of Microsoft 365 Apps for Windows. The **Add-in Health** pivot shows the health of COM and VSTO add-ins for Office, whether those are internally developed by your organization or obtained from public sources. Note that this pivot will only appear if you’ve enabled inventory collection.
- **Security Update Status** (currently in preview) shows you which machines are behind on security updates for the Microsoft 365 apps (not general Windows Update updates, sadly). This is based on the release date of the most recent Microsoft 365 apps security update; as of this writing, the most recent update was June 14, 2022, so any device that doesn’t have that update will be marked as “not up to date”. You can set goals for the percentage of devices you want to be up to date or the average time between patch release and deployment; these are useful to track.
- **OneDrive Sync** (currently in preview) shows reports and data about the sync health of OneDrive for Business clients, including how (or whether) the [known folders sync feature](#) is enabled, and how many clients have reported sync errors, and an aggregated view of issues by type.
- **Service Health** shows you whether the Microsoft 365 Apps portion of the Apps Admin Center is itself healthy.

There’s a lot of potential for Microsoft to improve and extend the features of this section of the Apps Admin Center, including linking “who’s not up to date” lists with the policy tools, or a notification mechanism, to help automatically notify users and/or bring their machines into compliance.

## Inventorying Your Application Estate

Keeping track of which devices have which versions of the Microsoft 365 apps on them is an enduring hassle for many Microsoft 365 customers. This hassle is magnified if you don’t have a solution such as Microsoft Endpoint Configuration Manager running; as individual machines can freely self-update, you can quickly end up with a confusing mess of different versions spread across your tenant. The **Inventory** section of the Apps Admin Center is intended to help address this by summarizing which devices, which builds, and which add-ins are present in your tenant. There’s a new client-side component of Office known as the Serviceability Manager (SM) that gathers this data and provides it to the service. SM only starts collecting this data after you do two



things. First, you must sign into the Apps admin center, and second, you must accept the onboarding prompt. Once you've completed both of those tasks, SM will start gathering data and passing it to the service and you'll see it appear.

Each of the sections in the inventory report (Figure 15-1) has links that allow you to see more detailed data. For example, the "Show all devices" link in the Devices section will show whether each device has macros or add-ins, what Office version and build it's running, what update channel it is on, and so on.

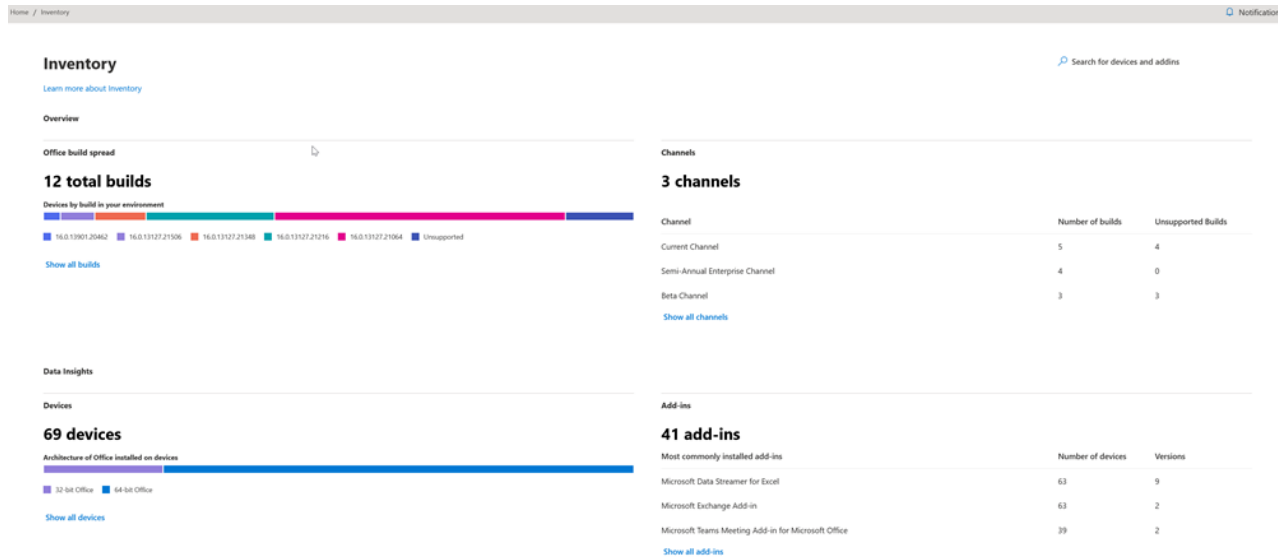


Figure 15-1: An example application inventory

## Managing the Outlook Client Family

Outlook is still the most common desktop email client used to access both on-premises and cloud Exchange mailboxes. It's also arguably the most important of the clients, and it's fair to say it has the broadest set of management features (in part because those management features are spread across several parts of the service). Microsoft is trying to create a unified user experience and identity for Outlook across multiple platforms, which can lead to some confusion when discussing "Outlook" in books like this. (In fact, they go so far as to [say that](#) "One of the Outlook design principles is to make Outlook feel native to your preferred operating environment"). Despite that confusion, we'll tackle managing the Outlook client family as a single topic. The family currently includes several members:

- Desktop Outlook, available for both macOS and Windows. There are [some differences in the functionality](#) between the two clients, and of course, there are differences between different versions, but in general, you can consider the two as equivalent. Both versions support multiple Exchange Online accounts in a single client session; both can work online or offline, and both offer extensive caching features.
- The new "One Outlook" client (codenamed "Monarch") is an ambitious effort to replace the Windows, web, and macOS clients with a single unified codebase.
- Outlook Web App (OWA) is the webmail interface for Exchange Online, providing users access to their mailbox, calendar, contacts, and tasks, as well as a launch pad for other services such as the Office Online apps, OneDrive for Business, SharePoint sites, Delve, and Stream. Although Microsoft officially refers to "Outlook on the web," this is a clunky name, so we won't use it, especially as most documentation for the product still refers to "OWA".

- Outlook mobile for iOS and Android is Microsoft's first-party mobile client for Exchange Online. It adapts the Outlook user experience to the smaller screens and touch-based interface of mobile phones and tablets, and it includes some mobile-only features such as Play My Emails (a voice-based system that will read your emails to you and let you respond via a connected Bluetooth or wired headset). As with the macOS-vs-Windows split, you'll find that features come to the Android and iOS versions of Outlook at different speeds.
- An [Outlook extension for the Edge browser](#). This extension shows up in your browser toolbar when installed and lets you do some simple tasks without having to open a separate window or tab to load OWA. Reviews are mixed on both the utility and stability of the extension so far, and Microsoft had to back away from its initial plans to advertise it to users (at least users who have enterprise licenses for Microsoft 365 apps).

**Real-world:** Outlook for Windows performs best when you configure it to use cached mode to connect to Exchange Online. In this configuration, you can decide how much of the server copy of the mailbox Outlook should keep locally on the user's computer for faster access. The local cache is known as the OST file and is on the hard drive of the user's computer in a default folder path within their local profile. You can use a slider control in the Outlook account settings to manage how much of the mailbox Outlook synchronizes to the OST file so that you can balance the availability of data and storage consumption. For instance, many users find that keeping the last year's email in the OST is the right balance. Outlook for Mac uses a similar approach to cache mailbox contents on the local drive; it uses a file very similar in concept to the OST file, but with a different structure. Unfortunately, there's no equivalent of the Outlook for Windows slider control, so you can't regulate how much email is synchronized with the client. This is a long-standing, and popular, feature request that Microsoft will hopefully get to sometime soon.

Features come to the three main flavors of Outlook at different times and in an unpredictable order. For example, Microsoft first developed and released "dark mode" for OWA, then rolled it out to macOS users, then to Outlook for Windows customers who were on Monthly Channel (Targeted)... after which the Outlook mobile team released their dark mode implementation. While there are certainly exceptions, in this case, it's fair to say that *most* new features come to OWA first. The individual versions of Outlook have a pretty high degree of independence for feature releases; in general, it seems that new features that require changes to the back-end services in the substrate usually appear first in OWA but that the mobile and desktop versions freely introduce UI-only changes on a more frequent schedule. Overall, the principle Microsoft's trying to apply is that the user experience for its multiplatform applications should be as consistent as possible whilst still fitting into the design language of the host platform—but those underlying implementation choices are to be minimized in favor of giving users a consistent *Microsoft-centric* experience on whatever device they're using. Common UI elements, affordances, color schemes, and so on are the vehicle Microsoft uses to deliver this consistency—if you look at the UI of OWA on Edge, Outlook for Windows, Outlook for macOS, and Outlook on iOS side-by-side, you'll notice the consistency.

Of course, there are always exceptions to this feature flow—starting in October 2021, Microsoft added a toggle labeled "Coming Soon" to Outlook for Windows that allows users to opt into seeing some prominent changes to the Outlook UI, including moving modules to a small left-side navigation rail instead of the bottom of the folder pane. If you want to turn this specific feature off, you can but it [requires a registry change](#).

When it comes to managing Outlook, it's important to differentiate between managing settings on the client (such as whether certain features are enabled or disabled on the client) versus managing them at the tenant level versus managing them at the mailbox or protocol level. For example, you might (and probably should) disable IMAP4 and POP3 in your tenant but that wouldn't necessarily stop an individual user from connecting Outlook using those protocols to another mailbox in another tenant (or another service). Several of the

newest features added to Outlook in late 2020, for example, are integrations with Cortana-powered services for things such as word suggestions. You may be able to control those features at the tenant level but there is no separate policy or setting that allows you to control the features' availability in Outlook itself.

## Outlook for Windows

The C2R version of Outlook for Windows is almost always delivered as part of a Microsoft 365 Apps installation. There are several ways to manage its behavior:

- Policies defined in an Office administrative template and distributed via Group Policy to domain-joined clients.
- Policies defined by the Office Client Policy Service and distributed when the client connects to the service.
- Local customizations made through the Windows registry.
- Local customizations made directly through the Outlook settings interface. These changes typically apply to the user's profile and thus only impact the specific user for whom the customizations are applied.

For the most part, the first three of these management techniques have an identical effect. Which one you use will depend on the infrastructure you deploy. User-specific customizations are a bit tricky, though, because they are usually chosen by the user herself, and users can be *very* protective of their customizations.

## Using Outlook for Windows in a Virtualized Desktop Environment

In virtual desktop environments, such as Citrix XenApp or XenDesktop, where multiple users log onto published applications or desktops hosted on a farm of servers, some issues exist with Outlook when it runs in cached mode:

- The OST files for multiple users accumulate on the hard drive of the server and can quickly consume all the available disk space.
- If a user logs on to a new session on a different server than their last session, the OST file may need to go through a major resynchronization, or even recreate its copy of the mailbox from scratch, which causes both a poor user experience as well as a server performance issue that can affect all other users on the same server.

Similarly, in virtual desktop environments that use non-persistent storage, Outlook needs to rebuild the OST file from scratch each time a user opens Outlook to start a new session.

Microsoft has somewhat belatedly realized that customers want to use Office 365 in environments that are built around virtualized desktops, and they have been making steady improvements to support those environments. For example, in 2019 they modified the default sync behavior for Outlook cached mode to sync the Inbox folder before the Calendar folder, and to reduce the number of folders that were synced (or resynced) by default. As an administrator, you also can tweak Outlook sync settings using the [Office 2019 Group Policy administrative templates](#) (ADMX files) to improve the sync experience in these environments:

- **Sync Settings** – this setting allows you to specify the amount of data (by age) that Outlook should cache in the OST file. For example, you could set this to 1 month for users logging on to the virtual desktop environment and then set it to 12 months for users who typically work from the same desktop or laptop each day. Users who frequently work offline in remote locations may even prefer *all* their email to be cached always.

- **Fast Access** – this setting allows Outlook to connect to Exchange Online in online mode while building the OST file at the same time, and then switches to cached mode when there is enough data available.
- **Cache File** – this setting can be used to specify a network location for storing the OST file so that it can be stored on persistent storage and accessed from any virtual desktop that the user is logging on to. Microsoft has [published guidance](#) and some specific requirements covering this configuration scenario.

There are also third-party solutions, such as [Liquidware's ProfileUnity](#) or [FsLogix Office Container](#), that can help Outlook cached mode perform well in VDI environments. (Technically, FsLogix isn't a third-party solution since it comes from Microsoft but, as it must be deployed separately from Office, I'll call it that).

**Real World:** Microsoft offers several interfaces that developers can use to extend the various Outlook clients. For the most part, there are three categories of plugins: those which work only in Windows Outlook, those which work only in Mac Outlook, and those which are written using the Office extension APIs and can be loaded in both versions of desktop Outlook, OWA, and even Outlook Mobile. For example, the familiar button that allows you to make a meeting into a Teams or Skype meeting with a single click is an "add-in" (the preferred term for desktop-native extensions to Outlook), while the FindTime tool is written with the Office APIs and can thus be used in multiple clients once it's enabled for the user.

## One Outlook/Monarch

We already have a single Outlook client for macOS, Linux, and Windows—it's called OWA. But, for a variety of reasons, not every user wants to, or can, use a web-based client for their everyday work. Microsoft has recognized this, but at the same time they have the quite reasonable desire to consolidate on the smallest possible number of unique client versions. Announced at Ignite 2020, "One Outlook" promises to satisfy both users who want a desktop client and Microsoft's desire for consolidation. For now, the Monarch client is only [available for Windows](#), and even then only to users who are in the Office Insider Beta channel. The current version is missing lots of features that both OWA and the desktop clients have. For example, there's no unified inbox, no support for PST files, and only limited support for offline access. Over time, Microsoft's goal is to reach feature parity with the existing versions, but they aren't quite there yet.

The earlier discussion in this chapter about WebView2 is important for Monarch because some parts of its functionality are enabled by a technology known as OPX ("OWA Powered Experiences"). Just as desktop Outlook embeds the OWA room finder in a separate subpane instead of duplicating its functionality, OPX allows Monarch to consume features originally developed for OWA. That's one big reason for the list of things that Monarch supports or doesn't—it's easier for Microsoft to plug in OWA features using OPX, so those features, in general, were the first ones delivered.

## Outlook for iOS and Android

For many years, Microsoft pursued a terrifically successful strategy of licensing Exchange ActiveSync (EAS) as widely as possible to third parties to enable their devices to connect to Exchange. EAS proved stable and both the cloud and on-premises versions of Exchange include basic security and management capabilities for mobile devices. The strategy of licensing ActiveSync to all and sundry worked in terms of making EAS the de facto protocol for Exchange mobile connectivity. The downside was that Microsoft ceded control over the user experience to the device vendors. Companies like Apple and Samsung incorporate EAS into the mail apps running on their devices while exerting absolute control over which features the clients expose to end-users. Even Microsoft's clients do not use the entire EAS protocol. Exchange ActiveSync was always intended to do

two different things: synchronize mail items between the device and a mailbox and provide mobile device management (including remote device wipe, PIN enforcement, and so on). Despite the popularity of EAS, there was never a real advantage for device vendors to implement all the functionality in the EAS protocol, instead of just whatever subset they felt was useful. A good-enough job allows a user to connect to Exchange and access their mailbox. No added value exists for Apple or Samsung or any other EAS licensee to extend their mail apps to offer more functionality to Exchange users than they do when the apps connect to Gmail or another email server, although some vendors grudgingly added support for a subset of EAS device management features. The native mail app is “good enough” if you can use it to process messages, meetings, tasks, and contacts. This attitude has existed for years and has resulted in situations like the famous series of bugs in the Apple iOS mail and calendar apps that caused excessive transaction log growth on Exchange servers or calendar appointment “hijacking”. The net result is that EAS is now the lowest common denominator protocol for mobile devices to connect to Exchange. The need to do better and to have control over functionality available through mobile clients drove Microsoft to come up with a new two-pronged strategy: first, they bought a company called Accompli and rebranded their clients as Outlook Mobile for iOS and Android, then they invested in improving device management by building mobile device management tools into Microsoft 365 and then releasing Microsoft Intune.

**Licensing Outlook for iOS and Android:** Not all plans include a license for the Outlook mobile clients. [Microsoft's FAQ](#) makes it clear that you need a suitable commercial license. “Suitable” really means that the license grants access to a desktop version of Outlook; Business Basic, Business Standard, and Enterprise E3 or E5 all count, as do or the corresponding versions of those plans for Education or Government. Exchange Online plans or the kiosk (front-line worker) plans may not include a license; these lower license tiers allow the use of Outlook on devices with screens smaller than 10.1”. Users who only have on-premises Exchange mailboxes with no service licenses aren't permitted to use the client at all. However, Microsoft does not enforce these licensing restrictions, so users can violate the license terms without realizing it.

Microsoft has kept up a rather astonishing cadence for feature improvements in the mobile versions of Outlook. The client is regularly updated every two weeks, with minor features and bug fixes appearing in each release and larger, more complex features arriving as soon as they're ready. It's always safe to assume that major Microsoft events will bring out a clutch of feature changes; for example, Ignite 2020 brought better separation between personal and corporate email and contacts, the ability to automate S/MIME certificate distribution, and policies for automatic signing and encrypting, along with several changes to “Play My Emails,” support for the Surface Duo, and more. However, Microsoft regularly releases features as a matter of course, too.

Feature support varies between platform and license. For example, dark mode is available on Android and iOS for both commercial and consumer customers. Support for sensitivity labels and shared mailboxes, on the other hand, are only available to commercial customers, and delegate access is available only to commercial customers whose mailboxes are in the service. As it's impossible to roll out new features to more than 100 million clients worldwide simultaneously, Microsoft uses a random selection process to decide who picks up new features first. Features for both commercial and consumer users go to a random selection of people, which means that it's possible to have some people within a tenant see a feature and some not. Microsoft releases commercial-only features on a tenant-wide basis. Successive waves of releases move the roll-out to 100% status over a period that can take some weeks to complete.

Another unique point about the Android and iOS Outlook clients: their support lifecycle is necessarily quite different from their desktop cousins. Microsoft typically supports two versions of iOS: whatever the latest major release is, and its immediate predecessor. For example, when iOS 14 was released, Microsoft stopped supporting Outlook on iOS 12. In September 2021, Apple released iOS 15, which triggered Microsoft's August

2021 announcement that they'd drop support for Outlook mobile on iOS 13. The rapid adoption of iOS versions means that this n-1 support policy won't generally be a problem for most enterprises, as users will self-update their devices to get the latest shiny goodness, but you should remain aware of it.

## Supporting Outlook Cloud Settings

As part of its steady output of new features, Microsoft occasionally delivers something from the long backlog of features that have been requested for years and years. One example: the ability to synchronize signatures between different machines, something that people have been requesting for nearly two decades! Signature synchronization was the first element of the plan to synchronize most Outlook settings through the cloud, a worthy goal. However, signature sync has been delayed multiple times (see MC305463 for the latest). For now, sync covers a subset of Outlook settings (the settings found in the Advanced, Calendar, Ease of Access, General, Groups, Mail, People, Search, and Tasks tabs), and settings are only synchronized for Outlook on Windows.

By default, this feature is enabled in your tenant, and you can't turn it off at the tenant level; you'll need to use a Group Policy template to disable it. Users can [control the setting themselves](#), though.

## Managing Outlook Web App

When Microsoft first shipped Outlook Web Access, it was primitive compared to desktop Outlook, and yet still revolutionary. It was the first mass-market enterprise product to use the technologies we now call AJAX (asynchronous JavaScript and XML). For several releases, OWA's feature set lagged Windows Outlook, but as time passed, we eventually reached a turning point where new features started to be delivered first into on-premises OWA before reaching the desktop rich clients, and now that trend has accelerated so that on-premises Exchange is basically in maintenance mode with very few new features being released. New client features are more likely to appear first in Outlook Web App in Exchange Online than in desktop Outlook, and it's become very clear that we will see few if any new features in the on-prem version of OWA.

The modern OWA user experience rivals that of the full Outlook client. In fact, for many users, OWA provides all the functionality they need, with the convenience that it can be used from anywhere that has an Internet connection without having to first install client software. A good way to start an argument is to ask a room full of heavy email users whether they prefer OWA or desktop Outlook—you'll probably hear some vigorous opinions pros and cons.

## OWA Browser Compatibility

Outlook Web App works with all the browsers Microsoft 365 supports: the latest versions of Microsoft Edge (including both the original version, officially deprecated in May 2020, and the newer version based on the Chromium browser engine), Firefox, and Chrome on Windows, Safari on macOS, and Firefox or Chrome on Linux. You may see some minor differences in your client as the OWA interface evolves rapidly.

OWA supports a "light" mode that is automatically used for older versions of web browsers or operating systems such as Linux that are not fully supported by Microsoft 365. Being HTML-based, OWA light is used as the accessibility version because it can work better with screen readers. The OWA light user experience is degraded when compared to the full version of OWA, with some features not working or simply not available in the interface.

With the pace of changes and new features appearing across the service, maintaining updates to web browser clients is important for customers. If you experience some unexpected behavior of OWA when running a

supported combination of web browser and operating system, you should check for recent updates to the browser. You can find the latest guidance on [OWA web browser compatibility](#) on the Office Support website.

## Managing Mailbox Settings

Over time, Outlook and Exchange have accumulated a large number of mailbox-level settings to provide fine control over things like whether week numbers are displayed in the calendar view or whether read receipts should be sent. Normally, users will configure these settings themselves, but in some cases, you may want to standardize them; for example, as part of user onboarding, you might want to preload a signature into every newly created mailbox. [Set-MailboxMessageConfiguration](#) controls mailbox settings such as signatures and the use of read receipts; [Set-MailboxCalendarConfiguration](#) covers how the calendar looks and how reminders are managed; [Set-MailboxRegionalConfiguration](#) controls regional settings such as the date and time formats, and [Set-MailboxSpellingConfiguration](#) governs the default spelling checker configuration and language. These cmdlets are covered in more detail in Chapter 6.

## Managing Features with Outlook Web App Mailbox Policies

Outlook Web App mailbox policies provide administrators with control over which OWA features are available to individual mailboxes. Examples of the features controlled by OWA mailbox policies include inbox rules, calendar access, mobile device controls, and social network integration. Although an OWA mailbox policy controls access to these features for OWA, it does not disable them on the mailbox itself. For example, a mailbox user prevented from accessing their calendar due to an OWA mailbox policy can still access their calendar using Outlook or a mobile device.

In the on-premises version of Exchange, the OWA virtual directories also provide a point of control for client functionality. You don't have the option of messing about with vdirs within Exchange Online, so you must work with OWA mailbox policies if you wish to disable specific OWA features. A default OWA mailbox policy is created for each tenant. You'll find OWA mailbox policies, including this default, in the **permissions** section of the classic Exchange admin center. The default policy enables all features for OWA.

The OWA mailbox policy that is assigned to a mailbox user can be viewed and changed in the classic EAC by selecting a mailbox and clicking **View details** under **Email Connectivity**. (The modern EAC shows any assigned OWA mailbox policies on the **Mailbox** pivot of the user details but you can't change policy assignments there yet.) You can also retrieve the OWA mailbox policy that is assigned to a mailbox user by running the `Get-CASMailbox` cmdlet.

```
[PS] C:\> Get-CASMailbox Kim.Akers | Format-List OwaMailboxPolicy
```

```
OwaMailboxPolicy : OwaMailboxPolicy-Default
```

The set of OWA mailbox policies defined in a tenant is accessible through the **Permissions** section of the classic EAC. After opening a policy, you can disable or enable features like autosignatures, the premium client interface, and inbox rules by checking or unchecking boxes. Although convenient to manage the set of features published to users through the EAC, it's important to realize that the EAC surfaces some, but not all, of the possible configuration options for OWA mailbox policies. In this example, the `Get-OWAMailboxPolicy` cmdlet is used to display the name of the OWA mailbox policy, as well as a setting that controls whether the LinkedIn contact sync is enabled within OWA.

```
[PS] C:\> Get-OwaMailboxPolicy | Select Name, LinkedIn*
```

```
Name                : OwaMailboxPolicy-Default
LinkedInEnabled     : True
```

Most organizations only need a single, default OWA mailbox policy to which they make one or two changes to suit organizational needs. Alternatively, you can create multiple OWA mailbox policies and assign them to different mailbox users to limit access to some features. For example, you might want to create an OWA mailbox policy that turns off Calendar access for temporary employees or people working on an assembly line, along with another OWA mailbox policy that allows Calendar access for other users. After configuring a new OWA mailbox policy it needs to be assigned to mailbox users by running the *Set-CASMailbox* cmdlet.

```
[PS] C:\> Get-OwaMailboxPolicy | Format-List name  
  
Name: Limited Access OWA Users  
Name: OwaMailboxPolicy-Default  
  
[PS] C:\> Set-CASMailbox Kim.Akers -OwaMailboxPolicy "Limited Access OWA Users"
```

The next time the user refreshes their session or logs on to OWA the new policy is applied. To revert the mailbox user to the default OWA mailbox policy, run the *Set-CASMailbox* cmdlet again.

```
[PS] C:\> Set-CASMailbox Kim.Akers -OwaMailboxPolicy (Get-OwaMailboxPolicy |  
Where {$_.IsDefault}).Name
```

An OWA mailbox policy is not mandatory for mailboxes (the default policy is applied if one is not specified), so you can also remove policies entirely from the mailbox using *Set-CASMailbox*.

```
[PS] C:\> Set-CASMailbox Kim.Akers -OwaMailboxPolicy $Null
```

## Customizing the OWA URL

The OWA URL for all Exchange Online customers is the same, <https://outlook.office.com>. There are some variations of that URL that will also work thanks to redirections that Microsoft has in place. For example, <http://outlook.office365.com> will redirect from HTTP to HTTPS and from the root of the domain to the /OWA virtual directory.

Although this is a simple URL some organizations would prefer a custom URL that includes their domain name, which may be easier for their end-users to remember. This can be achieved with a simple DNS CNAME record in the organization's public DNS zone. For example, a CNAME of *webmail* in the DNS zone for *office365itpros.com* with a target of *outlook.office365.com*, will redirect anyone browsing to <http://webmail.office365itpros.com> to <http://outlook.office365.com>, where Microsoft's redirection to /owa takes care of the rest.

Aside from the branding and ease of remembering, this is also a good strategy for retaining the same OWA URL that you may have previously used for on-premises Exchange, by simply changing it to a CNAME that resolves to the Exchange Online OWA URL instead. This avoids issues with your end-users still having web browser bookmarks for the on-premises OWA URL. However, this technique should not be used for Hybrid deployments.

## Using Microsoft 365 Groups with Desktop Outlook

Originally, Microsoft intended Microsoft 365 Groups to be used for email-based collaboration as "Outlook Groups." The focus for collaboration has since moved to Teams, but organizations can still use email-based collaboration with Groups through all Outlook clients. Group owners can also manage membership and other group settings through Outlook.

Both Outlook for Windows and Outlook for Mac support Microsoft 365 Groups. Apart from having a suitable client, there are two major prerequisites. First, the Autodiscover service must deliver information to clients about what groups are available in the tenant. This is a straightforward matter (you should not have to do



anything) for a cloud-only tenant but can be more complicated for a hybrid environment. Second, your clients must be capable of using Outlook cached mode.

Groups appear in two places in Outlook. First, Outlook includes any groups marked by a user as a favorite (with either OWA or Outlook) alongside other favorites such as folders and shared mailboxes. The mailbox holds details of favorite groups to make them available across all clients. Second, if the user clicks the **Groups** root in Outlook's navigation pane, Outlook expands the set of Outlook groups to which the user belongs with recently-access groups shown first. Groups do not have a hierarchy, but the full list is sorted in order of use, with the most-recent items at the top. You can browse the set of groups, but it can be easy to "lose" sight of an important group if it is at the end of the list. It's best to highlight important groups by marking them as favorites so that they show up in Outlook's list of favorite folders. The interface used to display the contents of group conversations is like the interface used to display the content of mailbox folders.

When you contribute to a conversation, Outlook posts the message to the group mailbox. Because you're emailing the group, copies of your messages are in the Sent Items folder of your mailbox. Given the way that people use Outlook, it is almost inevitable that many of the interactions with Groups are via normal email rather than opening a group to browse its content. If users have opted to do so, they will receive messages sent to groups in their inboxes and can then respond as they do to any other message.

## Groups Toolbar

After opening a group, Outlook switches its normal ribbon to a special group toolbar. Options available are to start a new conversation, switch to the group calendar, or open the shared group document library or notebook. Selecting **Files** or **Notebook** opens web pages to access the documents in the library or the contents of the shared notebook. Outlook calls the desktop version of OneNote if available on the PC. Otherwise, Outlook invokes the browser version to access the group notebook. The **Group Settings** button allows the user to choose what updates they receive for messages posted to the group (for example, all messages or just replies to your messages).

## Group Calendars

Group calendars appear under Outlook's "All Group Calendars" calendar group. To view items in a group calendar, click its checkbox to display it on-screen. Two ways exist to add an event to a group calendar. You can create an event in your calendar (in which case, you are the organizer) and add the group to the list of attendees. The event invitation goes to the group calendar as well as all members of the group. However, if you create the event in the group calendar, the group is the organizer. Email notifications for the new event will only go to group members who subscribe to the group. The difference in behavior is because a group member creates the item in the calendar instead of the group receiving a meeting request via email. In this case, the mail transport service expands the group membership and subscribers receive copies of the event.

## Groups in the GAL

Like all mail-enabled objects, Groups are in the Address Book, where they appear in the Global Address List (GAL) and the All Groups address list. Users can browse the GAL to find groups and examine their properties. After selecting a group, the user will see the group's properties. If the group is public, they can use the **Join group** button to join the group. If private, the **Request to Join** button emails a request to the group owners to add the person as a member. The group owners can then decide to accept or reject the request to join. There's also a **Send email** button that the user can use to send mail directly to the group, whether or not she has previously joined it.

Navigating through a large directory and scattering groups here and there between all the other objects that make up the GAL can create a search challenge for users. It is helpful when a tenant deploys a naming policy

for groups so that they are all gathered in a single place in the GAL. Once a user has found and joined a group, they can add it to their folder favorites.

## Discovering Groups

The prospect of reviewing hundreds of groups to find just the right one to join can be an intimidating task. To make things easier, users can find groups to join with **Browse Groups** (accessed in the Outlook ribbon or by right-clicking the **Groups** resource in the navigation pane). Like OWA's **Discover** feature, Browse Groups displays a set of suggested groups that the user might be interested in joining (Figure 15-2). Outlook uses the Microsoft Graph Insights API to calculate the set of suggested groups based on interactions between users. Simply put, if three other people often use a group and you interact with them often, a high chance exists that the group will show up as a suggestion.

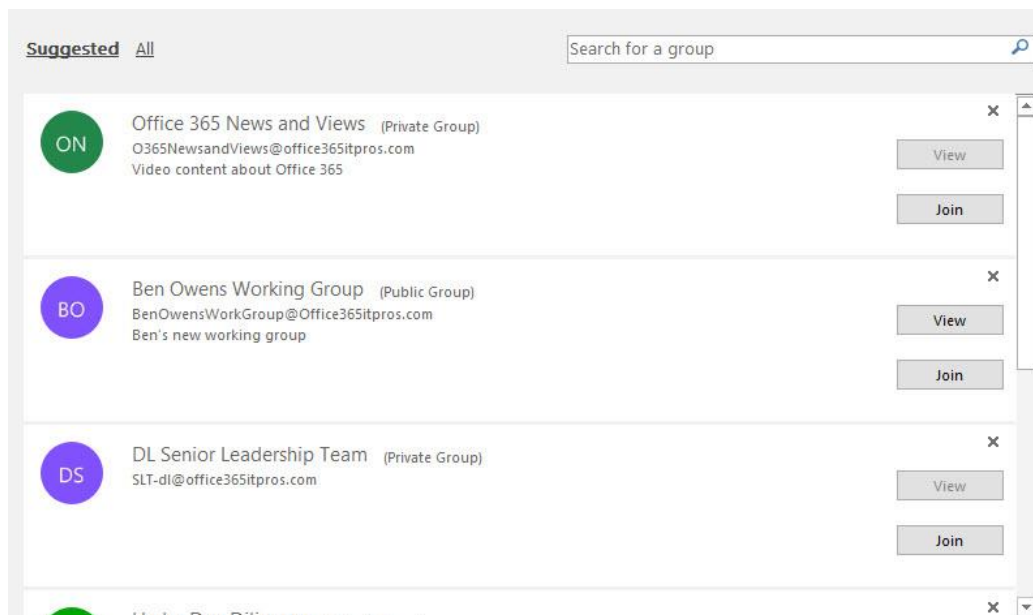


Figure 15-2: Browsing suggested groups from Outlook

Users can see what is happening in a public group by clicking **View**. This opens the group and shows the user the most recent conversations. Clicking **All** reveals a full alphabetic listing of available groups, including those that the user has already joined. It is the equivalent of viewing the *All Groups* address list.

## Offline Access for Groups

One of the biggest advantages of Outlook is its ability to enable users to continue working when a network connection is unavailable. Outlook stores offline group content in the GST (group storage table file). The GST is analogous to the OST used for offline mailbox content and is in the same folder as the OST with a .nst extension.

For performance reasons, Outlook synchronizes conversations and calendar data for groups accessed by the user into the GST. Outlook must run in cached Exchange mode for synchronization to occur, which means that Groups are not available when the client runs in online mode. Groups do not even appear in Outlook's resource list when the client works in online mode. This is an issue for deployments where Outlook runs on thin clients and local files are undesirable. A partial workaround exists by configuring Outlook to synchronize a minimal amount of data. However, given that much of the information shared by groups is only available online, OWA might be the better client choice as people in these situations don't need one of Outlook's strongest selling points (offline access).

Offline access to data in the group document library and shared notebook is not dependent on Outlook. You can synchronize the group document library with the OneDrive for Business sync client to have copies of those files offline, but you cannot get to this cache directly from Outlook. Instead, you can access the synchronized files through File Explorer and make changes there. The group notebook is also accessible offline if you have used it with the OneNote desktop application. OneNote synchronizes any change made in the desktop copy of the notebook back to the online version when the network is available again.

Like mailbox folders, Outlook synchronizes any updates for group conversations created when working offline to the server once a network connection is available.

## How Autodiscover tells Outlook About Groups

Outlook learns about new mailbox resources such as a new shared mailbox through the XML manifest returned by Autodiscover. Although they show some of the same characteristics as shared mailboxes, Autodiscover does not process group mailboxes in the same way. After it returns information about the user's mailbox and any other mailboxes that a user can access, Autodiscover issues a call to list all the Groups that the user belongs to. Details of each group are put in a set of XML files in the `C:\Users\<username>\AppData\Local\Microsoft\Outlook\16` folder (for Outlook 2016). Autodiscover generates a separate XML file, prefixed with "AutoD", for each group, and refreshes the data hourly.

## Group Visibility

Not all Microsoft 365 Groups appear in Exchange Online clients like Outlook. Since mid-2018, Microsoft hides the groups used by Teams from Exchange Online. The rule is:

- New Groups created by Teams set the *HiddenFromExchangeClientsEnabled* property to *\$True*. This stops Exchange clients like Outlook from listing team-enabled groups. Teams also sets the *HiddenFromAddressListsEnabled* property for its groups to *\$True* to prevent the groups from appearing in the GAL and other address lists.
- New Groups created by Outlook and Microsoft 365 administrative interfaces set the *HiddenFromExchangeClientsEnabled* property to *\$False*.

These are the default settings. As described in Chapter 11, you can set their values to control whether users can find team-enabled groups in the GAL.

## Using Microsoft 365 Groups with OWA

OWA was the first client to support Groups and is usually the first client to surface new functionality. If you expand the **Groups** section in OWA's left-hand navigation pane a list of the most used Groups appears. You can click **More** to expand the full list. In this case, OWA lists the groups that you have added to your Favorites list, followed by all groups of which you are a member. You can use the right-click menu to view information about a group or to add or remove it from your favorites list. An option to copy the email address of the group is also available in the right-click menu.

Click on the name of a group to access its contents. The first thing you see are the conversations in the group mailbox (Figure 15-3). A list of conversations in chronological order (last in, first out) appears to the left, and the items that make up the selected conversation appear in the viewing pane to the right. You can set the read status for individual conversations as read or unread, just like you can for messages in your mailbox.

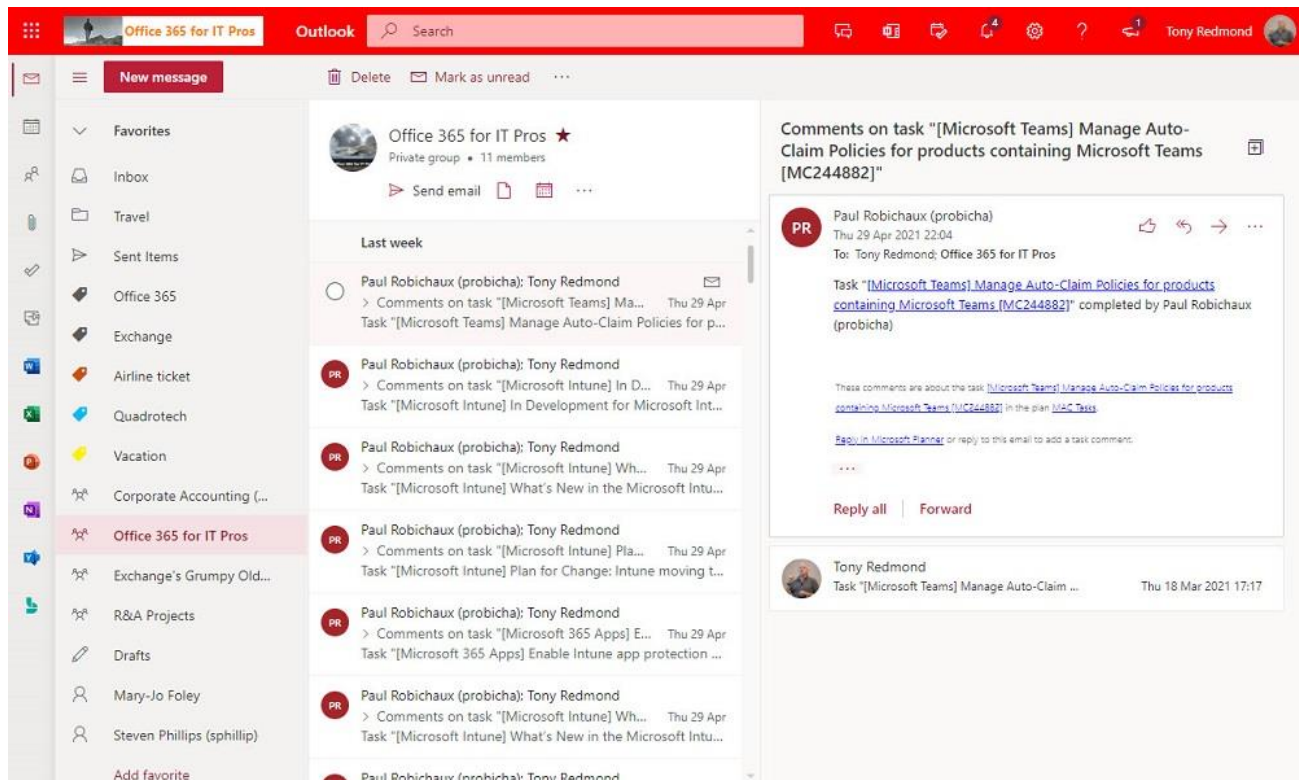


Figure 15-3: The OWA interface for group conversations

Users can reply to the selected conversation by hitting the corresponding button at the bottom of the list of contributions, which invokes the OWA editor. You can use the @ sign to address individual people (including those who are not members of the group). If you add someone outside the list to a conversation, they receive a copy of all earlier items in the conversation and Groups adds their email address to the conversation to ensure that they receive later contributions. When the reply is complete, click send. OWA submits the message for processing through the messaging system to allow Exchange Online to apply transport rules before delivering the item to the group mailbox. Once delivered, the item appears in the conversation.

In the menu bar under the group name, we see icons to switch between:

- **Group Files:** The group document library in SharePoint Online. This view shows the files and email attachments available to the group.
- **Group Calendar:** The group calendar. By default, the group calendar displays alongside the user's calendar using the normal OWA calendar interface.

The [...] menu reveals:

- **Planner:** Access the default plan for the group – if one does not exist, it will be created.
- **Notebook:** The group shared notebook. Other OneNote notebooks can be stored in the document library.
- **Site:** Open a new tab to access the home page of the SharePoint team site for the group.
- **Settings:** You can:
  - **Manage group email:** Decide what messages you receive from the group.
  - **Edit Group** (for group owners): Edit group settings, like whether it is public or private.
  - **Connectors:** Connect the group to various network data sources. We explore how to use Connectors with Groups in Chapter 13.

The settings menu includes the option to **Leave** the group. The menu bar also shows whether the group is private or public.

**Deleting conversation items:** Sometimes people post content to groups that should not appear. Perhaps a conversation includes some inaccurate information or maybe it is just a matter of formatting, spelling, or the choice of words that makes someone want to retract a conversation or a reply to a conversation. OWA allows a user to remove a complete conversation if they start that conversation. Group owners can remove conversations started by any user. OWA also allows users and group owners to remove individual items from a conversation and user who starts a conversation can remove an individual reply from that conversation, even if they are not its author. Removing items during a conversation is not something that should happen without thought as someone else might be replying only to discover that their reply does not make sense when posted. Users and group owners can also remove conversation items with Outlook, but Outlook does not support the deletion of individual items within a conversation. Deletion affects all replies, including those from other users. Unlike personal mailboxes, when a user removes a reply or conversation from a group mailbox, the deleted items do not move into the Recoverable Items folder. Instead, Exchange Online moves the deleted items into the Deletions sub-folder of Recoverable Items in the group mailbox. The deleted items stay in Deletions for 14 days after which the Managed Folder Assistant permanently removes them from the database. The exception is when the group is under the control of a hold, such as those placed by an eDiscovery case, or must be kept because a retention label exists for an item. In this case, the deleted items stay in Deletions for 14 days as normal and then move to the Purges sub-folder, where they stay until an administrator releases the hold or the hold expires.

## Discovering Groups

Naturally, you might want to find some more groups to join. To see what is available, click **Discover Groups** under the Groups section in OWA's resource list and OWA lists a set of cards displaying information about groups that you might want to join (Figure 15-4). Queries against the Microsoft Graph Insights API calculate the set of groups to display, based on the people with whom the user interacts inside the tenant and other signals. Frequent interaction with someone is a strong sign that you might like to be part of a group that the person already belongs to. You can click the X in the top right-hand corner of a card (exposed when you hover the cursor over the card) to hide it from view and show that you are not interested in the group.

Clicking a card displays more information about the group. The name and description are critical to conveying the intent and purpose of the group to a potential member. For example, "HR Working Group" is a reasonable title for a group as it is easy to understand that this group is probably associated with the Human Resources department and is to do with some item of work that the department needs to do.

The search box allows you to look for a group based on its name and the results include groups that you have yet to join as well as those you already belong to. One immediate shortcoming is obvious in that no facility is available for users to narrow their search. For example, you cannot look only for public groups that have a certain word in their title or use other search criteria to find the right group to join.

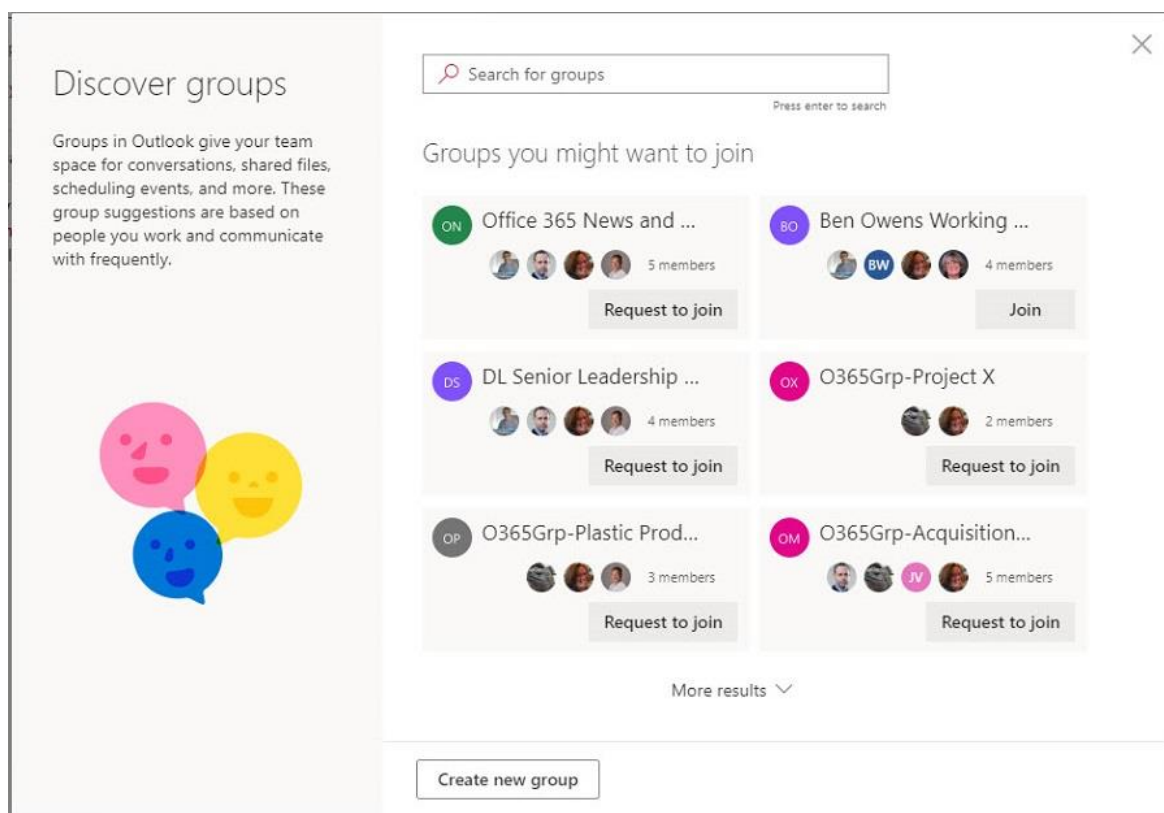


Figure 15-4: Viewing Groups a user might want to join

## Managing Client Access and Protocols

Managing which clients can use mailboxes, and what protocols they can use to do so, is not so common a task as it was in the days before the cloud, but you may still find that you need to apply some restrictions to block old clients, protocols you don't want on your network, and generally tailor connections to meet your security needs. Microsoft 365 can do this in multiple ways, including applying conditional access settings (as described in Chapter 3), using mobile application management (MAM) to control what mobile applications can do (Chapter 16), or controlling the protocols and client access on the service side.

### Managing Application Idle Timeout Periods

Since its introduction as part of Exchange 5.5, OWA has had an idle timeout feature that allows administrators to sign out users after a specified period of inactivity. SharePoint Online has a similar feature, and of course, Azure AD token expiration and revocation may force users to log in again at any time even if they've been active. In late June 2022, Microsoft made a [new idle session timeout feature](#) generally available for a selected set of Microsoft 365 web apps. For now, idle session timeout covers OWA, OneDrive for Business, SharePoint, Office.com, the Office web apps, and the Microsoft 365 admin center, but it should cover more apps in the future. Once you configure the idle session timeout feature in the **Security & privacy** tab of the **Org settings** item in the Microsoft 365 Admin center, the timeout you specify will take effect in all supported apps and supersedes the timeout values set for OWA and SharePoint Online.

### Managing Mailbox Client Protocols

A set of connectivity protocols is available to allow clients to connect to Exchange Online. By default, Exchange Online mailboxes can use all protocols and features, which include:

- **Mail Application Programming Interface (MAPI)** - Outlook clients use MAPI to connect to mailboxes. Older Outlook clients are still able to connect using RPC over HTTP, but Outlook 2010 and newer clients running the latest service packs and updates connect using MAPI over HTTP. RPC over HTTP is an unsupported protocol from October 31, 2017.
- **HTTPS:** the OWA webmail client for Exchange Online connects via the HTTPS protocol.
- **Exchange ActiveSync** - Most mobile devices and applications that connect to Exchange Online do so via the Exchange ActiveSync (EAS) protocol. Note that Microsoft has no control over how a vendor implements the EAS protocol in their email client, a fact that accounts for some "interesting" problems that have occurred with mobile devices over the years. Of course, because the Windows 10 Mail client uses Exchange ActiveSync, Microsoft does have control over its behavior!
- **POP3/IMAP4** - although these are now outdated protocols that do not support the advanced features found in most email clients, some users still use POP or IMAP to connect to mailboxes, and IMAP is also commonly used by external applications that need to ingest email items from mailboxes. POP and IMAP are covered in Chapter 9 of the companion volume to this book.
- **Exchange Web Services (EWS)** - EWS is the API for application access to Exchange mailbox resources, commonly used for integrating external applications. EWS is also used by Outlook to retrieve free/busy, out-of-office, and calendar sharing information.
- **SMTP** - Outlook can perform all its functions using MAPI, but POP and IMAP are access-only protocols. IMAP and POP clients need to use SMTP to send emails. SMTP is also used by external applications to send emails.
- **Microsoft Graph/REST API** - The Microsoft 365 APIs allow applications to access data such as email messages, contacts, calendars, and files. There are separate APIs for the various services (like Outlook and Teams), as well as the Microsoft Graph API.

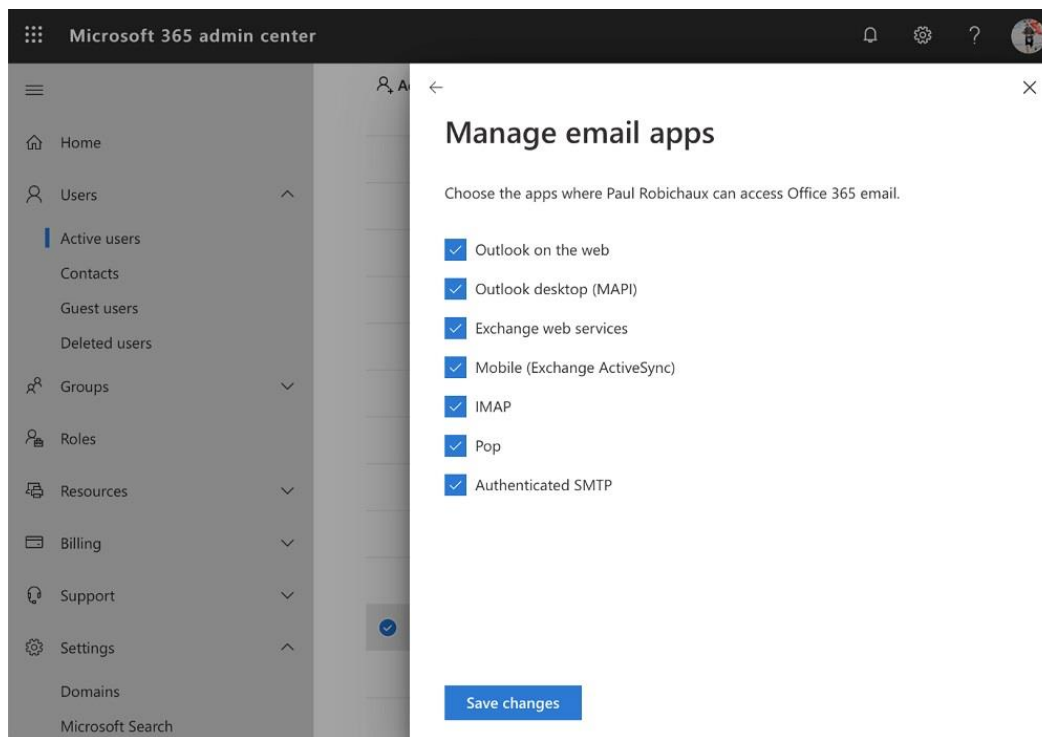


Figure 15-5: Accessing email apps settings in the Microsoft 365 admin center

You can manage mailbox client protocols through the Microsoft 365 admin center, the classic EAC, or PowerShell. Different capabilities exist through each interface, which we will explore in this section. First, the Microsoft 365 admin center includes an "Email Apps" section for user properties. Select a user, and then click

the **Manage email apps** link in the **Mail** pivot and you'll see a list of available protocols (Figure 15-5). Enabling or disabling access to a client protocol is merely a matter of checking the boxes accordingly.

To change the client protocols available to a user through the classic EAC, go to the **mailboxes** section under **Recipients**, open the properties of any mailbox, then **mailbox features**, and then scroll down the controls for the different protocols. The modern EAC doesn't support this feature yet.

You can disable a protocol by clicking the **disable** link. To enable the protocol for the mailbox user again click the **enable** link. Most protocols can be managed in the classic EAC, however, access to Exchange Web Services (EWS) can only be managed using PowerShell. Managing EWS is described in more detail later.

To view the status of each client protocol for a mailbox user via PowerShell, use the *Get-CASMailbox* cmdlet.

```
[PS] C:\> Get-CASMailbox -Identity Kim.Akers | Select *enabled
```

```
ActiveSyncEnabled      : True
OWAEnabled              : True
OWAforDevicesEnabled   : True
ECPEnabled              : True
PopEnabled              : False
PopMessageDeleteEnabled : False
ImapEnabled             : False
MAPIEnabled             : True
MapiHttpEnabled         : True
UniversalOutlookEnabled : True
OutlookMobileEnabled   : True
MacOutlookEnabled      : True
EwsEnabled              : True
OneWinNativeOutlookEnabled : False
```

To disable a client protocol in PowerShell, you could use the *Set-CASMailbox* cmdlet. For example, to disable access to the Monarch client for a mailbox, run the following command:

```
[PS] C:\> Set-CASMailbox -Identity Kim.Akers - OneWinNativeOutlookEnabled $False
```

You can repeat the same steps for each protocol. If you have a combination of enabled/disabled protocols you want to apply to mailboxes by default, then you will need to develop a custom script to apply those changes as part of your mailbox provisioning process. You may be able to avoid some of this tedium now that Microsoft has given us the [Set-CASMailboxPlan cmdlet](#). This cmdlet allows you to set up a template for protocol access, then apply the template automatically to newly created mailboxes. By default, you get four CAS mailbox plans in the service; you cannot create, rename, or remove them. The plan applied when a mailbox is created depends on the license assigned to the mailbox: Exchange Online Kiosk plans get the ExchangeOnlineDeskless plan, Exchange Online P1 gets ExchangeOnline, and Exchange Online P2 gets ExchangeOnlineEnterprise (which means that E3 and E5 subscribers get it too). Microsoft 365 Business Basic subscriptions get ExchangeOnlineEssentials. Suppose you wanted newly created mailboxes for your enterprise users to have IMAP and POP disabled—you could easily do this by running the following cmdlet:

```
[PS] C:\> Set-CASMailboxPlan -Identity ExchangeOnlineEnterprise -POP3Enabled $False -IMAPEnabled $false
```

**Note:** There are two protocols in the *Get/Set-CASMailbox* cmdlets that you can probably ignore for now. *MapiHttpEnabled* is only applicable to on-premises Exchange organizations, so changing it in a cloud environment will do nothing. The second, *UniversalOutlookEnabled*, controls whether the Mail app for Windows 10 (sometimes called "Universal Outlook" in news articles and blog posts) is enabled.



## Managing Exchange Web Services

The Exchange Web Services (EWS) API was originally introduced with Exchange 2007 as a replacement for WebDAV and, at the time, the planned successor to MAPI. MAPI lives on, but so does EWS. Applications can use EWS to retrieve information from Exchange Online services or to interact with data in Exchange Online mailboxes. For example, an EWS application can retrieve information about calendar items for room mailboxes to determine which items might have an organizer who no longer works for the company.

Since 2018, Microsoft has steadily been [moving away from EWS](#) as the preferred API for programmatic access to Exchange Online data, in favor of Microsoft Graph. This change has no impact on on-premises Exchange servers. However, for those who are using EWS to talk to Exchange Online, this change means that:

- EWS will not receive future feature updates.
- Microsoft will decommission basic authentication for EWS and exclusively use [OAuth 2.0](#). Originally planned for October 13, 2020, due to the worldwide COVID-19 pandemic, Microsoft temporarily paused its plan to block basic authentication for connectivity protocols where tenants actively use these connections. However, the plan is now back to full strength: Microsoft will decommission basic authentication fully in October 2022 and is already being turned off for select tenants when Microsoft detects that the protocol is unused in the tenant.
- EWS will remain supported in production for Exchange Online environments if you use OAuth 2.0.
- Microsoft deprecated 25 EWS APIs ([listed here](#)) on March 31, 2022.
- Beginning in September 2022, Microsoft will block the registration of new EWS applications. That means that migration vendors, as well as vendors that create client applications that leverage EWS, won't be able to register new instances of their applications—which will pose a problem for vendors whose applications depend on capabilities that aren't yet available in Graph.

If you create (or use) applications that use EWS features not yet present in Graph, you'll be stuck with using EWS unless and until Microsoft delivers feature parity between the two.

Despite Microsoft's plans to eventually deprecate it, there are lots of clients that still use EWS today. For example, the Windows Outlook client uses EWS for calendar free/busy information, Out of Office settings, calendar sharing, and other features such as MailTips. Other applications (including the macOS Mail application) use it for access to mailbox data as well. As with the other protocols used to access Exchange Online, there are controls available for administrators to use for a variety of scenarios. EWS controls can be managed at the mailbox level or the organization level. The EWS settings for a mailbox are retrieved by running the *Get-CASMailbox* cmdlet, and the organization level settings are retrieved by running the *Get-OrganizationConfig* cmdlet.

### Using the EWS Application Allow or Block List

Because EWS was the preferred protocol for mailbox access, naturally an ecosystem of third-party tools sprung up that used EWS to access mailbox data. For example, the Fantastical calendaring application for macOS and iOS uses EWS to access your calendar when you install it, and LinkedIn for years had an option to synchronize your mail contacts to their service using EWS. Thankfully, Microsoft has provided tools for controlling these applications without having to block *all* EWS traffic.

Disabling the entire EWS protocol because of one unapproved example of application access would deny your organization the many good things that EWS allows. Fortunately, we can be selective in what we block or allow for EWS by configuring an EWS application access policy. The EWS application access policy can be configured on a per-mailbox basis or configured for the entire organization. It's important to know that this

blocking mechanism depends on the User-Agent header sent by the client to the service, which is easily spoofed. Don't rely on this mechanism alone to protect critical data.

**Real World:** If your tenant has "K" (Kiosk) or Office 365 F3 licenses, the EWS allow/block controls for your tenant will not work at the organization level. For non-Kiosk mailboxes, apply the allow/block controls on a per-mailbox basis. For the mailboxes themselves, direct EWS application access to the mailboxes is not permitted anyway.

Suppose you wanted to block access to an application whose User-Agent string had a value of "LinkedInEWS." This would require two separate calls to the *Set-OrganizationConfig* cmdlet. First, set the *EWSApplicationAccessPolicy* to enforce the block list.

```
[PS] C:\> Set-OrganizationConfig -EwsApplicationAccessPolicy EnforceBlockList
```

Unless and until you do this, the EWS block list won't be used at all. Next, add the target user agent to the EWS block list.

```
[PS] C:\> Set-OrganizationConfig -EwsBlockList @{add='LinkedInEWS'}
```

The EWS block list is a multi-value attribute so it should be managed using add/remove methods to avoid overwriting existing values when you are making modifications. For example, to add the Outlook mobile user agent to the block list you'd run this command.

```
[PS] C:\> Set-OrganizationConfig -EwsBlockList @{add='Outlook-iOS/*',"Outlook-Android/*"}
Get-OrganizationConfig | Format-List ewsblocklist
EwsBlockList: {Outlook-iOS/*, Outlook-Android/*, LinkedInEWS*}
```

Similarly, to remove an entry you run this command.

```
[PS] C:\> Set-OrganizationConfig -EwsBlockList @{remove="LinkedInEWS*"}
Get-OrganizationConfig | Format-List ewsblocklist
EwsBlockList: {Outlook-iOS/*, Outlook-Android/* }
```

Unlike ActiveSync device access rules, the strings used for EWS block and allow lists can use wildcards for partial matches. However, unlike ActiveSync controls there is no quarantine action available, only allow or block.

If you only wanted to block access for an application on a single mailbox, you'd follow the same process, except that you'd use the *Set-CASMailbox* cmdlet instead of *Set-OrganizationConfig*. Enforcing a block list will permit any EWS application that is not in the block list to connect. A more restrictive approach is to enforce an allow list instead, which requires that any EWS applications be listed in the allow list before they can connect.

```
[PS] C:\> Set-OrganizationConfig -EwsApplicationAccessPolicy EnforceAllowList
```

Enforcing a block or allow list for EWS has no impact on the Entourage, Outlook for Mac, or Microsoft Outlook applications. These applications are controlled with different EWS settings which are discussed next.

## Blocking/Allowing Mac Clients

Organizations that standardize on Windows may want to block Mac clients from various parts of their network. In addition, organizations that allow Macs may still want to control how clients connect. For example,

an organization that supports BYOD may wish to require users to connect only through Outlook Web App or Outlook Mobile.

Even though some versions of the Mac Outlook client use EWS, Microsoft provides a separate way to prevent only Microsoft's Mac clients from connecting to the service. The Mac clients are allowed by default and can be blocked using *Set-CASMailbox* or *Set-OrganizationConfig*. For example, to block Mac Outlook for a mailbox user you would run the following command.

```
[PS] C:\> Set-CASMailbox Kim.Akers -MacOutlookEnabled $False
```

Apple's Mail and Calendar clients use EWS, and you can block their access by specifying their user agent strings using the EWS block list described in the preceding section. Note that the *EwsAllowEntourage* flag to these cmdlets only affects the ancient, and long-deprecated, Microsoft Entourage client for macOS. It doesn't affect any other client.

## Blocking/Allowing Outlook EWS Access

You can also block Windows Outlook's access to EWS. This will break free/busy information, Out of Office, and calendar sharing. Blocking Outlook is a rare requirement and typically only applies to organizations that want to limit specific users from being able to share calendars or see the availability of other users.

The *Set-CASMailbox* or *Set-OrganizationConfig* cmdlets are used to apply the block. For example, to block EWS access by Outlook for a mailbox user you would run the following command.

```
[PS] C:\> Set-CASMailbox Kim.Akers -EwsAllowOutlook $False
```

## Other Uses for EWS

Exchange Web Services is used by many organizations for custom application development, such as creating integrations between Exchange Online and their in-house line of business applications. It was also the API used for integration between Exchange Online and other Microsoft services such as Skype for Business and SharePoint Online. EWS applications can send and receive email messages, manage calendar items, and a whole lot more.

Many third-party migration tools use EWS to access Exchange Online mailboxes. In addition, an ecosystem of "cloud backup" products uses EWS as the access protocol to extract mailbox data from Exchange Online to copy to cloud storage and so meet the needs of many organizations who would like to make use of cloud systems but are concerned about backup and recovery.

Starting on September 30, 2022, Microsoft will block tenants from registering new EWS applications in Azure Active Directory. This will effectively prevent new deployments of any applications that use EWS, whether from commercial vendors or internal applications.

## Client Access Rules

Exchange Online supports the ability for a tenant to create and apply client access rules, which let you block or allow specific protocols, authentication types, and IP address ranges. The *New-ClientAccessRule* cmdlet creates [client access rules](#); as you would expect, *Get-ClientAccessRule*, *Set-ClientAccessRule*, and *Remove-ClientAccessRule* cmdlets are also available. You must use PowerShell to create and manage client access rules, as there is no support for them in the Exchange admin center.

To be more specific, you can use these rules to selectively block or allow access through the Exchange ActiveSync, IMAP, EWS, Outlook Anywhere, OWA, POP3, PowerShell, and REST API access, as well as controlling access to the Exchange Admin Center and PowerShell Web Services. You can also restrict or allow

connections using AD FS, basic, certificate-based, or OAuth authentication types. These selectors can be combined so that, for example, you can do things like block all OWA access that's not authenticated via AD FS.

By default, a tenant has no client access rules, so anyone with a valid mailbox can access Exchange Online from any IP address using any protocol not blocked at the tenant or mailbox level. A tenant can define up to 20 rules to control what protocols clients can use over what connections. You can even create a filter based on some Azure Active Directory account properties to control the scope of a rule (for instance, only apply this rule to users in Dublin). This example creates a new rule to block ActiveSync connections such as those from Apple iOS mail app clients unless they come from a specific IP address range corresponding to the corporate network.

```
[PS] C:\> New-ClientAccessRule -Name "Block ActiveSync" -Action DenyAccess -AnyOfProtocols ExchangeActiveSync -ExceptAnyOfClientIPAddressesOrRanges 203.0.113.0/24
```

Each rule must contain an action (specifying whether the rule should allow or block access when its conditions are met) and then a set of conditions. These conditions may either be positive or negative; that is, for a positive condition, the rule matches if the specified parameter *does* match the rule, and for a negative condition the rule matches if the specified parameter *doesn't* match. The rule above is an example of a negative condition: it matches any IP that is not in the specified range. You could get the same effect using this command:

```
[PS] C:\> New-ClientAccessRule -Name "Block ActiveSync" -Action AllowAccess -AnyOfProtocols ExchangeActiveSync -AnyOfClientIPAddressesOrRanges 203.0.113.0/24
```

However, this example wouldn't do anything unless you created other rules, because by default all clients can connect using all protocols from all addresses. To make this rule work, you'd need another rule that blocks everything else... so it's simpler just to use the first form unless you need a more complex set of rules to enforce your requirements.

Exchange evaluates client access rules in priority order, which you can alter to promote or demote rules. As soon as the conditions of a rule match, the rule is applied, and rule processing stops. When you create a rule, you can explicitly apply a priority (0 is the highest priority value). If you don't specify a priority, the first rule you add becomes priority 0, the second priority 1, and so on. Changing the priority of an existing rule bumps down the priority of all the rules after it; if you create a new rule with priority 5, the former priority-5 rule becomes priority 6, and all rules after it follow suit.

**Creating Client Access Rules Programmatically:** PowerShell is one of the protocols you can control with client access rules. Microsoft strongly recommends that the first rule you create is this:

```
[PS] C:\> New-ClientAccessRule -Name AllowPS -Action Allow -AnyOfProtocols RemotePowerShell -Priority 0
```

This rule guarantees that you'll be able to use Remote PowerShell even if you later create a rule that might block it. If you do manage to create a rule that prevents you from using PowerShell against your tenant, you'll have to contact Microsoft support to ask them to fix it.

You can create rules that apply to users instead of only to protocols and IP ranges, too. The *UserNameMatchesAnyOfPatterns*, *UserRecipientFilter*, and *ExceptUserNameMatchesAnyOfPatterns* conditions are useful for meeting requirements such as "Allow any users in the sales department to use Exchange ActiveSync but block it for everyone else." In that case, using *UserRecipientFilter {Department -eq "Sales"}* as a filter condition would do the trick.

You can also scope rules to control whether they apply to users only or all applications and clients. If you scope a rule using `-Scope All`, connections from services and applications (such as EWS connections from an on-premises hybrid Skype for Business server) will be affected by the rule.

Because you can create complex chains of rules, it is very important to validate the effect of the rules you build using a test tenant. Even with the [Test-ClientAccessRule](#) cmdlet, there's no substitute for a thorough test before you unleash these rules on your environment.

## Managing OneDrive for Business Clients

The OneDrive for Business client has come a long way since its 2016 debut. The first version of the client was based on the legacy Groove.exe client, then Microsoft introduced the Next Generation Sync Client (NGSC). NGSC morphed into what is now known as OneDrive and has by now replaced the legacy client. The current client offers full sync functionality for personal OneDrive folders, OneDrive for Business libraries, and SharePoint Online document libraries (including those associated with Microsoft 365 Groups). Sync works across tenants and Active Directory domains; for example, the files used to produce this book are hosted in the Office 365 for IT Pros tenant, but I synchronize them to several Mac and Windows PCs that are joined to various (or no) other Azure AD tenants.

### Deploying the OneDrive Sync Client

If you've already deployed a reasonably modern version of the desktop applications to your computers (say, Office 2016 or later), then the latest OneDrive client is installed. Note that the OneDrive client has its separate update mechanism and doesn't get updates through the same channel as either Office or the host OS.

If you have domain-joined computers, you should probably also download and configure the [OneDrive Group Policy template](#) for your organization. If you plan to use Group Policy to control sync settings for end-users, you will need to modify the copy of the Group Policy ADMX file that you place in your Central Store. You can edit the file using any text editor such as Notepad or Visual Studio Code. Within the ADMX file replace any instances of the placeholder text for the tenant ID with your real tenant ID, discovered by running the `Get-MgOrganization` cmdlet. The tenant ID is the GUID string displayed after a successful connection.

```
PS C:\> Connect-MgGraph
Get-MgOrganization | Select Id

Id
--
2b9bca49-687e-4e5f-8a52-21350b719b06
```

The Group Policy template also contains a [wealth of other settings](#), including controls that allow you to restrict what tenants users may synchronize with, control what the client should do on laptops when they're on battery power, apply default bandwidth limits for upload and download, block sync traffic when the computer is idle and no one is signed into it, and block or allow the client from synchronizing personal OneDrive accounts.

Besides these Group Policy-based settings, Microsoft supports [silently configuring OneDrive using the user's credentials](#) to provide single sign-on, but this requires you to configure the `SilentAccountConfig` registry key. If you don't do that, plan on communicating to your users how to sign into their OneDrive client manually.

## Managing OneDrive Sync Client Updates

Whether you install the OneDrive sync client with Microsoft 365 Apps or by using the OneDrive standalone installer, the sync client application files are installed in the `%localappdata%\Microsoft\OneDrive` folder. This means that OneDrive is installed by default as a per-user application, not as a per-computer application. It also means that OneDrive can be installed and updated by users who do not have local administrative rights on the computer. If you want to install the OneDrive client per machine, you can do so [with these instructions](#).

Microsoft releases OneDrive sync client updates through [three deployment rings](#):

- **Insiders** – this ring is an opt-in distribution channel for early adopters; releases into this ring go only to users who have enrolled in the Windows or Office Insider programs. Complete deployment of a new version into this ring takes about 3 days.
- **Production** – Once Microsoft releases a new update for this ring, updates are released to a small percentage of clients in this ring first, and the remaining clients receive the update within approximately two weeks. This is the default deployment ring for OneDrive installs, and there are no controls available for including or excluding your computers from that initial small percentage of clients that receive updates first. During the release of updates to the production ring, Microsoft uses its telemetry to measure the success of the initial rollout. If a problem is discovered, updates will be stopped while a fix is developed. The clients that received the problematic update will then be the first to receive the next update to fix the issues, and then updating will continue through the rest of the production ring.
- **Deferred** - Updates are released to the deferred ring only after they have been successfully deployed to the production ring without any major issues. The deferred ring receives updates over a 60-day window that starts after the production ring deployment window has ended. Critical updates might be released to clients early in that 60-day window, but otherwise, you can expect your clients to update any time within that period. You can use [Group Policy objects](#) to control this behavior.

The OneDrive sync client checks for newly available updates every 24 hours while it is running. To check for or apply sync updates, the client must be able to reach the `oneclient.sfx.ms` and `g.live.com` domains. If OneDrive has not been running within the last 24 hours, it will check for updates immediately the next time it starts. Windows 10 clients also have a scheduled task that will check for OneDrive updates every 24 hours regardless of whether the sync client has been running. These mechanisms are all designed to keep OneDrive updated with the latest bug fixes and feature enhancements. If you're tempted to block OneDrive from checking online for updates then you'll need to periodically check for an updated OneDrive standalone installer, download the updated application files, package them for deployment, and then manage the rollout using your enterprise software deployment tool. That's an unnecessary overhead for most organizations though. If you prefer to take a conservative approach to OneDrive updates then you can configure clients to use the enterprise ring, but you should be aware that this will slow the release of new features as well.

You can configure the update ring for OneDrive by using the [OneDrive Group Policy template](#). In **User Configuration, Administrative Templates, OneDrive**, there is a setting to *Delay updating OneDrive.exe until the second release wave*. Enabling that setting places the client in the deferred deployment ring.

## Managing Teams Clients

The Teams desktop app is available for:

- **Windows:** 32-bit versions are available for Windows 8.1 or later; a native 64-bit ARM version is available for ARM on Windows 10. You can also install Teams on Windows Server 2012 R2 or later. Teams requires .NET Framework 4.5 or later.
- **macOS:** running OS X 10.11.0 or later. There is no officially released Apple Silicon-native Teams client, although the Intel version runs fine on Apple Silicon machines.
- **Linux:** .deb and .rpm formats.

If you prefer to use MSI packages to control the roll-out of Windows desktop clients, you can [download 32-bit and 64-bit versions](#). A special Teams app is available for [Microsoft Surface Hub](#) devices. The [Teams client downloads page](#) has pointers for the supported platforms while the [Teams features by platform page](#) describes the features supported on the different platforms, including the mobile clients.

Microsoft uses a portable framework known as [Electron](#) to write the desktop clients. As part of the Windows 11 announcement, Microsoft said that they are drastically changing the Teams client architecture, moving away from Electron and the Angular user interface library. The new version, which is informally being called “Teams 2.0”, should dramatically improve the Teams client’s performance, as well as making it easier for developers to take full advantage of OS-native features such as notifications and external audio and video devices. Microsoft has not made any public statements about the exact schedule for when the Windows 11 integrated client will be usable with the Teams workload in Office 365 (as opposed to the “Teams for Life” consumer version), nor when the improved client will be available on macOS or Linux.

In most cases, the browser client is functionally equivalent to the desktop client, but as evident in the hints dropped when you use Teams with a browser, Microsoft prefers the desktop client. As with other Office Online applications, Teams supports modern browsers like Chrome, Edge, and Brave. The Safari 14 browser (and later) on macOS supports most features of the Teams browser app, including using video cameras and sharing in calls and meetings.

Mobile Teams apps are available for [Apple iOS](#) and [Android](#). Platform-specific technology is used (Swift for iOS and Java for Android) to create the user interfaces. Underneath, the mobile clients share the same basic functionality as their browser and desktop counterparts, including the ability to create and manage teams (membership and channels). As you’d expect, mobile clients are easier to use and more responsive than the other clients in some areas and less functional in others. For example, the mobile client often switches to a different tenant faster than the desktop or browser client do, but it is easier to compose a complex post with the desktop or browser client.

As detailed in Table 15-2, the Teams desktop and browser clients support limited offline capability to read and send personal chats and channel conversations. Almost every other piece of functionality needs an internet connection to work unless an app is designed to work offline. In some cases, like the calendar or documents, other tools are available to maintain local copies of data that permit offline activity.

<b>Feature</b>	<b>Offline capability</b>
Personal and group chat messages	Cached copies of messages are available offline for pinned and recent chats. Messages can be composed and queued for delivery locally when the network connection resumes. Messages are sent if a connection becomes available within 24 hours. If the delay is longer, the attempt to send the message will fail and Teams prompts the user to retry.
Channel messages	Cached copies of messages are available offline for pinned channels and channels recently accessed by the user, going back about 90 days. Messages for hidden channels are unavailable. Sending channel messages when offline works similarly to chat messages.

Calendar app	Cached copies of calendar events can be viewed. You can't schedule, initiate, or join meetings using Teams when offline. You can use Outlook to work with your calendar when offline.
Files	Both channel folders (SharePoint Online) and personal files in (OneDrive for Business) are unavailable offline. Files can be synchronized to the local drive with the OneDrive for sync client and accessed offline.
Tasks	None. You can use Outlook or the To Do app to work with personal tasks (but not tasks in Planner) when offline.
Wiki	None.
Lists	None.
Whiteboard	None. This may change now that Microsoft has moved Whiteboard storage to OneDrive for Business.
Yammer communities	None.
People card	None.
Third-party apps	Depends on the offline capability of the app.
Switch tenant	Not possible—when you go offline, you're stuck in the tenant you're currently in until you regain network access.
Manage team	None.
Calls	None unless the location deploys a <a href="#">Survivable Branch Appliance</a> (SBA).

Table 15-2: Teams offline capabilities

From a practical perspective, no matter what client you're using, the internet connection must be reasonably capable in terms of both latency and bandwidth. Even though Teams clients cache data for faster access, the high latency, and low bandwidth often available in airplane Wi-Fi services can make using Teams an excruciating experience. To help assess the impact of Teams usage on a network, the Teams admin center includes a network planner to help figure out what network capacity is needed. Your mileage will vary depending on how people use Teams and the clients used. Remember that although Teams usage increases over time, the usage of other applications like email might decrease to offset the extra network demand. In terms of security, Teams supports the same multi-factor authentication methods as other Microsoft 365 applications do.

## Teams Insider and Pre-Release Channels

Like Windows and the rest of Office, Teams supports pre-release channels for the Enterprise version to allow customers early access to pre-production features. Three channels are available:

- **Beta Preview Channel:** The earliest that non-Microsoft users can access new Teams features. Microsoft uses rings to describe the release of new functionality from the development group (ring 1) to general availability (ring 4). This channel was Ring 1.5.
- **Private Preview Channel:** Access to more developed forms of new Teams features. Microsoft previously called this Ring 3.
- **Public Preview Channel:** Tenants can enable public preview for selected users via an update policy (see later section) to allow those users to have early access to new features. This channel layers pre-release functionality on top of the general availability build (Ring 4).

Collectively, the three pre-release channels are known as Teams Insider. Restricted access to the Beta Preview and Private Preview channels is available to companies participating in Microsoft's Technology Adoption Program (TAP). Features released to the TAP are under Non-Disclosure Agreements (NDA).



Clients configured to use the Current Channel (Preview) release of Microsoft 365 Apps for enterprise automatically use the Teams Public Preview channel unless the Teams update management policy assigned to their account disables the link between Office preview and Teams (the link is on by default). To disable the link, update the policy in the Teams admin center or use the `Set-CsTeamsUpdateManagementPolicy` cmdlet (in the Teams module). For example, this command disables the link for any account assigned the default update management policy:

```
[PS] C:\> Set-CsTeamsUpdateManagementPolicy -Identity Global -AllowPublicPreview Disabled
```

The value of the `AllowPublicPreview` setting can be either Enabled or Disabled. Enabled allows users to switch the Teams desktop client into preview mode manually. Disabled hides the option to switch to preview mode. When a client runs in preview mode, the Teams desktop client signals this status by displaying a **P** beside the upper right-hand quadrant of the user profile photo in the menu bar.

## Teams Client Release Notes

Like most of Microsoft 365, Teams is developing rapidly. Microsoft publishes [release notes for Teams](#) online to give formal guidance about the introduction of new features. The release notes don't cover every new feature released for Teams, but they are a useful resource.

## Teams and Memory

The Teams desktop client has a reputation for being a performance hog on Windows which loves to grab lots of memory from the operating system. This well-earned reputation came about because Teams is an Electron application, and both the desktop and browser clients use the Chromium memory management model. This model uses free available memory to cache data and releases memory when it's needed by other applications (much as on-prem Exchange used to), so it's expected to see the set of Teams processes use what appears to be a large amount of memory, even when idle. As memory demand from other applications grows, the amount used by Teams should reduce, or at least not grow. For more information, see [this Microsoft article](#).

The Teams processes include:

- Main window: This process caches chat and channel data and uses the most memory.
- Experience renderer: This process drives the pop-out window for meetings and chats.
- GPU process: This is the Electron GPU process responsible for GPU interaction.
- Plugin Host: The Teams media processor (video and audio).
- Notification Manager: Handles notifications.

You can see the names of the processes by running the Windows process explorer and hovering over a process listed under Teams.exe. Note that on macOS these process names will of course be different.

## Teams Desktop Client Updates

For Windows computers, the Teams desktop client can be deployed automatically [using group policy or other distribution mechanisms](#). Microsoft automatically deploys the Teams client as part of the Microsoft 365 Apps installation for users who have Microsoft 365 Business or Enterprise subscriptions.

The Teams Windows client is self-updating and the mechanism used (Squirrel, an open-source component that is the default auto-update method for Electron-based apps) works differently from the other Office click-to-run desktop applications as users do not have to exit the client before an update proceeds. The combination of using this approach plus installing into a folder in the user area (`%LocalAppData%\Microsoft\Teams\`) rather than the program files structure (which requires administrative

permission) makes it simpler to ensure that Teams runs the latest software if you allow Teams to auto-update. Among the contents of the Teams directory, you'll find:

- **Current:** The files for the currently installed version of Teams.
- **Previous:** Files for the previously installed version of Teams. If the update process fails, the Squirrel utility can roll back using these files.
- **Packages:** Files downloaded by the Squirrel utility to update Teams.

Note that Teams user data is in `%LocalAppData%\Roaming\Microsoft\Teams\`. Here you'll find folders like `IndexedDB`, which store cached copies of Teams data for local (and offline) access. More information about the use of this file [is available here](#).

On macOS, the Teams client also installs its updates using Squirrel. The other Office apps use a macOS-native app called Microsoft AutoUpdate to download feature updates for both production and Insider rings, except for versions of the apps downloaded through the Mac App Store. Teams updates for old versions are delivered through Microsoft AutoUpdate, presuming it's installed. This covers the case where a previously installed version didn't receive updates and is now too old to be directly updated through Squirrel—the updates will be downloaded through AutoUpdate and applied, and, once the Teams client is up to date, the normal Squirrel-based mechanism will be used for future updates.

Microsoft usually releases Teams update packages on a two-week schedule. This schedule could vary if a problem discovered in a build is important enough to warrant an immediate fix, in which case Microsoft will release a new package. When signed into their home tenant, Teams desktop clients perform a check every few hours to detect new updates, and, if found, the client downloads and applies the update in the background when the workstation is idle.

Users can also update the client through the **Check for updates** option in the profile menu. This does no more than accelerate the process of finding and applying available updates, but it might be necessary if Microsoft releases an urgent update. Depending on the type and scale of the update, Teams might prompt the user to refresh the client (reload) to complete the procedure.

Several reasons can affect the ability of the Teams client to download updates, including:

- Antivirus software blocks the Teams update executable or the download of the update packages.
- Network infrastructure to specific locations is not capable of supporting software downloads.
- People only use Teams desktop client for calls and/or meetings. To ensure performance, Teams does not check for updates during these activities, and if the user closes Teams after their meeting finishes, the client will never download and apply updates.
- The files for the Teams client are in a non-standard path (not `%LocalAppData%\Microsoft\Teams\`). In this situation, users must update the client manually.

The fact that the Teams client self-updates is challenging for some organizations who like to control the software users have on their workstations. However, Teams is more aligned with the app model found on mobile devices than traditional software distribution channels and the aim is to deliver the best and most functional experience to users. For this reason, if you want to know what Teams delivers in client updates, you need to keep a close eye on the [release notes published by Microsoft](#). The release notes are also available by typing `/whatsnew` into the command box.

## New Features and Enablement Flags

When Teams ships a new feature, it uses a two-phase process to enable the functionality. First, updated software rolls out to backend services and for download to desktop, browser, and mobile clients (as

appropriate). To allow Microsoft to enable features separately from software updates, they use configuration flags to turn features on or off. Features don't become available to users until Microsoft sets the relevant flags.

Clients can receive software updates with the feature flags set off. When Microsoft testing and validation with early adaptors show that the feature is stable, Microsoft enables the flag to "light up" the feature. This process happens progressively as tranches of clients are enabled. Given the distributed nature of Office 365, the process of enabling a new feature can take some time. To ensure continual validation against a cross-section of workstation configurations and usage patterns, Microsoft does not organize users into tenants. Instead, tranches include users selected from multiple tenants and multiple countries. All of this means that some users within an organization may see a new feature before others.

## Mandatory Updates

Teams supports [Microsoft's Modern Lifecycle Policy](#), which requires users to run a recent version of software to maintain access to services. Once the Teams desktop client is more than a month behind the current release, the client displays a banner to advise that a software update is necessary. The banner includes a link to start an update. If the user doesn't perform a software update, Teams continues to prompt the user to update until the software is more than three months behind the released version. At this point, Teams displays a blocking page to give the user the choice to update the software now, looking for help from their IT administrator (to help with the update or apply a manual update), or to continue accessing Teams via the browser client.

New installations of the Teams client can deploy software more than three months old. In this situation, a 28-day grace period begins during which the user can continue to use the old software while Teams attempts to update the client in the background. At the end of the grace period, the user cannot use Teams, and an IT administrator must update the device.

## What Version of the Desktop Client is on a Workstation?

To learn what version of Teams is running on a workstation, select **About** and then **Version** from the profile menu. The client lists the current version number and the date of the last software update in a banner at the top of the screen.

You have Microsoft Teams Version 1.4.00.34266 (64-bit). It was last updated on 09/12/2021.

Sometimes it is useful to programmatically fetch information about the version of the Teams client running on a workstation. If you need to retrieve the information programmatically, on Windows, you can use PowerShell to report details about the Teams executable and updates from and the log file that records updates. With an eye on the requirement to keep the client software up to date, here is a script to report about the last update and age of the Teams executable on a Windows workstation.

```
[PS] C:\> $TeamsExecutable = Get-Item("${Env:LocalAppData}" + "\Microsoft\Teams\Current\Teams.exe")
$TeamsInstallFile = $Env:AppData+"\Microsoft\Teams\InstallTime.txt"
$TeamsLogFile = $Env:AppData + "\Microsoft\Teams\Logs.txt"
[datetime]$TeamsInstallDate = Get-Content $TeamsInstallFile
$TeamsVersion = [System.Diagnostics.FileVersionInfo]::GetVersionInfo($TeamsExecutable)
$R = Get-Content $TeamsLogFile | ? {$_.Contains("ring=") }
$LastRec = $R[-1]
$SplitRec = $($LastRec) -split "<"
$TeamsLastUpdated = (Get-Date($SplitRec[0].Split("+")[0]) -format d)
Write-Host "Teams Executable:" $TeamsExecutable
Write-Host "Teams Version:" $TeamsVersion.ProductVersion
Write-Host "Installed on:" (Get-Date($TeamsInstallDate) -format d)
Write-Host "Last Updated:" $TeamsLastUpdated
$DaysSinceUpdate = ($TeamsLastUpdated | New-TimeSpan).Days
If ($DaysSinceUpdate -le 30) { Write-Host "Congratulations. It is" $DaysSinceUpdate "days since
Teams was updated. It's nice to have updated software" }
```

```
Else {Write-Host "Your Teams software is" $DaysSinceUpdate "old. Time for an update!"}
```

To see details of the environment a Teams desktop client runs in, use the command:

```
[PS] C:\> Get-Content $env:UserProfile\AppData\Roaming\Microsoft\Teams\settings.json |
ConvertFrom-Json | Select-Object Version, Ring, Environment, CloudEnvironment
```

```
version      ring      environment cloudEnvironment
-----
1.5.00.17271 ring3_6 Production
```

Apart from the version number, the important information here is the ring. *Ring3\_6* means that the client runs the preview version of Teams attached to the production service.

## Teams and VDI

VDI, or Virtual Desktop Infrastructure, is a virtualization technology that allows desktop systems to be hosted on centralized servers. If you wish to run the Teams desktop client on VDI, make sure that [you read Microsoft's guidance](#), including advice about caching of Teams content, client installation, and VDI performance. If you use Azure Virtual desktop, [separate guidance is available here](#).

## Teams Phone Application

The Teams Phone application is a special version of the Android mobile client built for device manufacturers to build [Teams-enabled devices](#). These devices allow users to call other Teams users or PSTN numbers, join a Teams meeting, and retrieve voicemail. In effect, the phone application delivers the same functionality as other mobile clients with the chat (personal and channel), files, and extensibility features removed. Only devices specifically designed to work with Teams can support the Teams Phone application.

## A Few Teams Application Tips

Unlike, say, Outlook or Word, which have been around a very long time, Teams has lots of little features that aren't well understood by many users or admins. Here are a couple of things you might want to delve into within the app to help you, and your users, get the best experience.

### Notifications

Teams can notify users when different events occur in chats and channels, such as being @mentioned in a conversation. These are global settings that apply across all teams and channels that appear in your teams list. Teams doesn't notify you of events in hidden teams or channels or when you mute a channel or conversation. You can choose default settings such as notifications for all activity or just mentions and replies or customize each source of notifications.

A user can choose to receive missed activity emails. These messages are for events that happen when the user does not sign into Teams. The available intervals range from "as soon as possible" to "once a day." You can disable notifications by setting the interval to "Off." Teams can also monitor the status of priority users and flag when they appear online or sign in.

In terms of the type of notifications, you can choose:

- **Banner and feed:** Teams signals events in both the activity feed and a desktop notification. You can choose between Teams built-in notifications or native operating system notifications (Windows and macOS). For example, if you opt for native notifications on Windows, Teams posts its notifications to the Windows notification center.
- **Only show in feed:** Show events in the activity feed.

- **Off:** Don't show these kinds of events (for example, don't show reactions to messages).

Both the Windows and macOS clients allow you to choose between the Teams notification style and notifications that use the underlying host OS notifications. This is largely a matter of personal preference.

## Switching Between Organizations

Users might have accounts in multiple tenants. The desktop and browser clients show the current tenant as a clickable link in the title bar. To switch tenants, click the link and select the target organization from the list of available tenants to which the signed-in user has access. The list includes:

- Their home tenant.
- Other Microsoft 365 tenants where the signed-in user has a guest account.
- The consumer version of Teams accessed using a Microsoft Service account (MSA).

Each tenant in the list includes the number of unread items in the user's activity feed in that organization. If the user's credentials for a tenant have expired, you see a warning sign. Unread item counts are not available for these organizations until the user reauthenticates to gain new credentials. Switching from one tenant to another terminates any open calls and closes chat windows. Teams caches credentials for each tenant to make switching easy if the credentials are still valid. Settings are specific to an account in a tenant. If someone is a guest in another tenant, they must configure settings for Teams in that tenant if they want to see the same behavior everywhere.

On mobile devices which support multiple profiles, if you have notifications enabled, you will see notifications from all the organizations you're signed into. This can cause some confusion if you tap on a notification from tenant B whilst you're logged into tenant A, because tapping the notification will cause the client to load the item of interest, which means now you're seeing something from tenant B. It's easy to switch back to the desired tenant at any time.

You should also be aware that there are many subtle (and not-so-subtle) bugs still lurking around tenant switching. For example, while writing this book I often found that switching to the Office 365 for IT Pros tenant to work, then switching back to my primary work tenant would cause some chat participants in my work tenant to be labeled "Unknown User" until I signed out and signed back in.

## Switching to Preview Versions of Teams

Teams desktop and browser clients support two preview versions. If the policies assigned to user accounts allow switching to these versions, the options appear in the **About** menu.

- The **Public preview** allows users to test new features released in public preview. The idea is that early exposure to new features assist organizations to prepare (training, help desk support, documentation) for their general availability. Control over switching to [public preview](#) is through the Teams update policy assigned to an account, which must allow the user to select public preview. Using the public preview version of Teams is broadly equivalent to running the preview channel version for Microsoft 365 apps.
- The **Developer preview** allows app developers to test their code against beta versions of Teams. Control over switching to [developer preview](#) is through the App setup policy assigned to an account, which must allow the user to upload apps to Teams.

When a user switches to a preview version, they are asked to confirm the action. If confirmation is given, Teams terminates the current client and loads the preview version. The user might be asked to reauthenticate at this point. To switch back to the current client, the user opens the About menu and unticks the selected preview. Switching between tenants will usually switch back to the release version (that is, if you've enabled the preview in tenant A and switch to tenant B, you will normally see the release version in tenants A and B

until you switch back to a preview version). Microsoft does not support the use of preview versions, but users can provide feedback to the development group if they encounter problems. Bugs are expected in preview code.

To manage the set of users allowed to select public preview, it's a good idea to use some method to mark the accounts. For example, you could use:

- A value in a mailbox custom attribute.
- Membership in a Microsoft 365 group.
- Membership in a distribution list.

Once the preview members are marked, it's easy to use a Teams bulk policy assignment job to assign the necessary update policy to their accounts. See [this post](#) for details.

From a support perspective, it's important to know that the version number reported by the desktop client does not change when the preview or developer preview is enabled. The only trace that a preview is enabled is a tick against the chosen preview shown in the About menu.

## Managing Quiet Times

Teams has multiple ways to let users know when new information has arrived from banner notifications to the activity feed to email digests. If you belong to busy teams, you might like to have some quiet time when notifications aren't delivered unless they are urgent. Desktop and browser clients obey operating system settings for notifications. For instance, on Windows 10 clients, the Focus Assist setting controls when applications can send notifications. One of the automatic Focus Assist rules stops notifications between 23:00 and 07:00, and if the rule is enabled, you won't see notifications for channel or personal messages during that period. Even when notifications are suppressed, they still accumulate in the activity feed.

The Teams mobile clients take a different approach and have their own Quiet Hours settings to control when during each day notifications are allowed and Quiet Days to control on what days notifications are accepted. For example, you can disable notifications entirely at the weekend and define that you only want to see notifications between 09:00 and 17:00 during the working week. These settings are specific to each tenant—so if you've signed into multiple tenants on mobile, you will have to set quiet hours in *each* tenant or you may still receive notifications from any tenant where you haven't done so.

Incoming calls are not governed by the notification settings. If you are signed in and available when a call arrives, you'll have to decide if you want to answer or ignore it.

## Understanding the Teams User Profile

The settings chosen by a user are in their profile. Because Teams works across platforms, Teams stores the profile online and caches some information locally. For instance, on Windows, the *desktop-config.json* file in the *%appdata%/Microsoft/Teams* folder holds information like the user principal name, tenant GUID, user account GUID, and the list of tenants where the user is a registered guest. The source of this information is a table managed by the Teams Chat service called *users/ME/properties* (a Fiddler trace will expose the interaction between a Teams client and the table). Unfortunately, no management interface is available to allow tenants to control user settings for Teams clients and it's up to each user to choose whatever settings they like.

Another file called *settings.json* stores additional configurational information for the desktop client. Teams doesn't support updates of the information held in these files in the same way as you can update the system registry to affect how other desktop clients like Outlook work.

## Using the Command Box

The command box at the top of the Teams window gives users fast access to common Teams operations. For instance, to call someone, type `/call` in the command box and then select the person you want to talk to while the `/chat` command switches to a personal chat with the selected person. Unfortunately, the command box limits user lookup to tenant accounts so you can't select a federated or guest user in the command. The command box is also a great way to update your presence with commands like `/dnd` (Do not disturb) or `/busy`.

In addition to Teams commands, if you install apps into Teams, they can become available in the command box. For example, if you install the weather app, you can type `@weather` to find out the current weather in a location, or `@news` to use the News app to look for breaking stories about a topic, or `@YouTube` to look for a video. Once you find what you need, you can copy it to the clipboard and then paste it into a conversation.

# Managing Microsoft Authenticator

You can make a good argument that reusable passwords are bad. Giving users individual passwords for access control can lessen security—if you require long complex passwords, users will reuse them or write them down, but if you don't, users will pick short, simple passwords that can easily be brute-forced. In the same vein, passwords that expire are bad (because users will just recycle old ones) but passwords that *don't* expire are bad too (because once compromised a non-expiring password might be used for a long time). However, for many years, reusable user-specific passwords were less bad than the alternatives. Thankfully, we live in a world where the widespread availability of smartphones and hardware tokens has led to the increasingly broad deployment of multi-factor authentication (a topic discussed at some length in Chapter 3). Microsoft 365 and Azure MFA support several authentication factors, including phone-based authentication, one-time codes sent via SMS, and push notifications or codes generated by a mobile application. Although Google, LastPass, and other companies make mobile authentication apps that are compatible with Microsoft MFA, Microsoft's Authenticator app is best matched with Azure and Microsoft 365 MFA.

You can use Authenticator to set up MFA for both your Microsoft account and other accounts homed in Azure Active Directory, including accounts that are federated or synchronized from on-premises Active Directory. The MFA challenge/response process occurs after, and only if, the user presents valid credentials. Authenticator can operate in two modes: it can display a code (which changes every 60 seconds) that serves as the authentication challenge, or it can display a push notification prompt that allows you to accept or reject the authentication request. Push notifications only work with Microsoft and Azure Active Directory accounts, but you can use the Authenticator app to generate codes for other apps, including Google's mobile apps and the LastPass and 1Password mobile password managers.

One interesting thing to note about the Authenticator app: just as new features tend to be delivered in the service first, with on-prem versions coming later if at all, new Authenticator features tend to accrue to the consumer Microsoft account side first, before (or if) they are rolled out to Azure Active Directory. For example, code matching for passwordless login (which I'll discuss below) was first introduced for consumer Microsoft accounts before making its way to Azure AD. Some features (such as the ability for [Authenticator to show you alerts](#) when your Microsoft account password is changed) haven't made it to the Azure AD world yet, and they may never.

## Basic MFA with Authenticator

There are several steps required to use Authenticator as an MFA client with Microsoft 365, all well-covered in the [Authenticator app documentation](#). First, of course, you must license (if necessary) and enable MFA for your tenant. Once that's done, you can enable individual users for MFA. The app is one of the available

authentication methods to enable or block individual users, but the app is available to MFA-enabled users by default. To configure it for the user, here's what to do:

1. Download the Authenticator app itself from the Apple App Store or Google Play. If you have Microsoft Intune or another similar MDM solution deployed, you can use that solution to push the app to devices as well.
2. Visit <https://myprofile.microsoft.com> and make sure that "Microsoft Authenticator - notification" is selected as the default method. Until you set up the Authenticator app on one or more devices, you'll see an error message telling you that you need to set the app up.
3. Click the **+ Add method** link, then choose "Authenticator app" and click **Add** (Figure 15-6).
4. When prompted, set up the app as directed on your mobile device. As you follow the dialogs, the web page will eventually display a QR code.
5. In the mobile app, use the **+** icon to add an account by using the device camera to take a picture of the on-screen QR code displayed in step 4.
6. You might be asked to enter a code from the app to verify that you've got it set up properly.

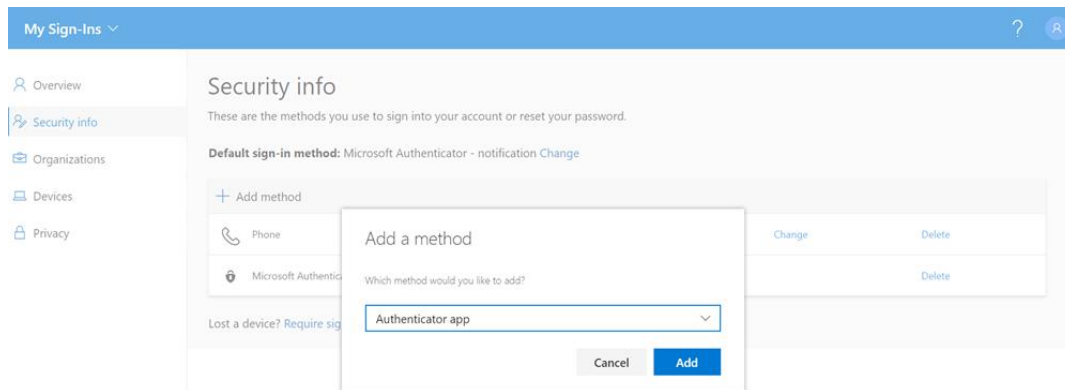


Figure 15-6: Setting up a user for MFA using the Authenticator app

Once the app is set up, when you log on to a service that requires MFA, you'll see a screen on your device like the one shown in Figure 15-7. (This version shows the authentication context features added to the service in late 2021, including the app name and rough IP-based geolocation of the authentication request.)

Authentication requests are presented on the phone as modal dialogs, so you must respond to them before doing anything else on the phone. If you've been using Authenticator for a while, you may have noticed a change made in late 2020; Microsoft used to display and allow action on, requests on the lock screen of iOS and Android devices. Now the App Lock feature of Authenticator is enabled by default, meaning that you must be logged into the Authenticator app, and have it active, before it accepts an authentication request.



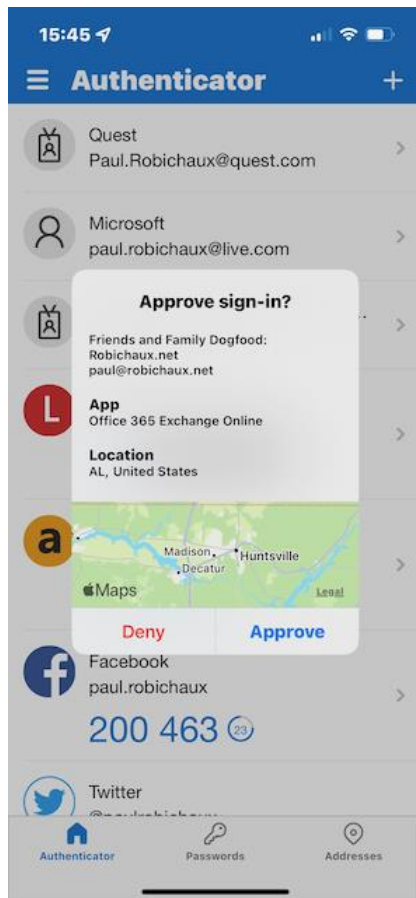


Figure 15-7: The iOS Authenticator app displays an MFA request

## Using Authenticator for Passwordless Authentication

In Chapter 3, we describe how passwordless authentication works in Azure AD. If you've enabled passwordless authentication in your tenant, you can optionally enable individual users to sign in without using a password. The way the process works is simple: when the user requests access to an app, the Microsoft authentication library will generate a two-digit challenge code and display it on screen. At the same time, the Authenticator app will prompt the user to select the matching code from a list of 3 choices. After properly selecting the matching code, she'll be prompted to use a PIN or biometric to authenticate to the device. This approach means that no user passwords are exposed at any point, nor do they need to be stored—with no password in use, there's no way for a user to accidentally compromise it, nor for an attacker to steal it.

The [instructions to set up passwordless](#) authentication for Authenticator are reasonably detailed, but the basic steps are as follows:

1. Enable the [combined security information registration](#) process in Azure AD. You do this once per tenant; note that tenants created after August 15, 2020, already have this done. In September 2022, the rest of us will have this change applied to our tenants automatically.
2. Configure Azure AD to allow passwordless sign-in with the Authenticator app. This setting applies at the tenant level.
3. Each user who wants to use passwordless authentication must add the Authenticator app as an authentication method for their account.

As of July 2022, Microsoft is rolling out the ability to have multiple accounts with passwordless authentication within Authenticator. Until this change hits your tenants, you'll only be able to add one passwordless account

to your device. The change is coming first to iOS; Microsoft promises it for Android but hasn't released a timeline.

As with all other Authenticator-based sign-in methods, if notifications don't make it to the Authenticator app the user won't be able to sign in unless they choose an alternate authentication method such as SMS.

## Authenticating from Your Wrist

If you have an Apple Watch, you can use the companion app included with the iOS Authenticator app to approve or deny authentication requests without using your phone. While this might seem like useless frippery, it's surprisingly useful because it allows you to respond to an MFA prompt without having to find and unlock your phone if it isn't already in your hand.

# Optimizing Microsoft 365 Client Network Access

Microsoft has published extensive guidance on [designing your enterprise network to work optimally](#) with Office 365. The primary principle they push is that you should design your network to minimize round-trip latency between the client endpoint (that is, the machine or device that the user's working on) and the "front door" service connection point that the client connects to. Traffic that reaches the front door will be carried over [Microsoft's internal network](#), which they manage with a ruthless focus on lowering round-trip latency.

To be more specific, Microsoft recommends that you consider the following when designing your network. First, you should know that Microsoft has separated Microsoft 365 service endpoints into [three categories](#): Optimize, Allow, and Default:

- **Optimize** endpoints represent roughly 75% of traffic to the service and include the *outlook.office365.com* and *sharepoint.com* name spaces. Microsoft says you should prioritize traffic to these endpoints by bypassing network inspection and proxies and ensuring the fastest possible DNS resolution for these namespaces.
- **Allow** endpoints, such as *protection.outlook.com*, are still important but are less sensitive to network latency. There are roughly 100 of these endpoints; because of this large number, many customers will choose not to do anything special with these endpoints to reduce the need for change management and control.
- **Default** endpoints have no special requirements for performance and can be treated as ordinary Internet traffic.

Second, Microsoft has some specific recommendations that you should apply to the design of your enterprise network:

- Don't use proxies, WAN accelerators, or other devices that sit between the client and the front door. Since you may very well need to do so anyway, see the section "Proxies and Network Devices" below for additional guidance.
- [Allow each office \(or client\) to connect directly to the Internet](#) instead of backhauling traffic from a remote office to a centralized data center and thence to the Internet.
- [Avoid network hairpins](#), or connections where a client's traffic exits your corporate network and then comes right back in again. Instead, you should always try to route all client traffic for Microsoft 365 endpoints directly to the Internet.

## Proxies and Network Devices

Client applications and servers within your corporate network that will be connecting to Microsoft 365 will need to have access to all the appropriate endpoints. The general recommendation from Microsoft is to bypass any web proxies or other network devices that shape or rewrite traffic such as WAN accelerators, for connections to the service. That recommendation is counter to the security policies of many organizations, with the adoption of cloud services making them more likely to want to push all their internet traffic through a security or performance device.

There are good reasons for customers to want to secure their internet traffic. These days, clients on their network will be connecting to more services over HTTPS, making the traffic invisible to traditional security measures. In response, organizations implement security devices that can intercept, decrypt, and inspect HTTPS traffic to detect viruses, data leakage, and ransomware. Similarly, with so much traffic now running over internet connections, organizations try to reduce the impact on their bandwidth with WAN accelerators and traffic caching devices.

Unfortunately, such measures often result in degraded performance for clients, and in some cases, they completely break connectivity to application endpoints, causing problems such as Offline Address Book download errors, OneDrive for Business sync failure, and poor audio or video performance in Teams calls. The specifics will vary depending on the vendors involved and their ability to provide guidance and templates for properly implementing their devices into your network without impacting client traffic. If you're going to attempt to implement such measures, approach it with caution and do ample testing to ensure you don't cause problems for end-users, particularly those on BYOD devices that will require special considerations for any SSL traffic inspection you plan on doing.

There may be a happy middle ground if your chosen security devices are not fit for use with Microsoft 365, but you want to continue using them to manage other internet traffic. Bypassing security devices for clients connecting to the service is a reasonable approach but can be poorly implemented if you're not careful. The natural inclination of firewall administrators is to apply controls based on IP addresses. But as Microsoft [explains in their guidance](#), Microsoft 365 IP address ranges are a constantly moving target. To keep up, you might be making changes to firewall configurations as often as every 14 days, which is a heavy burden to carry for very little real benefit. If your device supports it, Microsoft [publishes a web service](#) that your perimeter devices can query to get updates on IP address ranges and URLs for the service. Endpoint definitions are published at the start of each calendar month and take effect 30 days after their publication.

No matter the security devices at play, NATing of outbound connections to Microsoft 365 services is also a concern. Simply put, if too many clients are being NATed to the same public IP address, port exhaustion can occur which results in degraded performance for the clients. NATing is fully supported, but Microsoft recommends limiting the number of clients on one public IP to around 2000. This is mainly due to the number of connections a single client can have open. Outlook alone has 2-4 connections open for a basic user scenario, and more if they are also connecting to shared mailboxes, Microsoft 365 Groups, or are performing an OAB download. Once you add in other applications such as Teams or a web browser, the number of connections per client can easily be 10 or more. With a maximum of 64000 ports per public IP address, limiting NATing to 2000 clients per IP is a safe place to start. A pool of public IP addresses will be needed for NATing outbound connections if you are running a very large organization.

**Real World:** If you use a proxy, security, or performance device for client traffic, and you notice problems, bypassing the device is a good troubleshooting step. The goal is not to point the finger of blame at the network devices, rather it is to narrow down the likely cause of the issue. Ultimately if you choose to

implement such network devices in your environment, you are better off identifying issues and working to resolve them rather than pretending they do not exist.

## Split Tunnelling and VPNs

Some users may be required to connect to their corporate network through a virtual private network (VPN). This is especially common after the huge surge in telework brought on by the COVID-19 pandemic. For a lucky portion of the workforce, all the services they need are accessible directly through the Internet, but users who must use remote desktop solutions or access internal line-of-business systems are probably stuck with VPNs. The problem with combining VPN access and Office 365 is that the default VPN configuration will probably send all the user's network traffic from her computer over the VPN to the VPN endpoint on the corporate network, where it will be routed onwards. Traffic that's bound for the Internet (whether that's Office 365, Netflix, or something else) will thus have to traverse the VPN and then be sent onwards, with return traffic taking the same path. Microsoft calls this behavior "forced tunneling."

The solution to the performance problems caused by forced tunneling is to allow Microsoft 365 traffic to go straight to the Internet, a configuration known as *split tunneling*. When you [enable split tunneling on the VPN connection](#), traffic to select destinations is allowed to go directly to the Internet. Microsoft's guidance on how to enable split tunneling lists a set of "Optimize" endpoints that you should allow traffic interchange with over the Internet. The specifics of how you implement this are up to the specific VPN solution you use.

## Checking Office Client Network Connectivity

As described in Chapter 2, Microsoft now offers basic network connectivity monitoring in the Microsoft 365 admin center. This monitoring uses telemetry sent by the OneDrive for Business client on Windows computers to assess several aspects of your connectivity, including whether user traffic from a particular location is going to the optimal service front door location and whether there are excessive backhauls. As an example of the real-world problems this testing can catch, an early version of the connectivity tester identified that some traffic from a facility in Slovakia was being (incorrectly!) routed to Macomb, Michigan, more than 7,200 kilometers away. The unnecessary extra routing caused some additional delay for some types of traffic, but not enough delay to be routinely noticeable.

# Chapter 16: Managing Devices

## **Brian Desmond**

As workers become more mobile and security risks for corporate data increase, the management of mobile devices becomes an increasingly important factor in daily operations. Organizations have a choice of solutions for mobile device management (MDM) and mobile application management (MAM). Each of these solutions offers different features, strengths, and weaknesses.

Let's start with a quick definition to separate MDM and MAM. MDM refers to controlling the entire device, whereas MAM is concerned only with controlling the behavior of specific applications and the associated data on the device. This may seem like an obvious distinction, but there are some subtle points to it. For example, is it better to enforce a requirement to use a PIN on the entire device, or just on applications that contain corporate data? Your security team might prefer the former, but your users might prefer the latter.

Some of the considerations that come into play when planning for a mobility strategy include

- The devices and operating systems within scope.
- Who owns the target devices (BYOD versus corporate).
- The applications running on the devices (Microsoft, other vendor, and custom apps).

Some organizations can take a unified approach to mobility management, while others need to apply different policies and configurations to deal with different sets of use cases. Specific compliance requirements are also important, as some organizations fall under strict government or industry regulations.

At a high level, there are three solutions that you can choose from in Microsoft 365:

- Exchange ActiveSync (covered in the companion volume).
- Microsoft 365 Basic Mobility and Security (BMS).
- Microsoft Endpoint Manager (MEM), which includes Microsoft Intune.

In addition to the Microsoft solutions, there is an extensive range of third-party mobility solutions provided by other vendors. Here, we focus on MEM, and more specifically, Microsoft Intune.

## Comparing the Three Solutions

Before we dig into the details of Intune, let us take a quick look at the three options Microsoft offers. Consider this section somewhat of an executive summary to help you understand the basic capabilities and tradeoffs of each.

Exchange ActiveSync costs you nothing. It works on a very broad range of devices, and it offers basic device management functionality. It is only loosely integrated with the rest of Microsoft 365; it is very much Exchange-centric. It is not undergoing active development and it is increasingly being replaced by Microsoft's other solutions. We have coverage of Exchange ActiveSync in the Companion Volume.

Microsoft 365 Basic Mobility and Security (BMS) offers a [broader set of functions](#) than Exchange ActiveSync, including the ability to secure access to documents stored in OneDrive for Business and SharePoint Online. It also gives you more control over which devices you want to allow to connect, and what they can do when they do connect. There is no additional cost for BMS if you have an enterprise or business subscription. Since BMS offers, as the name implies, very basic functionality, we also have coverage in the Companion Volume.

Microsoft Endpoint Manager (MEM) is a full-fledged MDM and MAM solution. It does everything ActiveSync and BMS do, plus much more. For example, Endpoint Manager can also manage Windows PCs and macOS devices, and you can manage access to individual applications and their features even for devices that are not enrolled in Endpoint Manager for MDM. This is extremely useful in environments that have BYOD policies. MEM requires you to buy licenses, through Microsoft 365, the Enterprise Mobility + Security suite, or as a standalone license. Table 16-1 summarizes the differences between the three options.

<b>Feature</b>	<b>Licensing requirements</b>	<b>Manages data</b>	<b>Manages apps</b>	<b>Manages devices</b>
Exchange ActiveSync	Included with Exchange Online	Exchange only (online and on-premises)	No	Limited (PIN, device encryption)
Microsoft 365 Basic Mobility and Security	Requires Office 365 enterprise or business	Yes, for most Microsoft 365 workloads	Limited	More than Exchange ActiveSync
Microsoft Endpoint Manager/Intune	Requires EMS, Microsoft 365, or standalone license	Yes, including conditional access	Extensive	Extensive

Table 16-1: Choices for mobile device management

With that bit of perspective, let's dig into how you can use Microsoft Intune to manage mobile devices and apps.

## Getting Started with Intune

In this chapter, we will discuss how to use Microsoft Intune. More specifically, we will focus on managing Apple iOS/iPadOS and Android devices. Intune also supports Windows devices and macOS too, but we will not discuss those capabilities. If you are just getting started with Intune, there are a few tasks you will need to do to configure your tenant to work with Apple and Google services. You may also wish to configure branding to make the end-user experience more familiar to your users.

Unless we note otherwise, all the tasks discussed here are performed in the Microsoft Endpoint Manager admin center. You can access the admin center at <https://endpoint.microsoft.com>. You will need to be a member of the Global Administrator or Intune Administrator role in Azure AD.

### Company Portal

The Company Portal application is how you enroll most devices into Intune. If you will be doing MDM on iOS or Android devices, or MAM on Android devices, your users will need to install the Company Portal app from their device's app store. The Company Portal app is responsible for brokering Intune's MDM capabilities as well as providing a private app store, and a place to manage your device on iOS and Android. On Android, even if you choose not to use MDM, the Company Portal serves as an authentication broker. As an authentication broker, the Company Portal is responsible for enabling an integrated single sign-on experience across Microsoft's mobile apps. On iOS, Microsoft's Authenticator serves as the authentication broker.

You can customize the Company Portal to show your organization's name, logo, colors, and other branding elements. While these items are not mandatory, they do provide a familiar look and feel to end users. To customize the app, browse to **Tenant administration > Customization** in the admin center. Microsoft documents the specific requirements for branding elements, as well as certain best practices [in their documentation](#). We recommend that you work with stakeholders in your organization such as marketing

and/or internal communications teams to select branding elements that best represent your organization. Sometimes this process can take a while, so do not leave it until the last minute.

## Apple Device Enrollment

Before you can enroll your first Apple device, you must configure an important certificate in Intune. This certificate is used with the Apple Push Notification Service (APNs). The APNs does exactly what the name suggests: it lets Intune send push notifications to your devices. You can request this certificate for free from an Apple website with a few minutes of work. Once a year, you must renew the certificate. It is **extremely** important that you do not forget to renew the certificate. If you do, you must re-enroll your devices.

To setup your Apple MDM push certificate, browse to **Devices > iOS/iPadOS > iOS/iPadOS enrollment** in the admin center and click on **Apple MDM Push certificate**. The next screen takes you through the process step-by-step. Intune will provide you with a certificate signing request (CSR) in step 2 that you upload to the Apple Push Certificates Portal [linked](#) in step 3. If you have never used this portal, you must register for an Apple account first. It is very important that you do not use a personal email address for this account. Your organization's APNs certificate is permanently linked to the account you use. Instead, you should use a shared email address such as a distribution list or shared mailbox to complete the registration.

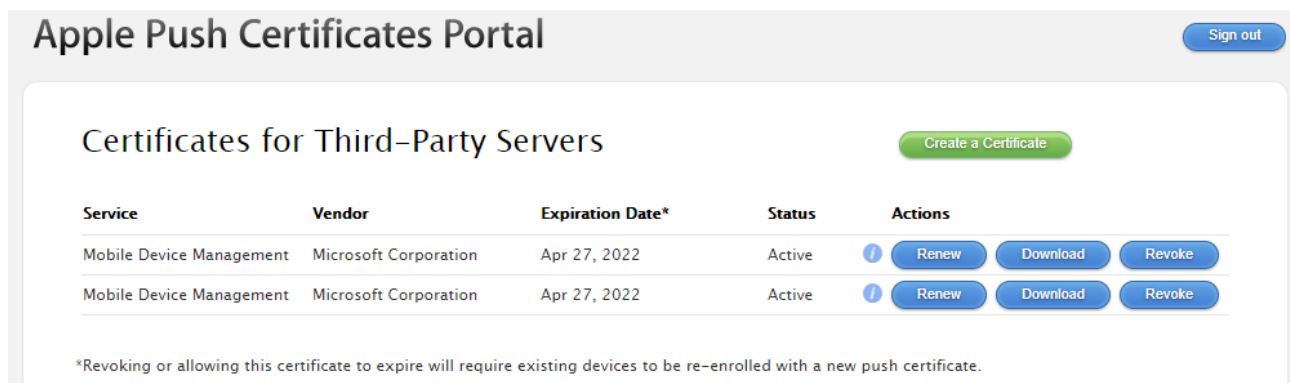


Figure 16-1 Apple Push Certificates Portal

Click **Create a Certificate** (Figure 16-1) and upload the CSR Intune generated. You will then be able to download the APNs certificate and upload it to Intune in Step 5. In Step 4, you must enter the email address of the Apple account you registered. Click **Upload** and you will be ready to enroll iOS devices using the Company Portal app.

You can also use the Apple Configuration application or Apple's Device Enrollment Program (DEP) to enroll iOS devices. However, we do not cover these topics here.

## Android Device Enrollment

While Apple makes device enrollment a universal process, Android offers a plethora of options for enrolling your devices. Android's enrollment options primarily revolve around Android Enterprise. With Android Enterprise, you have a choice of five MDM enrollment options:

1. **Android Enterprise personally-owned with a work profile** – use this option for BYOD devices that also need to access and store corporate data. Android separates corporate data and does not let the MDM system control personal data.
2. **Android Enterprise corporate-owned with a work profile** – use this option for devices your organization owns that you also allow end users to store/access personal data and apps from.

3. **Android Enterprise fully managed** – use this option for devices your organization owns that should be 100% controlled by Intune. Do not use this option if users are storing personal data on the device.
4. **Android Enterprise dedicated** – use this for devices that serve a singular purpose such as for a walk-up kiosk. Dedicated devices are restricted to a specific set of apps that can be used on the device.
5. **Android Open-Source Project (AOSP)** – use this for devices that do not interact with Google Mobile services. This feature is currently in public preview.

Regardless of the enrollment option(s) you choose, you will need to connect Intune to the Google Play store. To do this, browse to **Devices > Android > Android enrollment** in the admin center, and click on **Managed Google Play**. Click **Launch Google to connect now**. You will need to create a Google account or sign-in with an existing one. Much like Apple, you should not use a user-specific email address for this account. Use a distribution list or shared mailbox. Once you sign-in, supply your organization's name and agree to the terms and conditions and click **Confirm**. Once this is complete, you are ready to begin enrolling Android devices.

You can also enroll device using Android's legacy "device administrator" management model. While this may be necessary with certain older devices, Google is no longer investing in device administrator, and we do not recommend you use this model if it can be avoided. Microsoft has begun removing support for certain features in device administrator mode as well.

## Enrollment Restrictions

If you plan to use MDM to manage some (or all) of your devices, you may also be concerned with what devices are allowed to enroll. By default, any user in your Azure AD tenant that has been assigned an Intune license is permitted to enroll up-to five devices. You can alter the maximum number of enrolled devices per-user (up-to 15) or restrict users to only enroll certain operating systems (iOS or Android) or operating system versions, Android enrollment type (Enterprise or device administrator), Android device manufacturer, and/or block personally owned devices from enrolling.

To create or manage any of these restrictions, browse to **Devices > Enroll devices > Enrollment device limit restrictions** or **Enrollment device platform restrictions** in the admin center. The default profiles cannot be deleted, and they apply to everyone in your tenant. You can assign alternative restrictions that take precedence by clicking **Create restriction** on the toolbar. These alternatives can either be assigned to all users in your tenant, or to specific users or groups. If, for example, you need to permit certain users to participate in a BYOD program while other users can only use corporate devices, you can create two device type restriction profiles and assign them to separate groups.

## Device Categories

If you need to classify the types of devices used in the organization, device categories may be helpful. Device categories provide a list of options end users must choose from in the Company Portal when they enroll a mobile device. These can be anything you want. For example, you might create categories for departments like *Manufacturing* and *Sales* or you might categorize devices by purpose, e.g., *Point-of-Sale* and *Quality Control*. The options you choose are persisted on the device's record in Azure AD. This is helpful because you can create dynamic groups in Azure AD based on a device's category (amongst other options).

Azure AD will keep the group's membership up-to-date as the devices in your tenant change. You can then deploy Intune configuration items like apps, profiles, and policies to these groups. If you have a set of apps that are only used by Manufacturing devices, for example, you can see how this capability can be useful. To configure device categories, browse to **Devices > Device categories** in the admin center and then click **Create device category** on the toolbar.



## Device Updates and Intune

As Apple and Google evolve the iOS and Android operating systems (OS), Microsoft also evolves their minimum supported OS versions. Generally, this does not present an issue since devices typically update automatically or make an ongoing effort to ask the end-user to perform the update. However, as devices age, they may no longer be compatible with the latest OS versions.

When you have a BYOD program, and older devices are no longer supported by Intune, it can be difficult to inform end users that they must purchase a new device to continue accessing your organization's resources. Microsoft begins communicating planned OS support changes well in advance in the [Notices](#) section of the Intune docs.

For Apple devices, Microsoft supports for iOS 13 and better. For Android devices, Microsoft supports for Android 9 and better. For up-to-date support information, refer to the [supported operating systems](#) documentation. Microsoft supports the three most recent version of iOS and Android. As new versions are released, Microsoft will begin planning to deprecate support for earlier versions. While these supportability limits apply to device enrollment, individual applications (e.g., Microsoft Teams) may have different support lifecycles for iOS and Android.

## Managing Apps

Regardless of whether you choose MDM, MAM, or both, chances are you will need to manage some aspect of the apps on a device. For many organizations, Intune's MAM capabilities combined with the Microsoft Office apps (and select third-party apps) provides a good balance of control over corporate data while not putting controls on a BYOD device. If you use MDM, you will likely also use MAM to manage the apps on the device, but you might also want to use Intune to install apps on the device, or provide a curated app store in the Company Portal app.

### App Deployment

Intune has two app deployment options: "required" and "available". Depending on whether you are working in the context of MDM or MAM, the options you must install apps on mobile devices varies. If a device is MDM enrolled, you can use either option. When an app is deployed as "required", it will automatically be installed on the device. You can also choose to require uninstallation of an app you have previously deployed as "required". On the other hand, "available" publishes the app in the Company Portal. End users can click on "available" apps in the Company Portal, and they will be installed on their device.

Whether you choose required or available, you need to target your app deployment to a group of users or devices. Intune uses Azure AD groups so you can either use an existing group, or you can create a new one. If you simply want to target an app deployment to everyone, you can choose "all devices" or "all users". Be careful if you try to mix groups containing users and devices. Intune does not try to resolve the relationship between users and their devices so you may not get the results you want. As a best practice, we recommend that you stick with assigning to users *or* devices, but not at the same time.

To deploy an app, browse to **Apps > All apps** in the admin center. Click **Add** on the toolbar and choose **Android store app** or **iOS store app**. In this example, we will create a deployment for a Microsoft Outlook for Android. The process is very similar for iOS. You will need the URL of the application from the Google Play store which you can find by searching the store at <https://play.google.com/store/apps/>. Fill in the details for the application as shown in Figure 16-2. You can copy and paste all the details from the app's store listing.

While you do not need to provide optional details like the category or logo, they greatly enhance the user experience in the Company Portal. You can obtain the logo by saving the image in the app store listing as a PNG file and then uploading it to Intune. The category is used to help end users find apps that are available for installation in the Company Portal. There are nine built-in categories that come with Intune. You can add your own categories by browsing to **Apps > App categories** in the admin center.

## Add App ...

Android store app

1 App information   2 Assignments   3 Review + create

Name \* ⓘ

Description \* ⓘ

Publisher \* ⓘ

Appstore URL \* ⓘ

Minimum operating system \* ⓘ

Category ⓘ

Show this as a featured app in the Company Portal ⓘ

Information URL ⓘ

Privacy URL ⓘ

Developer ⓘ

Owner ⓘ

Notes ⓘ

Logo ⓘ [Change image](#)




Figure 16-2 App deployment for Microsoft Outlook for Android

Once you configure the app information, you can assign the app to users or devices. When you create an app assignment, you have several options to choose from. The assignment options are very similar across everything you will do in Intune. Fundamentally, you can assign apps to groups of users and/or devices by setting the **Group mode to Included**. If you need to exclude certain users or devices, you can change the **Group mode to Excluded**. Make sure you do not combine groups of users and devices, or you might see unexpected results. If, for example, you include a group of users and exclude a group of devices, Intune will not filter those devices as you would expect. If you need to only include or exclude certain devices in combination with users, use a Filter. We discuss Filters later.

For apps, you have a choice of assigning the app as **Required**, **Available for enrolled devices**, or **Available with or without enrollment**. Apps assigned as required will be automatically installed on devices. Apps assigned as available will be displayed in the Company Portal app for end users to install. The choice of whether to make apps available only to enrolled devices (MDM) or with or without enrollment (MDM or MAM) gives you additional control over when users can select the app for installation.

If you are not ready to deploy the app, this step is optional. If you do list one or more groups, the deployment will begin as soon as you click **Create** and complete the wizard. Find your app under **Apps > All apps** and click the app to monitor deployment status, change the properties of the app, or modify what users or devices the app is deployed to.

You can use the **App install status** report under **Apps > Monitor** to keep an eye on your app deployments. From here you can investigate the progress of a specific app deployment and see installation failures at-a-glance. To investigate a specific app, click on the row to go to the status dashboard for that app.

### Bulk Purchased and Private Apps

Apple and Google support managed app procurement as well as the ability for app publishers to advertise private apps to select customers. Apple manages this process through the Apple Business (or School) Manager, while Google uses the managed Google Play store. With both Apple and Google, once you approve or make a bulk purchase of an app, Intune will allow you to synchronize your app library and the apps will become available for deployment using the same process as we described earlier in this section.

The process with Android is straight forward. Browse to **Tenant administration > Connectors and tokens > Managed Google Play**. Click the **Open the managed Google Play store** link and approve the app(s) you want to deploy with Intune. Next, click the **Sync** button in the admin center. The apps you approved in the managed Google Play store will appear under **Apps > Android**. You can assign these apps the same way you assign any other app.

With Apple, you must enrol for an account in Apple Business Manager. This process requires several steps and data about your organization. Apple documents this process [here](#). Once you have obtained an Apple Business Manager account, you can browse to <https://business.apple.com> to login. Next, you must connect Intune to Apple Business Manager. To do this, go to **Settings > Apps and Books**. Click the **Download** link under *My Server Tokens*.

Next, in the Microsoft Endpoint Manager admin center, browse to **Tenant administration > Connectors and tokens > Apple VPP Tokens** and click **Create** on the toolbar. Give the token a name, provide the Apple ID you used to login to Apple Business Manager, and upload the file you downloaded in the previous step. Intune will automatically synchronize with Apple Business Manager the first time, but you can manually synchronize by selecting the token clicking the ellipsis (...) and clicking **Sync**.

To make apps available for Intune in Apple Business Manager, you must buy licenses (even if they are free) for the app in the **Apps and Books** area of Apple Business Manager. When you buy the license, you must be sure to assign it to the location that you downloaded the server token for in the previous step. If you have multiple locations configured, you must add a token for each location to Intune. If you are attempting to work with a private app and you do not see a Custom Apps area, go to **Settings > Enrollment Information** in Apple Business Manager and configure **Custom Apps** to **Enabled**.

Once you have bought licenses for an app, synchronize the token in Intune, and you will find the apps under **Apps > iOS/iPadOS**. You can assign these apps the same way you assign any other app. Occasionally Apple updates their terms and conditions. When this happens, the sync between Intune and Apple Business

Manager will stop working until you sign-in to Apple Business Manager and accept the updated terms and conditions.

## App Protection Policies

MAM is a core component of managing your organization's data on mobile devices. With MAM and Intune's app protection policies, you can apply controls to your data in apps that support MAM while leaving personal data untouched. This often creates a good balance of control versus autonomy in a BYOD program. MAM is primarily supported by Microsoft's Office applications on iOS and Android, but Microsoft publishes an SDK and an app wrapping tool that let third parties and in-house developers support MAM too., Zoom, Adobe Acrobat Reader, and Box are three examples of third-party apps that support MAM.

To implement MAM, you use app protection policies. These policies define what users and the supported mobile apps can do with your organization's data. For example, you can allow copy and paste between managed apps, but, if a user tries to copy managed data from Microsoft Word to Gmail, you can block this. You can also protect access to your data in managed apps by requiring a PIN to access the app rather than requiring a device PIN. The settings and capabilities are slightly different between iOS and Android, but fundamentally, the concepts are the same.

To create an app protection policy, browse to **Apps > App protection policies** in the admin center, click **Create policy** and select **iOS/iPadOS** or **Android**. For this example, we will create an iOS policy. Once you give the policy a name, you will need to select the devices and apps that the policy applies to. On the **Apps** tab, you can either select specific apps to apply the policy to, or use curated collections maintained by Microsoft. In the Target Policy to menu, you can choose from the following options:

- **All Apps** – automatically apply the policy to all publicly available apps that support app protection. Microsoft maintains [a list in the documentation](#).
- **All Microsoft Apps** – automatically apply the policy to any Microsoft-published app that supports app protection.
- **Core Microsoft Apps** – automatically apply the policy to a small set of "core" Microsoft-published apps including the Office apps, Microsoft Teams, and Edge.
- **Selected apps** – allows you choose the specific apps the policy applies to.

If you use any of the curated collections, Microsoft will update them from time-to-time as new apps become available that support app protection.

If you want to select specific apps, configure **Target Policy to** **Selected apps** and use the **Select public apps** link to target your policy to apps available in the app store. If there is an app that is available in the app store but not listed, or you develop an app that integrates with Intune, use the **Select custom apps** link to target the policy to the app.

Unlike app deployments, app protection policies can only be targeted to users. If you have a mixture of MAM and MDM devices, you might want to have one app protection policy for MAM-only devices that requires a PIN to access managed apps, while on MDM devices, you might require a PIN to access the device, rendering a second PIN to access managed apps unnecessary. If you need this level of control, change **Target to apps on all device types** to **No** and select **managed** (MDM) or **unmanaged** (MAM) from the **Device types** menu.

Next, select the apps you want to deploy the policy to. For traditional store apps, click **Select public apps** and select the apps you want to manage. If you have in-house developed apps that you want to manage, or MAM-capable apps that are not listed under **Select public apps**, click **Select custom apps**.

The bulk of the MAM controls are configured on the **Data protection** screen:

- **Backup org data to iTunes and iCloud backups** – controls whether your organization’s data can be backed up to Apple’s iTunes/iCloud services. We generally recommend selecting **Block**.
- **Send org data to other apps** – what apps on the device can the end user transfer data between using iOS’ sharing tools. We generally recommend selecting **Policy managed apps** or **Policy managed apps with Open-In/Share filtering**. These controls allow the end user to move organizational data between other MAM-managed apps, but not to unmanaged apps. If you have a special case that isn’t MAM-managed, you can use **Select apps to exempt** to enable sharing.
- **Save copies of org data** – can the user “save as” to an alternate location. The user interface (UI) for this control is confusing. You probably will want to select **Block** and then modify the **Allow user to save copies to selected services** and enable locations like OneDrive and SharePoint, Local Storage, or Box.
- **Transfer telecommunications data to** – controls whether the user can click on a phone number in an email, for example, to dial the number. By default, **Any dialer app** lets the user use any calling application (whether the native Phone app or otherwise), but, you can restrict this if you need to.
- **Receive data from other apps** – indicates whether managed apps can receive data from other apps. If you select the default, **All apps**, a user can transfer data from their personal email to their corporate email (for example). Like the **Save copies of org data** setting, if you want to choose specific data sources, you must select **Block**.
- **Restrict cut, copy, and paste between other apps** – this is an important data sharing control that determines whether an end user can copy organizational data and paste it into an unmanaged app, or vice-versa. We generally recommend selecting **Policy managed apps with paste in**. This enables full cut/copy/paste between managed apps, but also enables an end user to copy *from* an unmanaged app and paste it into a managed app.
- **Third party keyboards** – iOS supports alternative keyboards that can be installed from the iTunes Store. These keyboard apps can present a security risk since they have access to every keystroke.
- **Encryption** – enforces encryption of organizational data on the device. The downside of this is that it requires the user to have a device PIN which you may not wish to enforce.
- **Sync policy managed app data with native apps** – this is a roundabout way of saying if Outlook Mobile, for example, can add contacts to the native contacts app. There are more elaborate controls specific to Outlook Mobile that you should use if this setting is of interest.
- **Printing org data** – as the name implies, can an end-user use the native iOS printing functionality inside a managed app?
- **Restrict web content transfer with other apps** – if you want to require Microsoft Edge for browsing or to bring data into managed apps, you can use this control.
- **Org data notifications** – controls whether managed apps can create notifications on the device (or a companion device like an Apple Watch). These notifications can potentially disclose sensitive data such as in the subject of a meeting invitation.

Next, you can further restrict access to managed apps in the **Access requirements** tab by requiring the user to enter a PIN or biometric (e.g., TouchID) and/or login with their corporate credentials (the **Work or school account credentials for access** setting). Typically we find that these settings are useful if you are not using MDM. Without MDM, you cannot guarantee that users have setup security to unlock their device. By using the settings on the Access requirements tab, you can implement equivalent security controls to launch apps protected by the policy.

You may also wish to apply additional security precautions on the **Conditional launch** screen. These controls are divided into two sections: **App conditions** and **Device conditions**. App conditions control the behavior of managed apps (e.g., has the device been offline for too long, has the user entered too many invalid PINs, or

what to do if their account disabled in Azure AD) and what happens (e.g., restricting access or wiping data). Device conditions inspect the status and health of the device (e.g., is the device jailbroken, is it running a minimum OS version, and/or is it at an acceptable risk level as reported by a [mobile threat defense \(MTD\) solution](#)). Like app conditions, device conditions can take access such as blocking access to data or wiping the data entirely.

Finally, you must assign the policy to one or more groups of users on the **Assignments** tab. Once this is complete, click **Create** and complete the wizard. You can come back to the Apps > App protection policies screen later to monitor the status of the policy or make changes. New apps that support MAM periodically become available so you should occasionally check that you are including all the desired apps in the policy.

Once you deploy your policy, use the **App protection status** report under **Apps > Monitor** to keep track of your policy as shown in Figure 16-3. The dashboard in this report gives you an at-a-glance view of how many users are protected by the policy, the top apps, and any errors that you should investigate. To get more detail, you can use the download links on the toolbar, such as **App protection report: iOS, Android** to export the data to a CSV file.




Figure 16-3 App protection status report

## Email Client Configuration

Generally, we recommend choosing Outlook Mobile as the mobile email client for Exchange Online whenever possible. It provides the most complete combination of Microsoft 365 integration and administrative control. However, you may need to support the native email client on mobile devices as well. iOS in particular supports complete configuration of a profile in the native email client for MDM enrolled devices. Android devices support relatively limited pre-configuration of the Gmail and Nine for Work email apps. Outlook Mobile for iOS and Android can be pre-configured regardless of whether you use MDM or MAM.

### iOS Native Mail Client

To deploy settings for the native mail client to MDM enrolled iOS devices, you will need to create an email configuration profile. To begin, browse to **Devices > Configuration profiles** in the admin center and click **Create profile**. Select **iOS/iPadOS** for the **Platform** and **Templates** for the **Profile type**, and **Email** for the **Template name** and then click **Create**. You will need to complete the following fields to configure the profile:

- **Email server** – the ActiveSync endpoint to connect to. For Office 365 in the commercial cloud, use **outlook.office365.com**.
- **Account name** – this is shown in the iOS mail app as the name of the email account. You could use the name of your organization, for example.
- **Username attribute from AAD** – Select the user's username as represented in Azure AD. Generally this is the User Principal Name (UPN), but if you are connecting to an on-premises Exchange server, you might need to pick sAMAccountName, for example.
- **Email address attribute from AAD** – Select the user's primary email address as represented in Azure AD. If your UPN does not match the primary email address, then you will need to select Primary SMTP Address.
- **Authentication method** – select Username and password unless you are using client certificate authentication.
- **SSL** – select Enable.
- **OAuth** – this enables modern authentication. For Office 365, and any on-premises Exchange organization using modern authentication, select Enable.

**Note:** Microsoft will remove support for Basic Authentication in Exchange Online. You should only deploy profiles with modern authentication to prepare for this change. If you are using Basic Mobility and Security, discussed in the Companion Volume, Microsoft will automatically update your email profiles to use modern authentication.

- **Exchange data to sync** – generally you will select All data, however you can filter this to a subset. You might want to combine the **Calendar only** option with Outlook Mobile to provide access to the native Calendar app while using Outlook mobile for email and contacts, for example.

The remaining choices control behavior of the native mail app and are not critical for deploying an email profile. Like deploying an app, you must target one or more Azure AD groups of users or devices on the **Assignments** screen. Once this is complete, you can click **Create** to complete the wizard. You can monitor the status of your email profile deployment as well as modify settings and assignments by browsing to the profile under **Devices > Configuration profiles**.

One important thing to note is what will happen if you deploy an email profile to a device that already has a matching profile as determined by the Email server and Account name fields. Intune will not allow the profile to be installed. Instead, the end user will receive a notification that they must delete the profile that is already installed on their device. Intune will subsequently install the profile you have configured.

## Outlook Mobile

You can configure Outlook Mobile much more granularly than the native email client apps. In conjunction with an app protection policy, Outlook Mobile gives you a great deal of control over how the app behaves and what end users can do with your data. To configure Outlook Mobile, browse to **Apps > App configuration policies** in the admin center and click **Add > Managed apps**. Give your policy a name and click **Select public apps**. Select **Microsoft Outlook** for iOS and/or Android and then click to go to the **Settings** screen.

Microsoft documents all the Outlook Mobile configuration settings [here](#). Most of these settings are listed in Intune under **Outlook configuration settings** and are self-explanatory. Sometimes, though, Outlook Mobile introduces new settings before Intune catches up. You can use the **General configuration settings** in Intune to set these settings. For example, if you want to prevent end-users from adding a personal email account to Outlook Mobile, add a row under General configuration settings that looks like Table 16-2.

<b>Name</b>	<b>Value</b>
IntuneMAMAllowedAccountsOnly	Enabled

Table 16-2 Outlook Mobile custom configuration settings

All the settings for Outlook Mobile are documented and can be configured like the example in Table 16-2. Outlook Mobile custom configuration settings, however most are also listed in a much more friendly manner under **Outlook configuration settings**. Before you use General configuration settings, take a minute to confirm that is the only way. The raw settings are also useful if you are using a third-party MDM solution.

## App Selective Wipe

One of the benefits of MAM is the ability for Intune administrators to wipe organizational data without touching a user's personal data. This is useful when a user leaves the organization, for example. You can target a selective wipe to a specific device or all the user's devices. Once a selective wipe request is created for a user, it will remain until you delete the request.

To create a selective wipe request, browse to **Apps > App selective wipe** in the admin center. To selectively wipe organizational data on specific device, click **Create wipe request** on the toolbar. To selectively wipe organizational data on *all* a user's devices, open the **User-Level Wipe** tab and then click the **Add** hyperlink at the bottom of the list.

Select a user and/or their devices to create the request. You can open the request and monitor the progress of the wipe request for each managed app on each of their devices.

## App Configuration

Some applications support managing how the application works through Intune. For example, Outlook Mobile provides this capability as discussed above. Microsoft Edge is another example. These capabilities are not specific to Microsoft, however. You might need to specify a server for a line-of-business application to connect to, or you might need to change how the application behaves so that it can correctly interact with conditional access policies.

The availability of settings you can configure is completely up to the application developer. One common example is Salesforce. Salesforce [documents](#) a variety of settings that you can apply to their app through an MDM service like Intune. To do this, you must create an app configuration policy for each platform that you will deploy the app to.



To create an app configuration policy, browse to **Apps > App configuration policies** in the admin center and then click **Add > Managed devices** on the toolbar. Select the platform you are targeting (iOS or Android), and the app you will configure settings for. You must select an app that you have previously deployed in Intune. Next, you can choose how you enter the settings. In the **Configuration settings format** menu, **Use configuration designer** gives you a simple GUI for data entry, while **Enter XML data** is useful if the app developer has provided pre-configured settings that you can paste into Intune.

Finally, you must assign the app configuration policy to users or devices. The settings will be applied to the targeted app on any device that has the app installed.

## Managing Devices

Now that we have discussed managing apps on devices, what if you want to also manage the device itself? To do this, you will use the MDM capabilities of Intune. With MDM, you have substantial control of every facet of the device, limited only by what iOS and Android let you do. There are countless scenarios for how you deploy MDM, and thousands of settings that you can control on devices. We will not dive into all these settings or scenarios, but we will show you how to get started.

### Configuration Profiles

You use configuration profiles to control the behavior of devices, as well as install organizational settings such as wireless networks, virtual private networks (VPNs) or PKI certificates. Configuration profiles are specific to device platforms (iOS or Android) and must be configured individually for each. To create a configuration profile, browse to **Devices > Configuration profiles** in the admin center and click **Create profile** on the toolbar. You will need to select a **Platform** (e.g. iOS/iPadOS), **Template**, and **Profile type**. Most of the time, we expect you will use the **Device features** and/or **Device restrictions** template.

As you explore the configuration settings available in your profile, pay attention to the notes and information tips that Intune displays. Many settings are very specific about when and how they apply. For example, some iOS settings apply to any MDM enrolled device, while others only apply to *supervised* devices that are enrolled through Apple's automated device enrollment program.

Once you have configured the profile as you desire, assign it to one or more groups of users and/or devices. You can combine groups of users and devices and even exclude certain groups. As a reminder, be very careful if you decide to combine user and device groups in your profile assignment! You may find that it does not work the way you expect. If you need to only include or exclude certain devices in combination with users, use a Filter. Like other assignments, you can come back to your configuration profile later to monitor the status of the deployment or make changes.

### Device Inventory and Actions

Intune collects many data points about enrolled devices and makes it available to you in the properties of the device. You can learn about device hardware and installed apps, explore the status of configuration profile deployments, and check on app deployments. All this information is available by browsing to **Devices > All devices** in the admin center and clicking on a specific device. You can also use Intune's reporting features, which we will discuss later, to learn about this information for many devices at once.

In the same screen, you can also remotely take a limited number of useful actions on MDM enrolled devices. These actions are all accessible by. Actions you can take include:

- **Retire** – wipes managed app data and settings and removes the device from Intune. The user's personal data is retained.
- **Wipe** – removes all data from the device, resets the operating system settings, and unenrolls the device from Intune. Optionally, you can select to preserve user data on the device.
- **Remote Lock** – returns the device to the PIN lock screen.
- **Remove Passcode** – removes the device passcode, such as in scenarios where the user has forgotten their PIN.

These actions are primarily useful when someone leaves the organization or if a device is lost or stolen. From the **Devices > All devices** screen you can also perform these actions in bulk on many devices by clicking **Bulk device actions** on the toolbar. Be very careful using this tool! You could accidentally wipe many devices at once.

## Storing Custom Device Properties

Intune and Azure AD track various properties of devices such as their model, operating system version, etc. You might want to keep track of information that is specific to your organization too. For example, what department does a device belong to, or is the device approved for specific uses? To do this, you can use device extension attributes. You can use device extension attributes to populate dynamic groups or in device filters that are applied to a conditional access policy.

To populate device extension attributes, you must use the Graph API. Microsoft does not currently make an editor available in the admin center. You can use the Microsoft Graph [PowerShell module](#), the [Graph Explorer](#), or a custom script or program. In the example below, we will use the Graph Explorer to populate extensionAttribute1 with the value 'Finance Department'.

1. Obtain the Object ID of the device you want to update by logging in to the Azure Portal and navigating to **Azure Active Directory**. From there, select **Devices**, find the device you are planning to update, and click **Properties**.
2. Browse to the Graph Explorer at <https://developer.microsoft.com/en-us/graph/graph-explorer>.
3. Click **Sign in to Graph Explorer** on the left side of the screen.
4. Click **Modify permissions (preview)** and then click **Consent** for the *Device.ReadWrite.All* permission.
5. Configure your query with the following parameters:
  - a. **Method** – Patch
  - b. **Version** – v1.0
  - c. **URL** - `https://graph.microsoft.com/v1.0/devices/<device Object ID>` **URL** - `https://graph.microsoft.com/v1.0/devices/<device Object ID>`

6. Configure the body of the query with the following JSON:

```
{
  "extensionAttributes": {
    "extensionAttribute1": "Finance Department"
  }
}
```

7. Click **Run query**. A response of "No Content – 204" indicates success. You can modify Step 5 to use the Get Method to confirm the results of your update.

For more information on using Device Filters or dynamic groups, refer to Chapter 3.

## Custom Push Notifications

Intune can send push notifications with custom text to all the iOS and Android devices enrolled by a user. This can be useful for sending emergency alerts, such as for a weather event or a security situation. Custom notifications are shown on the lock screen of the phone or tablet so you should be careful not to include sensitive information. You can send notifications to individual users or to groups of users. Microsoft restricts you to sending up to 25 messages per-hour to groups, and up-to 10 messages per-hour to a specific user's devices.

To send a custom push notification, browse to **Tenant administration > Custom notifications** in the admin center. Give your notification a title (up-to 50 characters in length) and a body (up-to 500 characters in length). The title and body appear on the lock screen of the device when the notification is received. On the Assignments tab, select individual users or groups of users to send the notification to. Note that if the group includes devices, those devices will only receive the notification if the device's owner is *also included* in the group. Once you complete the wizard, Intune will immediately send the custom notification to the targeted users. It is not possible to cancel the notification or track its status.

## Security by Compliance

The configuration profiles and policies you deploy with Intune can be used as signals for access control decisions in Azure AD. This integration makes Intune even more powerful and an important part of your toolkit if you are implementing a zero-trust architecture. Whether you decide to use MAM, MDM, or both, you can make powerful decisions about whether a user and their device have access to applications or data using signals from Intune.

## Compliance Policies

If you use MDM, you can define policies that declare devices as compliant or non-compliant. Device compliance is determined by measuring certain settings on the device as well as risk indicators from MTD solutions. Based on the results of this evaluation, Intune sets a flag on the device's object in Azure AD that indicates if the device is compliant (or not). Subsequently, Azure AD can use this compliance flag to influence access control.

Like configuration profiles, compliance policies are created on a per-platform basis. To create a compliance policy, browse to **Devices > Compliance policies** in the admin center and click **Create Policy** on the toolbar. Select a **Platform** and give your policy a name. In this example, we will create a compliance policy for iOS. You'll find the settings you can choose from are relatively limited in comparison to a configuration profile, but the settings you can choose from are typically the most important indicators of a device's security posture:

- **Email > Unable to set up email on the device** – if you set this to require, the device must have an iOS Email configuration profile deployed to it that is successfully installed. If you are not using the native email app, this setting is not useful.
- **Device Health > Jailbroken devices** – this setting lets you mark jailbroken devices as non-compliant. We typically recommend that you select **Block**.
- **Device Health > Require the device to be at or under the Device Threat Level** – if you are using a mobile threat defense solution, this integrates the risk score from the MTD into the compliance indicator.
- **Device Properties > Operating System Version** – use these settings to require minimum (or maximum) versions of the device OS for the device to be considered compliant.

- **Microsoft Defender for Endpoint** – if you are deploying Defender for Endpoint to your mobile devices, this integrates the device’s risk score from Defender for Endpoint into the compliance indicator.
- **System Security > Password** – use this setting to require the device to have a PIN or passcode. We generally recommend you set this to Require. You can subsequently configure more specific details of the PIN such as its length or how often it must be changed.
- **System Security > Device Security** – list apps here that you want to block from mobile devices. If a listed app is installed, the device will be considered non-compliant.

When a device is considered non-compliant you can configure what Intune should do. By default, Intune will immediately mark the device as non-compliant which may have adverse consequences for the end-user’s access to organizational apps and data. You can take a more measured approach on the **Actions for noncompliance screen**:

- **Mark device noncompliant** – after how many days should Intune mark the device as noncompliant in Azure AD?
- **Send email to end user** – after how many days should Intune send an email to the end user (and optionally to others such as your helpdesk) informing them about their device?
- **Send push notification to end user** – after how many days should Intune send a push notification to the end user informing them about their device?
- **Remotely lock the noncompliant device** – after how many days should Intune lock the device, requiring the device’s PIN/passcode to be entered?
- **Retire the noncompliant device** – after how many days of noncompliance should Intune take the device retirement action, discussed earlier.

You might decide to first send the user an email immediately, wait three days to mark the device as non-compliant in Intune, and then retire the device if it is still non-compliant after 90 days. On the other hand, your security requirements might not allow a non-compliant device for any period so you might decide to both send an email and mark the device as non-compliant immediately (after 0 days).

Once you have configured the policy, you must assign it to groups of users or devices in the same way as you have assigned other profiles or apps. It is important to remember that compliance policies only *test* settings. They never configure settings on a device.

## Conditional Access

Conditional access policies are a feature of Azure AD Premium that let you make decisions about the who, what, when, where, and how of access to apps and data. We discuss the conditional access feature of Azure AD in detail in Chapter 3, including various examples of how to configure policies that integrate information from Intune. A classic example is organizations that want to make sure email is only accessed from devices that are either enrolled in Intune and policy compliant, or from end-users that are using a managed app like Outlook Mobile.

You can use two indicators from Intune to achieve these goals (and we show you how in Chapter 3). Since the compliance status of a device is reflected in Azure AD based on the results of your compliance policies, you can require a device to be compliant in your conditional access policy. Likewise, conditional access policies can determine if a user is using a managed app (e.g., Outlook Mobile) and if an app protection policy manages the app. If these factors are met, Azure AD will permit access. Otherwise, it can be configured to deny access.

## Terms of Use

Some organizations require end users to accept terms and conditions before they can enroll a device in Intune. Common examples are acceptable use policies for organization-owned devices and agreements for how personal mobile devices can be used to access organization data. There are two ways you can accomplish this with Intune:

1. Intune terms and conditions
2. Azure AD conditional access terms of use

We recommend using Azure AD's terms of use if you have Azure AD Premium. Azure AD terms of use support a much broader set of features including uploading PDFs that can have formatting and hyperlinks, multiple languages, requiring users to scroll through the entire terms document, and more.

You must first create a terms of use document in Azure AD. To do this, login to the Azure Portal and navigate to **Azure Active Directory**. From there, select **Security > Conditional Access > Terms of use** and click **New terms** on the toolbar. Configure the terms with the following settings:

- **Name** – a name for the terms that is useful to administrators. This value is not shown to end users.
- **Terms of use document** – upload a PDF that will be shown to end users for acceptance. You can upload additional PDFs in alternate languages. You must specify the **default language** for the first PDF you upload and provide a **Display name** that will be displayed to end users.
- **Require users to expand the terms of use** – if you select **On**, users must view the contents of the PDF before they can accept the terms.
- **Require users to consent on every device** – if you select **On**, users will be prompted to accept these terms on every device rather than just once.
- **Expire consents** – if you select **On**, users will be forced to re-accept the terms after the number of days specified in the **Duration before re-acceptance required** field.
- **Enforce with conditional access policy templates** – choose **Create conditional access policy later**.

Next, you will need to create a conditional access policy to apply the terms you created to Intune device enrollment. We cover conditional access policies in depth in Chapter 3. The policy that you create must be configured with the following minimum settings:

- **Cloud apps or actions** – choose **Select apps** and then select the **Microsoft Intune** and **Microsoft Intune Enrollment** apps.
- **Grant** – check the box next to the name of the terms of use you created previously.

Auditing information for when users accept terms of use is accessible from the **Audit logs** area of the Azure AD portal. Click the **Activity** filter below the toolbar and select **Accept Terms of Use** to filter the audit data.

## Intune Management

As you plan to use Intune in your organization, you will need to determine how to operationalize management of the service. In a small organization, this may be simple, but in larger organizations, you will probably need to grant access to different teams such as your service desk or endpoint management team to manage different parts of your Intune configuration. You may also need to give management insight into how the service is being used and what your mobile device and application footprint looks like. In this section we will explore some of these requirements and how you can accomplish them with Intune.

## Privileged Access and Role Based Access Control

Members of the Global administrators and Intune administrator roles in Azure AD are automatically granted full access to Intune. While these roles grant broad access, you probably will want to delegate more granular access to administrators in your IT organization. There are several [built-in roles](#) that you can start with by browsing to **Tenant administration > Roles > All roles** in the admin center.

If your organization is more complex, you can define custom roles to meet your specific needs. Custom roles define two components: the permissions the grant and the objects (e.g., devices, policies, apps, profiles, etc.) the permission are applicable to. There are hundreds of individual permissions that you can explore as you create a role. The objects that the role grants access to use a capability called scope tags. We have not discussed scope tags thus far, but you may have noticed that you can set them anytime you created a policy, profile, or app.

If you do not need to segment access to individual devices or policies/profiles/apps, you can use the built-in Default scope tag. If you do need to segment access, scope tags are for you. You might want to create individual scope tags for different departments, regions or geographies, or categories of users (e.g., executives). Scope tags are assigned to devices based on group membership. You can either put devices in a group manually, or you can use Azure AD dynamic groups to automatically place devices in certain groups. All other objects in Intune (e.g., configuration profiles, apps, etc.) are manually assigned one or more scope tags.

To create a scope tag, browse to **Tenant administration > Roles > Scope (tags)** in the admin center. Click **Create** on the toolbar and give your new scope tag a name and description. You can choose anything you want – for example “Headquarters Devices” or “Executives”. If the tag will apply only to certain devices, select a group that meets your needs.

Next, when you create a custom role by browsing to **Tenant administration > Roles > All roles** in the admin center and clicking **Create**, you can choose one or more scope tags in the wizard. The permissions you select will *only* be granted to items that are tagged with the scope tag(s) you select. To apply a scope tag to almost anything in Intune, go to the Properties of that item and click **Edit** in the **Scope tags** section of what you are editing.

## Complex Assignments with Filters

One of the limitations of traditional assignments in Intune is the inability to combine inclusion and exclusion groups that mix users and devices. For example, you cannot include a group of users in a specific department but then exclude a group of devices that are personally owned. Filters allow you to create reusable rules for when an assignment should (or should not) apply to certain devices.

Filters can be used with most types of assignments, but there are a [few exceptions](#). To create a filter, browse to **Tenant administration > Filters** in the admin center and then click **Create** on the toolbar. Give your filter a name and select the **Platform** that it will apply to. You can either manually specify the filter in the **Rule syntax** editor or use the expression builder. For example, if you wanted to create a filter that targets only devices that are personally owned, configure the following expression:

- **Property** – deviceOwnership
- **Operator** – Equals
- **Value** – Personal

Under the **Rule syntax** editor, there is a **Preview devices** link. Clicking this link allows you to test the results of the filter in real time before you deploy it.

After creating the filter, you can apply it to an assignment for an App, Configuration Profile, Enrollment Restriction, or Compliance Policy. You will configure the **Filter mode** to **Include** or **Exclude** in the assignment and select the filter you previously created.

## Simplifying Assignments

Throughout the examples described here, we assign various policies, apps, and profiles to groups. As your Intune deployment grows in scale and complexity, this can become difficult to accurately manage. To address this, Intune has a feature called *Policy Sets*. You can use policy sets to create a single assignment for any combination of apps, app configuration policies, app protection policies, configuration profiles, compliance policies, device type restrictions, and more. For example, you might create a policy set for all of your users in North America and assign everything they need at once.

To create a policy set, browse to **Devices > Policy sets > Policy sets** (or **Apps > Policy sets > Policy sets**) in the admin center and click **Create** on the toolbar. Give your policy set a name and then click through the **Application management** (apps, app configuration policies, and app protection policies) and **Device management** (configuration profiles and compliance policies) tabs. Add the relevant items that you want to deploy on each tab. Finally, on the assignments tab, add the groups of users or devices that you want to include (and, optionally, exclude) from the deployment. Note, again, that you cannot combine inclusions and exclusions of users and devices in the same assignment.

Once you create the policy, all the items you selected will be deployed according to the assignment. In the future, if you need to deploy a new item such as an app or profile to the same set of users or devices, you can simply add it to the policy set without recreating potentially complicated assignment logic in yet another location.

## Maintaining a Clean Admin Center

Over time, your Intune console will probably become cluttered with old devices. People get new phones all the time, but they do not call you to ask for their old device to be removed from Intune. You might remember that you can configure a compliance policy to retire non-compliant devices from Intune which will remove them from the console next time they check-in. This does not solve all the clutter, though.

Fortunately, you can configure Intune to automatically delete devices that have not checked-in for some time. To do this, browse to **Devices > Device clean-up rules** in the admin center. Set **Delete devices based on last check-in date** to **Yes**, and then set **Delete devices that haven't checked in for this many days** to a number of days ranging from 30 to 270 days. This configuration applies to every device in Intune (except Macs managed by Jamf), and it is not possible to choose different settings for different types of devices. You will need to choose a value that makes the most sense for your organization. For example, if you have a set of tablets that are only used for special events that happen quarterly, you might configure this feature to 120 days to allow time for those tablets to be used periodically.

## Reporting

The **Reports** section of the admin center contains a [set of reports](#) that highlight specific aspects of your Intune deployment. One thing you may notice when you click on a report is that it does not show any data. You must manually refresh the Intune reports before you can see any data in them. From time to time, Microsoft makes changes to the reports available here, or the contents of existing reports. The list of reports in this section is short, but, in the individual areas of the admin center, you will find a **Monitor** section with more reports. For example, if you browse to **Devices > Monitor**, you will find twenty additional reports that you can generate.

While the built-in reports are useful, the most powerful way to take advantage of the data in Intune is to access Intune’s data warehouse. The data warehouse is available as an OData feed that you can integrate into any tool you want. Microsoft supplies a [sample app](#) for Power BI, shown in Figure 16-4, that is a very useful getting started tool if you have access to Power BI.

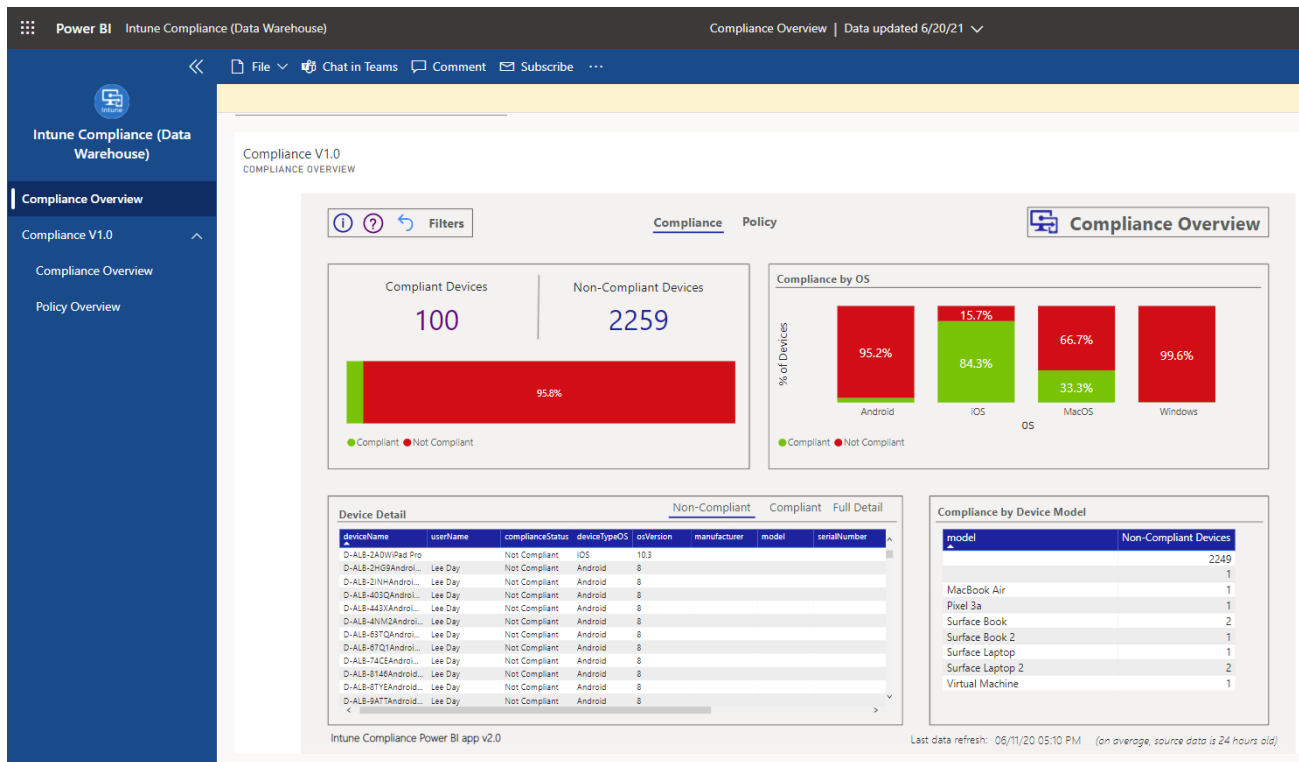


Figure 16-4 Power BI Intune data warehouse dashboard



# Chapter 17: Managing Data Governance and Compliance

## **Tony Redmond**

The Microsoft Purview suite covers a range of risk and compliance and data governance solutions available to Microsoft 365 tenants. For Information Technology, compliance is how organizations comply with established guidelines, regulations, and legislation. It can also refer to the practical measures taken by an organization to ensure compliance. Legislation varies from country to country. Well-known examples in the United States include the Sarbanes-Oxley Act (SOX 2002) and the Health Insurance Portability and Accountability Act (HIPAA – 1996). Companies working in these industries must ensure that their IT systems store and keep data in compliance with the regulations laid down in legislation. The [General Data Protection Regulation](#) (GDPR) and the California Consumer Privacy Act (CCPA) are good examples of how regulations exist to compel companies to protect personal data properly.

Microsoft began its journey to deliver compliance features in Exchange 2010. Since then, Microsoft has delivered an array of features. Although not every workload currently supports every aspect of compliance functionality, coverage is steadily spreading. The aim is to cover all workloads where users generate information that organizations might need to keep for business or regulatory purposes. Collectively, these categories give organizations a framework for “data governance.”

A data governance framework helps organizations satisfy regulations and applicable standards. Technology helps users to be compliant, but only if that technology is easy to use and unobtrusive. Experience proves that trying to implement difficult-to-use or complex compliance technology is not a recipe for success. Building compliance into applications like SharePoint Online or Teams intuitively and unobtrusively is challenging, especially as the volume, complexity, and sources of data requiring governance increase. Without good governance, organizations cannot meet compliance and regulatory requirements and expose themselves to the consequences of data leakage or external threat.

## Data Governance

Originally, workloads implemented different compliance capabilities limited to the extent of the workload’s functionality. Because each took a different approach to compliance, bringing the on-premises functionality to the cloud resulted in a fragmented and disjointed situation. Some data have protection, and some do not. Some data are subject to retention policies, and some are ignorant of policies. Some data come within the scope of eDiscovery, and some are invisible to searches. And so on.

A complication is that some applications have been slow to support the compliance framework. Another arises through the introduction of new applications like Teams and Planner, which create new sets of data that might be of interest from a compliance perspective. In short, the old approach of depending on workload-specific compliance functionality was obsolete. A new cross-service architecture was necessary to handle compliance for cloud applications running inside Microsoft 365. Microsoft rolled out the new Data Governance framework in April 2017. Over the next five years, many add-on solutions joined the framework, renamed the Microsoft Purview suite in April 2022.

Depending on the business requirements that exist within an organization to preserve and protect its information, the individual compliance solutions in the Purview suite are less or more important. Fitting them together into a coherent plan is a major task that needs input from Information Technology, Human Resources, and Legal personnel in addition to overall direction and support from senior management. As the

regulatory and legal frameworks that govern compliance vary extensively from country to country and industry to industry, the treatment of the topic here is by necessity at a reasonably high level. However, it should be enough to give guidance as to what is and what is not possible.

## Licensing Data Governance and Compliance Services

Many of the data governance and compliance services require premium or advanced licenses. This isn't a problem if you have the necessary licenses, but if you don't, you should remember that:

- Microsoft might enable a service for every account in the tenant. This can incur a potential licensing liability if you are not careful to restrict access to functionality to the accounts that need to use the service.
- Microsoft does not enforce licensing for many services. This might change in the future and if the accounts do not possess the correct licenses, they will lose access to a service once "targeted capabilities" are available.

Among the Purview solutions you might want to use are:

- Microsoft Purview Information Protection and Sensitivity Labels.
- Microsoft Purview eDiscovery (Premium).
- Microsoft Purview Customer Lockbox.
- Microsoft Purview Data Loss Prevention.
- Microsoft Purview Data Lifecycle Management.
- Microsoft Purview Information Barriers.
- Microsoft Purview Insider Risk Management.
- Microsoft Purview Communications Compliance.
- Microsoft Purview Records Management.
- Microsoft Purview Audit (Premium).

Other Microsoft Entra and Microsoft 365 components assist with data governance, including:

- Privileged Access Management.
- Azure AD Conditional Access Policies.
- Microsoft Defender for Cloud.
- Office 365 Message Encryption.
- Office 365 Advanced Message Encryption.
- Microsoft Priva.

These lists are not complete and will change over time as Microsoft introduces new services and capabilities.

Some inconsistencies in Microsoft licensing deserve mention. For example, Data Loss Prevention policies for Exchange Online, SharePoint Online, and OneDrive for Business require Office 365 E3 while DLP and retention policies for Teams need Office 365 E5. Finally, make sure you know the rules for which accounts need additional licenses. For instance, anyone who is a custodian in a premium eDiscovery case because they "control" documents or messages needed by the case must have an Office 365 E5 or Microsoft 365 E5 Compliance license.

Because detail is important, before deciding to invest in any functionality, you should first consult [Microsoft's licensing guidance](#) and review the feature chart available on that page in both PDF and spreadsheet format.

With that information in hand, you can work with the organization from which you purchase licenses (Microsoft or partner) to clarify exactly what licenses you need for the features you want to use.

## Keeping Content in Place

Traditionally, when a company needed to preserve email or documents for compliance purposes, they move copies of the items to a separate archiving system. Veritas Enterprise Vault is a classic example of such an archival system. When necessary, administrators run eDiscovery or retrieval activities to find content held in the archival system. This approach was acceptable when products like Exchange and SharePoint did not have archiving capabilities. Today, the focus is on leaving data in place within its native repository and integrating compliance technology inside applications. Advantages of this approach include:

- No need exists to transfer copies of messages, documents, and other content to external systems for searching, analysis, or otherwise interrogation for legal reasons.
- Avoidance of additional cost for software and hardware by not having to use external systems. In addition, the IT infrastructure is simplified because the extra systems are unnecessary.
- It is easier to prove the “chain of custody” and show that no one can tamper with emails and documents before these items become evidence in legal cases.
- Keeping content in its native repository means that no opportunity exists to compromise content during transfer to a new repository.

Removing the need for copying files from one system to another makes it easier to meet chain of custody requirements and prove that no one could have interfered with an item to compromise its contents. Guaranteeing the chain of custody is important when data proves facts during litigation. In-place archiving means that information stays in its original location and is protected against interference for as long as the hold persists. The data retained in place is immutable as neither administrators nor users can interfere with it after a hold is in place.

Keeping everything in place inevitably means that repositories become larger. The software therefore must be more intelligent and more capable to deal with higher volumes of data. Searches must be able to find information quickly and accurately when needed for compliance purposes. The work done by Microsoft Research to improve and make search technology smarter together with the acquisitions of FAST (otherwise known as the Search Foundation) and Equivio are examples of how Microsoft responded to the need to have better search technology. Today, Microsoft Search indexes information as soon as documents, files, messages, chats, and other items are added to repositories via user activity or bulk imports while Equivio has evolved to become Microsoft 365 Advanced eDiscovery and is capable of processing millions of messages or documents to find the desired results very quickly. Another advantage is that when you have a unified view of data such as the Microsoft 365 substrate, it is much easier to build the necessary infrastructure around that data to secure and audit its use. For instance, the audit log captures details of user activity from multiple workloads which administrators can then search using the audit log search in the Microsoft Purview Compliance portal, Office 365 Cloud App Security, or third-party tools such as [Quest On Demand Audit](#).

## Principles of Data Governance

Microsoft 365 takes a cross-service approach to data governance based on common policies that apply across multiple workloads. This approach is obvious in the way that content searches, eDiscovery, and Data Loss Prevention work, and is rolling out for data governance in the form of retention labels, retention policies, and associated features. The Microsoft Defender and Microsoft Purview Compliance portals are important parts of that strategy because they bring together functionality designed to work across applications rather than just a single workload.

Data Governance is a policy-driven framework to control the retention and removal of content across all applications. We can break the framework down into four parts:

- **Import:** Bring information from multiple sources, including from non-Microsoft sources like Facebook and Bloomberg messaging, into Exchange Online through the Import Service or third-party products. The aim is to allow customers to gather all the information they need to manage in repositories such as archive mailboxes or SharePoint sites to replace older methods like PSTs and file servers. Once the service holds the data, it becomes possible to manage the data through the application of policies.
- **Retain:** Provide customers with the ability to state what data they need to retain for compliance purposes, for whatever period is necessary. The policies used to enforce retention must apply across all workloads (or as many as possible).
- **Delete:** It is equally important for customers to be able to remove data that they no longer need. Methods need to exist to remove information from all repositories according to policy.
- **Classify:** Some data is more important or secret than others. The framework must allow users to classify information as they work. Applications should highlight information classified by users to raise awareness of the importance of the information. The system should recognize the sensitivity of the information and potentially restrict what users can do with that content.

Of course, when a company retains data for data governance, the data should be immutable. The Office 365 Import Service is the cornerstone for “Import.” Retention policies and retention labels make sure that tenants can keep or remove information as they need and classify items to mark important data for retention. Other functionality, like Data Loss Prevention, handle the need to protect against the misuse or inadvertent disclosure of data, while sensitivity labels (applied manually or automatically) restrict access to confidential information to authorized users.

## The Influence of Privacy Laws

Companies with business operations inside the European Union must follow the [General Data Protection Regulation](#) (GDPR) to safeguard how they process personal data while companies operating in California must follow the California consumer privacy act (CCPA). In both cases, heavy penalties for breaches ensure that companies must take this and similar legislation seriously. According to Microsoft, over 90% of an organization’s data stored in Exchange Online, SharePoint Online, and OneDrive for Business is in Word, Excel, PowerPoint, OneNote, and Outlook (Exchange). It is easy to imagine examples of where applications hold personal information, including:

- Annual employee reviews that are stored in a SharePoint or OneDrive for Business site.
- A list of applicants for a position is in an Excel worksheet attached to an email message.
- Lists holding data (names, employee numbers, hire dates, social security numbers, salaries) about employees in files in SharePoint Online sites.
- Discussions about potential new hires in a Teams chat.
- Word documents holding applications for employee work visas.

Given that personal information can be found in almost any application, the work to locate the 10% of personal information stored outside Office documents is likely to take the most effort (for example, see this [guide](#)). Organizations must know what applications hold personal data and how they process personal data in compliance with legislation. Knowing where to look is only the start of the process and it might not be possible to automate many of the steps involved in responding to data subject requests.

Fortunately, the compliance features available in Microsoft 365 assist tenants to satisfy legislative requirements. These features include:

- Sensitivity labels and policies to mark and potentially protect (encrypt) documents, messages, and other objects, like Power BI reports.
- Auto-label policies to find and classify documents holding personal data. Microsoft 365 includes sensitive information definitions for many country-specific personal identity cards and passports. You can create an auto-label policy to find and label documents holding these sensitive information types.

Retention processing can remove items stamped with a suitable label after a defined period, perhaps after including a manual disposition review.

- Data Loss Prevention (DLP) policies can use the same sensitive information types to stop people from sharing confidential data outside the tenant. A General Data Protection Regulation DLP template makes it easy to deploy protection for all the defined European Union sensitive information types.
- Content searches to find personal data stamped with retention labels used to mark items holding personal data.
- Alert (or advanced alert) policies to detect actions that might be privacy violations. For example, multiple downloads of documents from a SharePoint site holding HR information. You can also search the audit log to discover and report potential issues.
- The Microsoft Privacy Subject Rights Request solution makes it easy to retrieve information stored in Microsoft 365 when necessary to satisfy a request to provide information held about an individual.

Microsoft publishes information about how various aspects of Microsoft 365 align with GDPR (as an example, see this post covering [SharePoint Online and OneDrive for Business](#)). Having technology available to help satisfy the regulatory requirements is helpful. However, it is only one step to achieving full compliance. Tenants still need to protect data through a mixture of user education and technology. See [this paper](#) for information comparing the terms used in GDPR with those used in U.S. eDiscovery contexts.

This chapter covers retention labels and retention label policies. Other relevant chapters are:

- Chapter 18: Managing eDiscovery.
- Chapter 21: Managing auditing.
- Chapter 19: Managing Data Loss Prevention.
- Chapter 20: Managing Information Protection.

## Compliance Manager

Microsoft Purview [Compliance Manager](#) is a tool to guide compliance managers through the often-confusing array of rules involved in privacy frameworks. Compliance Manager is a static tool in that it does not have the necessary features to organize the work needed to achieve compliance with complex regulations like GDPR. However, many tools are available to help tenants to organize and manage the work set down in Compliance Manager. For example, Planner tasks can track progress to completion for responsibilities assigned to those working on ensuring compliance with regulations. To gain some collaboration capabilities, including document management, the tasks in plans are accessible on SharePoint Online sites, Teams, and Groups.

# Compliance Permissions

Microsoft 365 Purview compliance functionality uses role-based access control to allow or deny access to individual features. When users access the compliance portal, Purview evaluates the permissions held by their account to decide what they can do. Microsoft Purview uses a separate set of permissions to control access to compliance functionality than those used for general administrative purposes or in Exchange Online.

Compliance permissions are grouped into roles to reflect the access needed to perform specific tasks, like Data Investigator or Compliance Data Administrator. Roles also make it easier to assign relevant permissions to users. Only members of the Compliance Administrator role group can create or manage retention policies.

Compliance roles are managed through the **Permissions** section of the Microsoft Purview Compliance portal, where users holding the Organization Management role can manage the set of permissions assigned to a role or who holds a role. Some compliance roles can also be assigned by editing user accounts in the Microsoft 365 admin center. After you assign the necessary role to a user, they must sign out and back into Microsoft 365 to ensure that the new permissions are respected. Sometimes it takes a little while before the new role assignments are acknowledged. If this happens, just wait for a few minutes, and then retry.

# Retention Policies and Publishing Label Policies

The retention strategy for an organization usually contains a mixture of removal and retention. Organizations want to remove items after their useful lifetime or to stop the ongoing accumulation of data that has no value. On the other hand, keeping high-value content is important because these items form the collective recorded memory of the company. Documents and emails relating to policies and procedures, strategy discussions, board minutes, reports, research papers, and so on must remain available for as long as the organization needs them. Sometimes laws or regulations define the retention period, and sometimes the organization sets the period. Within Microsoft 365, retention processing operates by applying retention labels to items and retention policies to locations (mailboxes, sites, or groups). The settings in retention policies and retention labels can do the following:

- **Retain-only:** Microsoft 365 retains items in their home locations for a specified period (or forever), called the retention period. When the retention period expires, users decide how to dispose of items. The retention period can be forever.
- **Delete-only:** Microsoft 365 permanently deletes items after their age reaches a specific period (set by the creation or last modified date). Because retention is not enforced, users can delete the items beforehand.
- **Retain and Delete:** Microsoft 365 keeps items for a specified period and then deletes the items permanently after the retention period expires.

Other settings exist to enable more granular processing, but the basics of retention boil down to defining a retention period and action.

A retention label is more precise because it applies to a single item instead of to every item found in a location. For instance, a user who works on a new corporate policy will know what retention label is most appropriate for that kind of content. Retention policies apply to all the items in containers coming within the scope of the policy, such as all the documents stored on a site. A retention label always takes precedence over a policy because of its specificity. Two types of retention policies exist to satisfy the specific requirements of a tenant. A tenant can deploy both types of policies.

- **Publishing Label policies** make retention labels available to users by publishing the labels to workloads. Settings within the retention labels control what happens to content marked with the labels. Multiple publishing policies can make labels available to a user. Workloads are responsible for combining all the labels available to a user and presenting the complete set through client interfaces.
- Assigning individual retention labels to items is more precise than container-based assignment through **Retention policies** because users select and apply the labels to specific items. Retention policies allow organizations to achieve broader coverage by applying retention settings to all items stored in selected containers or locations (such as mailboxes, sites, or groups). Background processes such as the Exchange Managed Folder Assistant (MFA) process the retention settings defined in policies against the items stored in the target locations. Tenants can use retention policies to preserve or remove content for the entire organization or specific groups of up to 1,000 accounts. Typically, tenants use retention policies to ensure that a company meets the compliance requirements set out in legislation such as the Sarbanes-Oxley Act. Management of retention labels and policies is in the **Data lifecycle management** section of the Microsoft Purview Compliance portal.

It can be confusing to understand the scenarios where it's best to use retention policies and where retention labels are a better choice. Usually, tenants end up using a mixture of both to ensure broad coverage through policies and precise retention of specific items through labels. For instance, you might apply a retention policy

to a mailbox or SharePoint site that mandates the removal of items after six months and then provide some retention labels for users to assign to items they wish to keep for longer periods.

Table 17-1 lists several features that you should consider when planning how and when to use retention policies and retention labels.

<b>Feature</b>	<b>Retention Policy</b>	<b>Retention Label</b>
<i>Workloads:</i>		
<i>Exchange Online</i>	Yes	Yes (except for items in public folders)
<i>SharePoint Online</i>	Yes	Yes
<i>OneDrive for Business</i>	Yes	Yes
<i>Yammer</i>	Yes	No
<i>Planner</i>	No	No
<i>Teams</i>	Yes (chats and channel messages)	No
<i>Microsoft 365 Groups</i>	Yes (group mailbox)	Yes (items in group mailbox)
<i>Skype for Business</i>	Yes (conversation history)	No
<i>Automatic application</i>	Yes	Yes (by auto-label policy, DLP, etc.)
<i>Manual application by a user</i>	No (policies can only be set and applied by administrators)	Depends on the client's UI
<i>Labels displayed in client UI</i>	Partially (by Outlook clients for email)	Depends on the client's UI
<i>Persistence</i>	Policies are location-dependent, so items move out of scope if moved out of location	Persistent within supported Microsoft 365 locations. For instance, if you move a message from one folder to another, it keeps its assigned label
<i>Mark item as a record</i>	No	Yes
<i>Event-based retention</i>	No	Yes
<i>Require manual disposition review</i>	No	Yes
<i>Office 365 audit</i>	Audit records not generated for application of policies to locations	Audit records are generated when labels are applied, changed, or removed to/from items
<i>Find items subject to retention</i>	No	Via content searches (select retention label as a search condition), content explorer, and activity explorer

Table 17-1: Comparing retention policies and retention labels

Over the remainder of this chapter, we explore the details of how to create and manage retention labels, the policies used to publish retention labels and to auto-apply, and general retention policies. We will also

examine how Exchange Online mailbox retention policies work because of the ongoing need to support hybrid environments.

# Rules or Principles of Retention

The many workloads and types of data in use across Microsoft 365 create a complex environment for retention management. Multiple ways exist to mark content for retention or deletion. Therefore, a need exists for a mechanism to resolve the conflict that can occur when several policies apply to a mailbox or site, especially when Microsoft 365 retention policies remove items after retention periods expire. Microsoft applies the following four rules of retention (sometimes called the *principles of retention*) to decide how to process content. The rules work from top to bottom. Workloads use the rules as tie-breakers to set precedence when processing items if multiple policies apply.

1. **Retention wins over deletion:** You could call this the “keep safe” principle. In practical terms, it means that when multiple retention policies apply to content, retention wins over deletion. Take the example of where a mailbox comes under the scope of two retention policies. The first removes all messages after they are four years old. The second, perhaps applied to a subset of mailboxes belonging to senior managers, retains all messages for seven years. The solution is to move messages into the Recoverable Items structure after four years and to keep them there for three more years. Both policies are respected because items seem to be deleted after four years (users do not realize that the messages are still available) while keeping the items in the background ensures that the messages stay indexed and discoverable for the full seven years.
2. **Longest retention period wins:** If multiple policies specify different retention periods, items are always kept for the longest period. This principle ensures that content is kept for as long as it might be needed. If you want to deliberately remove content after a certain period, deploy a policy that explicitly removes the content after that period elapses and make sure that the location holding the content does not come within the scope of any other policy (Exchange mailbox retention policy or retention policy). Note that an item will be kept for longer if a user applies a label with a longer retention period to it.
3. **Explicit wins over implicit:** Explicit means that a user or administrator has selected specific content for special retention. This can happen when a user applies a label with a retention action to an item in a mailbox or site, or when a location is within the scope of a “non-org wide” policy (one that only applies to some content within the tenant). The logic here is that the user or administrator has made an explicit decision about retention for a specific item (a label applied manually) or location (non-org wide policy created by the administrator). Manual application of a retention label always takes precedence over a catch-all retention policy that applies across all locations, including when an auto-label policy applies retention labels automatically. This principle has long existed in Exchange retention policies where a personal retention tag applied manually by a user always has precedence over a tag applied to a folder or a default tag applied to a mailbox.
4. **Shortest deletion wins:** Retention policies allow administrators to actively remove content after a certain period (still called the retention period). If multiple deletion policies apply with different retention periods, Microsoft 365 applies the shortest retention period and removes the content when that period expires. The logic here is that an administrator decided to remove content after a certain period. The presumption is that good reason guided this decision. It would therefore not make sense if another policy, perhaps created by another administrator, interfered with the decision to remove content after that period. Again, this is a reason to consider how the retention policies in a tenant interact with content.



Remember that holds always take precedence over deletion. If an eDiscovery case places a hold on content for a set period, Microsoft 365 cannot remove items within the scope of the hold until the hold period expires or an administrator releases the hold, even if the retention periods of policies applied to the content expire.

It often takes time and some experience in working with retention policies across different workloads to understand the effect of the retention principles and how to best deploy these principles to support the data governance strategy for the organization. Microsoft [publishes a helpful flowchart](#) to explain how retention decides to keep or remove items (Figure 17-1). It's much easier to follow and understand the steps in the flowchart after gaining some experience working with retention policies and labels.

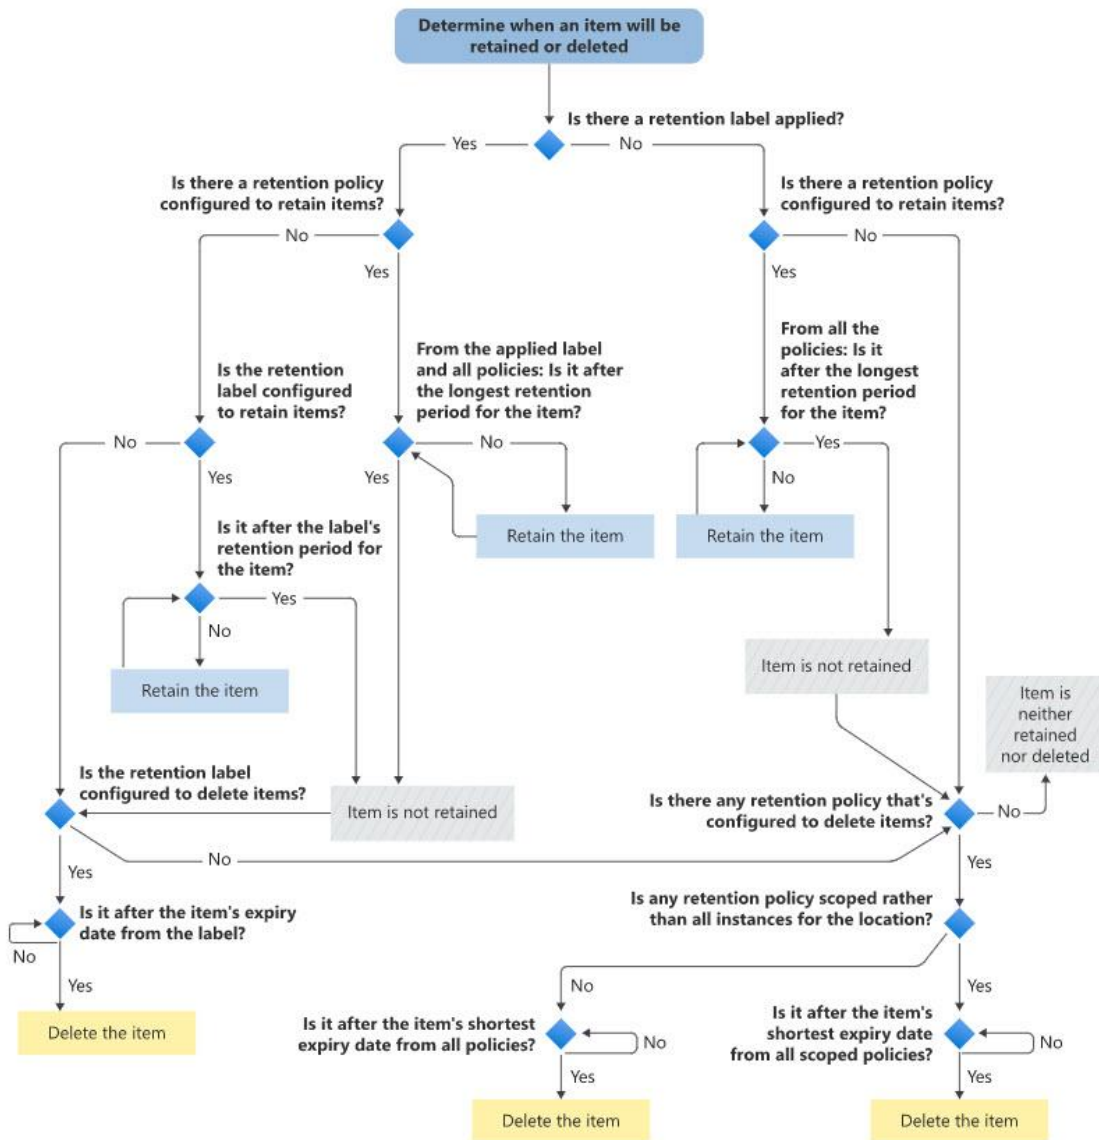


Figure 17-1: Microsoft's flowchart to determine how retention removes or keeps items

## Retention Policies

Retention policies apply retention rules to target containers (locations). The supported containers include:

- Exchange (mailboxes and cloud-based public folders).
- SharePoint Online (including files owned by Yammer, Teams, and Groups).
- OneDrive for Business.
- Teams compliance records for chats and regular and shared channel conversations. Compliance records for chats are in personal mailboxes, while regular and shared channel compliance records for

conversations in group mailboxes. Because shared channels aren't associated with a group mailbox, Microsoft 365 creates a special version called a *SubstrateGroup* mailbox to hold their compliance records.

- Teams private channel conversations.
- Yammer (when configured in Microsoft 365 network mode, user messages and community conversations can be subject to retention).
- Groups (Outlook conversations in the group mailbox).

Retention labels contain settings to tell Microsoft 365 how the organization wishes to retain or remove items in a location. The basic settings are a retention action or rule, defining what should happen when a retention period expires, and a retention period, which defines when the action should occur.

The Managed Folder Assistant processes retention policies for Exchange locations:

- User and shared mailboxes (primary and archive, if enabled).
- Public folders (you can't select specific sections of the public folder hierarchy for processing).
- Group mailboxes.

A separate Microsoft 365 retention assistant processes all the other workloads, including the compliance records stored in Exchange Online.

## Broad and Narrow Retention

Retention policies can be broad or narrow in scope. A broad retention policy applies the same retention settings to a set of containers. Each workload has its locations: Exchange has mailboxes (user, shared, and group), SharePoint has sites, OneDrive for Business has accounts, and Teams has chats and conversations. The broadest retention policy is one that applies the same retention settings to all the locations managed by all workloads in a tenant. As described below, this is an example of an org-wide policy.

Narrow retention policies make retention labels available to workloads. Each retention label has its retention settings and a retention policy for labels can have just one label or include many labels. Publication means the process of making workloads like Exchange and SharePoint aware of the existence of the labels in a retention policy. Following publication, the workloads expose the retention labels in client interfaces so that users can then apply the labels to the selected content. The application of a retention label to a message or document is a much more precise way to keep information. As such, retention labels trump the settings applied by broad retention policies. For instance, the organization uses a broad retention policy to assign a five-year retention period to all SharePoint sites. All documents stored in a document library inherit the policy and SharePoint will keep documents until they are five years old. If a user then assigns a retention label with a ten-year retention period to a set of documents, SharePoint will keep those documents for that period.

## Org-Wide and Scoped Retention Policies

Two types of retention policies exist:

- **Org wide:** The purpose of an org-wide policy is to apply a single retention policy to all supported workloads. The scope of an org-wide policy is the tenant, but the scope can be amended to exclude or include specific workloads. Currently, org-wide policies can apply to Exchange mailboxes, Exchange public folders, Groups, Yammer private and community messages, and documents stored in SharePoint Online sites and OneDrive for Business accounts. Separate policies handle the retention of Teams chats and channel messages (use of these policies requires E5 licenses). There can be up to 10 org-wide policies in a tenant. You should use org-wide policies sparingly as it is easy to create a policy that has unintended consequences. For instance, if you create a policy to keep all content for two years and then remove the content afterward, the policy will remove everything from the tenant that is more than two years old and does not come under the control of another policy. That could

lead to the removal of most content across workloads, which might not be what you want to do. Org-wide policies are “entire location” policies because they cover all the locations within the selected workloads (for example, all Exchange Online mailboxes).

- **Non-org wide:** These policies apply retention to a subset of the locations available within a tenant. The scope of these policies is set by picking specific mailboxes, sites, or groups to include in the policy or by applying a query based on keywords or sensitive information types to find the content to which the policy applies. Each workload has its restrictions for the number of locations in a non-org-wide policy. For example:
  - Exchange Online and Teams: 1,000 mailboxes (accounts).
  - Microsoft 365 Groups: 1,000 groups.
  - SharePoint Online sites and OneDrive for Business accounts: 100 of each.
  - Teams chat or channel messages: 1,000 accounts.
  - Yammer community or private messages: 1,000 accounts.

[This article](#) details the current limits for Microsoft 365 retention policies.

The number of org-wide and non-org-wide policies available for deployment means that organizations can get very creative with their retention strategy. In general, it is best to simplify retention processing by limiting the number of active policies as this will help to avoid situations where the actions and retention periods of policies clash. See the *Rules of Retention* section to understand what happens when several policies apply to an item.

**Getting Round Policy Limits:** The limits for the target locations processed by retention label policies divide into SharePoint locations (sites and accounts) and Exchange locations (mailboxes and groups). SharePoint locations have a 100 limit while Exchange can deal with 1,000 locations. In either case, these limits are often too small for large tenants. To get around the limits, you can create multiple policies that have the same settings but different target locations. For example, if you need to apply retention settings to 500 SharePoint sites, create five examples of the same policy and include a different set of 100 target sites in each policy. With the necessary licenses, adaptive scopes (see below) allow organizations to solve the problem by creating scopes to find target locations and using the scopes in retention policies. Adaptive scopes don't have the same numeric limitations as static policies do.

## Adaptive Scopes

Adaptive scopes are a way to find target locations by applying a filter (query) against the set of available locations in a tenant. In some respects, adaptive scopes work like dynamic Azure AD groups or dynamic distribution lists, both of which use a query to find a set of objects. Adaptive scopes must be created first before they can be used in a retention policy.

Three kinds of adaptive scope are available:

- **Users:** Filter applied against selected account properties like job title, city, state or province, department, and the Exchange custom properties. For example, an adaptive scope can find all the users located in France whose job title starts with “Manager.” Mailbox states and types are also usable in filters, meaning that you can create scopes to look for inactive mailboxes (a state) or shared mailboxes (a type). Figure 17-2 shows the query builder in use to build the scope for an adaptive scope for users. You can also see the set of properties available to build the query. Adaptive scopes of this type apply retention to mailboxes and OneDrive for Business accounts.
- **Sites:** Filter against the site URL, name, or the 100 refinable string properties (refinablestring00 to 99) available to customize the SharePoint Online search schema. For example, find all sites where the value of a custom property (represented by the RefinableString99 refiner) is “Secret.” To update custom properties for a site, its administrator must write values into the site property bag. [This article](#)

explains how to update the site property bag using PowerShell. Another example is to use an adaptive scope with a query to find sites with URLs starting with a certain value to find all OneDrive for Business accounts in a tenant.

- **Microsoft 365 groups:** Filter against group properties like name, description, and Exchange custom properties.

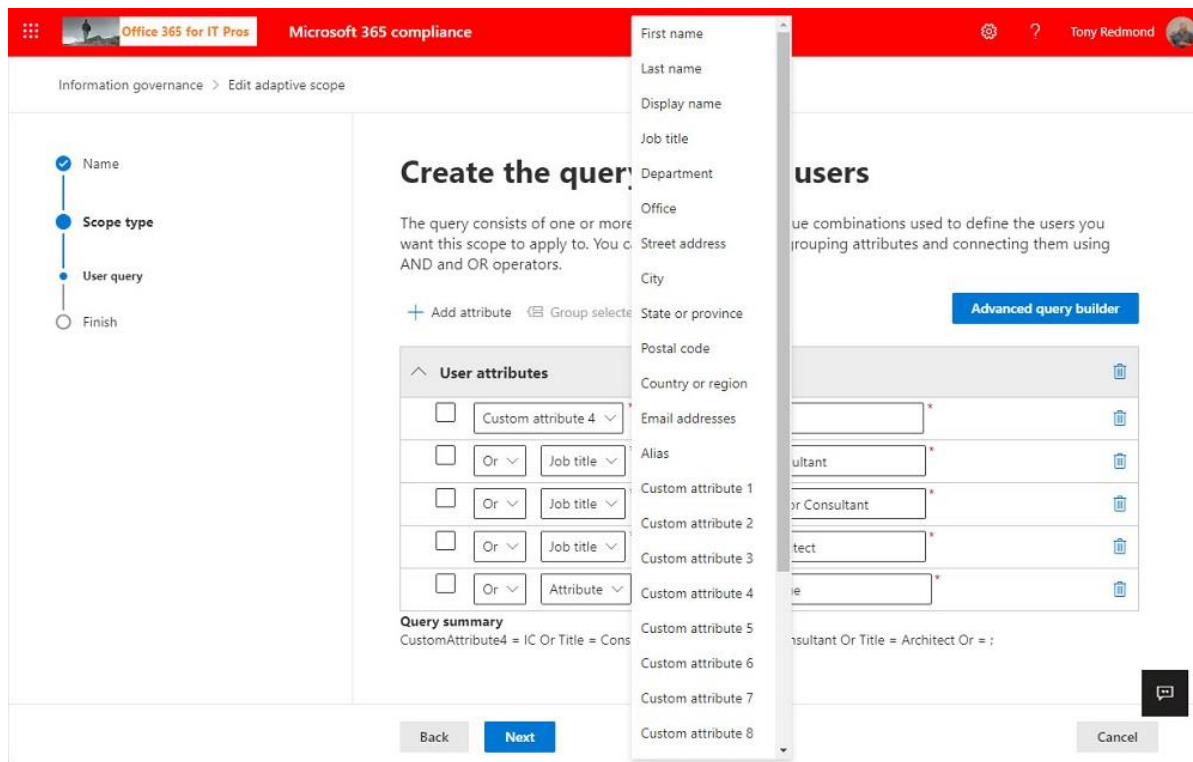


Figure 17-2: Defining an adaptive scope

The operators available to build an adaptive scope query are:

- is equal to.
- is not equal to.
- starts with.
- not starts with.

You can combine *And* or *Or* checks against properties. For instance, the query shown in Figure 17-2 contains several *Or* tests for different values in the Job title property. If the query builder can't construct the query you need, you might be able to build an advanced query in OPATH syntax (users and groups) or KQL (sites). For instance, an adaptive scope to find inactive mailboxes uses the OPATH query:

```
IsInactiveMailbox -eq 'True'
```

While the query to find shared mailboxes is:

```
IsShared -eq 'True'
```

However, although advanced queries can build user adaptive scopes using mailbox types and states and employ access to more operators (such as *like* or *notlike*) than available in the query builder, advanced queries can only use the same set of Azure AD account properties presented by the query builder.

The results of adaptive scope queries are not available immediately. Instead, a background process that runs daily uses the queries in adaptive scopes to identify the set of target locations identified by the scope query. You can view the current set of locations by accessing the scope details (Figure 17-3). If you see *No data*

*available*, it means that either the background process has not yet resolved the query, or the query doesn't find any objects.

Information governance >

Scope

### Individual contributor mailboxes

6 items Refresh Search

Filters: Filters

Display name	Location type	State
Marc.Vigneau@office365itpros.com	User	Added
Chris.Bishop@office365itpros.com	User	Added
Oisin.Johnston@office365itpros.com	User	Added
Eoin.Redmond@office365itpros.com	User	Added
Ben.James@Office365itpros.com	User	Added
Brian.Weakliam@office365itpros.com	User	Added

#### Details

**Name**  
Individual contributor mailboxes

**Description**  
An adaptive scope to find the mailboxes belonging to individual contributors

**Type**  
User

**Query summary**  
CustomAttribute4 = IC Or Title = Consultant Or Title = Senior Consultant Or Title = Architect

**Last modified by**  
Tony Redmond

**Last modified**  
Nov 1, 2021 2:50 PM

Figure 17-3: Viewing the set of locations calculated using an adaptive scope query

A retention policy can use one or more adaptive scopes to find the locations to which it applies. Figure 17-4 shows that a retention policy uses two user adaptive scopes. Because these are user adaptive scopes, the set of workloads that the retention policy can cover doesn't include SharePoint Online. However, it does include OneDrive for Business because these accounts are personal and linked to individual users. When a workload processes a retention policy with adaptive scopes, it resolves the queries from all scopes to find the up-to-date set of locations and applies the settings in the retention policy to those locations.

Adaptive scopes are only available to tenants with Office 365 E5 or Microsoft 365 E5 compliance licenses. Office 365 E3 tenants are limited to retention policies with static scopes. These policies process a defined set of locations. To change the set of locations, an administrator must amend the policy. This approach is satisfactory when an organization doesn't need to change its retention framework very often. It is less good when an organization wishes to apply retention policies based on some criteria, such as the people who work in a certain country or department, or SharePoint sites containing a certain kind of information.

Because adaptive scopes depend on settings such as Azure AD account properties or custom site properties, it's easy to add new locations to policies by updating the account, mailbox, or site settings. Workloads then pick up new locations or remove locations no longer within the scope the next time they process locations to enforce retention policies.

Retention policies with adaptive scopes require an Office 365 E5 or Microsoft 365 E5 compliance license for every account coming within their scope. Administrators can manage adaptive scopes through the Data lifecycle management or Records management sections of the Microsoft Purview Compliance portal.

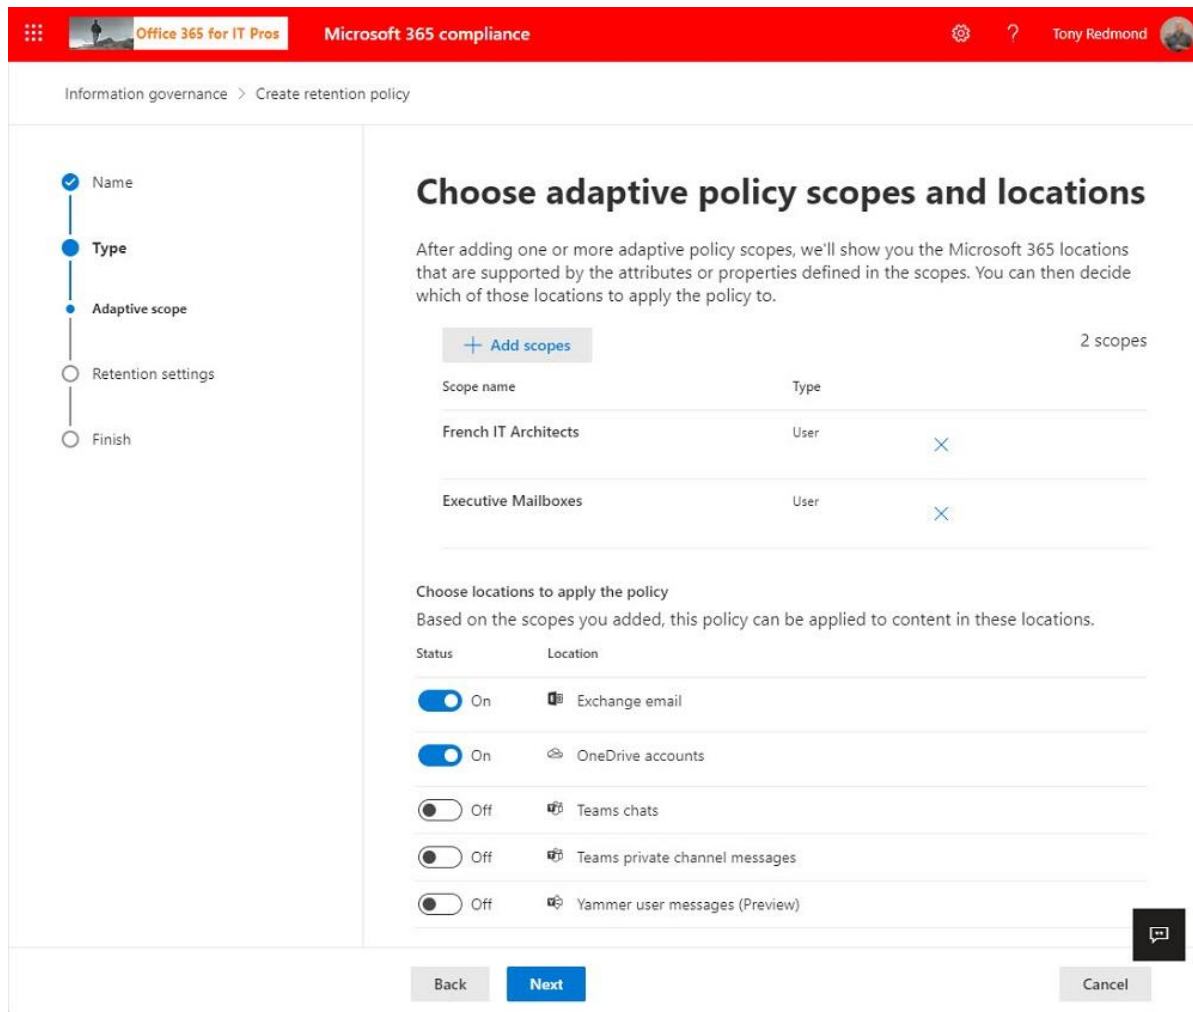


Figure 17-4: Using an adaptive scope in a retention policy

## The Effects of Retention Policies

It's important to understand how content might come under the influence of retention settings through broad policies or the more precise application of retention labels:

- Broad retention publishing policies that define retention settings for locations (selected or org-wide defined through static or adaptive means) are invisible to users. These policies run in the background and users are unaware that retention is in force. This assignment is implicit because an item inherits the retention policy because of its location instead of having a retention label assigned by a user.
- Narrow retention publishing policies make labels available to users when they work with content in the covered locations. After publication, a user can assign one of the published retention labels to a document or mail message. This assignment is explicit because it takes a deliberate action by a user to assign the label.

Organizations can apply retention to an item in these ways:

- A user assigns a retention label to an item. This explicit assignment has the highest priority. Remember that a label used purely as a visual marker does not have a retention action or period.
- A location comes within the scope of a broad org-wide or non-org-wide policy, including policies that use advanced processing to find content. The items take the retention settings from the policy unless someone assigns a retention label to the item.
- The owner of a SharePoint site defines a default retention label for a document library. New items created in the library inherit the label. Library settings can also dictate if existing items receive the

same label. A label inherited from a document library overrides any org-wide or non-org-wide retention policy assigned to the location.

- An auto-label policy assigns the label contained in the policy to items. An auto-label policy never replaces a label if a user has already assigned a label.

If their account has the appropriate permissions, users can change the retention label assigned to an item (unless the label is a regulatory record). They do not see or cannot affect the retention policy assigned to a location.

If multiple retention policies apply to an item, an explicit assignment always takes precedence over an implicit assignment. It is worth emphasizing that the presence of any type of retention affects the ability to remove a mailbox, site (including OneDrive for Business accounts), or group. Mailboxes holding retained items or under the scope of a retention policy become inactive when deleted until the retention expires. Sites and accounts cannot be removed until retention on individual documents, lists, and files lapses, and no retention policies are in effect for these locations.

## Planning a Retention Policy

Before beginning the process to create a new retention policy, it is sensible to write down the basic structure of the policy, broken down into several headings:

- **Broad or Narrow:** Do we need to apply default retention settings to some or all locations, or publish retention labels to allow users to dictate how to retain content? If the policy uses retention labels, what are those labels? Organizations often implement org-wide (broad-scope) policies to cover all workloads and all users with selective deployment of non org-wide (narrow-scope) policies to target specific workloads, users, or sensitive information types.
- **Scope:** What locations will this policy cover? If the policy targets a subset of locations within a workload (for instance, ten mailboxes or five SharePoint sites), write down those locations. If you want to exclude a subset of locations from the policy's scope, you should also note those locations. If the policy covers inactive mailboxes, this fact should be noted. If an adaptive scope will be used, some thought must be given to how the set of target locations can be found.
- **Purpose:** Will the policy keep or remove content? If so, how long will the retention period be?
- **Records:** Is the content that comes under the scope of the policy considered to be a formal record for the company and if so, must the content be immutable? (See the sections on Preservation Locks and Records).

For the example used here, the conditions listed in Table 17-2 apply.

<b>Heading</b>	<b>Policy setting</b>
Scope	The policy applies to mailboxes and OneDrive for Business sites belonging to senior managers defined as members of the Senior Leadership Team (SLT) distribution list. It also covers the content held in the SLT group and team.
Purpose	The policy keeps all information in the targeted locations for ten years and then removes the content.
Type	This is a "non-org wide" policy. The group membership dictates the accounts to which the policy applies.
Lock	A preservation lock is not required.

Table 17-2: Planning a retention policy

It is bad to find yourself in a situation where you create and deploy a retention policy only to discover that the policy removes needed information. For example, it is very easy to apply an aggressive retention policy to “All Exchange mailboxes” that removes items after 30 days. This is the equivalent of a default delete tag in an Exchange retention policy to remove content after 30 days. When stamped on a mailbox, the Exchange Managed Folder Assistant (MFA) applies the “Delete and Allow Recovery” retention action to all items in the mailbox that do not have a more explicit tag. After items reach 30 days old, the MFA moves them to the Recoverable Items folder and preserves them for a further period (the deleted items retention period). When that period expires (usually 14 days), MFA permanently removes the items from the database. Exchange Online does not use backups, so you cannot recover the items at this point. Because retention policies affect the content stored in user mailboxes, it is only sensible to consider and understand exactly what will happen when a retention policy is active.

It is also possible that a retention policy will clash with an Exchange mailbox policy or another retention policy. It is a good idea to take some of the target locations and work out what policies have those locations within their scope to figure out if a clash occurs.

### No Change for Retention Action or Period

A crucial factor to consider when planning the implementation of retention policies and labels is that you cannot change some of the important settings that control how the policy functions after creation. For example, you can alter the retention period for a policy, add new locations to its scope, and alter the KQL query to find content for the policy to apply. However, you cannot change its basic operation. For instance, you cannot change a policy that keeps content into one that simply removes content.

The logic is that users expect consistency in the processing of their data. If you can change the fundamental operation of how retention works inside a tenant, users will not know whether their data will be kept or removed or when this will happen. For this reason, it is wise to take time to chart out how retention will work across the tenant for all workloads before you create any policies. Fools rush to implement retention without thought!

### Naming a Retention Policy

With our structure in mind, we can go to the **Data lifecycle management** section of the Microsoft Purview Compliance portal, select **Retention policies**, and then **New retention policy**. The first step is to assign a name and description (only visible to administrators) to the new policy. Some tenants insist that administrators include their name and a pointer to supporting documentation in the description of a new policy. It's more useful to include notes about what the policy does. For example:

*“This is a retention label policy to publish a set of general-purpose labels to every location in the tenant.”*

*“This policy publishes the Highly Confidential label to the Senior Leadership Team locations.”*

*“This policy searches for content with the “Kazaa” keyword and applies the Ten Year retention label.”*

### Setting the Scope for a Static Retention Policy

The simplest form of a retention policy is one that includes every available location, but quite often the need exists to focus on a select group of individuals and the data with which they work. Retention policies allow you to include or exclude subsets of locations. When you enable a retention policy for a location, you can choose the scope to be:

- **Org-wide:** Cover Exchange mailboxes, public folders, Groups, SharePoint Sites, and OneDrive for Business accounts. The policy applies to all content in all locations. As the tenant adds new locations, they come under the control of the policy.
- **Non org-wide** (choose specific locations): You can select individual mailboxes, sites, and so on. You can select everything from a workload, like all Exchange mailboxes, or you can select a subset of the



locations available within a workload, such as only a few SharePoint sites. You can also exclude a selected subset from the policy. This means that the policy will not apply to the mailboxes or sites that you select.

Two special processing cases exist for retention policies. First, you can't exclude or include specific public folders, all of which are either processed by the retention policy or not. Second, you cannot mix inclusions and exclusions for a location in a policy. If you exclude some sites or mailboxes from a policy, it means that the policy applies to all other sites or mailboxes but not to those selected for exclusion.

A retention policy for Teams can only cover Teams content (chats, channel messages, or private channel messages). The same is true for Yammer policies, which can only process messages posted to Yammer communities in networks configured in Microsoft 365 mode. You cannot include another workload in a retention policy for Teams. However, you can have multiple retention policies for Teams that cover different subsets of users, or separate policies for channel messages and personal chats. Table 17-3 explains the various location types and how you input the selected locations.

<b>Location type</b>	<b>Identified by</b>
Exchange email	Select <b>All recipients</b> to include all items stored in Exchange user mailboxes, otherwise, select the mailboxes to apply the policy. You can also input the name or alias of a distribution list or mail-enabled security group.
SharePoint site	Select <b>All sites</b> to include all the sites in the tenant or input the URLs for selected sites. Example: <i>https://tenant.sharepoint.com/Projects/</i> .
OneDrive for Business accounts	Select <b>All accounts</b> to include all OneDrive accounts in the tenant or input the URLs for selected accounts. Example: <i>https://tenant-my.sharepoint.com/personal/kim_akers_office365itpros_com/</i> .
Microsoft 365 Groups	Select <b>All groups</b> to cover the mailboxes used by all Microsoft 365 groups in the tenant (but not the compliance items stored in mailboxes for Teams and Yammer) or the names of the selected groups.
Exchange public folders	Select <b>All</b> to extend the policy to cover every public folder in the hierarchy. The default is "None." You cannot select a subset of public folders.
Teams channel messages	Select <b>All teams</b> to cover messages posted to all channels in every team in the tenant or select the individual teams to come within the scope of the policy.
Teams chats	Select <b>All users</b> to include all personal chats sent by users in the tenant or select the users to come within the scope of the policy.
Teams private channel messages	Select <b>All users</b> to include the messages sent to private channels in all user mailboxes or select the users to come within the scope of the policy.
Yammer community messages	Select <b>All communities</b> to include all messages sent to Yammer communities or select the communities the policy will apply to. This option is only available when the Yammer network runs in Microsoft 365 mode.
Yammer user messages	Choose <b>All users</b> to include all private messages sent between Yammer users or select the users to which the policy will apply. This option is only available when the Yammer network runs in Microsoft 365 mode.

Table 17-3: Microsoft 365 workloads and locations supported by retention policies

## Choosing Specific Locations

Figure 17-5 shows the user interface to add locations to a policy. If your policy covers a subset of workloads and locations, some up-front work is necessary to list the locations and gather the information to input each location. The easiest way to add a set of mailboxes to a retention policy with a static scope is to use a distribution list or a mail-enabled security group. Each mailbox counts against the 1,000 location limit for the policy. In addition, Microsoft 365 includes the owners of the distribution list in the mailboxes added to the set of locations (some consider the inclusion of distribution list owners as a feature; I believe it to be a bug). The population of the Exchange locations in the policy is a one-time operation and any future additions or removals to the membership of the distribution list do not synchronize with the locations in the retention policy. You must edit the policy to ensure that it continues to cover the correct individuals.

### Choose locations to apply the policy

The policy will apply to content that's stored in the locations you choose.

Status	Location	Included	Excluded
<input checked="" type="checkbox"/> On	Exchange email	All recipients <a href="#">Edit</a>	None <a href="#">Edit</a>
<input checked="" type="checkbox"/> On	SharePoint sites	All sites <a href="#">Edit</a>	None <a href="#">Edit</a>
<input checked="" type="checkbox"/> On	OneDrive accounts	All accounts <a href="#">Edit</a>	None <a href="#">Edit</a>
<input checked="" type="checkbox"/> On	Microsoft 365 Groups	All groups <a href="#">Edit</a>	None <a href="#">Edit</a>
<input type="checkbox"/> Off	Skype for Business		
<input type="checkbox"/> Off	Exchange public folders		
<input type="checkbox"/> Off	Teams channel messages		
<input type="checkbox"/> Off	Teams chats		
<input type="checkbox"/> Off	Teams private channel messages		
<input type="checkbox"/> Off	Yammer community messages		

Back **Next** Cancel

Figure 17-5: Choosing a set of target locations for a retention policy

While mailboxes are relatively simple to add to policies with static scopes, the same is not true for SharePoint Online and OneDrive for Business locations. This is because you must input the URLs for each location that you want to add to the policy. It can sometimes be difficult to know what the URL is for a location, so it's a good idea to collect the URLs beforehand in a text file and cut and paste the URLs from the file into the policy. For example, you can use the *Get-SPOSite* cmdlet to output the URLs and export the information to a CSV file. It's relatively easy to automate adding SharePoint Online and OneDrive for Business locations to a retention policy with PowerShell. You can add up to a hundred individual sites in a non-org wide policy.

The processing of Teams chat and channel messages and Yammer messages rely on compliance records stored in personal and group mailboxes. The rules are:

- Retention policies for chats and channel conversations cannot include any other non-Teams locations.
- Retention policies for Teams private channel conversations cannot include any other location (even other Teams locations). It's important to note that retention policies for private channels operate

based on individual accounts rather than teams. This is because the compliance records for private channel conversations are in user mailboxes and it's not necessarily true that all the members of a team are members of a private channel. To be sure of processing all messages for a private channel, add the mailboxes of all the members of the channel to the policy. Also, make sure that the SharePoint content for the channel comes within the scope of a separate SharePoint Online retention policy as the Teams private channel policies process only channel conversations.

- Retention policies for Yammer content can process only Yammer locations.

After adding the target locations to the policy, click **Next** to continue.

**Be Careful with Inclusions and Exclusions:** Retention policies allow you to include or exclude specific locations. For example, you might create a policy to process a single mailbox or SharePoint site. If you edit the policy and remove the exclusion or inclusion, the target locations revert to All. This might be what you want (for example, you've tested the effect of the policy and are now happy to apply it to all locations in a chosen workload), but if it's not, it's easy to end up applying a retention policy inadvertently to all locations, which might remove items that the organization wants to keep.

## Keeping or Removing Content

The final step is to define what the policy does to keep or remove content (Figure 17-6). When a retention policy removes items, it uses a "delete and allow recovery" action, to allow users to recover items later if needed from Exchange's Recoverable Items structure or the SharePoint recycle bin. In either case, we must know the length of the retention period and how to calculate the age of an item. You can keep content forever, but it is more common to set a period like seven or ten years. For mail messages, the creation date is used, but when a policy spans both documents and other items, it is best to choose the last modification date as shown here as this accommodates the fact that documents are often changed well after their creation date.

## Decide if you want to retain content, delete it, or both

Retain items for a specific period  
Items will be retained for the period you choose.

Retain items for a specific period

7 years

Start the retention period based on

When items were created

At the end of the retention period

Delete items automatically

Do nothing

Retain items forever  
Items will be retained forever, even if users delete them.

Only delete items when they reach a certain age  
Items won't be retained, but when they reach the age you choose, we'll delete them from where they're stored.

---

Back **Next** Cancel

Figure 17-6: Defining the actions a retention policy takes to process content

Administrators accustomed to working with Exchange mailbox retention policies see an immediate difference here. Exchange retention policies let you remove or archive items after the retention period lapses. While retention policies do not support an archive action, they let you say that you want to keep information for a set period. When the retention period lapses, you can choose to leave the content alone (in which case

another policy could apply) or remove it. Alternatively, if you are more concerned about cleaning out locations to remove information that the organization no longer needs, you can instruct retention policies to remove items after they reach a certain age.

**Immediate Evaluation:** After you publish a retention policy, workloads make it active. This means that the workload processes that evaluate items know about the new policy. If you're not careful, the introduction of a new policy can have a significant effect on users. For example, let's say that you create a policy to remove items after three years and set the scope of the policy to be all OneDrive for Business accounts. After publication, the retention policy begins to evaluate items in all accounts and will remove anything more than three years old. If users aren't prepared for this clean-up to happen, the sudden removal of many items could come as a nasty surprise.

## Reviewing Policy Settings

The last step is to review its settings. If all looks to be in order, click **Submit** to begin the publication process to the locations covered by the policy. You can opt to save the policy for later, meaning that the policy is in a draft state that you can update later (perhaps to add some extra locations) before making the policy live.

It can take some time before a retention policy becomes fully effective across all locations. The assistants must process each location coming under the scope of a policy, including the compliance records stored for Teams and Yammer. You can check the distribution status of the policy by selecting the policy and viewing its properties. The following values are available:

- **Enabled (Success):** All locations know about the policy and are processing content as per the policy. However, the policy might still not be effective everywhere because background jobs might not have completed their processing to enable all locations.
- **Enabled (Pending):** The policy is being enabled in the target locations.
- **Enabled (Errors):** The policy is active, but some errors have occurred in its distribution.
- **Disabled (Success):** The policy is disabled and is not being applied to content. You can reenable the policy at any time.
- **Disabled (Pending):** Microsoft 365 is in the process of disabling the policy. All locations are stopping the processing of content. However, the process is incomplete.
- **Disabled (Errors):** The policy is off, but due to some errors, it might still be active in some locations. You can click the error to get more information. In many cases, the fault will disappear if you leave it alone for an hour or so.

Note that you can't disable a policy that applies a preservation lock.

## The Actions Taken to Deploy Our Retention Policy

In terms of what you might have done previously with Exchange retention policies or SharePoint document delete policies, the retention policy we just created has the following effects:

- The policy imposes a seven-year retention period on every item in the target mailboxes that do not already have an explicit mailbox retention tag (directly assigned on items or inherited from folders) or retention label. In addition, Exchange places an in-place hold on the mailboxes (both primary and archive) for seven years. During this period, Exchange preserves any item that a user removes from the mailbox in the Recoverable Items structure. The preserved items are indexed and discoverable. When an item's retention period elapses (in this case, seven years after creation), the Managed Folder Assistant applies a "Delete and Allow Recovery" action and moves the now-expired item into the Recoverable Items folder. The user can recover the item until the deleted items retention period set for the mailbox expires (usually 14 days).

- An equivalent seven-year in-place hold applies for documents and lists stored in SharePoint or OneDrive for Business sites. When users delete or remove items, SharePoint Online captures copies in the preservation hold library for the site.
- The Exchange Online Managed Folder Assistant processes the contents of group mailboxes to apply retention policies. SharePoint Online treats the sites belonging to groups (teams) like other non-group connected sites.
- Although our policy does not cover public folders, if they were, any copies kept by policy stay in the public folder mailboxes to which they belong until the retention period expires. A process like that used to assess content removed from user mailboxes ensures that any attempt to remove held content from a public folder cannot succeed.

## SharePoint Online's Preservation Hold Library

To preserve copies of information required for retention purposes, SharePoint Online uses a special retention destination called the **Preservation Hold Library**. SharePoint Online creates the preservation hold library automatically when first needed by a site to retain information about changes and deletions to files and lists. Once a site comes within the scope of a retention policy or individual items receive retention labels, user actions to remove or change content cause SharePoint to copy items to the preservation hold library. A single preservation hold library handles retention for the complete site, no matter how many document libraries and lists exist on the site.

User actions that cause SharePoint Online to capture information in the preservation hold library include:

- **Update or deletion of a document under a retention policy:** SharePoint Online captures a copy of the document. If versioning is enabled for the site (the default), SharePoint Online retains copies of all changes made to the document.
- **Deletion of labeled files (with unexpired retention periods):** Until a change deployed in November 2021, SharePoint Online stopped users from deleting labeled files and displayed a message saying that the label applied to the file blocked its removal. This behavior differed from OneDrive for Business and exposed a weakness in that members of a group could remove the label and then delete the file. With the change, both SharePoint Online and OneDrive for Business allow users to delete labeled files and capture the deletion in the preservation hold library. Users cannot remove files assigned a record label from either SharePoint Online or OneDrive for Business.
- **Unlocking of a retention label (record) on a file:** Users must unlock files assigned retention labels marked as a record before editing. SharePoint Online captures a copy of the file when it is unlocked.
- **Deletion of a OneNote section:** A copy of the section is captured.
- **Deletion of a OneNote notebook:** A copy of the notebook is captured.
- **Update of document metadata:** A copy of the update is captured.
- **Update or deletion of a list item:** A copy is captured in an Excel XLS. If the list item has an attachment, SharePoint Online captures the list and the attachment separately. The same happens if an attachment is removed from a list item.

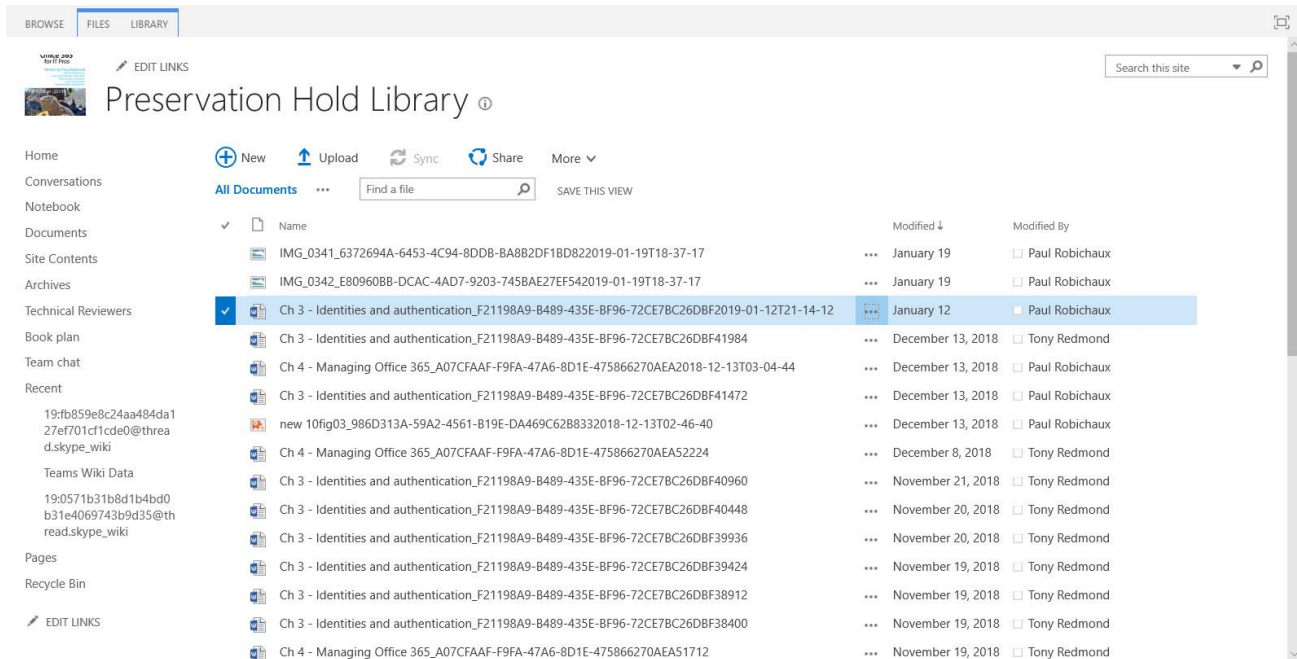
Not every change generates an update to the preservation hold library. The aim is to retain the original content before the imposition of retention controls plus the current content. When a user attempts to update or delete an item, SharePoint Online checks if this is the first action since retention for the item became active (or the policy changed). If it is, SharePoint Online captures the original content and allows the user to update it. Any subsequent updates to the content are available in the versions that SharePoint Online captures automatically, meaning that the original content plus the complete change history are available through the data in the preservation hold library plus the online content. If a user deletes an item coming within the scope of a retention policy, SharePoint Online removes the item from the library or list and stores it in the preservation hold library until its retention period expires.

Users are unaware of the preservation hold library because it is only visible to site administrators, who can access the preservation hold library by going to Site Contents and then selecting the library, or they can append */PreservationHoldLibrary* to the site URL (because users administer their own OneDrive for Business account, they can access the preservation hold library for their account). Site administrators cannot remove or change items kept in the preservation hold library, but they can copy items from the library to retrieve a file for a user.

## Using the Preservation Hold Library

When users add new documents to a library, several operations happen to create the document, apply updates, add metadata, and so on. SharePoint Online captures none of these actions in the preservation hold library. However, once the initial creation is complete, SharePoint Online monitors and captures subsequent file updates.

Every week, a background job runs to clean out old items from the preservation hold library. The job looks for items in the library for more than 30 days and compares them against the retention settings applicable to the site to find and remove items with expired retention periods. If the file or list item is still present in its original location, the background job removes it too. Items removed from the preservation hold library go into the second stage recycle bin from where SharePoint Online removes them permanently after the normal 93-day period in the recycle bin expires. Figure 17-7 shows items in a preservation hold library. SharePoint Online generates the names of the retained items by combining the original name with a GUID (and sometimes the date and time of the change) to create a unique value.



The screenshot shows the SharePoint interface for a Preservation Hold Library. The top navigation bar includes 'BROWSE', 'FILES', and 'LIBRARY'. Below the navigation bar, there are options for 'EDIT LINKS' and a search box. The main content area displays a list of documents with the following columns: Name, Modified, and Modified By. The list contains several items, including images and documents, with their names containing GUIDs and dates. The item 'Ch 3 - Identities and authentication\_F21198A9-B489-435E-BF96-72CE7BC26DBF2019-01-12T21-14-12' is highlighted in blue.

Name	Modified	Modified By
IMG_0341_6372694A-6453-4C94-8DDB-BA8B2DF1BD822019-01-19T18-37-17	January 19	Paul Robichaux
IMG_0342_E809608B-DCAC-4AD7-9203-7458AE27EF542019-01-19T18-37-17	January 19	Paul Robichaux
Ch 3 - Identities and authentication_F21198A9-B489-435E-BF96-72CE7BC26DBF2019-01-12T21-14-12	January 12	Paul Robichaux
Ch 3 - Identities and authentication_F21198A9-B489-435E-BF96-72CE7BC26DBF41984	December 13, 2018	Tony Redmond
Ch 4 - Managing Office 365_A07CFAAF-F9FA-47A6-8D1E-475866270AEA2018-12-13T03-04-44	December 13, 2018	Paul Robichaux
Ch 3 - Identities and authentication_F21198A9-B489-435E-BF96-72CE7BC26DBF41472	December 13, 2018	Paul Robichaux
new 10fig03_986D313A-59A2-4561-B19E-DA469C62B8332018-12-13T02-46-40	December 13, 2018	Paul Robichaux
Ch 4 - Managing Office 365_A07CFAAF-F9FA-47A6-8D1E-475866270AEA52224	December 8, 2018	Tony Redmond
Ch 3 - Identities and authentication_F21198A9-B489-435E-BF96-72CE7BC26DBF40960	November 21, 2018	Tony Redmond
Ch 3 - Identities and authentication_F21198A9-B489-435E-BF96-72CE7BC26DBF40448	November 20, 2018	Tony Redmond
Ch 3 - Identities and authentication_F21198A9-B489-435E-BF96-72CE7BC26DBF39936	November 20, 2018	Tony Redmond
Ch 3 - Identities and authentication_F21198A9-B489-435E-BF96-72CE7BC26DBF39424	November 19, 2018	Tony Redmond
Ch 3 - Identities and authentication_F21198A9-B489-435E-BF96-72CE7BC26DBF38912	November 19, 2018	Tony Redmond
Ch 3 - Identities and authentication_F21198A9-B489-435E-BF96-72CE7BC26DBF38400	November 19, 2018	Tony Redmond
Ch 4 - Managing Office 365_A07CFAAF-F9FA-47A6-8D1E-475866270AEA51712	November 19, 2018	Tony Redmond

Figure 17-7: Items in the Preservation Hold library for a SharePoint site

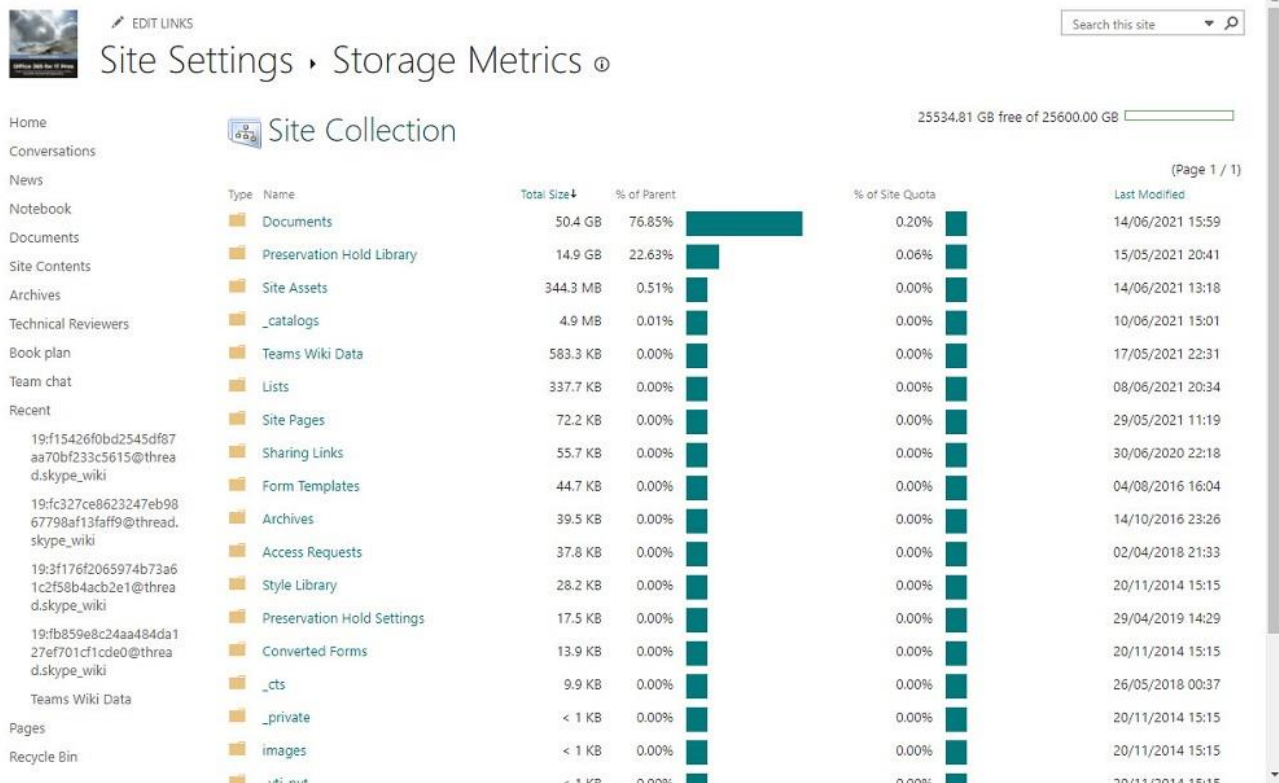


Figure 17-8: Site Settings include storage metrics

Files in the site recycle bin and those retained in SharePoint Online because of a retention policy count against the overall storage quota for the tenant. When a retention policy is in place, the SharePoint admin center can help you understand the effect of retention through the storage metrics available through site settings (Figure 17-8). The fact that retained files occupy 22.63% of the total storage for the site underlines how much tenant quota retention can consume.

Your mileage might vary in terms of storage consumption for retained files. The number of files retained by policy differs from site to site and is highly dependent on user activity and the size of the files. If people create and save files and never edit them again, no extra versions exist and no deletions happen, so few files end up in the preservation hold library. On the other hand, if you constantly edit files stored in a library, the versions accumulate (especially through Office's Autosave feature) to consume storage. This is especially evident in libraries that store large files with frequent updates, such as the Word documents for the chapters for this book.

### Purging the Preservation Hold Library After Removing a Retention Policy

When an administrator removes a retention policy from a site, a 30-day grace period starts to stop the release of the hold on the site. During the grace period, SharePoint Online preserves the items in the preservation hold library because the hold is still in place. Once the 30-day grace period elapses, the hold lapses and SharePoint Online proceeds to delete the items. Items deleted from the preservation hold library go into the second-stage recycle bin instead of an immediate purge. Items stay in the second-stage recycle bin for up to 93 days after their deletion before permanent removal. During this time, site administrators can recover items.

This mechanism gives administrators the ability to recover from the consequences of an error in removing a retention policy that results in data loss. Instead of having background processes purging content soon after the removal of a retention policy, administrators can access the preservation hold library and second-stage recycle bin to recover data during the grace period. Recovering data from these locations is a manual process that will take time and effort, but at least it's better than losing valuable documents.

## Knowing That a Retention Policy Works

It can be difficult for a user to know that a retention policy is in force. It can also be difficult for an administrator to know when retention policies work as expected. Here are some ways to verify that all is well:

- For **user mailboxes**, you can check to see whether the retention tag assigned by a policy is stamped on folders and messages as expected. Clients display these tags like older Exchange Online retention tags, so if Outlook or OWA displays the tag, you know that the policy is working. Another method is to check the Purges sub-folder of Recoverable Items to see if retained items accumulate there.
- Conversations in **group mailboxes** do not display retention tags. You must check that the items in the Inbox folder are what you expect based on the retention policy in force. You can also check that the Managed Folder Assistant is processing the group mailbox and removing items.
- **Teams** and Yammer compliance records are in special hidden folders in group and user mailboxes. You can use the MFCMAPI utility to check these folders to ensure the expected retention of compliance data. Another (easier) method is to use the `Get-ExoMailboxFolderStatistics` cmdlet to check the number of items in the folder before the retention assistant processes it and compare it to the number of items reported afterward. For example, the command below shows that the oldest item in the folder used to hold Teams compliance items (for group mailboxes, items from all channels are in the same folder) is from 7 October 2020. You know the retention period set in the policy, so it is easy to calculate what the date of the oldest item should be if the retention policy is working.

```
[PS] C:\> Get-ExoMailboxFolderStatistics -Identity O365ITPros -FolderScope NonIPMRoot -
IncludeOldestAndNewestItems | ? {$_.FolderType -eq "TeamsMessagesData"} | Format-Table Name,
ItemsInFolder, NewestItemReceivedDate, OldestItemReceivedDate
```

Name	ItemsInFolder	NewestItemReceivedDate	OldestItemReceivedDate
TeamsMessagesData	513	02/02/2022 16:10:41	7/10/2020 08:17:09

- Documents stored in **SharePoint Online** and **OneDrive for Business** sites, including the sites used by Groups, show no trace of retention because everything remains in place until user actions trigger the need for SharePoint to capture copies of documents. If a policy dictates that an item should be kept, and an attempt is made to remove it, a copy of the item will be in the special Perseveration Hold library for the site. The capture happens even if a retention label stops the user from removing an item. The site collection administrator can access the Preservation Hold Library through Site Contents to verify the capture of changes or deletions for files under the control of retention policies are there.
- Check the audit log to look for *FileDeleted* events logged for deletions of SharePoint and OneDrive items because their retention period expired, and the retention action forces a deletion. *FileDeleted* events capture information about what account removes an item. In the normal course of events, the data captured in the user property is the User Principal Name of the account. When a retention policy or label causes a deletion, the audit event captures the name of the policy or label. For example, this command searches for all *FileDeleted* events logged during a month:

```
[PS] C:\> Search-UnifiedAuditLog -Operations FileDeleted -StartDate 4-Nov-2018 -EndDate 5-Dec-2018 -
ResultSize 2000
```

After retrieving the audit events, you can use the techniques explained in Chapter 21 to interpret them and generate data for analysis. A review of the information reveals if retention processing deleted files. In this example, a user deleted two files and a retention policy deleted the next three files.

TimeStamp	User	File
2018-11-28T15:15:22	tony.redmond@office365itpros.com	Board Meeting Agenda 12 Sept 2018.docx
2018-11-28T15:22:17	tony.redmond@office365itpros.com	Privacy Policy for Website.docx
2018-12-04T03:51:45	Preservation Lock - Mailboxes and Sites	Encoding Time.csv



2018-12-04T03:51:46 Preservation Lock - Mailboxes and Sites TeamsNotebook(Shared).onetoc2  
 2018-12-04T03:51:43 Preservation Lock - Mailboxes and Sites Ch25Skype for BusinessRewriteV1.docx

## Retention Policy Lookup

Sometimes many different retention policies affect an individual location (user, site, or group) and it's difficult to understand what influence these policies have on the location. For example, the SharePoint admin center might inform you that a retention policy is blocking a site from deletion, and you want to find out what policies might be the block. To help understand the set of policies that apply to a location, use the Policy Lookup option under Data lifecycle management in the compliance portal. To lookup the policies for a location, enter the:

- Primary email address of a user.
- URL for a SharePoint Online site or OneDrive for Business account.
- Primary email address of a Microsoft 365 group.

The response is the full set of retention policies for the selected location. These can be:

- Label publishing policies for which the location is a target. In other words, policies that make labels available for use by the user or in the site or group.
- Auto-label retention policies that process the location.
- Retention policies that include the location in their scope.

Figure 17-9 shows a typical result where a user account is subject to seven policies, one of which has an adaptive scope. The others are all org-wide policies created for different reasons, including to process messages in Teams channels and Yammer communities. Two of the policies are label publishing policies (for example, the manual disposition label policy publishes labels used for manual disposition), but the only way to know this is to recognize the policy name. Disabled policies are included in the list, so the need exists to check policies to understand the full coverage of retention,

**Data lifecycle management** Remove from navigation

Overview Retention policies Labels Label policies Adaptive scopes Policy lookup Import Archive

Search for a specific user, SharePoint site, or Microsoft 365 Group to find out which data lifecycle management policies (retention, label, and auto-labeling) they're included in.

---

Find policies that include a

User

7 items

Policy name	Scope types	Applications	Last modified	Date created
<input type="checkbox"/> Retention Policy for French IT Architects	AdaptiveScope	Exchange, OneDriveForBusin...	3 Nov 2021 13:04	3 Nov 2021 13:04
<input type="checkbox"/> Formal Company Records	OrgwideScope	Exchange	31 Dec 0000 23:35	31 Dec 0000 23:35
<input type="checkbox"/> Company Confidential Policy	OrgwideScope	Exchange	31 Dec 0000 23:35	31 Dec 0000 23:35
<input type="checkbox"/> Manual Disposition Label Policy	OrgwideScope	Exchange	18 May 2021 10:15	18 May 2021 10:15
<input type="checkbox"/> Teams Private Channels	OrgwideScope	MicrosoftTeamsChannelMes...	11 Nov 2021 15:42	30 Jun 2021 19:24
<input type="checkbox"/> Yammer messages	OrgwideScope	Yammer	11 Nov 2021 15:43	11 Nov 2021 15:43
<input type="checkbox"/> Label Customer Invoices	OrgwideScope	Exchange	18 Jan 2022 11:38	25 May 2020 15:58

Figure 17-9: Looking up the retention policies applying to a user

The results of a retention policy lookup only include Microsoft 365 retention policies. They do not include Exchange Online mailbox retention policies or litigation holds.

## Retention Policies and Inactive Mailboxes

An inactive mailbox is a mailbox kept for compliance or some other reason that belonged to a now-removed Microsoft 365 user account within the scope of a hold placed before deletion. Including a mailbox in the locations processed by a retention policy ensures that a mailbox becomes inactive if its account is removed, but only if the policy settings keep content. A policy configured to remove content does not transform a mailbox into inactive status upon deletion. This is logical because if you deploy a policy to remove content from mailboxes, it doesn't follow that you want to keep mailboxes after their accounts are deleted.

It is a good idea to have a retention policy to hold the complete contents of user mailboxes for a period after deletion. To make this work, you must add the mailboxes to the policy (up to 1,000 mailboxes can be included in a single policy) before removing the account. Alternatively, you can use an org-wide policy that covers the content in all mailboxes. In either case, when the account is deleted, Exchange Online recognizes that one or more holds exist on the mailbox and makes the mailbox inactive. See [this page](#) for more information on how retention policies process inactive mailboxes.

## Hybrid Governance

Retention policies do not apply to on-premises locations such as Exchange mailboxes, public folders, or SharePoint sites. The same issue occurs for content searches and eDiscovery cases. If you have a hybrid environment, it's a good idea to try to have the same retention policies (or as close as possible) apply in both on-premises and cloud locations and be prepared to perform eDiscovery processing on both platforms to capture all information necessary to satisfy eDiscovery searches. Unfortunately, synchronizing retention policies across on-premises and cloud environments is a manual process.

## Groups, Teams, and Retention

When a group is deleted, it enters a 30-day soft-deleted state. Following this period, Azure AD permanently removes the group, and it becomes irrecoverable. Deletion of the group also removes all the linked resources associated with the group. However, if content in the group mailbox or the SharePoint site belonging to the group (or team) comes within the scope of a retention policy or retention label, the group mailbox and/or site is retained until the retention period expires. The group mailbox becomes an inactive mailbox while the SharePoint site becomes a deleted but retained site.

### Retention Processing for Teams Compliance Records

When we discuss retention for Teams content, it is important to understand that retention policies process the compliance records that the Microsoft 365 substrate creates to record messages posted in channel conversations and chats. The substrate also creates compliance records for messages posted in channels by Office connectors. The compliance records are in a folder called *TeamsMessagesData* in the system part of user, group, and cloud-only mailboxes that are invisible to email clients. Retention policies do not directly process any data in the Teams message store in Azure Cosmos DB. Synchronization with Exchange apply the effects of retention to that store.

Retention policies can process all Teams accounts with a valid Office 365 license. However, a difference exists between accounts with enterprise licenses (Office 365 E3 and E5) and other licenses. Accounts with enterprise licenses can use retention periods as low as one day while the other accounts have a minimum retention period of 30 days.

When a Teams retention policy is in place, the substrate ensures the retention of changes made to Teams messages. If a user modifies or deletes a message in a personal chat or channel conversation that comes within the scope of a retention policy, the substrate moves (for deletions) or copies the message to the

*SubstrateHolds* subfolder of the Recoverable Items folder. Messages remain in *SubstrateHolds* until their retention period expires.

The retention assistant processes mailboxes covered by Teams retention policies to remove compliance records according to the criteria set in those policies. Different policy settings cover chat and channel messages. Messages in chats are “owned” by all the chat participants, so each participant has a copy of each message in their mailbox. As chats are removed from mailboxes by retention processing, the number of references to a message drop until it eventually reaches zero. At that point, Teams removes the message from its data store in Azure Cosmos DB. Only one copy of channel messages is kept in the group mailbox belonging to the team. When this copy is removed by retention processing, the resulting synchronization removes the message from the Teams data store.

MFA processes Teams compliance records in the same way that it deals with other mailbox items (for example, a mailbox holding compliance records is not processed if the mailbox size is under 10 MB). The date used to determine whether items exceed their retention period are the creation dates of compliance records in the *TeamsMessagesData* folder.

When MFA removes Teams compliance records, it first moves the messages from the *TeamsMessagesData* folder to the *SubstrateHolds* folder. Messages remain in *SubstrateHolds* for a day before they are permanently removed from the mailbox. During this period, a background job removes the source items in the Teams data store in Azure Cosmos DB and later, through server-to-client synchronization, from client-side caches. Depending on the load on different components across the service and how often clients connect to the Teams service, the end-to-end removal process usually takes a couple of days to complete and can take up to a week. The minimum period for Teams compliance policies is one day. Great care should be taken when setting such policies as any error might lead to an irrecoverable data loss like [that suffered by KPMG](#).

System messages posted to channels (for example, the addition or removal of a member) are not removed by retention policies. In addition, some older messages posted by guest or hybrid users might not be removed.

If you want retention policies to apply to the content posted in the SharePoint document libraries used by Teams, you must include those sites in the **SharePoint** section of the retention policy. A retention policy cannot process data stored in other locations used by Teams such as third-party applications accessed through tabs or bots.

Retention policies for Teams messages cannot use the advanced features to search for content based on keyword queries or sensitive information types. The minimum retention period for Teams is one day.

**Be Careful with Short Retention Periods:** The reason why Teams supports a 1-day minimum retention period is that some organizations don't like the idea of keeping chats around for any longer. Although there can be good business reasons for such a stance, it's important to understand the downside. If you configure a very short retention period for items, the intended recipient might never see some messages. For instance, a message sent on Friday might be removed before the recipient checks for new messages on Monday. Short retention periods are really to cover scenarios where messages do not need to be persistent. In most Teams scenarios, removing chats and conversations quickly can mar business effectiveness if you're not careful and users understand just how quickly items are removed.

## Org-wide Policies and SharePoint Online Sites

When an org-wide retention policy covering SharePoint sites is deployed, SharePoint Online keeps deleted sites until their retention period expires. The policy spans all SharePoint sites in the tenant, including the team sites belonging to Groups and Teams as well as the hidden sites used by Teams private channels. This is expected behavior: SharePoint is told to retain deleted sites for a period and that's what it does.

If you want to permanently remove a deleted SharePoint site before the retention period expires, you must:

- Use the SharePoint Admin Center to restore the site from the Deleted sites list. If the site was connected to a group and more than 30 days have elapsed since the group was deleted, you won't be able to restore the group and reconnect the site.
- Wait a few minutes after the site restore finishes and then edit the retention policy to exclude the now-restored site from the locations covered by the policy (you can't exclude a deleted site from a retention policy because it's already deleted). The delay allows the Microsoft Purview Compliance portal to recognize that the site exists.
- Wait another little while for the amended retention policy to be effective and then delete the site from the Active sites list in the SharePoint Admin Center.
- Finally, permanently remove the site from the Deleted sites list.

## Applying Retention Policies to Microsoft 365 Groups

By default, when retention policies cover Microsoft 365 groups, the same retention settings apply to content in both the mailboxes and the SharePoint team sites owned by the groups. You cannot, for instance, remove conversations after a month and keep documents for a year. If you want retention policies to process either mailboxes or sites, you can update the policy with PowerShell by running the *Set-RetentionCompliancePolicy* cmdlet. For example:

To process only the contents of group mailboxes covered by the policy:

```
[PS] C:\> Set-RetentionCompliancePolicy -Identity "General Retention Policy" -Applications "Group:Exchange"
```

To process only the contents of the SharePoint team sites for groups covered by the policy:

```
[PS] C:\> Set-RetentionCompliancePolicy -Identity "General Retention Policy" -Applications "Group:SharePoint"
```

To reset the policy so that it covers both mailboxes and team sites:

```
[PS] C:\> Set-RetentionCompliancePolicy -Identity "General Retention Policy" -Applications $Null
```

The Managed Folder Assistant (MFA) processes the group mailbox to remove or keep items based on the settings in a retention policy. MFA will not process mailboxes unless they hold more than 10 MB of data, so MFA might never process some groups, even if their mailboxes hold several hundred conversation items. The information in the group team site does not count against the 10 MB threshold. You can use the *Get-ExoMailboxStatistics* cmdlet to check the current storage for a group mailbox:

```
[PS] C:\> Get-ExoMailboxStatistics -Identity Office365TenantServiceHealth | Format-Table DisplayName, TotalItemSize, ItemCount
```

DisplayName	TotalItemSize	ItemCount
Tenant Service Health	11.29 MB (11,836,393 bytes)	154

In this case, MFA will process the group mailbox. See the section about Logging the Managed Folder Assistant later to understand how to see a summary of the actions taken by the MFA to remove items from a mailbox per the settings of a retention policy.

## Removing Retention Policies

To remove a retention policy, select it in the Microsoft Purview Compliance portal and take the **Delete policy** action. When you remove a retention policy, Purview notifies the affected workloads that the policy no longer exists so that they cease to implement it. Retention labels applied by automatic label policies stay in place after the removal of that policy.

The next time a background process reviews content, it might apply a new retention policy to items. The time taken to switch retention policies depends on how quickly workloads cease processing the original policy and how soon afterward the content covered by the original policy is reevaluated.

If you delete a label publishing policy, the labels published by the policy are no longer available in the locations covered by the policy. However, any labels applied when the policy remain assigned to items.

## Preservation Locks

Some regulatory regimes require that after an organization implements a retention policy, administrators cannot turn the policy off or make it less restrictive. To meet this need, you can apply a preservation lock to a retention policy. After applying a preservation lock to a policy, an administrator cannot disable the policy or remove locations from the policy. The lock remains in force and active for all locations under the scope of the policy until the retention period expires. Users cannot remove or update content within the scope of the policy during this period either. The only option open to an administrator is to add locations to the policy or extend its duration.

To lock a retention policy, set its *RestrictiveRetention* property through PowerShell. For example:

```
[PS] C:\> Set-RetentionCompliancePolicy -RestrictiveRetention $True -Identity "Management Preservation Policy"
```

Microsoft can remove the preservation lock from a retention policy. If you get into a situation where you need this to happen, the tenant administrator must open a support incident and supply Microsoft will the information to justify unlocking. If Microsoft concurs, they will remove the preservation lock from the policy. Because the potential exists that Microsoft might not agree to unlock a policy, it is wise to pause and think before enabling preservation lock on a retention policy. You might need to implement such a policy to satisfy a legal or regulatory need, but in most cases, tenants do not need to lock down content in this way. In short, make sure that you need a preservation lock before turning it on for a policy. Apart from waiting for the policy to expire, there is no way back if you make a mistake and enable the lock-in error.

To discover if any policies include preservation locks and the workloads they cover, view its details through the Microsoft Purview Compliance portal or run this command:

```
[PS] C:\> Get-RetentionCompliancePolicy | ? {$_.RestrictiveRetention -eq $True} | Format-Table Name, Workload
```

Note that if you move a mailbox that is subject to a preservation lock back to an on-premises Exchange server, Exchange Online keeps a copy of the mailbox to satisfy the lock. The copy held in Exchange Online is a point-in-time copy and no mechanism exists to synchronize the two copies after the mailbox moves.

## Auto-Label Retention Policies

Publishing labels for people to apply to documents and messages is certainly one way to solve the need for data governance. The problem with this approach is that it depends on human beings to be very precise, consistent, and persistent in how they classify material. As noted in this book, we know that humans are very inventive, but they also tend to lose interest in boring technicalities after a while. On the other hand, computer systems are all about processing the same steps time after time until told to stop. Auto-label policies help organizations meet their data governance needs by finding content to retain and applying the right label to that content. Auto-label policies work alongside standard label policies. People apply the labels published to their accounts to explicitly label important or specific items while most items that need labels receive them automatically. Auto-label policies require Office 365 E5 or Microsoft 365 E5 compliance licenses for any account coming within the scope of the policies.

Take the “Business Critical” label for example. We know that we want users to use the label to classify any content that the organization needs to keep for audit purposes. We have reasonable confidence that people will classify new documents and messages, but tens of thousands of pertinent documents exist in SharePoint Online and OneDrive for Business that the business should retain. If we can build rules to tell Microsoft 365 how to find the content, an auto-label policy is the right tool for this job.

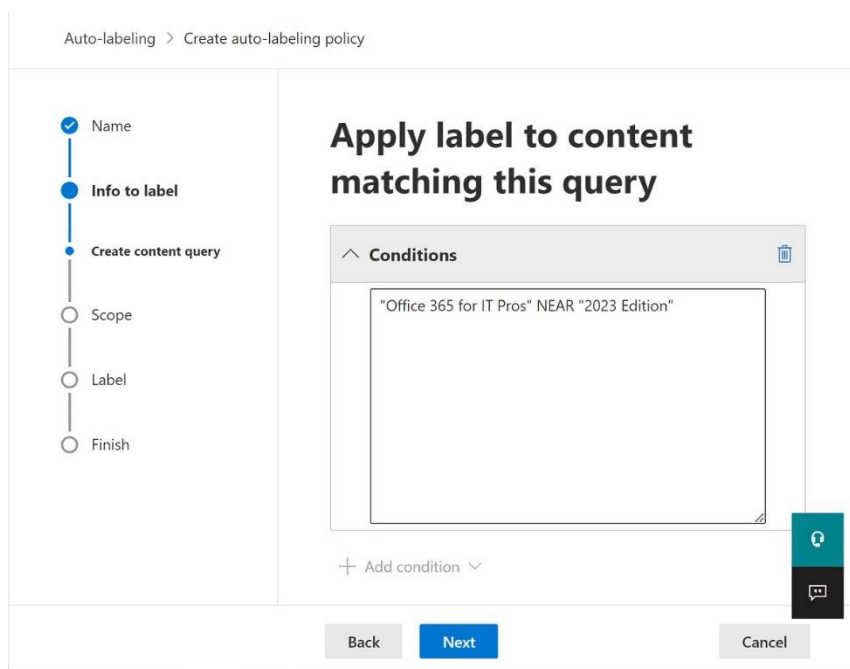


Figure 17-10: Using keyword queries to find content to keep through a retention policy

Auto-label policies use conditions to find matching items in the locations covered by the policy and assign a retention label to those items. For instance, you might want to use an auto-label policy to protect items identified with matching criteria. Three options exist to find content for an auto-label policy:

- **Sensitive Information:** Find items containing sensitive information types such as credit cards, bank account numbers, taxpayer identification numbers, and social security numbers. Any sensitive information type known within a tenant is usable, including custom types defined by the tenant or created using digital fingerprinting.
- **Content (KQL) Queries:** Find items using a content query containing keywords or search terms. Good examples of this kind of policy inaction are those used to [find and label Teams meeting recordings](#) and [apply retention labels to documents marked with sensitivity labels](#). Like content searches, the query is expressed in KQL syntax (this [page is helpful](#) to understand how to construct KQL queries). A good way of testing a query before using it in an auto-label policy is to use it with SharePoint search or a content search. Either way, by checking the items returned by the search, you’ll know if the content query finds the correct items for the auto-label policy. See Chapter 18 for more information about running content searches.
- **Trainable classifier:** Find items matching a classifier created by Microsoft or created by the tenant. Auto-label policies that use trainable classifiers can’t process content that’s more than six months old.

To create a new auto-label policy, go to Label policies and then select auto-apply a label. After naming the new policy, you select the kind of criteria to match items against. In Figure 17-10, we see a simple query to look for one phrase close to another. You can specify as many keywords as necessary to create the best chance of matching items for retention using the same approach as for eDiscovery searches.

Content queries can filter out certain kinds of files for retention processing. For example, let’s assume that you only want to assign retention labels to Word documents and PowerPoint presentations stored on specific sites. The keyword query to find documents of these types must include a reference to their file extensions.

Some search experts recommend adding a term to have SharePoint Online only search items it considers documents (*isDocument* is true). However, an argument also exists that it's sufficient to pass the desired file extensions. This content query uses both file extensions and an explicit term to limit the search to documents:

```
filetype:doc* filetype:ppt* isDocument:1
```

If you change a policy to update the keywords used in the content query, the target locations evaluate the updated content query to decide whether to keep or remove content.

Microsoft 365 publishes auto-label policies to the target workloads in a similar manner to label publishing policies. Auto-label policies work against data at rest for both SharePoint and Exchange locations. Policies detect matches against conditions when the indexer processes items, leading to the assignment of labels to matched items. When an Exchange mailbox comes within the scope of an auto-label policy, the locations include both the primary and archive mailboxes.

Figure 17-11 shows how to include a sensitive information type in a retention policy. When you choose to use sensitive information types, you select a template like the way you create a new DLP policy. In this case, we chose the U.S. Financial Data template, which imports three sensitive information types into the policy. You don't have to use all the imported sensitive information types and can remove the unneeded types.

Auto-labeling > Create auto-labeling policy

**Define content that contains sensitive info**

Choose a category of industry regulations to see the classification groups you can use to classify that information or create a custom group.

**Content contains**

Sensitive Information Types Any of these

**Sensitive info types**

U.S. Individual Taxpayer Identification Number (ITIN)

Medium confidence 1 to Any

ABA Routing Number Medium confidence

Instance count 1 to Any

Back Next Cancel

Figure 17-11: Using sensitive information types in a retention policy

Figure 17-12 shows how to use a custom trainable classifier as the basis for an auto-label policy. You can use either the set of classifiers created and published by Microsoft or build a tenant-specific classifier. Classifiers are “trainable” because their creation involves a training process where artificial intelligence technology processes sets of sample documents to recognize their essential characteristics and create the classifier. In this instance, the trainable classifier is *Customer Billing* constructed by analyzing a set of 400 customer invoices. When the auto-label policy is active, each time the indexer meets a matching document (one with the same characters as learned from the sample set), Microsoft 365 applies the label to the matched item.

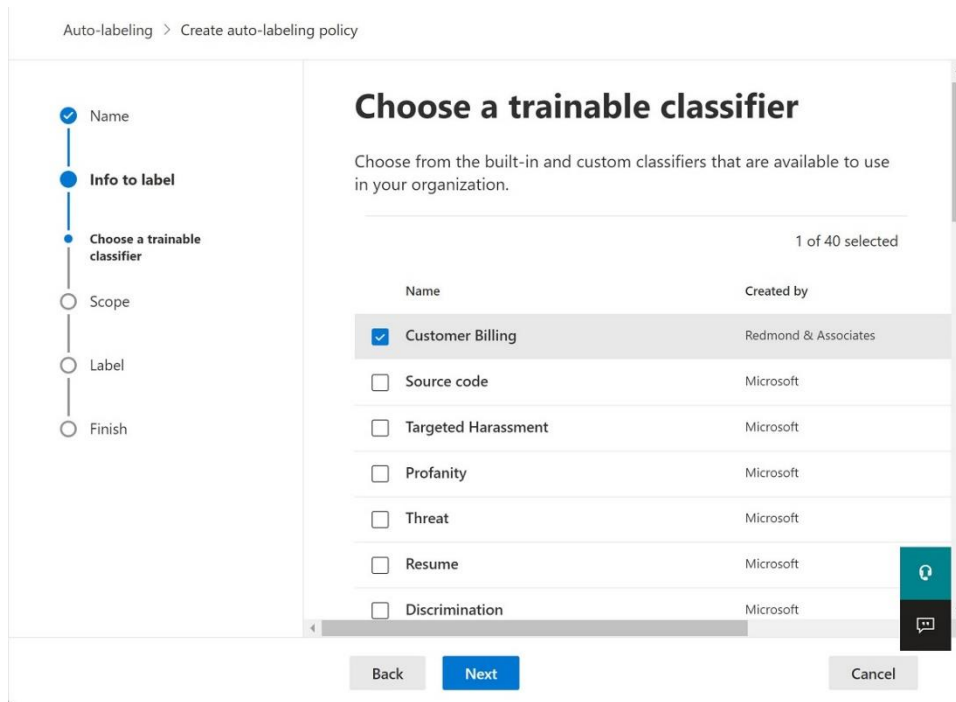


Figure 17-12: Choosing a trainable classifier for an auto-label policy

After choosing the criteria to find matching items, you select an adaptive scope or select static target locations from:

- Exchange mailboxes. Note that when using sensitive data as the criterion for auto-labeling, the policy applies to all mailboxes. You cannot include or exclude individual mailboxes.
- SharePoint sites.
- OneDrive for Business accounts.
- Microsoft 365 Groups. These locations include the group-connected SharePoint Online sites, including those used by Teams. When using sensitive data, you cannot include or exclude individual groups. Also, you must include Exchange mailboxes in the policy to ensure that the policy processes the complete content for the groups.

Finally, you choose the retention label to apply. After a last check of the policy settings, submit the policy for publication to the target workloads. The Compliance portal validates the settings and if everything is OK, go ahead and publish.

Auto-applying labels are an effective way to achieve broad coverage of content that a query can find using a keyword or because the content holds some sensitive data. Later, users can review items with the label and decide whether the label is correct, or if another label is a better fit. Unlike other labels, the labels included in auto-apply policies do not show up in SharePoint, OneDrive, or Exchange as labels that users can apply manually.

## Validating the Progress of Auto-Labeling

It can take up to a week before an auto-label policy becomes fully effective. Unfortunately, there's no easy way to measure the effectiveness of auto-labeling without some degree of manual checking. You can:

- Check the properties of individual items that you expect to receive labels to see if the auto-label policy has assigned a label to those items.
- Use the Activity explorer in the compliance portal to see what items have received the label applied by an auto-label policy.



- Analyze the *TagApplied* events from the Office 365 audit log to track the application of the label applied by an auto-label policy. An example [script in the Office 365 for IT Pros GitHub repository](#) shows how to approach the task.

If you deploy several auto-label policies, documents can end up with labels that you think are incorrect. Apart from using bad search queries to find the items, the problem might be because more than one policy applies to documents. When this happens, the oldest policy (by creation date) wins and once a policy applies a label to an item, other policies will ignore that item because it already has a label. Because auto-labeling works in this manner, you need to be careful with the queries used to find content and with the order in which you deploy policies.

## Replacing Labels Assigned by Auto-Label Policies

Auto-label policies are an effective way of applying retention labels to large quantities of items. Microsoft 365 won't process matching items in the target locations if they already have labels assigned manually (by users) or automatically (by policy). Auto-label policies ignore labeled items to avoid the potential that overwriting the existing label with another might cause data loss.

Good as this logic is, it does create a problem if an organization's retention strategy changes. You can update a retention label to change its retention period, but you cannot change its retention action. For instance, if a label's retention action is to delete items when the retention period expires, you can't change the action to leave items in place. Microsoft is aware that this lack of flexibility can cause problems for organizations forced to change their retention strategy due to government or other regulations. In the future, you might be able to alter an auto-label policy to instruct it to replace one label with another but for now, you should remember that once you apply retention labels through policy, those labels will stay assigned to items unless replaced by a label explicitly assigned by a user.

## Other Methods to Automatically Assign Retention Labels

Although auto-label retention policies are a good way to find and label items automatically, several other methods exist to apply retention labels to content. These include applying labels by:

- Defining a default retention label for a SharePoint Online or OneDrive for Business document library. This method requires Office 365 E5 or Microsoft 365 Compliance E5 licenses.
- Using Power Automate to assign a label to a document (here's [an example](#)).
- Applying a retention label to a SharePoint Syntex document understanding model.
- Programmatically using the [SetComplianceTag method](#).

Automatic assignment of retention labels usually requires a specific license. For example, if you define a default retention label for a document library, users accessing that library need Office 365 E5 or Microsoft 365 Compliance licenses. Applying retention labels using mail flow rules doesn't require any additional licenses.

## Retention and Sensitivity Labels

Two forms of compliance labels exist within Office 365:

- **Retention labels:** Mark content like documents and messages that the organization wants to retain, like keeping certain documents because they hold valuable information, such as accounting records, sales records, and so on. Users apply labels to messages or documents to mark the items as a visual marker of the importance of the content or to make sure that Microsoft 365 keeps the item for a set period. Two types of retention labels exist: standard and record. Standard labels are the more common type while the major use of record labels is in the Records management solution, covered later.

- **Sensitivity labels.** Apply marking with optional protection to content. [Sensitivity labels](#) are how Microsoft has extended the functionality of the original Azure Information Protection labels as part of their Microsoft Information Protection initiative. They refer to these labels as “unified” because they bring together the work done by Azure Information Protection to make it easy for users to self-classify content by applying labels in the Office applications with the data governance framework. Applying a sensitivity label to an item can protect it through encryption or add visual indicators such as watermarks to show users the importance or sensitivity of the information. See Chapter 20 for more information about how to define sensitivity labels and deploy them to users through sensitivity label policies.

When we refer to the two types of labels in general, we say “labels” or “Office 365 labels.” Otherwise, we refer to the specific type. The information presented here covers retention labels.

## Retention Label Concepts

A retention label can be passive, meaning that it serves as a marker for a certain type of content but takes no further action because it doesn’t have any retention settings. Passive labels are useful to mark content that someone will need to process in the future or to find items belonging to a project. Mostly, retention labels are active, meaning that the label has a defined retention action and period. The action tells the owning workload to keep content for the period or to remove content after that period. For instance, the organization might decide to retain all documents marked with the “Confidential” label for five years and removed afterward. Auto-label policies remove the need for users to apply labels manually. The combination of manual (explicit) and auto-applied (implicit) classification gives tenants great flexibility in how they manage important content.

Items such as a document or message can only ever have one retention label (it can also have a sensitivity label). Apart from the precedence of manually-applied retention labels, the other rules are:

- Anyone with write access to content can change the label assigned to that content whether the item receives a label manually or automatically. For example, if an email has the “Confidential” label, the user can go ahead and change that label to any other available label. The exception is retention labels that mark items as formal company records. Once a user applies a record label to an item, the item can’t be edited, updated, or removed until its retention period lapses (for instance, Exchange Online keeps deleted records in the Recoverable Items structure).
- Retention labels applied automatically can never overwrite labels manually applied by users.
- If multiple auto-label policies match an item, the workload uses the label belonging to the oldest policy. On the creation of a label policy, the policy receives a priority number incremented from 0 (zero) as more policies are created. Thus, the policy with the lowest number is the oldest. You can discover the priority order for policies by running the *Get-RetentionCompliancePolicy* cmdlet as shown below:

```
[PS] C:\> Get-RetentionCompliancePolicy | Format-Table Name, Priority, WhenChanged
```

Name	Priority	WhenChanged
Management Preservation Policy	0	13 Apr 2018 13:02:36
Company Confidential Policy	1	29 Apr 2017 03:07:38
Clean up Groups with Connectors	9	29 Apr 2017 03:07:40
Formal Company Records	13	29 Apr 2017 03:07:39

You cannot change the priority order of retention policies.

## Retention Label Publishing Policies

Broad retention policies assign retention settings to locations. Retention label policies make one or more retention labels available to end-users through a publication process to inform workloads about the labels and their settings. It is then up to the workload to decide how best to reveal the labels through the different

clients supported by the workload. A label policy is composed of one or more retention labels. A retention label can be in multiple label publishing policies. For instance, the “Confidential” label referred to above to keep content for five years could be in the policies assigned to different departments along with other labels that meet the specific needs of each department. The members of the legal department might have a policy that includes labels called “Case Review”, “External Counsel”, and so on while people working in Accounts might have labels for “Collections”, “Audit Records”, and “Tax.” Being able to create label publishing policies with a mixture of generic labels and work-specific labels supports flexibility in data governance. After all, not everyone who works in a company needs to deal with information in the same way.

When complete, administrators publish the policy to make the labels available to end-users. After a period, the labels included in the policy become available to clients and users can then apply the labels to email, documents, and other content to mark those items as being of interest for business reasons. Behind the scenes, background processes make sure that workloads respect the instructions contained in the label settings. For instance, items stamped with the “Archive Retention” label might be kept for 10 years before being removed.

## Planning Retention Labels

The first thing to decide is what labels the organization needs to build out the retention strategy. Broadly speaking, you can divide retention labels into the following categories:

- **Targeted:** Retention labels needed by certain departments and used for specific purposes. For example, “Project Documentation,” “Board Minutes,” or “Patent Material.” These labels usually keep information that is of high importance to the organization and might be published to a select group of locations.
- **Generic:** Retention labels used anywhere in the business. Often, these labels are named after the length of the retention period, as in “Keep Five Years.” They might also have names that describe the business purpose, like “Commercially Sensitive” or “Required for Audit.” These labels are usually published to all locations in an org-wide policy.
- **Special-Purpose.** Retention labels that the organization creates for a specific well-defined purpose. For example, to mark and keep a set of information needed for a merger and acquisition project.

After consulting with business units (including the IT department) to gather suggestions for retention labels, you can rationalize the set to a manageable number. Each label should serve a distinct and obvious purpose definable in clear and easily understandable terms. In addition, you should be able to say where the records marked by labels are stored. For example:

*“We are required to preserve financial records for five years because we can be audited during this period. We need a label to mark these records and ensure that they are retained for at least five years. Items needed for audits include messages and documents across all mailboxes and sites.”*

It is sensible to write down each of the retention labels that you plan to use before creating anything. It is much easier to delay the release of a label and the training of users to use the label properly than it is to launch a label into general circulation only to discover that you later need to withdraw it. Another thing to consider is how easy it is for users to decide between different retention labels when the time comes for them to apply a label. Too many labels, misleading names, or too many choices can lead to frustration and bad decisions.

Other points to consider include:

- An item can only ever have a single retention label. To change a label, you must remove the first label and replace it with another. Sometimes, you might have to remove a label from an item before you can remove the file.

- A retention label applied by a user always has the highest priority. A workload never applies a retention label to content if a label applied by a user already exists. On the other hand, a user can always replace a label applied by an automatic policy with one that they choose. The logic here is that the user understands the full import of the content and can therefore make the best decision about its importance.
- If you use automatic label policies to apply labels to content found using a keyword query or because the content has sensitive data, you might find yourself in a situation where some documents come within the scope of multiple policies. a tiebreaker decides which policy to apply. The tiebreaker is the age of the policy, and the workload always complies with the oldest policy. The logic here is that companies usually create their most important policies first, so it makes sense to apply those policies first. This is a key factor in the planning process as you must decide what your most important policy is and make sure to create it first. No way exists to reorder the priority of automatic label policies as the only factor is a policy's creation date. Therefore, if you make a mistake with your priority, you must remove policies and recreate them in the correct order.
- Although auto-label policies cannot replace labels assigned by users, they can replace labels previously automatically assigned to items by other policies.

In summary:

- Make sure that every retention label has an obvious purpose.
- Try to have a small number of retention labels so that it is easier for users to make good choices about for how long they should keep content.
- Create retention labels and policies in priority order.
- Deploy auto-label policies after users have had a chance to apply retention labels manually.

**Removing a Site:** Before you can remove a SharePoint Online site, you must remove any documents that have retention labels from the site. Normally, this means that you must remove the label from the documents and then delete them.

## Creating New Retention Labels

After understanding the labels necessary to implement the data governance policy, we can create them through the Data lifecycle management section of the Microsoft Purview Compliance portal. Click **Labels** and then **Create a label**. You can then start by entering the name of the label and some text to describe the purpose of the label for administrators and a separate description that is visible to users when they browse labels as they classify material.

### Naming a Retention Label

The name given to a retention label is important because this is what users see when they use the label to classify a message or document. One issue for multilingual tenants is that no facility exists to translate labels. Whatever name you give to a label shows up in clients, no matter what language they use. For this reason, it is best to give labels names that are simple to understand and unambiguous in their intent as this will make it easier to communicate how to use the labels to classify information.

### Create retention label

**Name**

Retention settings

Finish

#### Name your retention label

This is the name of the label your users will see in the apps where it's published (like Outlook, SharePoint, and OneDrive). So be sure to come up with a name that helps them understand what it's used for.

Name \*

Strategic Planning

Description for users

Documents and other files connected to corporate strategic planning. We'll keep these documents for ten years.

Description for admins

Documents and other files connected to corporate strategic planning. We'll keep these documents for ten years.  
Created 1-Oct-2020

Next Cancel Up

Figure 17-13: Entering the name for a new retention label

load ... 1 selected DocViewModified

Office 365 for IT Pros 8.pdf

None Clear the label

Approved Retain for 10 years

Audit Material Retain for 10 years

Board Records Retain forever

Commerically Sensitive Retain for 1 year

Confidential Retain for 5 years

Contractual Information Retain for 7 days

Draft

Locked Down Retain for 15 years

Manual disposition Retain for 6 months

Office 365 for IT Pros eBook Content Retain for 10 years

Patent Materials Retain for 20 years

Record (Legal) Retain for 1 year

Regulatory Record (Legal) Retain for 7 years

Name \*

Office 365 for IT Pros 8.pdf

Title

Enter value here

Thumbnail

Enter value here

Image Tags

Enter value here

Apply label

None

All our files for the best book

Edit columns

Figure 17-14: Selecting a retention label to apply to a SharePoint document

For example, if we publish the label shown in Figure 17-13 to OneDrive for Business and SharePoint Online, people can use the label to classify a document stored on a site by amending the document properties and selecting the label. As you can see in Figure 17-14, the user gets two visual hints about the meaning of the label: its name and the descriptive text that appears when they hover over the label. You can also see that it's

possible to include Windows emojis in the display name of retention labels to deliver another visual hint to users about a label's purpose.

## Retention Actions and Periods

The next step is to define how workloads will process items with a retention label. You must choose what action the label will apply and when the action occurs. A retention label can:

- **Retain** items for a **specific period**. The retention period is the number of days that elapse before workloads process the retention action.
- **Retain** items **forever**. If users attempt to delete a labeled item, the workload will keep a copy of the item in a safe location where the user cannot access it. For example, Exchange Online moves the copy to the Recoverable Items folder.
- **Enforce actions** after a specific period. Once the item reaches the set age (retention period), it is removed.
- **Just label** items. These labels act as visual markers for the type of information in an item (sensitivity labels with no encryption that apply watermarks, headers, and footers might be a better choice). For instance, you could create a "Draft" label to allow users to mark items that have not yet reached the point where the content is interesting or valuable enough to justify its retention. Another way of using labels without actions is as a convenient way to find information with content searches (for instance, find all documents for "Project X"). Labels that do not have a retention action show up with a "Never" expiration date when viewed through Outlook or OWA.

If you choose to keep items for a specific period, you can choose between date-based retention or event-based retention as the baseline for the retention period. For date-based retention, you select one of the three predefined periods (5, 7, or 10 years) or a custom period in years, months, and days. The GUI allows the selection of a retention period of over 100 years, but a shorter period is best. You then choose the basis for the retention period:

- The **creation date** for items: This is the default and works well for most items.
- The **last modification date** for items: This is a better retention base for documents that go through multiple review cycles before finalization.
- The date when a user or policy applies a **label** to an item.

If you use event-based retention instead of date-based retention, the retention period starts when the event occurs. Some of the standard event types are listed below. You can create additional event types to meet business needs. We cover event-based disposition in the Records Management section.

- Employee activity.
- Expiration or termination of contracts and agreements.
- Product lifetime.

The user interface presented to create a retention label depends on its settings. Figure 17-15 shows the how to configure a label that the label sets a seven-year retention period based on the creation date for items. When the retention period for an item expires, Microsoft 365 deletes the items. This means that the next time that the item goes through retention processing, its expired status causes the workload to apply its normal deletion process. For instance, a document removed from a SharePoint Online site goes into the recycle bin.

**Create retention label**

Define the period

Choose how long the period is and when it begins.

How long is the period?  
7 years

When should the period begin?  
When items were created

**Create retention label**

Choose what happens after the retention period

These settings determine what happens to items when the retention period ends.

Choose what happens after the period

Delete items automatically  
We'll permanently remove labeled items from wherever they're stored.

Change the label  
You can extend the period by choosing an existing label to replace this one with. [Learn more about relabeling items](#)

Figure 17-15: Defining retention periods and actions for a new label

You can also choose to trigger a disposition review (change the label). When this happens, the item stays in place until a designated person reviews it to decide to keep or remove the item, or do nothing (deactivate retention settings).. The item remains in place and it's up to its owner to decide what to do with it.

When all the details are complete for the new label, click **Next** to review the settings and then take the **Create label** option.

**One chance to get label settings right:** Be careful with the settings you specify when creating a new label as you can only change the retention duration afterward. The logic is that allowing other changes after label creation will disrupt how the label behaves when applied to content. For example, if you change a label that keeps content for ten years and does not remove it afterward to now remove the items, users might lose information that they expect to have.

## Publishing Retention Labels Through a Label Policy

To make the new label available to users, you must publish it in a label policy. You can publish the new label in a separate policy, but it's usually better to create policies to publish sets of labels to users. For instance, you might have a label policy called "Finance Department Labels" that includes all the labels needed by workflow processes in the finance department. The department users can apply the labels when an administrator publishes the policy to them.

Start the publication process by going to the **Labels policies** section under **Data lifecycle management** to view the set of label policies defined in the tenant. You can then select an existing policy and edit it to add or remove labels from the policy or use the **Publish labels** or **Auto-apply a label** option to create a new label policy. The first type of policy makes the label available to workloads, the second sets the conditions for the automatic application of a label. You can select several labels and take the action to publish the set of selected labels in a single policy (bulk publish). We will explain how to create an auto-applied policy in the next section.

After deciding on the set of labels to include, you select the locations where the labels are available to users. The options for static policies are:

- **All locations:** This means that the labels in the policy are available to all users in the following workloads. This is an org-wide policy.
  - **Exchange Online mailboxes.** Labels appear in clients in the same way as Exchange retention tags and behave like retention tags. For example, you can create a rule in Outlook desktop to apply a label to messages, such as all those that come from a specific address.
  - **SharePoint sites.** Users assign labels to classify individual documents or items inherit a label because it is the default for a document library. When a default label is present, new items created or uploaded in the library inherit the label.
  - **OneDrive for Business accounts.** Users assign labels to classify individual documents stored in OneDrive folders.

- **Microsoft 365 Groups.** Users apply labels to conversation items in Outlook Groups. Yammer communities don't support the labeling of conversations. Because SharePoint supports retention labels, they can classify information in the document libraries belonging to Outlook Groups or Yammer communities.
- **Let me choose specific locations:** For each of the supported workloads, you can opt to include or exclude certain mailboxes, sites, OneDrive for Business accounts, or groups. For mailboxes and groups, you enter the names of individual mailboxes or distribution lists you want to include or exclude. For site collections, you enter the URL in the form `https://mytenant.sharepoint.com/`. For OneDrive for Business accounts, enter the URL for the site. Figure 17-16 shows the locations available for label publication. We can see that a single Exchange recipient is selected together with all SharePoint sites and OneDrive accounts, and five groups.

## Edit retention policy

**Choose locations**

We'll publish the labels to the locations you choose.

All locations. Includes content in Exchange email, Office 365 groups, OneDrive and SharePoint documents.

Let me choose specific locations.

Status	Location	Included
<input checked="" type="checkbox"/> On	Exchange email	1 recipient <a href="#">Edit</a>
<input checked="" type="checkbox"/> On	SharePoint sites	All sites <a href="#">Edit</a>
<input checked="" type="checkbox"/> On	OneDrive accounts	All accounts <a href="#">Edit</a>
<input checked="" type="checkbox"/> On	Office 365 groups	5 groups <a href="#">Edit</a>

Back **Next** Cancel

Figure 17-16: Selecting locations to publish labels in a retention label policy

## Scoping with Non Org-Wide Policies

Label policies with excluded or included locations are known as non-org wide policies. You can have up to 1,000 non-org wide policies in a tenant. This figure includes both retention and label policies.

When you exclude or include locations in a label policy, certain limits exist in the picker used to select target locations:

- For **Exchange**, if you don't enter anything in the search box, the picker shows the first 50 mailboxes and distribution lists in the directory. Enter a search phrase to find the mailboxes to include or exclude. The easiest way to add many mailboxes at one time to a policy is to use distribution lists. The membership of the list is expanded, and the mailboxes are added to the policy.
- For **SharePoint**, you can include or exclude up to 100 sites. These are sites, not libraries within sites. Use this location to publish labels to traditional SharePoint sites that are not connected to Groups.
- For **OneDrive for Business**, you can include or exclude up to 100 accounts by specifying the URLs for the target accounts.



- For **Groups**, the picker shows the first 100 groups in the tenant, including those hidden from Exchange clients (used by Teams). You can add up to 1,000 groups to a label policy. Labels published to Groups apply to the chosen group mailboxes and the group team sites. Use this location type to publish labels to modern team sites connected to Groups.

Workloads like Teams and Planner do not currently support retention labels, so there is no need to include them in the publication process. Teams does support retention policies.

Click **Next** after selecting all the target locations. You now name the policy and give some optional information to explain its purpose. Click **Next** to review the settings for the policy. If any issues are detected at this point (for instance, you enter a SharePoint site instead of a site collection), the policy cannot be published, and you must fix the problem before you can continue and save the policy.

When everything is ready, click **Publish labels** to begin the provisioning process that makes the policy available to the target locations. You know when the provisioning process is complete when the policy status changes from "Pending" to "Success." Once published, the new label is available to the target workloads defined in the policy. This process can take up to one day to complete as, in some cases, clients must find out about the new labels. For example, the XML data used to inform Outlook desktop clients about retention policy settings must receive an update with details of the new label. The clients then need to download the label information before the new label appears in Outlook's user interface. Web clients typically pick up new labels faster, but it is reasonable to expect that the entire end-to-end publication and provisioning process might take one or two days before a new or updated label is available throughout the tenant.

## Comparing Retention Labels and Retention Policies

A retention policy (not a label publishing policy) applies a single retention setting to everything in a container or location (mailbox, site, group, or team). A retention label applies retention settings at the item level. Together, the combination of retention policies and retention labels gives administrators a lot of flexibility in planning a retention strategy for content. However, policies and labels support different settings, and this can be confusing at times.

Let's summarize what retention labels and retention policies can do:

- You can apply retention actions and periods through both policies and labels. Some workloads (Teams and Yammer) only support the application of the same settings to everything in the targeted containers. Others (Exchange, SharePoint, and OneDrive) support label assignment to individual items and targeted containers.
- A retention label can act purely as a visual marker. In this case, the label has no retention settings. Because it can add visual markings to documents and messages, a sensitivity label might be a better choice for this purpose.
- Both policies and labels support the ability to retain content for a set period, remove content after a set period, or retain for a period and then delete it. Both can use the created date or the last modified date as the date when the retention period begins.
- A retention label assigned explicitly to an item can override the retention period imposed by the policy assigned to a container. For instance, if a site has a retention policy that removes items after three years and an individual document in the site is assigned a retention label with a retention period of five years, SharePoint will remove all the unlabeled items on the site after three years and keep the labeled document for five years.
- After the retention period lapses, both policies and labels allow items to remain in place for the user to decide what to do with them. An analogy is government papers that have restricted access for ten years after creation. When the ten-year period expires, the papers are not dumped. Instead, they become available for public access. This subtle differentiation is important for records management.

- Because they are more explicit, retention labels offer extra control over what happens when a retention period is over. A retention label can invoke manual disposition or use event-based retention. We'll get to these topics later.
- A retention label can classify an item as a formal record or regulatory record. Labels created for this purpose are managed through Records management – see later section.
- Labels assigned by auto-apply policies can be replaced by explicitly (user-assigned) labels. Explicitly-assigned labels can never be replaced by automatic assignments.
- Labels inherited from a container (like the default label for a SharePoint site) can be replaced by labels assigned by auto-apply policies.
- Unlike Exchange Online retention tags, neither retention policies nor retention labels support a “move to archive” action. If moving mailbox items to the archive is important, you can continue to use mailbox retention policies in conjunction with retention policies.

Figure 17-17 shows the retention settings for a retention policy (right) and retention label (left). These are old UI screens (2021), but they show the extra flexibility available in retention labels better than the current GUI. The ability to set a label to mark items as records (label classification) is only available when editing retention labels through the File Plan section of Record Management in the Microsoft Purview Compliance portal.

### Define retention settings

When this label is applied to items, the content is retained and/or deleted based on the settings you choose here.

Retain items for a specific period  
Labeled items will be retained for the period you choose. During the retention period, Users will be able to edit items and change or remove the label. If they delete items, we'll keep copies in a secure location. [Learn more](#)

Retention period of  years  months  days

Retention period

Start the retention period based on

+ Create new event type

At the end of the retention period

Delete items automatically

Trigger a disposition review

Do nothing  
Items will be left in place. You'll have to manually delete them if you want them gone.

Retain items forever  
Labeled items will be retained forever, even if users delete them. Users will be able to edit items and change or remove the label. If they delete items, we'll keep copies in a secure location. [Learn more](#)

Only delete items when they reach a certain age  
Labeled items won't be retained, but when they reach the age you choose, we'll delete them from where they're stored.

**Label**

### Decide if you want to retain content, delete it, or both

Retain items for a specific period  
Items will be retained for the period you choose.

Retain items for a specific period of  years  months  days

Start the retention period based on

At the end of the retention period

Delete items automatically

Do nothing

Retain items forever  
Items will be retained forever, even if users delete them.

Only delete items when they reach a certain age  
Items won't be retained, but when they reach the age you choose, we'll delete them from where they're stored.

**Policy**

Figure 17-17: Retention settings in a retention policy (right) and retention label (left)

In summary, retention labels offer more flexibility, but they must be assigned manually or through auto-label policies. The retention settings in policies are simpler, but they are easier to apply because you can deploy retention across the entire organization or for a selected set of locations with just a few policies.

## Using Retention Labels

We now know how to set up, publish, and manage labels and understand the basic principles that only one label can exist on an item at any time, whether a user assigns the label to an item manually or a background process assigns the label automatically. With those thoughts in mind, we can now go ahead to discuss how to use labels within the different locations.

## Using Retention Labels with Exchange Online

Publication of a label means that workload-specific mechanisms make the new label available to users of an application. For Exchange Online, that mechanism is to insert the label into the set of retention tags available to a mailbox through its assigned mailbox retention policy. The Managed Folder Assistant (MFA) creates a unified set of retention tags and retention labels and publishes the set to user and group mailboxes (group mailboxes only see retention labels and not tags). MFA includes only labels with a retention action in its published set.

MFA operates a workcycle to process mailboxes at least once every seven days. Therefore, it might take up to a week before new retention labels become available to mailboxes. See the “Logging the Managed Folder Assistant” section later to understand how to extract and interpret MFA diagnostic logs for a mailbox to know when MFA last processed the mailbox.

Mailboxes must hold more than 10 MB of content before the MFA processes the mailbox to make retention labels available. This restriction exists to ensure that MFA does not waste processing cycles on unused mailboxes. When it opens a mailbox for processing, MFA makes sure that the current set of labels is known to the mailbox before it begins to check items against their retention status.

Integrating retention labels with retention tags allows OWA and Outlook desktop to handle the two kinds of tags consistently. In effect, OWA and Outlook treat labels in the same way as personal retention tags and make the labels available to users to tag individual messages or folders. Figure 17-18 shows how Outlook presents a mixture of retention tags and retention labels to the user when they want to apply a policy to a folder.

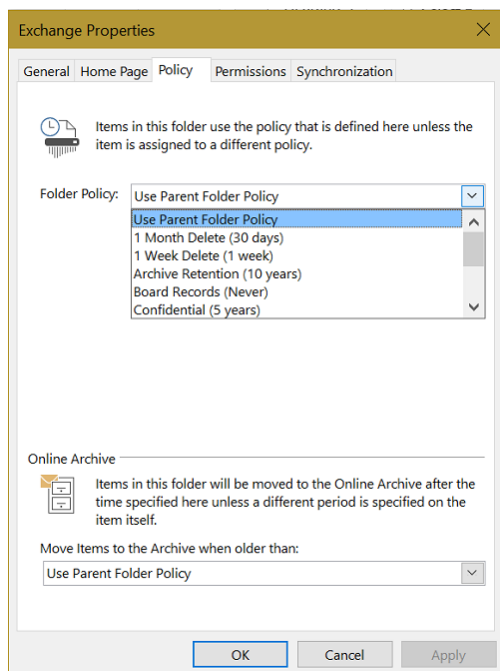


Figure 17-18: Outlook lists retention tags and labels

Because Outlook clients display retention labels alongside personal retention tags, users should not see any difference between the two types. If a user assigns a label to a folder, all items in the folder inherit that label, unless the item already has a personal tag or another retention label. Likewise, if you change a label at the folder level, the next time MFA processes the mailbox, it updates the items in the folder with the new label unless they already have an explicit label or tag.

Because Exchange public folders do not support retention policies, they also do not support retention labels.

## Integration with Exchange Retention Policies

Although OWA and Outlook present retention labels to users in the same way as they see personal retention tags from mailbox retention policies, we already know that labels function differently from retention tags.

Table 17-4 lists some of the differences between the two methods used to mark mailbox items for retention processing.

<b>Feature</b>	<b>Mailbox Retention Tag</b>	<b>Retention Label</b>
Remove content from mailboxes	Yes, when the retention period expires.	Yes, when the retention period expires.
Archive content from mailboxes	Yes, when the retention period expires.	No. Labels do not support an archival action.
Expiration of control	No. A retention tag stays with an item until it is removed from the mailbox.	Yes. The effect of a label ceases when it expires.
Policy-driven tagging of default folders	Yes. Retention policies can include folder tags for any of the Exchange default folders (like Inbox).	No. A retention label functions as a personal retention tag and can be applied to any other folder except default folders.
User-driven tagging of items and folders	Yes. Retention policies can include personal tags for users to stamp on non-default folders and individual items.	Yes. A label can be used in the same way as a personal tag to mark items or folders.
Policy-driven default retention	Yes. Retention policies can include default tags that apply to all items in the mailbox that are not stamped with a more explicit tag. A policy can include default tags to define removal and archival actions and you can have a specific default tag for voicemail.	No. It is possible to create an auto-label policy that applies to all items in the mailbox, but that is different from a default tag.
User-selectable tags	Yes. Users can select personal tags (through OWA options) that are not in the assigned retention policy and use them to tag items and folders.	No. Users do not have access to labels not published to their mailbox.
Target	Limited to Exchange mailbox folders, items within a conversation, individual items, or complete mailboxes.	As for retention tags, with the addition that labels can be used in other workloads.

Table 17-4: Differences between Exchange retention tags and retention labels

The last point is the most important. Retention policies and tags only cover Exchange content. OWA and Outlook desktop clients combine retention labels with the tags published through Exchange mailbox policies.

Meaning users can apply labels to email items in the same way as they use retention tags. Of course, the big difference between labels and tags is that labels are available in other workloads.

## Using Retention Labels with Groups

When you publish retention labels to a group, they become available in both the group mailbox and the group document library. At the time of writing, OWA is the only client that supports the application of labels to conversations in Groups. Only group owners can use labels to classify conversations as OWA hides the labels from ordinary group members.

However, all group members can assign retention labels to files and folders in the group document library and any group member can overwrite a label previously assigned by another group member, except if that label is a formal record (in which case only the site collection administrator can update the item). The reasons why labels behave differently in a group's mailbox and document library are because of the different ways that Exchange retention tags and SharePoint permissions work. An Exchange mailbox is typically under the sole control of its owner while a SharePoint site is designed to support multiple levels of shared access.

## Using Retention Labels with SharePoint and OneDrive for Business

Any user in the default members group for a SharePoint site (with the Contribute permission level) can apply labels to documents, folders, or items in a list within the site. If the site belongs to a group, any of the members can apply labels because they all have equal access to the site. The owner of a OneDrive for Business account can apply labels to content within their account. When you apply a retention label to a folder, all items in the folder inherit the same label, unless some of the items in the folder already have labels. Applying a retention label to a folder holding thousands of files can take a little time to complete. A retention label inherited from a folder stays with a document even when the document moves to another folder or another site. Also, if you upload a file to replace an existing document that has a retention label, the uploaded file inherits that label.

To apply a label to an item in OneDrive for Business or SharePoint, select the document or folder and then open the Details pane. Go to the **Apply Label** section and select a label from the set published to the location. Placing a label with a retention action on an item has some consequences:

- Depending on the Retention labels settings (available in the Records management section of the Microsoft Purview Compliance portal), users might be unable to remove a labeled item. See the section covering Record management settings.
- Users cannot edit documents marked with a label that classifies items as a formal record. They can update documents marked with labels not classified as formal records. Site administrators can remove the label or replace it with another if users need to edit the item.
- SharePoint records updates and deletions for documents in the site's Preservation Hold library.
- If a document is subject to a retention policy, users can remove it from the library, but because the document comes within the scope of a retention policy, SharePoint must keep a copy. In these instances, SharePoint Online moves the deleted document into the Preservation Hold library and keeps it there until the retention period expires.

The behavior is different for files held in OneDrive for Business because these sites are personal rather than shared and the owner of the site has the right to remove files from the site. When a user removes a file from their OneDrive for Business site, the file goes into the site recycle bin. Thereafter, when the retention period for the recycle bin expires (93 days), a timer job examines expired items and moves copies of any item with a label into the site's Preservation Hold library.

A label is a managed property that SharePoint indexes along with other document attributes. Indexing occurs when the crawler accesses new or updated documents. It can take up to an hour before a newly-assigned

label is in the index. When indexed, you can search for documents tagged with a retention label by using the *complianceTag* property. For example, you can input *complianceTag:"GDPR Personal Data"* in the SharePoint search box or a content search to find documents stamped with the *GDPR Personal Data* label.

## Displaying Retention Labels Used in a Document Library

The Outlook clients include the necessary user interface to make retention tags available to users, including those published by retention policies. This is valuable because users get a clear visual understanding of for how long Exchange will keep an item. Retention labels are not one of the default fields shown for SharePoint document libraries or lists. You can make the labels more obvious by customizing the view of items in the library or list to include the “Labels” and (if necessary), “Item is a record” fields. Figure 17-19 shows how information about assigned labels appears in a document library. Note that the column only displays the label name and doesn’t show anything else such as the outstanding retention period. This points to the need to either include the retention period in label names or to coach users about what each retention label does.

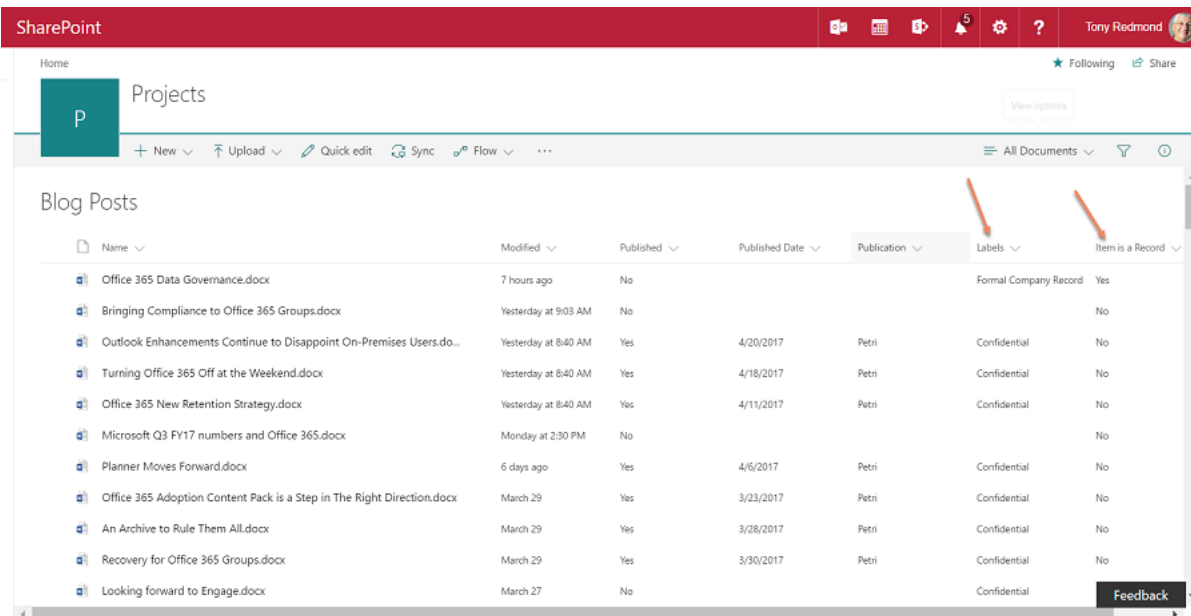


Figure 17-19: Viewing label information in a SharePoint document library

Unlike when you update properties like a document’s name or title, SharePoint does not treat applying a label as a modification. Instead, it is more of an administrative event. SharePoint, therefore, does not update the “Modified By” property with the name of the person who applies the label.

An organization might have many retention labels in active use at any time. It can be confusing for users to have to choose between multiple labels when they classify documents. In addition, some labels might be inappropriate for the content of certain sites. With these points in mind, it is sensible to consider what sites should receive a label when you publish it. Some labels are general purpose and are useful across all sites while others are better if restricted to specific sites.

## Finding Items Marked with Labels in Content Searches

You can create content searches (either standalone or part of an eDiscovery case) to find content marked with specific labels. You do this by using the “ComplianceTag” keyword in searches. For example, this search finds any items stamped with the “Draft” or “Approved” labels.

*(ComplianceTag:Draft) OR (ComplianceTag:Approved)*

If the label name contains spaces, enclose the name in quotation marks. See Chapter 18 for more information about eDiscovery searches.

## Audit Records Generated for Retention Label Actions

When someone assigns, changes, or removes a retention label from a document or folder in a SharePoint Online or OneDrive for Business library or an item in a list (including a list created in the Lists app), a *ComplianceSettingChanged* audit record captures the event.

- For retention labels applied by a user or retention labels assigned because a default label exists for a document library, the *UserIds* property holds the user principal name of the account that executed the action.
- For retention labels applied by the background job for auto-label policies, the *UserIds* property holds the GUID of the background job. You will see a value like *1b1c17be-a6f8-4691-9fca-9b6ac128c9e1*.
- SharePoint Online does not generate audit events for retention labels inherited by new documents because a site has a default label. This is an acknowledged gap in Microsoft's compliance story.

If a user removes a retention label from a document or list item, SharePoint Online captures the *TagUnApplied* audit event.

Audit records captured for retention label actions in a library include the site, the document name, the user, and the name of the label. If a label previously existed for a document, the audit record captures both the original (*SourceLabel*) and new (*DestinationLabel*) classifications. For label actions involving list items, SharePoint Online writes the item number into the audit event, so you see values like:

[https://office365itpros.sharepoint.com/sites/GDPRPlanningMarkII/Lists/Things to do/2\\_000/2](https://office365itpros.sharepoint.com/sites/GDPRPlanningMarkII/Lists/Things%20to%20do/2_000/2)

Exchange Online does not capture audit events when messages receive retention labels. This is because Exchange treats retention labels in the same way as mailbox retention tags. Exchange does not capture the assignment of mailbox retention tags as mailbox audit events.

### Analyzing Retention Policy Updates

Retention policies exercise control over user information, so it's wise to keep an eye on changes made to policy settings to ensure that administrators don't make mistakes that result in inadvertent loss of data. Microsoft 365 captures audit records for additions, updates, and removals of retention policies and retention label policies. [This article](#) explains how to analyze audit events for retention policy changes to report change details. The situation is more complicated than it first appears due to the information captured in audit events for different types of retention policies.

## Default Assignment of Retention Labels

If you have Office 365 E5 licenses, one solution is to classify items by using an automatic label policy to find documents using searches based on keywords or sensitive information types. This is a good way to make sure that documents with certain characteristics (like HR personnel files) are classified no matter where they are stored within SharePoint.

Another method to ensure that all items in a SharePoint or OneDrive for Business list or library are assigned a retention label is to assign a default label. Microsoft considers this to be an advanced feature, so it also requires an Office E5 license. To assign a default retention label, go to a library and select **Apply label to items in this list or library**, and then choose the default label from the set of available retention labels (Figure 17-20). In addition to assigning a default label to new items, you can choose to have SharePoint or OneDrive for Business apply the selected label to existing items. If you need to add default retention labels to multiple libraries, it's possible to [use PowerShell to script the assignment of a default label](#) to a SharePoint library.

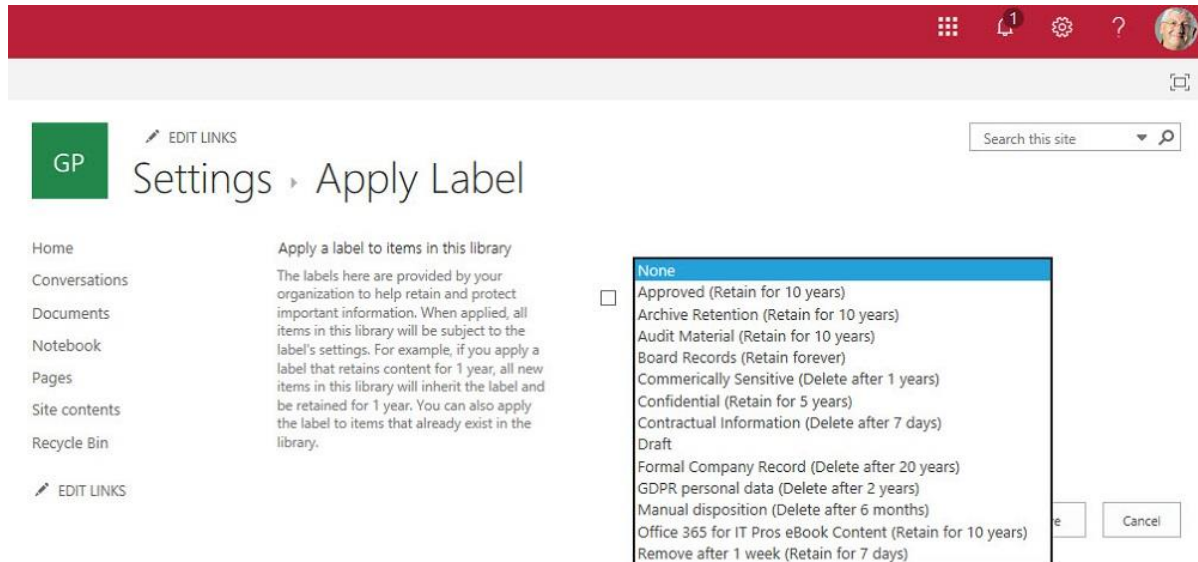


Figure 17-20: Applying a default retention label to a SharePoint library

**OneDrive for Business Library Settings:** If you want to set a default retention label for a OneDrive for Business account, you must either enter the URL for the library settings (it's the settings.aspx page) or use the old-style interface. Library settings are available when you expose the ribbon, select the library tab, and then library settings. You can then select a retention label as described for SharePoint above.

## Applying Retention Based on Sensitivity Labels

Given that files with sensitivity labels often hold confidential information that the organization wishes to keep (or wants to remove after a set period), it makes sense to use auto-label policies to find documents with certain sensitivity labels to make sure that they have appropriate retention labels. In this example, we'll create an auto-label retention policy to assign a retention label to documents and messages protected by the Highly Confidential sensitivity label. To do this, you:

- Connect to the compliance endpoint with PowerShell by connecting to Exchange Online and then running the *Connect-IPPSSession* cmdlet.
- Find the unique identifier (GUID) for the selected sensitivity label by running the *Get-Label* cmdlet. The *ImmutableId* property contains the GUID.

```
[PS] C:\> Get-Label | ? {$_.DisplayName -eq "Highly Confidential"} | Select-Object
-ExpandProperty ImmutableId
```

Guid

```
-----
9ec4cb17-1374-4016-a356-25a7de5e411d
```

- Use SharePoint search to test the KQL query for the auto-label policy. The search term is in the form *InformationProtectionLabelId:9ec4cb17-1374-4016-a356-25a7de5e411d* where the managed SharePoint property that holds the GUIDs of sensitivity labels (*InformationProtectionLabelId*) is combined with the GUID identifying the sensitivity label you want to search for. Run the search and open one of the documents returned by the search to check that it has the correct sensitivity label. If no documents are found, it might indicate that the GUID is incorrect or that your account has access to no documents that have this sensitivity label.



- If the search term finds the correct documents, create an auto-label retention policy that uses the same search term as the content query to find the target documents and apply a suitable retention label to keep the documents for the desired period.
- Configure the policy to find documents in the desired target locations. Remember to use Microsoft 365 Groups to cover SharePoint sites owned by groups and teams. Publish the policy when everything is complete.
- After ten days or so, check that documents with the sensitivity label have the correct retention label, remembering that if a user assigns a retention label to a document, an auto-label policy won't replace it.

The ten days mentioned above are an estimate rather than a guarantee. It can take SharePoint Online anything from seven days to two weeks for a new auto-label retention policy to become operational and start to apply retention labels.

## Records Management

Microsoft Purview Records management is a solution managed through the section of the same name in the Microsoft Purview Compliance portal. Records management makes management of retention labels and policies easier in large organizations and includes these features:

- **File plan:** Enterprises typically use more retention labels than smaller organizations and often have retention labels created for specific purposes. A file plan gives extra flexibility in grouping and managing retention labels based on different properties, such as labels used for manual disposition. Administrators can assign a set of file plan descriptors to retention labels, such as the *authority* for the label (the intended business category requiring the use of the retention label – default values are business, legal, or regulatory). In addition, the file plan allows tenants to mark retention labels as formal company records and regulatory records.
- **Retention label policies:** The same functionality is available to create and publish labels as exists under Data lifecycle management.
- **Events:** Create the events used for event-driven disposition, such as the termination of a project.
- **Disposition:** Process items that have reached the end of their retention period and require a manual check to decide what should happen to the item (deletion, further retention, and so on).

These features all require Office 365 E5 or Microsoft 365 Compliance E5 licenses.

## Records Management Settings

Three important controls for retention labels are in the Settings section of Records management:

- **Deleting content labeled for retention:** Until early 2022, SharePoint Online blocked users from deleting labeled items while OneDrive for Business allowed them to do so. To achieve consistency across the two applications, Microsoft changed SharePoint Online to behave like OneDrive for Business, meaning that users can delete labeled items and SharePoint Online will store the items in the site's preservation hold library until their retention period expires. Some organizations prefer the previous behavior because they believe that users should not remove labeled items. If this option is set, users see an error if they attempt to remove a labeled item. To proceed, a site administrator or user with permission must remove the label to allow deletion to happen or replace the label with one that does not have a delete action.
- **Configure record versioning:** By default, record versioning is on, meaning that users can unlock items assigned a record label and edit their content. If off, items assigned a record label remain locked and updates are not possible after the creation of these items. In effect, record labels then act like regulatory record labels.

- **Allow editing of record properties:** Apart from its content, an item has metadata like its title and other attributes. By default, users can edit items assigned a record label to update metadata. If this control is off, users cannot update item metadata after creation.

Although the Records management solution requires Office 365 E5 or Microsoft 365 Compliance licenses, administrators can set these controls without those licenses. The controls apply to all sites in a tenant.

## Record Labels

Retention policies and labels play a key role in the ability of Microsoft 365 to meet the requirements of Rule 17A-4 of the U.S. Security and Exchange Commission (SEC), stating that companies that employ brokers, dealers, and other workers in the financial industry must keep records of their electronic communications for between three and six years, depending on the type of communication. Among the requirements set out are that the records should not be amendable or erasable by an administrator.

Microsoft 365 supports two forms of record retention labels to help organizations satisfy regulatory requirements. Unlike regular retention labels, you cannot create or edit record labels through the Data lifecycle management section of the Microsoft Purview Compliance portal. Instead, these special forms of retention labels are managed in Records management. The two types of record label are:

- **Record:** After a user applies a record label to an item, only administrators can remove the label or change it for another label. Anyone with write access to an Exchange mailbox can apply a record label to an item in the mailbox. Any member of a SharePoint Online site can apply a record label to a file or list item. Once applied, a record label stops any attempt to delete the item. Items with record labels stored in SharePoint Online and OneDrive for Business can have a locked or unlocked status. Users cannot update the content of a locked item, but they can update its metadata (like the title) subject to the tenant Record management settings for updating items marked as records. Any site member can unlock an item to permit editing.
- **Regulatory record:** This is a stricter form of record label. Not even a tenant administrator can remove a regulatory record label after its creation, and the only changes allowed to the label settings are an increase in the retention period or publishing the label to additional locations. After applying a regulatory record label to an item, no one can remove the label or delete the item until its retention period expires. Site administrators cannot change the locked status of an item, so no one can edit an item's content. However, users can open documents in review mode and save their content as a new file.

Not every organization needs to implement the strict retention regime implied by regulatory records. For this reason, before you can create new regulatory record labels, you must expose the UI to allow the compliance portal to manage regulatory record labels. Do this by connecting a PowerShell session to the compliance endpoint and running the *Set-RegulatoryComplianceUI* cmdlet.

```
[PS] C:\> Set-RegulatoryComplianceUI -Enabled $True
```

The command is effective immediately. To disable the UI to manage record labels, run:

```
[PS] C:\> Set-RegulatoryComplianceUI -Enabled $False
```

When you create a new retention label through Records management, you can choose to make the label act as a:

- Regular retention label.
- Record label.
- Regulatory record label.

The latter two options depend on enabling the necessary GUI. You can't select an existing retention label and transform it into a record label. Just like normal retention labels, the tenant must publish record labels

through a label policy to make the labels available to users. If you want to withdraw a record label from active use, you should disable the policy used to publish the label.

After assigning a record label to a document, folder, or list item, users cannot delete the item until the retention period specified in the label expires. Any site member can lock or unlock records. SharePoint indicates the locked status of records with a small padlock on the item or folder icon. You can also see the locked status in the item properties (Figure 17-21). Assuming the record management settings for the tenant allow, after unlocking a record, users can edit its content. If someone subsequently updates a record after the last unlock action, SharePoint Online captures a copy of the item (before editing) in the *Records* folder of the site Preservation Hold Library. SharePoint Online highlights the version it captures by writing **Record** into the *Comment* field in the version history. You can't unlock an item assigned a regulatory record label, meaning that you shouldn't assign these labels until after you are sure that the record contains the final content.

Someone might lock a record while another member is editing the file. When this happens, the file contains anything saved (by autosave or the last explicit save) up to the point the user executes the lock action. To generate a complete copy including the changes made since locking, the person editing the file will have to save it under a different name and then exit the edit session. After a short period, SharePoint frees the file to allow an unlock to proceed. It will then be possible to merge any differences between the versions into the file.

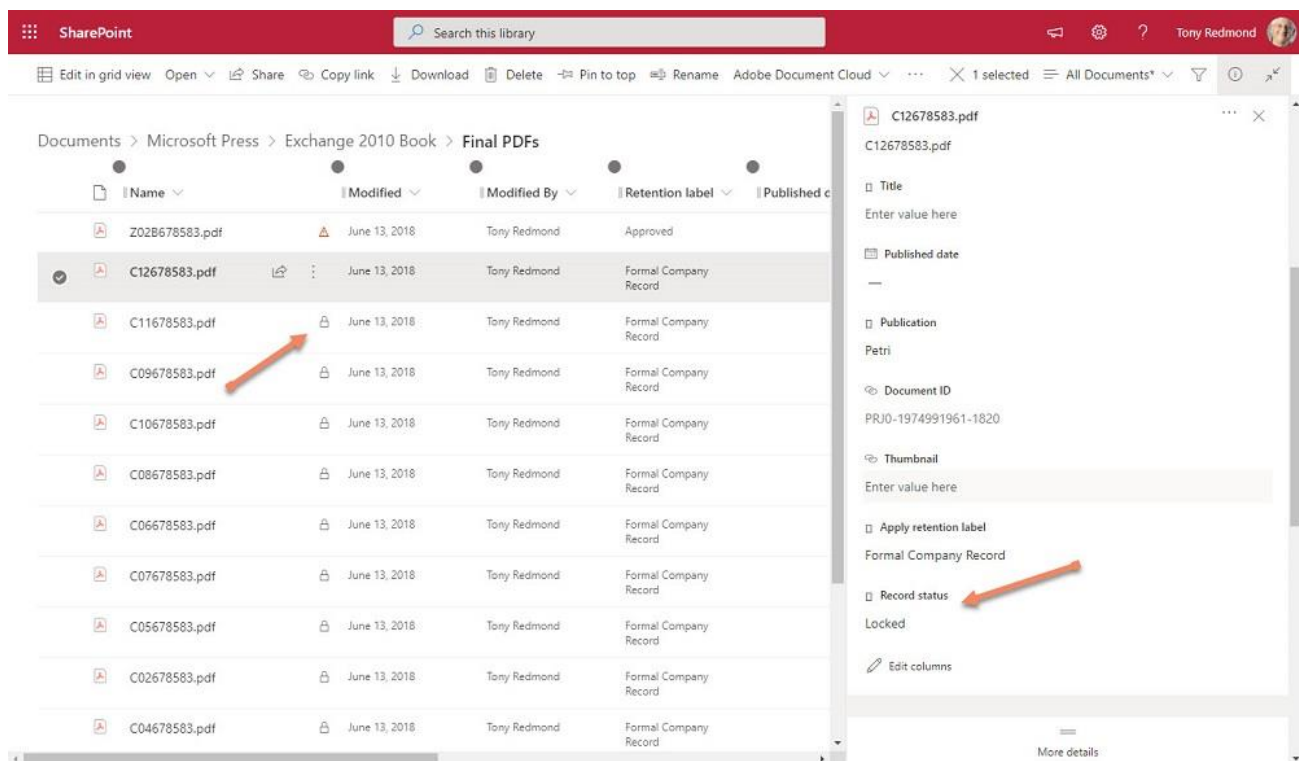


Figure 17-21: Documents in a SharePoint library marked as records

Although users need Office 365 E5 or Microsoft 365 E5 licenses to access the Records Management section in the compliance portal, any SharePoint Online or Exchange Online plan allows users to apply record labels to items. From a user perspective, after an item has a record label, it cannot:

- Permanently remove the item. If you try to remove an Exchange item labeled as a record, Exchange copies the item to the Recoverable Items structure and keeps it there until its retention period expires. If you try to remove an item marked as a record in SharePoint, an error is flagged, and the item is unchanged. If someone tries to remove a record from a OneDrive for Business account, OneDrive for Business moves the item into the preservation hold library where it remains until its retention period expires.

- Edit a locked item. Users can unlock an item assigned a record label and edit its content, but those assigned a regulatory record label cannot.
- If the assigned label is a regulatory record label, item metadata cannot be updated.
- Change or delete the label. A site administrator can remove a record label, but no one can remove a regulatory record label.

Only users can assign record labels to items. You cannot use an auto-label policy to apply regulatory record labels to items.

Owners of OneDrive for Business sites can apply and remove record labels to any item under their control. They can also remove items marked as records. However, for SharePoint content, while site members can apply a record label to an item, only the site administrator or a global administrator can change the label or update the properties of the item when it has a record label.

The implementation of regulatory record labels within Exchange differs from that used by SharePoint Online and OneDrive for Business. Browser interfaces interact directly with the server while Exchange must support the synchronization model that enables Outlook desktop clients to work offline for extended periods. After applying a regulatory record label to a message, a certain window of time is available to change the label. The window accommodates Outlook's synchronization model and the need to update the new label status across multiple clients. After a few minutes, the window closes, and no further change is possible. Also, when you apply a record label to an Exchange folder, all the items stored in the folder automatically become records, even if the user later moves some or all the items out of the folder. When an Exchange item is tagged as a regulatory record, Outlook clients block deletion of the item. However, users can move messages tagged as records between folders in the mailbox.

See this [white paper](#) for more information about the ability of Microsoft 365 to meet various regulations applicable to financial bodies through the use of record labels and other mechanisms.

**Audit records for Records Management:** When you apply a record label or a regulatory record label to an Exchange item, Exchange writes an *ApplyRecord* event into the audit log. This event comes from mailbox auditing, so it's not captured for the application of labels to SharePoint or OneDrive content using the browser interface. In these situations, SharePoint Online captures a *TagApplied* event, as it does for the application of any retention label. Note that SharePoint Online does not generate *TagApplied* events when new documents receive the default label assigned to a library. To find events for the assignment of record labels and regulatory record labels to documents and list items, you must search the audit log for *TagApplied* events and then examine the *AuditData* property of the events to look for those associated with record labels. See Chapter 21 for more information about how to find data in the audit log. As [explained in this blog](#), it's also possible to use Microsoft Search to find items tagged as records in both a locked and unlocked state.

## Processing Manual Dispositions

Retention labels are good for marking content for retention or removal, however, sometimes you do not want an automated process to function without supervision. For instance, you might have a label to mark documentation for customer projects. Usually, projects finish in a few months, and it is certainly safe to remove the associated documentation after five years. However, in some more complex or extended projects, you need to keep files for longer. A label that removes all documents classified as project documentation after five years would not work. The same might be true for items that the company might need for litigation or audit purposes.

As described in this [SharePoint blog from 2006](#), manual disposition is not a new idea. Microsoft 365 manual disposition allows retention labels to mark content for manual disposition from all workloads covered by data

governance instead of just SharePoint. For now, Microsoft has enabled “Disposition Review” for items in SharePoint Online, OneDrive for Business, Exchange Online, and Microsoft 365 Groups. Support for other workloads might come in the future.

Manual disposition means that human intervention is necessary to check expired content to decide if the business still needs the items or if the deletion (disposition) can happen. A workflow notifies one or more expert reviewers, nominated because they have the skills needed to recognize content that the organization should retain for longer when they examine items with expired retention periods. The expert then decides to approve the removal of items or to extend their retention. Disposition can happen through a single review, or items can progress through up to five review stages with different reviewers processing items at each stage. The reviewers defined for each stage are individual accounts or mail-enabled security groups.

Background processes for each workload scan to detect expired items for manual disposition. When items await disposition, the compliance portal sends mail notifications to the reviewers defined in the retention label settings to tell them that content awaits their decision. Items remain in the awaiting disposition stage until a reviewer decides about their future. Reviewers also receive weekly reminders about items waiting for disposition. If a retention label uses multi-stage disposition, items processed in one stage pass to the next stage and so on until the last stage is complete.

Reviewers must have accounts with cloud mailboxes in the tenant. They should also be members of two compliance role groups:

- **Records management:** Members of this role group (containing the Disposition management role) have access to the disposition section of records management within the compliance portal.
- **Content explorer content viewer:** Members of this role group (containing the Data classification content viewer role) can view the content of items awaiting disposition. Members of the record management role group can see the details and history of items, but not the content.

If you wish, you can create a new role group and assign the necessary roles and members to that group. For instance, you might decide that it's a good idea not to grant access to the other features enabled for members of the Records management role group, like being able to customize the messages sent to reviewers. In this situation, you create a new role group and include only Disposition management and (optionally) the Data classification content viewer roles.

In addition to making decisions about keeping or removing items, the review process helps organizations understand whether people apply labels correctly. For instance, if you see documents stamped with inappropriate retention labels, you might ask why people use labels in error and then take steps to update procedures or change behavior.

## How to Dispose of Items

When items marked with a label that triggers the manual disposition process reach the end of their retention period, background processes mark the items as being available for manual disposition. For example, when the Managed Folder Assistant processes mailboxes, it detects messages that need manual disposition and copies these items to a special hidden mailbox. Reviewers then process the items in this mailbox. When a reviewer decides how to dispose of an item, Exchange replicates the action taken for the item in the hidden mailbox and back to the source mailbox. The need for workloads to process items before the compliance portal recognizes them as being ready for manual disposition means that it can take a little time between the retention period for an item elapsing and that item showing up in the list of items awaiting disposition.

Those specified as reviewers in the label settings receive email notifications about items waiting for disposition. The reviewers can go to the **Disposition** section under **Records management** in the Microsoft Purview Compliance portal to see the items they should review, grouped into the items tagged with each label. Reviewers only see the items they can review while compliance administrators can see all items waiting

for review. To process the items, select a retention label and then the *Open in new window* icon to see the waiting items. In the review window, the compliance portal selects *SharePoint and OneDrive* as the source and loads the items awaiting disposition. If necessary, you can apply a filter (like a date range for item expiration) to refine the set of items shown. If email messages are awaiting disposition, you can see the messages by choosing *Exchange* in the Source filter. Figure 17-22 shows a set of messages waiting for disposition. Note that you can see messages shown in a set of SharePoint and OneDrive items. For instance, this can happen when messages are in SharePoint document libraries because someone sent them via email to a Teams channel.

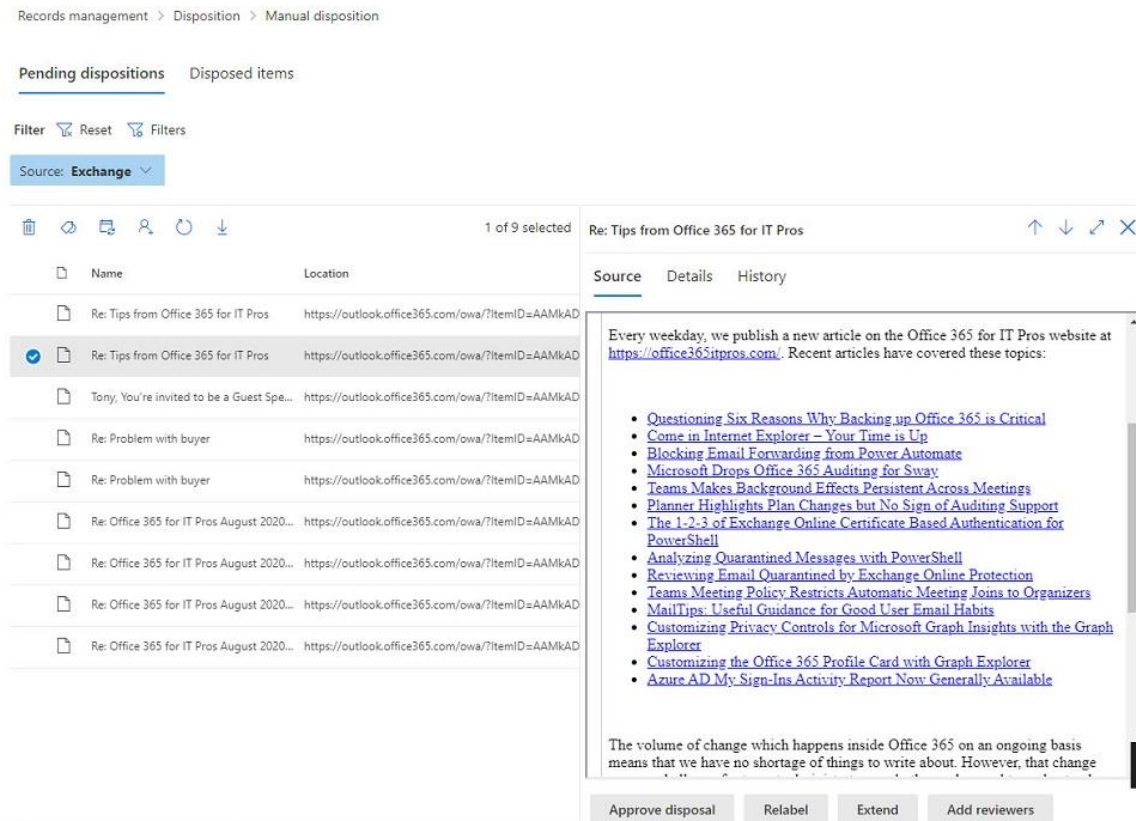


Figure 17-22: A set of messages marked for manual disposition

After selecting an item waiting for disposition, you can:

- **Approve disposal (Delete permanently):** The organization no longer needs this content. When a reviewer approves an item for final deletion, the compliance portal releases the block on the deletion and logs the action in an *ApproveForDelete* audit record. The workload-specific jobs that remove expired items delete the items the next time they process the host location. If the item is in a document library when its retention period expires, marking it for deletion means that the item goes into the first stage Recycle Bin within 7 days of the disposition decision. On the other hand, if the item was already deleted and is in a site's Preservation Hold library, it is eligible for immediate deletion and will be removed following the normal disposal cycle. The audit log captures the actual file deletion in a *Deleted file* or *Deleted file from the second-stage recycle bin* audit record. The delay in deletion is because of the need to run background jobs to process the disposition decisions.
- **Relabel:** The organization should keep this content by applying another retention label to the item. The other label might not have a retention action or a longer retention period.
- **Extend:** Leave the original label on the item but extend the retention period to a specific date (a one-year extension is the default period) after which the item goes through the review process again. This action overwrites the computed retention date for the item with the retention date selected by the

reviewer. The compliance portal captures details of the extension in an *ExtendRetention* audit record. The audit record does not include the new retention date.

- **Add Reviewers:** Add other people who need to review the content to decide on its disposition. These people must have the necessary permissions to access record management.

You can also export the set of items waiting for disposition to a CSV file.

If the reviewer has access to a document's location, they can use the link to view the content. The compliance portal creates the link when the item becomes eligible for manual disposition, and if the item is subsequently deleted or moved the link is invalid. A document might be in a different folder or a preservation hold library, while a message might be in Deleted Items, Recoverable Items, or another folder. A reviewer might want to understand the full context of an item before authorizing its deletion, and if they don't have access to its location (or the link doesn't work), they must consult with the location's owners to decide on its disposition.

Items that a reviewer extends or relabels remain in their original location, or if they were moved into the Recoverable Items folder or preservation hold library, they are moved from there back to their original location.

A busy tenant can generate a heavy workload of review items if disposition review is the norm rather than an exception. For this reason, users should receive training about when they should apply labels that trigger reviews. Reviewers also need the training to understand how to deal with items awaiting their attention so that they know when they can authorize deletion or when they need to seek further guidance from the business about how to handle items.

## Disposed Items and Proof of Disposal

Reviewers can see details of Items that they or other reviewers previously authorized for deletion by selecting the **Disposed items** tab. This view only shows items that reviewers approved for deletion. It does not show items where the reviewer decided to apply a different label or extend the retention period. Items don't show up disposed until the underlying workload has processed the deletion. Because of the reliance on background processes, it might take a couple of days between a decision to delete an item and its appearance in the dispositions list.

Sometimes organizations need to show evidence of the removal of items. To meet this need, Office 365 retains details of all delete dispositions. You can download a CSV file containing details of these operations using the Export option under Disposed Items, using the filters to decide to download documents or messages and the time range for the download. Each line in the file notes the deletion of an item, including:

- Location.
- Title or Subject.
- Retention Label.
- The user principal name of the user who authorized the deletion.
- The timestamp when the deletion occurred.

According to Microsoft, the compliance portal can log up to one million disposal operations.

## Records Management Settings for Disposition

The settings on the Record management page allow members of the records management role group to:

- Define a mail-enabled security group for record managers allowed to see awaiting dispositions for all retention labels. Normally, a member of the records management role group sees only the dispositions assigned to them.
- Define additional text to include in the email notifications reviewers receive to let them know when items await their attention. You can't change the default text and can only add text that appears after

the default text. For instance, you could add text to inform reviewers how to find company guidelines for disposition reviews.

You can also define the mail-enabled security group for record managers to see all disposition reviews using the *Enable-ComplianceTagStorage* cmdlet. Connect to the compliance endpoint and pass the email address of the security group as the parameter:

```
[PS] C:\> Enable-ComplianceTagStorage -RecordsManagementSecurityGroupName Compliance.Records.Managers@office365itpros.com
```

Allowing people to see all items waiting for disposition review is a highly permissioned capability. As such, the compliance portal doesn't support changing it through the GUI. Instead, if you need to make a change, run the *Enable-ComplianceTagStorage* cmdlet again to update the setting with the email address of a different mail-enabled security group.

The *Get-ComplianceTagStorage* cmdlet returns details of retention label (aka compliance tag) management settings. To see the email address of the group defining record managers with full access to disposition reviews, type:

```
[PS] C:\> (Get-ComplianceTagStorage).RecordsManagementSecurityGroupName
```

## Event-based Retention

Labels normally use age-based retention periods and invoke retention actions based on the creation or last modified date of the content. Event-based retention takes a different approach and waits until a specific event occurs before starting the retention countdown for items. For instance, let's assume that you want to preserve all project documents for seven years after a project completes. The event is the project completion, which the project manager might have to sign off. The retention period begins as soon as the event occurs.

Because it depends on something happening rather than just the passing of time, event-based retention is more complex than date-based retention. Here is the general flow of what happens:

1. The administrator creates a new label and selects "an event" as the decision point for the retention period rather than the usual "date created" or "date modified" as used with other labels.
2. The administrator selects an event type (which must already exist) to associate with the label. An event is something like, the expiration of a contract or the departure of an employee, or any other common occurrence in the life of a business. A set of pre-packaged event types are available, but you can create new event types if needed.
3. After saving the label, the administrator includes it in a label policy and publishes the policy to make it available to end-users. After an hour or so, the label is available to SharePoint and OneDrive for Business. It takes a little longer for the label to appear in Exchange.
4. Users apply the label to content that they want to link with the event. For example, they might look for the set of documents belonging to a contract and apply the label to those documents.
5. When they apply the label, users also give a value to a field called "Asset ID," which is part of the standard SharePoint Online schema. A label for an event type is reusable across many different events, so a mechanism is necessary to isolate the content belonging to a specific event. The Asset ID is used to identify individual projects, tasks, or other entities. For instance, if the event deals with the departure of an employee, the Asset ID might hold the employee's number. The Asset ID must be populated correctly because this is the value used to find content associated with an event. You can find out what items are stamped for a specific event by using SharePoint search or content searches to look for the *complianceassetid* tag. For example, find items with *complianceassetid:PK1*.
6. When an event occurs, like an employee leaving or a contract reaching its end, the administrator goes to **Events** under **Records Management** in the Microsoft Purview Compliance portal and creates a



new event to trigger compliance processing. They select the event type to use or choose an existing label configured for event-based retention used to classify items. To find the items for the event, they input the associated Asset ID (for SharePoint and OneDrive items in the form *complianceassetid:<value>*) and/or keywords to locate Exchange items. If other keywords are necessary to find the relevant items, they can be added at this stage (for instance, a tenant might already have assigned a different form of asset identifiers to project documentation). Finally, they select the date the event happened and save the event.

7. Background processes now start looking for content matching the event in SharePoint, OneDrive for Business, and Exchange (in effect, a content search is run). As the search finds matching items, the retention action and period specified in the label are applied to the items. Once the items are stamped, normal retention processing begins. For instance, if the label states that items should be kept for five years, the items are kept for five years after the event date. It can take up to a week before the background processes find all matching content across workloads.

Once created, you cannot update an event, so you should be sure that everything is ready to find the items relating to an event when you create it. In addition, once you associate a label with an event, you cannot change the label to associate it with a different event. For these reasons, it is important to have a good understanding of the business events that occur within the tenant and how people working in the business can use labels to aid the processing of content associated with using event-based labels. For more information about event-based retention, see the [Microsoft documentation](#).

## Removing Retention Labels

Four scenarios occur for removing retention labels from a tenant:

- **A label is created, but not published.** Because the label is not in use, you can edit the label in the Microsoft Purview Compliance portal and remove it with the **Delete label** option. Alternatively, run the *Remove-ComplianceTag* cmdlet.

```
[PS] C:\> Remove-ComplianceTag -Identity "Bad Label"
```

- **A label is published in one or more policies but has never been assigned to items.** If you try and remove the label, you'll see an error that the label is in use. This is technically correct because the label is in a policy even a user or policy never assigned the label to any items. To remove the label, you first remove it from all the policies it is included in (or remove a complete policy if the label is the only one in that policy). After removing the label from all policies, you can remove the label as described above.
- **A label is in active use and applied to content.** You can run a content search to find all the items that have the label (find a compliance tag equal to the label name) and remove the label from the items. This works when a label is recent and only applied to a small number of items but is unreasonable in a tenant of any size. Instead, you can follow the procedure as if the label was never assigned to any items by removing it from all policies and then removing the label in the Microsoft Purview Compliance portal or using PowerShell. The label then goes into a pending deletion state, meaning that some background processes in the different workloads must run to remove the label from items. A background timer job removes labels from SharePoint Online and OneDrive for Business documents and the process can take several hours to complete. For Exchange, the Managed Folder Assistant removes labels from mailbox items; it might take up to a week before the assistant processes a mailbox. Once the Managed Folder Assistant removes a label from an item, the item becomes a potential target for the assignment of another label by any other retention policy applying to the mailbox.

- **A label is a record.** As noted above, an item assigned a record label cannot be changed. You cannot remove any label marked as a record (even if the label has never been assigned to an item), but you can stop people from using it again by removing the label from any policies that it is in.

Removing labels is not something to do at a whim. The complexities involved in removing labels that are applied to content underline the need for planning and preparation for the deployment of labels as part of your data governance strategy. Removing labels from items can result in their deletion by background processes because their retention period has expired, so if you remove a label from content, you might need to replace those labels with different labels to ensure the retention of the items.

## Data Classification Dashboard

Retention labels have existed for several years. Sensitivity labels were the next step. Over time organizations are likely to accumulate a reasonable amount of labeled material. But how does a compliance administrator know that their data governance strategy is effective, that users apply labels as intended, and that the right information is protected? Microsoft's response is the Data classification dashboard (Figure 17-23) in the Microsoft Purview Compliance portal. The dashboard contains some useful statistics about where and what labels are in active use and has the following sections:

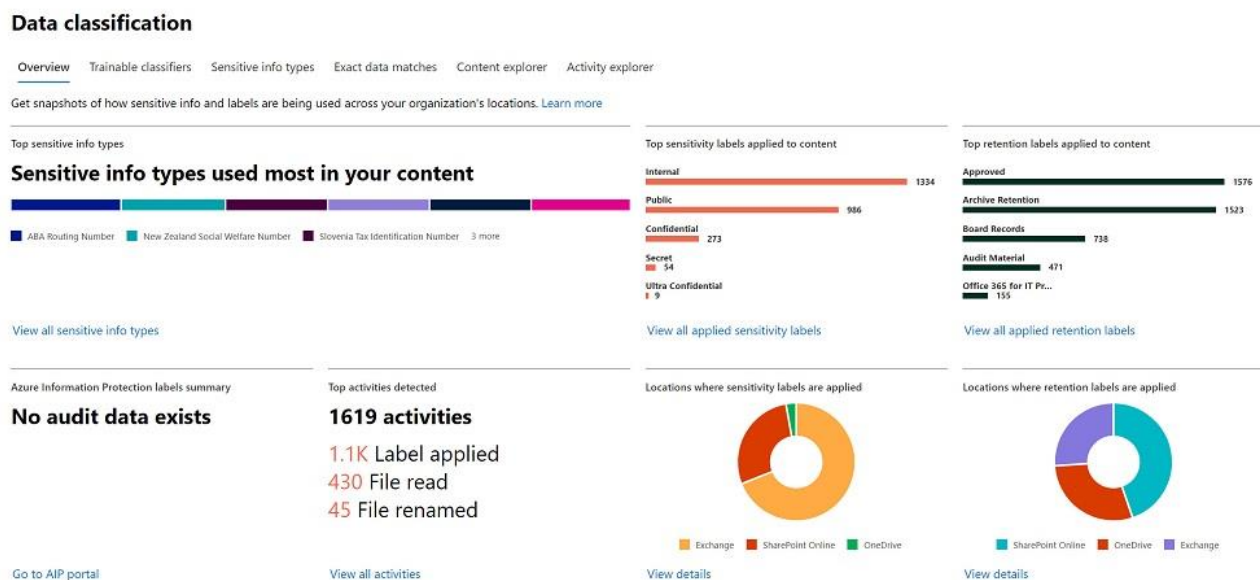


Figure 17-23: Data classification dashboard

- **Overview:** Displays the sensitivity and retention labels are in use and in what workload use the labels. The dashboard also highlights sensitive information types found in documents and messages.
- **Trainable classifiers:** This feature allows tenants to build custom sensitive information types by [using AI training based on a set of examples](#). Microsoft's default classifiers include document types like resumes and source code and classifiers used to detect objectionable behavior like profanity and threat.
- **Sensitive Info types:** Lists the sensitive information types known in the tenant, including the default set (over 200) created by Microsoft and those created by the tenant through digital fingerprinting, dictionaries, or simple rule matching.
- **Exact data matches:** The regular sensitive information types used in Microsoft 365 match items using generic patterns, an approach that works well for document-centric information. As the name implies, [exact data matches \(EDM\)](#) use more precise matching and are more suitable for structured data. This

option allows administrators to define EDM schemas and use those schemas to build custom sensitive information types.

- **Content Explorer:** Allows compliance administrators to see where retention and sensitivity labels are applied in Exchange Online, SharePoint Online, and OneDrive for Business. To view the locations where labels are applied, the user must hold the *Content Explorer List Viewer* permission; to view the source content of items in those locations, their account must hold the *Content Explorer Content Viewer* permission.
- **Activity explorer.** Shows usage data for both sensitivity and retention labels.

The overview is available with an Office 365 E3 license. Organizations need to have Office 365 E5 or Microsoft 365 E5 compliance licenses to use the content explorer, activity explorer, and trainable classifiers.

## Content Explorer

The content explorer is especially useful in validating the effectiveness of label usage. The explorer allows administrators to see how users and policies apply retention and sensitivity labels and the type of content users apply labels to. For example, let's assume that you define a retention label for formal company records with a ten-year retention period. You probably don't want this label applied to run-of-the-mill documents and messages, and the content explorer allows you to see what's in the labeled items to check if the content is as valuable as the label implies.

The downside of the content explorer is its access to any user information in a SharePoint Online site, OneDrive for Business account, or Exchange Online mailbox. Although compliance administrators can use other features to access user data (a content search, for instance), because the content explorer organizes items by label, it's easy to find information labeled as being the most confidential or valuable to the organization and then peruse the content (items protected with sensitivity labels with encryption can't be opened by content explorer). For this reason, it's wise to make sure that the only people with the permissions needed to use content explorer are those who must have it to perform specific operations, and then remove those permissions once the need passes.

## Activity Explorer

The Activity Explorer displays information about the application of retention and sensitivity labels to content in Exchange Online, SharePoint Online, and OneDrive for Business to help compliance administrators understand the use of labels within the organization. Data flows through to the Activity Explorer when sensitivity and retention labels are applied using Office Online, desktop (Microsoft 365 apps for enterprise), and mobile apps. Using a 30-day sliding window, the explorer gives an insight into:

- How labels were applied (manually by a user or by policy).
- Who applied, changed, or removed a label, including labels applied by an auto-labeling policy. If a user or policy changed a label on an item, you see details of the old and new labels. Details of labels applied by endpoint monitoring are also available.
- The location (Exchange Online, SharePoint Online, including OneDrive for Business).
- The target document or folder (the data does not show labels inherited by files from a label assigned to a folder).

You can filter by date range, label activity (all actions or just label changes, removals, or applications), user, workload, or specific label to focus on a specific activity. To see added information, click on an individual record to open the details pane. Sometimes it's good to create your view on this data. Chapter 21 includes an example of using PowerShell to extract and interpret audit records logging how users assign retention labels. Chapter 20 covers the capture of audit records for sensitivity labels.

## Trainable Classifiers

A [trainable classifier](#) is a digital map of a defined category of documents such as a customer invoice, order form, or personnel review. Microsoft 365 includes a set of built-in classifiers (created by Microsoft) such as HR, Healthcare, Legal Affairs, and Threat for use out-of-the-box. The pre-trained classifiers are available in English, Spanish, Japanese, French, German, Portuguese, Italian, and Chinese (simplified).

Organizations can create custom trainable classifiers by going through a training process. The major steps in the process are:

- Define the kind of information you want the trainable classifier to recognize. For example, you might want to create a trainable classifier that recognizes financial reports in a specific format.
- Assemble a set of sample documents for the training process to scan and build a digital picture using machine learning. The set acts as seed content for the classifier and should include between 50 and 500 “positive” samples available on a SharePoint Online site. Leave the content on the site for a day or so to allow indexing to happen before creating the classifier. By analyzing the seed content, the classifier builds a prediction model (the digital picture) to allow it to recognize documents of the same type when it processes them in the future. The more representative the seed content is of real-life documents, the more accurate the prediction model is likely to be.
- Test the classifier to see how it performs when it processes real content and tweak it if necessary. The seed set might require a refresh to introduce new examples to help the classifier refine its prediction model. During this phase, compliance administrators help the refining process by marking files identified by the classifier as good or bad matches. The aim is to increase the percentage of good matches to as close to 100% as possible.
- Publish the classifier to make it available for use with Microsoft Purview compliance solutions. If you discover that the classifier does not work well in production and retraining the prediction model with feedback does not improve its accuracy, you might need to withdraw (delete) the classifier and restart the training process with a new set of sample documents, including samples not detected by the previous iteration of the classifier. For this reason, use new trainable classifiers in limited policies (for instance, against a small number of SharePoint sites) and monitor its results closely until the classifier works as expected.
- Over time, monitor the items matched by the classifier and mark them as good or bad matches. The prediction model incorporates received feedback from administrators into new, hopefully, more accurate, versions. The caveat here is that feedback is only valuable when compliance policies use a trainable classifier to detect content. If this is not the case, retraining to take account of feedback is useless.

Machine learning takes time, and you cannot rush the creation of the prediction model for a trainable classifier. The entire end-to-end process from assembling the sample document set to publication might take several weeks before the classifier is accurate enough for production use (98% or better). After publishing the classifier, it becomes available to find items in:

- Communications compliance policies.
- Data Loss Prevention policies.
- Auto-label policies for sensitivity labels and retention labels.
- The Content explorer.

Trainable classifiers require Office 365 E5 or Microsoft 365 E5 Compliance licenses. See [this page](#) for more information.

# Ingesting Items for Review from External Sources

Companies that need to supervise employee communications need coverage over more than just email. Employees can conduct business using a variety of consumer and business services including Twitter, Facebook, Bloomberg, HipChat, Thomson Reuters, and BlackBerry messaging. To cater to the problem, Microsoft has signed deals with [third parties](#) who specialize in the extraction of data from different communication systems to create connectors to extract data from those systems and ingest the data into Microsoft 365. The basic approach is:

- A connector from a selected partner (like Actiance or ArchiveSocial) connects to the source data using whatever API is available. The connection runs on a scheduled or ongoing basis to find and extract data of interest.
- The connector uses Exchange Web Services to connect to an Azure endpoint for the ingestion of data into Exchange Online.
- Data flows across the connector to either:
  - User mailboxes, if a match exists between the identifier used by Exchange Online (usually the User Principal Name) and the identifier used by the source service. For instance, if the corporate Twitter account logs in as TwitterService@tenant.com and an Azure AD account exists with the same identifier, a match exists, and the data extracted from Twitter goes to that mailbox. Because you do not want someone to be able to access the information brought in via the connector, the items go into the Purges folder within Recoverable Items. The items are indexed and discoverable but invisible to anyone who logs into the mailbox.
  - Connector mailboxes are set up explicitly as a target for data ingested into Exchange Online through a connector. In this case, the items go into the Inbox folder because someone usually needs to check the items and decide where to keep the items over the long term.
- As items flow into Exchange Online through the connector, a separate set of agents watch the Exchange Web Services traffic to apply communications compliance policies.

When the data reaches Exchange Online, the imported items are indexed as normal and the content they hold is discoverable and usable by other data governance features. You can apply the data governance policies that exist within the tenant. In other words, you can assign retention policies to the mailboxes used by the connectors to ensure that you keep the ingested content for the desired retention period, including the mailboxes in content searches and eDiscovery cases, and so on.

## Microsoft Connectors for Third-Party Data

Microsoft supports methods to import information from third-party sources into Exchange Online where the third-party information is then subject to data governance functionality. Many connectors are currently available including popular social media feeds like:

- **LinkedIn:** Import information from a [company's LinkedIn page](#).
- **Facebook:** Import information from a [Facebook business page](#).
- **Twitter:** Import tweets sent and received from a [company's Twitter account](#).

Preview versions of other connectors are also available. In general, connectors work by extracting information from the target source and creating items in an Exchange Online mailbox as described above.

# Using PowerShell with Retention Labels and Policies

A set of cmdlets is available to manage labels and retention policies. To access the cmdlets, you must connect a PowerShell session to the compliance endpoint after connecting to the Exchange Online endpoint:

```
[PS] C:\> Connect-ExchangeOnline
Connect-IPPSSession
```

It is not the intention to have a detailed and in-depth discussion of all the cmdlets here as it would occupy too much space. In any case, given the complexity of some of the operations involving compliance, it is usually best to create and update policies and labels through the GUI.

## Working with Retention Labels

The *\*-ComplianceTag* cmdlets manage retention labels. In this example, we extract the set of retention labels defined in a tenant and list the most important properties of each tag to understand the purpose of the tag: to mark an item as a record, has a retention action, the retention duration (in days), the retention action, and if it is in use.

```
[PS] C:\> Get-ComplianceTag | Format-Table Name, IsRecordLabel, HasRetentionAction,
RetentionDuration, RetentionAction, Mode -AutoSize
```

Name	IsRecordLabel	HasRetentionAction	RetentionDuration	RetentionAction	Mode
Confidential	False	True	1825	Keep	Enforce
Remove after 1 week	False	True	7	Delete	Enforce
Patent Materials	False	True	7300	Keep	Enforce
Board Records	False	True	Unlimited	Keep	Enforce
Formal Company Record	True	True	7300	KeepAndDelete	Enforce

Labels marked as regulatory records can be found by examining the *Regulatory* property.

```
[PS] C:\> Get-ComplianceTag | ?{$_.Regulatory -eq $True} | Format-Table Identity, Notes
```

Identity	Notes
Regulatory Record (Legal)	A legal regulatory record

The *New-ComplianceTag* cmdlet creates a new retention label. In this example, we create a retention label to keep content for ten years. This retention label does not mark content as a formal record.

```
[PS] C:\> New-ComplianceTag -Name "Patent Information" -IsRecordLabel $False -RetentionDuration 3650
-RetentionAction Keep -Comment "Items marked with this classification are associated with patents"
-RetentionType ModificationAgeInDays
```

The *Set-ComplianceTag* cmdlet updates the properties of a retention label while the *Remove-ComplianceTag* cmdlet removes a label from a tenant. For example, this command sets the retention duration for the "Patent Materials" label to 15,000 days (the maximum is 24,855):

```
[PS] C:\> Set-ComplianceTag -Identity "Patent Materials" -RetentionDuration 15000
```

The audit log captures changes made to retention tags with the *SetComplianceTag* event. Updates to retention policies generate *SetRetentionCompliancePolicy* events for updates to policy settings and *SetRetentionComplianceRule* events for changes to rules belonging to policy rules, like updating the retention period.

## Working with Retention Policies

Retention policies use a parent/child structure. The parent is the policy itself and the child is the set of rules that implement the policy settings. Another way of thinking about this is that the parent for a policy defines the overall data to which the policy applies while the rules govern the application of a policy. It is a little simpler than it seems because a 1-to-1 relationship usually exists between retention policies and retention rules (including label policies). The Compliance portal hides the links between retention policies and rules, but the connection needs to be understood when we manage retention policies through PowerShell. Two cmdlet sets are used:

- The *\*-RetentionCompliancePolicy* cmdlet set manipulates retention policies. Use these cmdlets to manipulate the workload locations to which a policy applies or to enable or disable a policy.
- The *\*-RetentionComplianceRule* cmdlet manipulates the rules for retention policies. Use these cmdlets to work with properties such as the retention duration of a policy.

When you fetch details of retention policies with the *Get-RetentionCompliancePolicy* cmdlet, the set returned includes policies used to:

- Apply retention settings to workloads and policies.
- Publish retention labels to workloads.
- Publish sensitivity labels to workloads.

Not all the retention policies in a tenant publish retention labels to workloads. The *Get-RetentionCompliancePolicy* cmdlet returns all the policies used for retention labels. Here is some code to find the set of retention policies used to publish retention labels and then tell us what retention labels are in each policy:

```
[PS] C:\> $Policies = (Get-RetentionCompliancePolicy -RetentionRuleTypes | ? {$_.RetentionRuleTypes
-eq "Publish"})
ForEach ($P in $Policies) {
    Write-Host "Processing" $P.Name
    $Tag = $Null
    $Rules = (Get-RetentionComplianceRule -Policy $P.Guid)
    ForEach ($R in $Rules) {
        If (-Not [string]::IsNullOrEmpty($R.PublishComplianceTag)) {
            $Tag = $R.ComplianceTagProperty -Split(",")
            $TagValues = Get-ComplianceTag -Identity $Tag[0]
            Write-Host $P.Name "includes the retention label" $TagValues.Name }
    }
}
Processing Company Confidential Policy
Company Confidential Policy includes the retention label eBook Content
Company Confidential Policy includes the retention label Contractual Information
Company Confidential Policy includes the retention label Confidential...
```

*Get-RetentionCompliancePolicy* does not return all retention policies. Microsoft regards the retention policies for Teams private channels and Yammer conversations as “app specific” policies and the cmdlet does not return details of these policies (compliance records for conversations from Teams shared channels are in the mailboxes of channel members). Instead, you must fetch details of app specific policies using the *Get-AppRetentionCompliancePolicy* cmdlet, which behaves in the same way as *Get-RetentionCompliancePolicy* does. There’s some obvious inconsistency here because Microsoft doesn’t treat the retention policies for Teams (regular) channels and chats as app specific.

### Retention Rule Types

The *RetentionRuleTypes* property of a policy tells us what kind of policy an individual retention policy is. To see this information, you must return the policy type by including the *RetentionRuleTypes* parameter in the call to *Get-RetentionCompliancePolicy*. For example:

```
[PS] C:\> Get-RetentionCompliancePolicy -RetentionRuleTypes | Sort RetentionRuleTypes | Format-Table
Name, RetentionRuleTypes
```

Three values report the types of retention policies:

- **Apply:** The policy uses advanced settings (like a keyword search) to find the content to which it applies labels. These are also known as auto-label policies.
- **Default:** The policy publishes retention settings to workloads (static or adaptive policies).
- **Publish:** The policy publishes labels to workloads.

## Fetching Retention Policies

To fetch the properties of an individual retention policy, run the *Get-RetentionCompliancePolicy* cmdlet:

```
[PS] C:\> Get-RetentionCompliancePolicy -Identity 'GDPR Personal Data' -DistributionDetail | Format-List
ExchangeLocation, SharePointLocation, OneDriveLocation, ModernGroupLocation, TeamChatLocation, TeamChannelLocation, Workload, Enabled, Mode, RestrictiveRetention, DistributionStatus
```

```
ExchangeLocation      : {All}
SharePointLocation    : {All}
OneDriveLocation      : {}
ModernGroupLocation    : {}
TeamChatLocation      : {}
TeamChannelLocation   : {}
Workload               : Exchange, SharePoint, OneDriveForBusiness, ModernGroup
Enabled                : True
Mode                  : Enforce
RestrictiveRetention  : False
DistributionStatus     : Success
```

We can interpret the output as follows:

- **ExchangeLocation:** Lists the names of the mailboxes covered by the policy. If All, it means that all mailboxes are covered (the same is true for the other locations).
- **SharePointLocation:** Lists the SharePoint Online sites covered by the policy.
- **ModernGroupLocation:** Lists the aliases for the Groups covered by the policy.
- **OneDriveLocation:** Lists the OneDrive for Business locations covered by the policy.
- **TeamChatLocation:** Lists the locations for personal and group chats covered by the policy.
- **TeamChannelLocation:** Lists the Teams (all channels) covered by the policy.
- **Workload:** Lists the workloads where the policy is active. In this case, Exchange Online, SharePoint Online, OneDrive for Business, and Microsoft 365 Groups.
- **Enabled:** Tells you whether the policy is active or not. The default is *\$True*, but if a policy has been disabled for some reason, this value is *\$False*.
- **Mode:** If "Enforce", the policy is active.
- **RestrictiveRetention:** The default is False, meaning that the policy can be changed. If True, a preservation lock exists on the policy, meaning that administrators can only make limited changes to the policy.
- **DistributionStatus:** *Success* means that the different workloads have the information needed to enforce the policy settings. *Pending* means that Microsoft 365 is distributing details of a policy change to workloads. Any other value shows that a problem has occurred in the policy distribution. Sometimes this happens for good reason, such as a change occurring in a data center. If you see problems when distributing labels to certain locations, it might be possible to figure out where the problem lies by examining the *DistributionResults* property for the policy. You won't see the current distribution status or any errors unless you include the *DistributionDetail* parameter (see example below).

If you see "All" listed as a value for a workload, it means that the policy covers every location supported by that workload. For example, "All" listed in *ExchangeLocation* means that the policy covers every Exchange



mailbox in the tenant. You can exclude specific mailboxes or sites in a workload from the policy. If this is the case, you will find a list of those locations in the properties *ExchangeLocationException*, *SharePointLocationException*, and so on.

To find retention policies using adaptive scopes, check the value of the *IsAdaptivePolicy* property:

```
[PS] C:\> Get-RetentionCompliancePolicy -Identity "Retention Policy for French IT Architects"
-DistributionDetail | ? {$_.IsAdaptivePolicy -eq $True} | Format-List Name, IsAdaptivePolicy,
AdaptiveScopeLocation
```

```
Name                : Retention Policy for French IT Architects
IsAdaptivePolicy     : True
AdaptiveScopeLocation : {French IT Architects, Executive Mailboxes}
```

Retention policies for Teams private channel messages use the *ExchangeLocation* property to indicate the set of locations (compliance records for these messages are in Exchange mailboxes). The *Applications* property is set to **User:MicrosoftTeamsChannelMessages** to indicate the type of information the policy covers. The same approach is used for retention policies for Yammer messages. In this case, the *Applications* property holds **Group:Yammer** to indicate that the policy processes community messages and **User:Yammer** for user messages.

### Including Distribution Details

When you look at retention policies with the *Get-RetentionCompliancePolicy* cmdlet, you won't see details of the individual mailboxes or sites specified for locations, the up-to-date distribution status, or any error information unless you include the *DistributionDetail* parameter. To see details of locations or errors, you need to expand the relevant property for the location you want to examine. For example, here's how to examine details of the mailboxes to which a policy applies.

```
[PS] C:\> Get-RetentionCompliancePolicy -Identity "Senior Leadership Team" -DistributionDetail |
Select -ExpandProperty ExchangeLocation
```

```
DisplayName        : Kim Akers
Name               : Kim.Akers@Office365itpros.com
ImmutableIdentity  : f120e18f-8305-41e3-abd4-de93d4a2a493
Type               : IndividualResource
Workload           : Exchange
SchemaVersion      : 2
```

```
DisplayName        : Brian Weakliam
Name               : Brian.Weakliam@office365itpros.com
ImmutableIdentity  : aae8332a-6832-4c00-b873-6ec443c36395
Type               : IndividualResource
Workload           : Exchange
SchemaVersion      : 2
```

Here's an example of looking at retention policies that have encountered problems when distributed to workloads. Often the issue is a transient problem caused by a recipient selected for the policy being unavailable for some reason that you can fix by editing the policy to remove the recipient and then add them back again.

```
[PS] C:\> $Errors = (Get-RetentionCompliancePolicy -DistributionDetail | ? {$_.DistributionStatus -
eq "Error"})
ForEach ($E in $Errors) {
    $Results = ($E | Select -ExpandProperty DistributionResults)
    Write-Host "Policy:" $E.Name "Issue:" $Results }
Policy: Company Confidential Policy Issue: [Exchange]SMO-
Academy@office365itpros.onmicrosoft.com:Recipient not found: f120e18f-830
5-41e3-abd4-de93d4a2a493
Policy: Black Matter Policy Issue: [ModernGroup]'ModernGroup' Resources:Policy deployment has been
interrupted by an unexpected Office
365 data center issue. Please contact Microsoft support to fix the deployment issue.
[ModernGroup]BlackMatterTeam@office365itpros.com
```

```
[Recipient not found: 6661b878-83b5-41bb-aad4-ba14e8879b90]
```

## Retention Policy Rules

Returning to our discussion about rules, when you create a new retention policy, Microsoft 365 creates the underlying rule for the policy is created to instruct workloads on how to process content. We can see the rule for a retention policy by using the *Get-RetentionComplianceRule* cmdlet. Because policies can have similar names, passing the GUID identifying the policy makes sure that the correct information is returned. To discover the identifiers for retention policies, run the command:

```
[PS] C:\> Get-RetentionCompliancePolicy | Format-Table Name, Guid
```

If you don't want to use a GUID, pass the policy name and hope it's unique:

```
[PS] C:\> Get-RetentionComplianceRule -Policy 'Patent Materials' | Format-List ContentMatchQuery, RetentionDuration, RetentionComplianceAction, ExpirationDateOption, Workload
```

```
ContentMatchQuery      : Patent NEAR(10) claim
                        : Patent NEAR(10) prosecution
                        : Patent NEAR (10) application
RetentionDuration      : 3650
RetentionComplianceAction: KeepAndDelete
ExpirationDateOption   : ModificationAgeInDays
Workload                : Exchange, SharePoint, OneDriveForBusiness, ModernGroup
```

Some of the output (Workload in this case) matches what you see when examining the policy. The interesting pieces here are:

- **ContentMatchQuery:** The Keyword Query Language (KQL) query to determine whether items come under the scope of the retention policy. In this case, three separate tests are used.
- **RetentionDuration:** The length of time in days to retain items. You can also use "unlimited". In this case, the items remain held indefinitely until the hold set by the policy lapses or the policy is removed from the tenant. The hold duration is calculated using the created date for email items and the date last modified for files in SharePoint Online or OneDrive for Business sites.

Rule settings can be changed using the *Set-RetentionComplianceRule* cmdlet. For example, this command sets the retention duration for a rule to 3,600 days (the maximum duration is 24,855 days):

```
[PS] C:\> Set-RetentionComplianceRule -id 86f67249-74b9-48bf-8fe6-9e0c58416dfb -RetentionDuration 3600
```

Changes made to a rule will not be effective until the publication of the policy update becomes known to all the workloads. This might take a few hours.

To publish labels and make them available to workloads, we create a retention label policy and associate a rule for each label with that policy. In other words, every label published by the policy has a separate rule. Labels can be in multiple label policies and the connection between label and rule is through the label GUID. If you run the *Get-RetentionComplianceRule* cmdlet to find all the rules belonging to a policy, you can find the different labels by looking at the *PublishComplianceTag* property, which holds the GUID pointing back to the label. The rule for an auto-label policies specifies its label in the *ApplyComplianceTag* property.

## Setting Retention Policies

You can force the republication of a policy to the workloads by running the *Set-RetentionCompliancePolicy* cmdlet. For example:

```
[PS] C:\> Set-RetentionCompliancePolicy -Identity 'Patent Materials' -RetryDistribution
```

Republishing a policy to workloads only works if an error previously prevented a policy from reaching a workload. If a policy cannot be published to workloads for some reason, you should file a support incident with Microsoft.

The *Set-RetentionCompliancePolicy* cmdlet can also add or remove workload locations to the policy. In this example, we remove a mailbox and add a mailbox to the set of Exchange locations. The same approach is taken remove and add some SharePoint sites. Note that you must give the URL for each site.

```
[PS] C:\> Set-RetentionCompliancePolicy -Identity 'Patent Lock-down'
-AddExchangeLocation 'Frank Clonan' -RemoveExchangeLocation 'Rob Young'
-RemoveSharePointLocation 'https://office365itpros.sharepoint.com/Projects/'
-AddSharePointLocation 'https://office365itpros.sharepoint.com/Exchange Connections'
```

## Reporting Retention Policies Applied to SharePoint

As an example of using information about distribution detail and retention rule type to analyze or report on retention policies, let's say that you want to know what retention policies apply to SharePoint sites. The code in [this GitHub script](#) fetches information about the retention policies including their distribution detail, excluding retention policies for Teams, those that don't process SharePoint, and policies used to publish retention labels. We then examine each policy to extract the locations within the policy scope and figure out whether the retention settings are simple or advanced (using a keyword query or sensitive information type). Some policies apply to every SharePoint site, so the location is "All." Others have specific SharePoint sites defined, and some policies process everything except a set of excluded sites. The output is an ordered array. We can look at the data in different ways with PowerShell (see below) or export it to a CSV file to load into Excel or Power BI.

Policy	Site	Duration	Action
SharePoint Online Retention Policy	*Exclude* Interesting Patent		
SharePoint Online Retention Policy	*Exclude* Frank's Italian Job		
Label Customer Invoices	All SharePoint Sites	2555	Keep
SharePoint Online Retention Policy	All SharePoint Sites	2555	Keep
Preserve Office 365 for IT Pros Files	Company Communications	Unlimited	Keep
Preserve Office 365 for IT Pros Files	GDPR Planning Mark II	Unlimited	Keep
Senior Leadership Team (SLT) Retention Policy	Office 365 2020 Speakers	3600	KeepAndDelete
Office 365 for IT Pros eBook Content	Office 365 for IT Pros	3650	Keep
Preserve Office 365 for IT Pros Files	Office 365 for IT Pros	Unlimited	Keep
Preservation Lock - Mailboxes and Sites	PL Test Group	3650	KeepAndDelete
Management Presevation Policy	Projects	Unlimited	Retain

## Viewing Teams Retention Policies

Remember that a Teams retention policy can only cover Teams personal chats and channel conversations and a general retention policy applied to other locations cannot cover Teams. If you only want to work with retention policies that affect Teams, use the *TeamsPolicyOnly* parameter when fetching retention policies:

```
[PS] C:\> Get-RetentionCompliancePolicy -TeamsPolicyOnly
```

Likewise, to exclude the Teams policies, use the *ExcludeTeamsPolicy* parameter:

```
[PS] C:\> Get-RetentionCompliancePolicy -ExcludeTeamsPolicy
```

## Remove a Retention Policy

To remove a retention policy, run the *Remove-RetentionCompliancePolicy* cmdlet. Remember that you will not be able to remove a policy if a preservation lock is in place.

## Tracking Retention Holds for Mailboxes

When a non-org-wide retention policy applies an in-place hold to a mailbox, Exchange Online notes the fact by updating the *InPlaceHolds* property of the mailbox with the GUID for the hold. Thus, you can get a quick

view of what mailboxes are on hold by checking the mailbox properties for *InPlaceHolds* using PowerShell. For instance:

```
[PS] C:\> Get-ExoMailbox -RecipientTypeDetails UserMailbox -Properties InPlaceHolds | ?
{$_ .InPlaceHolds -ne $Null} | Format-Table DisplayName, InPlaceHolds
```

DisplayName	InPlaceHolds
Tony Redmond	{skp748f77b020124e6e8304e66021fb297b:3, mbx748f77b020124e6e8304e66021fb297b:3}
Kevin A. Laahs	{UniH26c5d797-0fd3-496d-92ac-4f405700c917}
Kim Akers	{mbx748f77b020124e6e8304e66021fb297b:3, skp748f77b020124e6e8304e66021fb297b:3, UniHec6...}

One reason why you might want to check the holds set on mailboxes is to find out what holds are keeping inactive mailboxes alive.

```
[PS] C:\> Get-Mailbox -RecipientTypeDetails UserMailbox -InactiveMailboxOnly -Properties
InPlaceHolds | Format-Table DisplayName, InPlaceHolds
```

DisplayName	InPlaceHolds
Holly Holt	{}
James Gangley	{mbx29550d04cffd42109bdd94cc56c65041:2}
Jodie Smith (Program Manager)	{}
Rob Young	{d9eb7052cc0f4200b6a1ad0d6f2171ed...}
Mary Smith (Customer Support)	{}
Ed Banti	{skp748f77b020124e6e8304e66021fb297b:3}

When you see an inactive mailbox shown with a null value in *InPlaceHolds*, you know that the hold on the mailbox comes from an org-wide retention policy (or an Exchange Online litigation hold) instead of a non-org-wide policy. More on this topic in a little while.

To get a full view of the holds that apply to an individual mailbox, use the *Get-ExoMailbox* cmdlet, and expand *InPlaceHolds*:

```
[PS] C:\> Get-ExoMailbox -Identity "Ben Owens" -Properties InPlaceHolds | Select -ExpandProperty
InPlaceHolds
```

```
mbx748f77b020124e6e8304e66021fb297b:3
skp748f77b020124e6e8304e66021fb297b:3
UniH47f67751-1036-4621-80d6-d25837adf813
UniHec6163be-6ed6-4b16-afe8-1b2165b9359f
UniH84dea76f-c845-4101-b066-a8b10c13c210
```

One of the holds listed applies to Skype for Business conversions (skp). Another is for mailbox contents (mbx). If you see a minus sign before an "mbx" hold, it means that a retention policy explicitly excludes the mailbox. The numeric notation following the hold identifier tells you the kind of hold it is:

- 1: The retention policy deletes items. Microsoft also uses this value for label publishing policies for retention labels and sensitivity labels.
- 2: The retention policy holds items. Microsoft 365 doesn't remove the items after the hold lapses.
- 3: The retention policy holds items and then deletes them after the retention period expires.

The holds with a "UniH" prefix are "unified holds" and belong to:

- A hold placed by an eDiscovery case created through the Microsoft Purview Compliance portal.
- A hold placed by an old preservation policy now upgraded to a retention policy.

If a hold shows up as a GUID without a prefix or it has a "cld" prefix, it belongs to an Exchange in-place hold managed from the EAC. These include holds created for mailboxes included in eDiscovery cases managed in the SharePoint eDiscovery Center. While you might come across these GUIDs, they should begin to disappear over time as workload-specific holds expire.

The GUID for a hold placed by a retention policy can be used to find which policy the hold belongs to. Take the value stored in the *InPlaceHolds* property, remove the prefix (like "mbx") and suffix (like :3), and use the value to check against the set of retention policies for the organization.

```
[PS] C:\> Get-RetentionCompliancePolicy 748f77b020124e6e8304e66021fb297b
```

Name	Workload	Enabled	Mode
Senior Leadership Team (SLT) Retention Policy	Exchange, ModernGroup	True	Enforce

We now know that the mailbox comes under the scope of the Senior Leadership Team (SLT) Retention Policy. This is a non-org wide policy that applies to selected Exchange locations (mailboxes). If we look at the *ExchangeLocation* property, we see a list of the mailboxes the policy applies to, or "All" if the policy applies to all mailboxes. For example, here's a typical entry for a mailbox.

```
[PS] C:\> Get-RetentionCompliancePolicy 748f77b020124e6e8304e66021fb297b -DistributionDetail |
Select -ExpandProperty ExchangeLocation
```

```
DisplayName      : Tony Redmond
Name             : Tony.Redmond@office365itpros.com
ImmutableIdentity : d446f6d7-5728-44f8-9eac-71adb354fc89
Type            : IndividualResource
Workload         : Exchange
SchemaVersion    : 2
```

The *ImmutableIdentity* property reported for each mailbox on hold equates to the GUID identifying the user account in Azure AD and is the same value that you see if you run the *Get-ExoMailbox* cmdlet and examine the *ExternalDirectoryObjectID* property.

## Group Mailboxes and Holds

To discover if any holds apply to a group mailbox, run the *Get-UnifiedGroup* cmdlet and examine the contents of the *InPlaceHolds* property:

```
[PS] C:\> Get-UnifiedGroup -Identity "Office 365 for IT Pros" | select InPlaceHolds
```

```
InPlaceHolds
-----
{grp6a9f7bf0507b4a4983c301a701958d11:2, UniH26c5d797-0fd3-496d-92ac-4f405700c917}
```

To find what groups a hold applies to, specify the *DistributionDetail* parameter when calling *Get-RetentionCompliancePolicy* to return details of the locations covered by the policy. For example:

```
[PS] C:\> Get-RetentionCompliancePolicy -DistributionDetail | ? {$_.ModernGroupLocation -ne $Null
-and $_.ModernGroupLocation -notlike "*All" } | Format-List Name, ModernGroupLocation
```

```
Name                : Clean up Groups with Connectors
ModernGroupLocation : {office365tenantervicehealth, office365roadmapupdates, askhr}

Name                : Senior Leadership Team (SLT) Retention Policy
ModernGroupLocation : {Senior Leadership Team}
```

This technique only works to report details of the holds placed by retention policies. We need to process the other types of holds differently to uncover their secrets. More on this topic soon.

## Org-wide Retention Policies

Org-wide retention policies applying to all mailboxes do not stamp hold information on individual mailboxes. Instead, the Exchange Online organizational configuration for the tenant stores details of the holds belonging to retention policies applicable to all mailboxes. This approach avoids the need to write the hold information into every mailbox and then check for individual holds when evaluating items for deletion. Here is how to retrieve information about the holds for org-wide retention policies.

```
[PS] C:\> Get-OrganizationConfig | Select -ExpandProperty InPlaceHolds
```

```
mbx15382841af9f497c83f9efe73e51888d:1
mbx9696959111f74ecda8a40aef97edd2c2:1
mbx703105e3b8804a1093bb5cb777638ea8:1
mbxf6a1654abdba4712a43c354e28a4d56c:2
grp6a1654abdba4712a43c354e28a4d56c:2
```

“mbx” refers to policies applied to user mailboxes and “grp” means policies applied to group mailboxes.

## Phantom Holds for Sensitivity Labels

If your tenant uses sensitivity labels, sensitivity label policies exist to publish the labels to workloads. As explained earlier, sensitivity policies appear like retention policies and must be excluded if you want to focus solely on retention policies. Even though the policies only publish labels, the policies insert entries into the set of org-wide retention holds held in the Exchange organization configuration. These holds can be ignored.

## Discovering Holds in Force for a Mailbox

Taking all that we know about the different forms of holds into account, we can come up with a PowerShell script ([downloadable from GitHub](#)) to report the set of holds that are in force for a mailbox. The script accepts the name of a mailbox and reports any organization-wide holds followed by holds where the mailbox is in scope.

```
Enter User to check: Ben Owens
```

```
The following organization-wide mailbox holds are in force...
```

Name	Workload
GDPR Personal Data	Exchange, SharePoint, OneDriveForBusiness, ModernGroup
General sensitivity policy	Exchange, SharePoint, OneDriveForBusiness, ModernGroup
Office 365 for IT Pros ...	Exchange, SharePoint, OneDriveForBusiness, ModernGroup
Formal Company Records	Exchange, SharePoint, OneDriveForBusiness, ModernGroup
Company Confidential Po...	Exchange, SharePoint, OneDriveForBusiness, ModernGroup

```
The following specific holds are in place on the Ben Owens (Business Director) mailbox...
```

```
Exchange In-Place Hold: New EAC Search
```

```
Hold Applied by: Senior Leadership Team (SLT) Retention Policy on Exchange, SharePoint, OneDriveForBusiness, ModernGroup
```

```
Hold Applied by: Management Preservation Policy
```

```
Litigation hold is enabled on the mailbox Ben Owens (Business Director)
```

If your tenant uses sensitivity labels, you need to amend the code to exclude the holds set by sensitivity label policies in the Exchange organization configuration. See the earlier discussion.

## The ComplianceTagHoldApplied and DelayHoldApplied Properties

If a user applies a retention label that’s configured to retain content to an item or folder in their mailbox, Exchange Online updates the *ComplianceTagHoldApplied* property for the mailbox. Because a specific hold is in place for one or more items, Exchange regards the mailbox as being in the same state as if an administrator assigned it a retention policy or placed the mailbox on litigation hold. For this reason, if the account for the mailbox is removed, Exchange puts the mailbox into an inactive state and keeps it there until the longest retention period for any label assigned to the mailbox expires. To see a list of mailboxes with applied retention labels, use the command:

```
[PS] C:\> Get-Mailbox -Filter {ComplianceTagHoldApplied -eq $True} | Select DisplayName
```

If a mailbox is not listed, it does not mean that its contents are not on hold due to a retention policy, in-place hold, or litigation hold. It simply means that the mailbox owner has never applied a retention label to an item or folder.

## Delayed Holds

After any type of hold is removed from a mailbox, Exchange puts the mailbox into a delayed hold state and keeps it there for 30 days. The extra retention step exists to give administrators the chance to recover information released by the hold after the hold is removed. In effect, this is a backstop to stop Exchange purging data in case a mistake is made in releasing a hold. While the mailbox is on delay hold, Exchange treats it as if it were on litigation hold and everything is kept. To mark the mailbox as being on delayed hold, Exchange updates the *DelayHoldApplied* property to *\$True*, so we can check to see what mailboxes are in this situation:

```
[PS] C:\> Get-Mailbox -RecipientTypeDetails UserMailbox -Filter {DelayHoldApplied -eq $true} |  
Select DisplayName
```

It takes about an hour after releasing a hold before the delay hold status shows up for a mailbox. Exchange does not expose any information about when the delayed hold status starts or ends, nor does Exchange give any information about the removed holds. When the 30-day delay hold period elapses, Exchange sets the property back to *\$False* to tell the Managed Folder Assistant that it can now purge the items that were on hold.

If the need exists to release the delayed hold, you can do so by running the *Set-Mailbox* cmdlet to set the *RemoveDelayHoldApplied* switch. For example, this command clears the delay hold for any mailbox where it is set.

```
[PS] C:\> Get-Mailbox -RecipientTypeDetails UserMailbox -Filter {DelayHoldApplied -eq $true}  
| Set-Mailbox -RemoveDelayHoldApplied
```

You can remove a delayed hold for an inactive mailbox to speed its removal:

```
[PS] C:\> Set-Mailbox -Identity Peter.Jenkins -InactiveMailbox -RemoveDelayHoldApplied
```

## Exchange Retention Holds

Apart from in-place holds, Exchange Online supports retention holds. This is not a hold of the type generally referred to as Exchange in-place or litigation holds, holds due to retention policies, or eDiscovery case holds, all of which stop the removal of information by a user or a system process (except retention processing). An Exchange retention hold stops the Managed Folder Assistant from processing the retention policy for a mailbox for a period and is usually applied when a mailbox owner cannot manage their mailbox because they are ill, on vacation, or absent for some other reason.

Exchange retention holds are still in use. Because they affect how MFA processes mailboxes, they influence the retention policies applied to mailboxes. In other words, if a retention hold stops MFA from processing a mailbox, MFA will not process either Exchange mailbox retention policies or retention policies for that mailbox while the retention hold exists. To set a retention hold, set the *RetentionHoldEnabled* property for a mailbox as follows:

```
[PS] C:\> Set-Mailbox -Identity "Kim Akers" -RetentionHoldEnabled $True
```

When the user can resume working with their mailbox, it is usual to give them a week or so to allow them to process new messages awaiting their attention. You can then disable the retention hold by setting the *RetentionHoldEnabled* property to *\$False*. MFA will then restart applying retention policies to the mailbox.

## Moving Data Between Tenants

Content can have multiple policies controlling how it is retained or disposed of. These policies exist within a tenant but have no application elsewhere. Therefore, if you move content from a tenant, even to another

tenant, the effect of retention policies and labels disappear. This problem already exists when users move from on-premises environments to the cloud and the situation is similar in merger/acquisition scenarios when tenants join or split. The migration technology available from Microsoft does not take retention policies into account because the focus is on extracting information from one tenant to bring it to another. It is unlikely that Microsoft will change its approach due to the complexities of reconciling retention demands as content moves from one organization to another. Instead, they'll leave the problem to ISVs to solve in third-party migration products, some of which do a better job in this respect.

Some basic steps can be taken to help:

- Understand how data governance functions within the source tenant (or tenants, in the case of amalgamations). Know what retention policies exist and how those policies treat content.
- Prepare for the migration by understanding where data is stored after the move.
- Create a new data governance strategy for the target tenant.
- Execute the strategy as data is moved from source tenants so that it is managed from day 1. In other words, have the retention policies in place in the target tenant so that newly arriving content is preserved as soon as it is transferred. This task might take some PowerShell scripting.

Data subject to holds remains in the source tenant until its retention period expires. That is, the data is kept if the owning tenant is funded (licensed). In the case of tenant mergers, the source tenants will likely eventually close. At this point, any data kept due to retention policies will be removed from workloads.

# Understanding the Exchange Mailbox Lifecycle

Items held in Exchange Online mailboxes follow a lifecycle from their creation by the user (calendar, posts, tasks, contacts, and so on) or when delivered as new messages. The major points in the lifecycle of Exchange Online items are:

- **Mailboxes** store and process incoming and outbound emails. Mailboxes can be archive-enabled to enable long-term retention of less-frequently accessed items for extended periods.
- The **Deleted Items** folder is the equivalent of the Trash or Wastebasket folder used by other email systems. The Deleted Items folder holds items that the user removes from other folders. Items stay in the Deleted Items folder until the user empties the folder or a retention policy dictates the removal of the items.
- Items removed from the Deleted Items folder move into **Recoverable Items**, a set of sub-folders used to preserve deleted items. The Recoverable Items folder also holds system items, such as audit entries to record actions taken within the mailbox. The Recoverable Items structure is not part of the normal IPM\_SUBTREE folder tree exposed by clients. It is in the hidden part of the mailbox that clients never synchronized to local copies. The Recoverable Items folder structure is also known as the "dumpster". This term goes back to its first implementation allowing users to recover deleted items without the need for administrator intervention.
- The **Managed Folder Assistant** (MFA) is a background assistant that processes mailboxes regularly to remove items no longer needed from the Recoverable Items structure, remove items from other folders, and move items into folders in archive mailboxes. Retention policies applied to mailboxes control MFA's processing. Holds that might exist for those mailboxes constrain the removal of items until their retention period lapses.
- Administrators create **Retention policies and tags** to define how the MFA automatically manages mailbox items on behalf of users. A mailbox can have a single retention policy, consisting of a set of tags that apply to default folders (such as the Inbox and Sent Items folders), other folders and items



as dictated by the user, and every other item in the mailbox (including the archive if available) that is not under the control of another tag.

- Administrators can apply different forms of **holds** (retention hold, legal or litigation hold, or in-place hold) to mailboxes to control whether users can remove information. Holds usually come about through a legal or other authorized action that forces a company to keep data for some purpose. When a mailbox is on hold, MFA will not permanently remove items from the mailbox, and Exchange Online copies any item that the user attempts to remove or update.

Collectively, these components interact with each other to form the lifecycle of Exchange Online mailbox items. Let's discuss what happens during the deletion of items to gain a better understanding of how the different parts of the lifecycle fit together.

## Cleaning Mailboxes

Over time, users remove items from various mailbox folders. Some remove messages immediately after reading, and some keep everything and leave items to accumulate. Exchange uses a two-stage removal process. First, the item is "soft-deleted" and moves to the Deleted Items folder, a default mailbox folder that acts as a convenient collection point for any item removed from another folder. Then, if the user empties the Deleted Items, the items are "hard deleted" (sometimes called "purged") and moved to the Recoverable Items\Deletions folder. You can also assign a folder retention tag to the Deleted Items folder to govern how long items stay in the folder.

A user can force a "hard delete" for an item by using the *Shift+Delete* key combination in either Outlook or OWA. This command instructs Exchange to ignore the Deleted Items folder and move the item directly to the Recoverable Items\Deletions folder. Exchange also moves items into the Deletions folder when a user moves an item from a mailbox folder to a PST. The reason here is that the move is a combination of an item deletion from the mailbox folder and an item creation in the PST.

When a user soft- or hard-deletes an item, Exchange writes a pointer to the original folder into the item's Last Active Parent Folder Identifier property (the MAPI property is *LastActiveParentEntryID* or LAPFID). Knowing which folder an item originally came from allows OWA to restore the item into that folder using the **Recover Deleted Items** feature, even if the mailbox owner renames or moves the folder. Exchange Online has recorded LAPFID information since 2016, but some older items might not have a value stored in this property. If this is the case, OWA uses a scheme called "folder type origin" to figure out where to recover items. Calendar items go back to the Calendar folder, Task items into Tasks, Contacts into Contacts, and mail and any other items go into the Inbox. It can be a little strange to find recovered items in the Inbox, especially if they are not mail items. For instance, if you recover a Word document or Excel spreadsheet (many people store these files in their mailbox), it shows up in the Inbox.

Outlook clients use a different mechanism to recover items. First, Outlook does not use the LAPFID to restore items. Second, Outlook recovers all items into the Deleted Items folder, the idea being that users can then decide where to move the items to as a permanent location. Microsoft might upgrade Outlook to support LAPFID in the future. Table 17-5 summarizes the target recovery location used by Outlook and OWA for different types of items.

	<b>Outlook</b>	<b>OWA</b>
Mail item (message)	Deleted Items	Original folder or Inbox
Contact item	Deleted Items	Original folder or Contacts
Calendar item	Deleted Items	Original folder or Calendar
Task item	Deleted Items	Original folder or Tasks

Any other non-mail item	Deleted Items	Original folder or Inbox
-------------------------	---------------	--------------------------

Table 17-5: The recovery destination for various Exchange item types

If you remove a complete folder, you can recover the individual items within the folder, but you cannot recover the complete folder as a single entity.

## Deletions, Purges, and Versions

The Recoverable Items structure uses sub-folders to organize items that it needs to keep. Some items stay until their deleted items retention period elapses. Others stay in the structure for a lot longer because one or more holds exist on the mailbox.

Items reach the Recoverable Items\Deletions folder when:

- They are hard-deleted.
- They are deleted from the Deleted Items folder.
- The Deleted Items folder is emptied.

Items stay in the Deletions folder for the period set in the deleted items retention period for the mailbox. In Exchange Online, the default period is usually 14 days. You can increase the deleted items retention period to a maximum of 30 days by running the *Set-Mailbox* cmdlet (and while you will probably use whole days, you can specify a retention time down to the second). For example:

```
[PS] C:\> Set-Mailbox -Identity 'Sanjay Patel' -RetainDeletedItemsFor 29.23.57.03
```

An exception exists for calendar items, which Exchange keeps for 120 days.

**Recoverable Items Only Online.** Because the Recoverable Items folder is only present on the server, you can only use Recover Deleted Items when Outlook can make a network connection to Exchange Online. In addition, if you ever need to use [the MFCMAPI utility](#) to see the items in the various sub-folders under Recoverable Items, you must configure MFCMAPI to open the message store online. Do this by going to the **Tools** menu, selecting **Options**, then setting the checkbox "Use the *MDB\_ONLINE* flag when calling *OpenMsgStore*".

When the deleted item retention period expires for an item, MFA removes it from the database and the user can no longer recover the item. The exception to the rule is when a mailbox is on hold or has single item recovery enabled as Exchange then moves items into the Recoverable Items\Purges folder and keeps the items there until the hold or the single item recovery period lapses. During this period, administrators can recover items with a content search.

A user can try to remove an item permanently by using the Recover Deleted Items feature to select the item and then select **Purge**. What happens next depends on the *SingleItemRecoveryEnabled* setting for the mailbox. Exchange uses the Single Item Recovery (SIR) feature to ensure that a deleted item can be recovered during the longest possible time set by the deleted item retention period:

- If *SingleItemRecoveryEnabled* is *False* and the mailbox is not subject to a hold, Exchange removes the item from the database at once and it is irrecoverable.
- If *SingleItemRecoveryEnabled* is *True* (the default value) and the mailbox is not subject to a hold, MFA moves the item to the Recoverable Items\Purges folder. The item stays there until its deleted item retention period (between 14 and 30 days as defined for the mailbox) expires. At that point, MFA removes the item from the database and the item is irrecoverable. However, as discussed below, while this description is true, a retention policy might cause something different to happen.
- If *SingleItemRecoveryEnabled* is *True* and the mailbox is subject to a legal hold, Exchange moves the item into the Recoverable Items\Purges folder and keeps it there while the hold applies.

- If *SingleItemRecoveryEnabled* is *True* and the mailbox is subject to an in-place hold, part of MFA known as the Email Lifecycle Assistant (ELC) determines whether a copy needs to be retained to satisfy the hold criteria (query and date range). If this is the case, MFA moves the item to the Recoverable Items\DiscoveryHolds folder, where the item stays until the hold lapses. ELC examines items that have exceeded the deleted items retention period in mailboxes at least once a day.

See Chapter 18 for more information on how to apply in-place holds to mailboxes.

**Nothing moves to DiscoveryHolds:** If you review items in the Recoverable Items structure, you might see that nothing ever moves into the DiscoveryHolds (or Purges) folder in the primary mailbox, even if those items are subject to an in-place hold. This can happen when the mailbox is under the control of the Default MRM Policy and is archive-enabled. Here is why.

When the Deleted Items folder is emptied, or items are hard-deleted by users, they end up in the Deletions sub-folder. Items in any folder under Recoverable Items are subject to the folder tag contained in the MRM policy, which instructs the Managed Folder Assistant to move the items to the archive after 14 days. If the mailbox is archive-enabled, the items are moved. If not, they stay in the primary mailbox.

ELC processes items when they reach the deleted items retention period. By default, this is between 14 and 30 days. ELC moves items subject to a hold to the DiscoveryHolds folder when their deleted items retention period expires. However, no items reach the 30-day deleted item retention period in the primary mailbox because they have already been moved to the archive. ELC processes the archive and will move the held items to the DiscoveryHolds folder in the archive, but no trace is seen of an item in the DiscoveryHolds folder in the primary mailbox. This can be confusing at first, but it is quite logical when you consider the age limits that control where items stay for different periods in a mailbox.

No client can permanently remove items under hold, including low-level utilities such as MFCMAPI. Exchange integrates checks for holds into how it manages data in its databases, and no one can circumvent the effect of a hold. Preventing the unauthorized removal of data from mailboxes allows Exchange to preserve items in an immutable fashion when needed by an organization.

Exchange monitors mailbox items that are subject to an in-place hold to detect if the user or any other process changes the item. If this happens, a copy of the original item is moved into the Recoverable Items\Versions folder to ensure that everything that happens to an item is fully recorded. This action is called a "copy on write". When the hold elapses, MFA removes the items from the Versions folder along with items held in the DiscoveryHolds and Purges folders.

**The SearchDiscoveryHoldsFolder folder:** This is a sub-folder of the DiscoveryHolds folder that is used by ELC when it processes the DiscoveryHolds folder to decide if any items can be removed. If such an item is found, it is moved to *SearchDiscoveryHoldsFolder* and kept there until it is eventually removed by the Managed Folder Assistant.

## Recoverable Items Quota

The Recoverable Items structure has a separate storage quota from that given to the primary mailbox. If a mailbox is archive-enabled, a separate quota is available to a similar Recoverable Items structure in the archive mailbox. The default quota for recoverable items is 30 GB, and Exchange Online increases the quota to 100 GB automatically when the first hold is applied to a mailbox. If a mailbox is created when org-wide holds exist in the organization, its recoverable items quota is set to 100 GB. Between the primary and archive mailbox, Exchange can keep up to 200 GB of recoverable data for a user. Items stored in the Recoverable Items folder in the archive behave in the same manner as if they were in the primary mailbox and are discoverable by eDiscovery searches.

You cannot change the *RecoverableItemsQuota* for a mailbox with the EAC or PowerShell, but you can log a support case with Microsoft Support if a mailbox is likely to exhaust its recoverable items quota. To find out how much data exists in Recoverable Items for a mailbox, use either of the *Get-MailboxStatistics* or *Get-MailboxFolderStatistics* cmdlets:

```
[PS] C:\> Get-ExoMailboxFolderStatistics -Identity Kim.Akers -FolderScope RecoverableItems | Format-Table Name, FolderSize, ItemsInFolder, FolderAndSubFolderSize -AutoSize
```

Name	FolderSize	ItemsInFolder	FolderAndSubFolderSize
Recoverable Items	0 B (0 bytes)	0	563.3 MB (590,650,854 bytes)
Audits	523.1 MB (548,504,424 bytes)	149324	523.1 MB (548,504,424 bytes)
Calendar Logging	18.37 MB (19,261,189 bytes)	734	18.37 MB (19,261,189 bytes)
Deletions	10.72 MB (11,236,121 bytes)	115	10.72 MB (11,236,121 bytes)
DiscoveryHolds	0 B (0 bytes)	0	0 B (0 bytes)
SearchDiscoveryHoldsFolder	0 B (0 bytes)	0	0 B (0 bytes)
Purges	5.707 MB (5,984,254 bytes)	753	5.707 MB (5,984,254 bytes)
Versions	5.402 MB (5,664,866 bytes)	53	5.402 MB (5,664,866 bytes)

Add the *Archive* switch to the cmdlet parameters if you want to see data for the archive mailbox.

If the mailbox exceeds the Recoverable Items quota:

- The user cannot remove items from their mailbox.
- The Managed Folder Assistant cannot move (hard-delete) items into the Recoverable Items.
- Write-on-protect cannot take copies if users alter items subject to a litigation or in-place hold.
- Exchange cannot save audit items if mailbox auditing applies to the mailbox.

You can take three actions to restore the Recoverable Items folder to normal working order. First, you can ask Microsoft to increase the recoverable item quota for the problem mailboxes. This might take a day or so to be effective and during that time the mailbox might experience problems such as those listed above. The second solution is to run the Managed Folder Assistant and instruct it to clean up duplicate items that might be present in Recoverable Items. This step should have an immediate effect and restore the mailbox to good health. To run the Managed Folder Assistant in clean-up mode, use a command like this:

```
[PS] C:\> Start-ManagedFolderAssistant -Identity TRedmond -HoldCleanup
```

Check the reported data for the Recoverable Items folder after the Managed Folder Assistant completes and hopefully you will discover that the overall size goes down. For more information about how to track what the Managed Folder Assistant does, see the section “Logging the Managed Folder Assistant” later.

The last action is to follow [the procedure laid down by Microsoft](#) to use PowerShell to clean up Recoverable Items.

## Exchange Mailbox Retention Policies

As we know, retention policies allow administrators to control how long data workloads retain or remove data. Both Exchange Online and Exchange on-premises servers support mailbox retention policies. Although these policies only cover Exchange mailboxes, they are still important where tenants want to impose the same retention regime for both online and on-premises mailboxes. In addition, mailbox retention policies can apply control at the level of individual folders instead of applying the same settings to a complete mailbox. Finally, mailbox retention policies can move items to archive mailboxes, something that is still not possible with Microsoft 365 retention policies.

Mailbox retention policies instruct the Managed Folder Assistant (MFA) about how long users can keep mailbox content (the retention period) and what to do once the retention period lapses (the retention action). The choice for retention action is to either remove the item (recoverable or permanent) or move it to the archive mailbox. A mailbox can have only one retention policy, and often the mailbox receives this policy

when created through the application of mailbox plan settings (see the Managing Exchange chapter). Many mailbox retention policies can exist within a tenant to ensure that the needs of different user groups are met.

A mailbox retention policy consists of one or more tags. These are:

- **Tags for the default mailbox folders (folder tags):** Among the default folders are the Inbox, Sent Items, and Deleted Items folders. Folder tags control how long items remain in these folders and what happens once the retention period elapses. Items in folders with folder tags inherit these tags unless the mailbox owner assigns them a personal tag.
- **Personal tags:** Users can apply these tags to any mailbox item and any folder except an Exchange default folder. A personal tag always has precedence over any other tag.
- **Default tags:** A mailbox retention policy can include a default delete tag and a default archive tag. The first determines the deletion of items not under the control of a folder or personal tag. The second determines when MFA moves items into the archive.

A mailbox retention policy commonly contains several folder tags, some personal tags, and one or both default tags.

Details of how to create and manage Exchange mailbox retention policies are in Chapter 8 of the companion volume.

## How MFA Processes Retention Policies

MFA processes both Microsoft 365 retention policies and mailbox retention policies against mailboxes. When MFA processes a mailbox to apply a retention policy, it examines the mailbox to “stamp” folders and items according to policy settings. MFA updates several MAPI properties for mailbox items with information, including:

- A GUID for the retention policy. Administrators can update the display name of a name, but its GUID (the retention identifier) is immutable.
- Retention period. The number of days that an item can stay in the folder before MFA removes it.
- Retention expiry. The calculated date when MFA will remove an item.
- Retention flags, including whether the item inherits the retention tag from the parent folder, or a user assigned the tag to the item.
- Archive period. The number of days that an item can stay in its folder before MFA moves it into the archive.
- Archive date. The calculated date when MFA will move an item to the archive.

In addition to processing normal messages, MFA can process compliance records for Teams, Yammer, and Planner. If a mailbox only has a mailbox retention policy, MFA ignores the compliance records. If it has a Microsoft 365 retention policy, MFA processes the compliance records for Planner along with other mailbox content. Special Microsoft 365 retention policies apply to Teams compliance records and Yammer compliance records.

### Logging the Managed Folder Assistant

You can use the *Export-MailboxDiagnosticLogs* cmdlet to find out the last time that the Managed Folder Assistant processed a mailbox and what happened during the run. The “Elc” (Electronic life cycle) properties in the report contain details of MFA activity. For example, the *ElcLastRunSuccessTimeStamp* tells you the date and time that MFA last successfully processed the selected mailbox while *ElcLastRunDeletedFromDumpsterItemCount* holds the total number of items removed from the Recoverable Items folder (the famous “dumpster”). In this truncated version of the output, we can see the last successful run of the assistant was at 01:03 on 12 April 2020 and that MFA removed 554 items (from all folders under the mailbox root), archived 16 items, and moved 556 items from the Recoverable Items structure to the archive.

```
[PS] C:\> $Log = Export-MailboxDiagnosticLogs -Identity James.Ryan -ExtendedProperties
$xml = [xml]($Log.MailboxLog)
$xml.Properties.MailboxTable.Property | ? {$_.Name -like "ELC*"}

Name                                     Value
----                                     -
ElcLastRunTotalProcessingTime           97103
ElcLastRunSubAssistantProcessingTime    61121
ElcLastRunUpdatedFolderCount           9
ElcLastRunTaggedFolderCount             0
ElcLastRunUpdatedItemCount              881
ElcLastRunTaggedWithArchiveItemCount    0
ElcLastRunTaggedWithExpiryItemCount     881
ElcLastRunDeletedFromRootItemCount      554
ElcLastRunDeletedFromDumpsterItemCount  0
ElcLastRunArchivedFromRootItemCount     16
ElcLastRunArchivedFromDumpsterItemCount 556
ELCLastSuccessTimestamp                  12/04/2020 01:13:48
```

## MFA Workcycle

MFA runs on a workcycle basis where a goal is set and the server hosting the mailbox figures out how to meet the goal, taking system load and available resources into account. In an on-premises deployment, the MFA workcycle aims to process every mailbox at least once daily. However, in Exchange Online, the workcycle used for the MFA aims to process every mailbox at least once weekly. The larger quotas granted to Exchange Online mailboxes is one reason why it is safe to extend the workcycle. However, although the formal workcycle goal is weekly, experience (and observation of mailbox statistics as described above) proves that MFA can process mailboxes up to five times weekly. If system resources are available, Exchange Online releases the resources to background processes like MFA, and this accounts for any discrepancy between formal workcycle goals and what happens in practice.

Knowing how to check whether MFA has processed a mailbox, we can write some code to check all user mailboxes. It's easy to amend this code to select a different group of mailboxes for processing, such as all those belonging to a department. In addition, the script only reports two of the ELC properties and you could add others as needed, such as the count of items moved to an archive mailbox.

```
[PS] C:\> $Mbx = Get-Mailbox -RecipientTypeDetails UserMailbox -ResultSize Unlimited
$Report = [System.Collections.Generic.List[Object]]::new()
ForEach ($M in $Mbx) {
    $LastProcessed = $Null
    Write-Host "Processing" $M.DisplayName
    $Log = Export-MailboxDiagnosticLogs -Identity $M.Alias -ExtendedProperties
    $xml = [xml]($Log.MailboxLog)
    $LastProcessed = ($xml.Properties.MailboxTable.Property | ? {$_.Name -like
"*ELCLastSuccessTimestamp*"}).Value
    $ItemsDeleted = $xml.Properties.MailboxTable.Property | ? {$_.Name -like
"*ELcLastRunDeletedFromRootItemCount*"}
    If ($LastProcessed -eq $Null) {
        $LastProcessed = "Not processed"}
    $ReportLine = [PSCustomObject]@{
        User           = $M.DisplayName
        LastProcessed  = $LastProcessed
        ItemsDeleted   = $ItemsDeleted.Value}
    $Report.Add($ReportLine)
}
$Report | Select User, LastProcessed, ItemsDeleted
```

## Forcing MFA to Process a Mailbox

If you need the Managed Folder Assistant to process a mailbox urgently, for instance, to publish the retention policy information to a mailbox, you can run the *Start-ManagedFolderAssistant* cmdlet to request Exchange Online to process the retention settings for the mailbox.

```
[PS] C:\> Start-ManagedFolderAssistant -Identity "Ben Owens"
```

If all goes well, MFA refreshes the retention settings for the mailbox and new tags and labels are available to users. Sometimes kicking the MFA convinces it to do the right thing and process a mailbox, sometimes it does not. The problem is that Exchange Online throttles background processing and does not allow mailbox assistants like MFA to run on demand. Microsoft wants to smoothen server load to reduce the risk that background processing interferes with the ability to be responsive to clients, which is why a seven-day workcycle exists for mailbox assistants. Telling MFA to process a mailbox will have an effect if MFA considers that it is reasonable to go ahead because it has not processed the mailbox recently. Running *Start-ManagedFolderAssistant* several times to convince MFA to start processing is a fruitless exercise.

The *Start-ManagedFolderAssistant* cmdlet supports an *-InactiveMailbox* switch to force immediate processing of inactive mailboxes. In the past, Microsoft has swapped between MFA processing inactive mailbox and not. The current state is that MFA processes inactive mailboxes and applies Exchange mailbox retention policies and retention policies against their content. One difference exists in the processing done for active and inactive mailboxes in that MFA does not process archive tags when it deals with inactive mailboxes.

## Stopping MFA Processing a Mailbox

Times occur when you might want to stop the Managed Folder Assistant from processing one or more mailboxes. For example, when you want to remove items from a mailbox when it is on hold because that information should not be in the mailbox as when someone forwards a message with confidential information to mailboxes that should not receive them. To remove the items from a mailbox that is on hold, you must release the holds temporarily, remove the items, and then replace the holds. During this time, MFA might process the mailbox and remove other items. To stop this from happening, you can disable MFA temporarily by running the *Set-Mailbox* cmdlet to set the *ElcProcessingDisabled* flag.

```
[PS] C:\> Set-Mailbox -Identity "Kim Akers" -ElcProcessingDisabled $True
```

When the holds are in place again, you can reverse the process and release MFA by switching the value to *\$False*. You cannot disable MFA processing if a preservation lock applies to the mailbox.

## Moving from Exchange Retention Policies

Exchange mailbox retention policies have been in use since Exchange 2010. Microsoft's long-term direction is to move away from workload-specific processing, like that done by mailbox policies, to use retention policies instead, whenever it makes sense for a tenant. If you are a new tenant starting with a retention strategy, the right decision is to use retention policies. However, if you have been using mailbox retention policies for many years and have a hybrid tenant, that decision is not quite so clear-cut.

To make retention policies available to other workloads, Microsoft evolved and expanded the core principles behind Exchange retention policies. In doing so, they have dropped some Exchange-specific features, like the ability to move items to archive mailboxes. Losing the ability to archive items automatically is regrettable but this action is only valid for Exchange and does not apply to other applications. A more fundamental problem for some is the loss of granularity in that an Exchange retention policy can include tags for default folders, personal tags, and a set of default tags to control deletion, archival, and voicemail. By comparison, a retention policy applies the equivalent of a single default tag (to remove or keep content) to mailboxes.

Retention policies create in-place holds when they include a retention action. The holds apply for all locations covered by a policy and last until the policy's retention period expires, meaning that users cannot remove content from Exchange, SharePoint, OneDrive for Business, Teams, or Groups if that content comes within the scope of the policy. The integration of in-place holds into retention policies is an advantage over Exchange mailbox retention policies.

Significant differences in functionality exist between the two types of retention policies. The most attractive feature of retention policies is that they apply to many locations rather than just Exchange Online. Another

good point in their favor is that you can combine Exchange retention policies with retention policies and labels, even if that might be a double-edged sword in terms of the resulting complexity in retention processing. Being able to apply a single consistent policy across multiple workloads is a huge advantage, but that advantage might be reduced because of the loss of some of the granular processing available in Exchange retention policies. The decision to move to retention policies will therefore need considerable thought and preparation on the part of tenants.

Every tenant is different and although it might be easy for a cloud-only tenant with relatively simple retention needs to go ahead and embrace retention policies, the situation is probably very different for large and complex tenants that already have a well-defined retention strategy in place. Things become even more complicated for hybrid tenants, who often want to use the same processes on-premises and in the cloud.

Experience and time will allow us to develop better answers. In the meantime, new tenants should start with retention policies and labels while older tenants test, compare, and contemplate what is their best course of action. For instance, you might decide on a strategy based on four stages:

1. Remove personal tags from Exchange Online mailbox retention policies and replace the tags with Microsoft 365 retention labels with the same name and retention settings. Users will begin using the retention labels to retain new items while the older retention tags will age out over time.
2. Remove folder tags from Exchange Online mailbox retention policies and replace them with Microsoft 365 retention policies. This means that no folder-specific retention processing is available because the retention policies apply the same settings and actions to the complete mailbox.
3. Replace default deletion tags in Exchange Online mailbox retention policies with Microsoft 365 retention policies. This is a natural consequence of step 2 and the only action required once a Microsoft 365 retention policy processes mailboxes is to remove the default deletion tag from the mailbox retention policy.
4. Eventually, limit the use of Exchange Online mailbox retention policies to moving items to archive mailboxes. When Microsoft 365 retention policies support this action, you can remove the default archive tags from Exchange Online mailbox retention policies.

This is an outline of tactics to move to Microsoft 365 retention policies that need to be adjusted to meet the unique circumstances of individual organizations. For example, you should run Exchange Online mailbox policies (with all tags intact) and Microsoft 365 retention policies alongside each other for a few weeks or so to make sure that no interregnum happens when retention processing does not occur for mailboxes. MFA will resolve any inconsistencies (such as having a personal tag and retention label with the same name).

The availability of the older workload-specific functionality allows organizations time to make the transition. It's a wise approach because the nature of retention is that items often require retention for long periods, and no one wants software to force them to change their data governance strategy in such a way that it might affect terabytes of retained content.



# Chapter 18: Managing eDiscovery

**Tony Redmond**

The Microsoft Purview Compliance portal is the administrative interface for cross-service compliance and protection operations, including the tools Microsoft creates to help tenants to find information needed for investigations, now gathered under the Microsoft Purview banner. In this chapter, we cover:

- **Content searches:** How to search through locations (sites, mailboxes, teams, etc.) to find information. Information found by content searches can be exported for review by investigators or subject matter experts.
- **Microsoft Purview eDiscovery (Standard) – previously Core eDiscovery:** How to run eDiscovery operations spanning searches, holds, and exports. Standard eDiscovery is available to tenants with Office 365 E3 licenses to organize content searches (including exports) and in-place holds into cases to make it convenient for investigators to process information. A single eDiscovery case can span multiple searches and holds. The holds are *in-place*, meaning that the items that come within the scope of the hold remain inside the host repositories (for instance, an Exchange mailbox or a SharePoint document library).
- **Microsoft Purview eDiscovery (Premium) – previously Advanced eDiscovery:** Uses completely different technology to process high-end eDiscovery cases which typically cover much higher volumes of data than Standard eDiscovery cases do. The organization of Premium eDiscovery cases is different too. eDiscovery Premium requires Office 365 E5 or Microsoft 365 Advanced Compliance licenses.

eDiscovery is a specialized subject and not everyone is interested in how to run searches, create holds on different resources, or export results. However, given the massive increase in data stored by companies and the litigious nature of the world, it is likely that many tenants will meet the need to run some form of search in any given year, even just to find a document or message that a user thinks they have lost. With that thought in mind, we begin by looking at content searches.

## Content Searches

The design goal for eDiscovery technology is to meet the needs of investigation professionals who need to pursue a compliance query from start to finish, potentially to the point where a company produces evidence in court to prove wrongdoing or other legal points. Searching for evidence is where investigations begin, and that is why content searches are so important.

Before beginning the process to create a content search, you should know:

- **What locations to search:** You can search Exchange Online user mailboxes (including shared and inactive mailboxes), group mailboxes, Exchange Online public folders, SharePoint Online and OneDrive for Business sites, tasks (To-Do), Sway, and Forms (owned by Microsoft 365 Groups). Any information held in mailboxes, including compliance records for Teams channel and private conversations or Yammer messages, is available. Planner tasks are discoverable if compliance records have been captured when tasks are created or edited (see the Planner chapter). Yammer messages are discoverable (and can be found in Exchange Online mailboxes) if the network is configured in native mode.

- **The keyword query to use to find information:** Content searches depend on the content indexes created and updated by Microsoft Search. As users create information using a supported workload, Microsoft Search automatically processes the information and adds it to its content indexes. Queries can be all-encompassing (“find everything”), basic (“search for this term”), or very complex.

Content searches use the Microsoft 365 server fabric, which means that a tenant can run multiple concurrent searches. No limitation exists on the number of mailboxes or sites that you can include in a search as Microsoft designed this generation of search technology to be able to handle the demands of the largest tenant. While content searches take care of finding and exporting information from Microsoft 365 locations, if you need to apply a hold to keep email and documents until investigators no longer need that information, you create an eDiscovery case. The holds processed by eDiscovery cases depend on content searches to find the information to hold. Content searches can also be part of an eDiscovery case, where investigators can use a set of searches to interrogate various locations using different queries to retrieve the information that they need. Table 18-1 lists some scenarios where content searches are useful.

<b>Scenario</b>	<b>Search action</b>
Assess the risk to the business from data shared by users with external people from SharePoint Online or OneDrive for Business sites	Include the <i>ViewableByExternalUsers</i> keyword in the search and set it to True. In order to exclude the aspx files used by SharePoint Online, the full query is <i>ViewableByExternalUsers:true AND ContentType:document NOT FileExtension:aspx</i> .
Find documents that hold more than a certain number of sensitive information types (as defined for Data Loss Prevention – see chapter 19)	Include the <i>SensitiveType</i> keyword in the search and specify the sensitive information type to look for and the number of instances of that type that must be in a document for the search to retrieve it. For instance, <i>SensitiveType:"Credit Card Number 5.."</i> finds documents that have more than five credit card numbers. See <a href="#">this page</a> for more information.
Investigate whether anyone sent emails holding confidential information to a specific address outside the organization	Include the <i>Recipients</i> keyword and specify their SMTP email address in a search. For example, <i>Recipients:"TRedmond@Yandex.com"</i>
Investigate whether mailboxes have received a specific phishing message (or not)	Include the <i>Sender</i> and <i>Subject</i> keywords with values to search. For example <i>Sender:"SomeGuy@FortuneForYou.com" AND Subject:"I have transferred \$100million to you"</i>
Check that a message that has some specific text in its body is not circulating within the organization	Include the <i>Body</i> keyword in the query. For example, <i>Body:"We are about to be taken over by MegaCorp"</i>

Table 18-1: Content search scenarios

From July 2022, Microsoft limits the number of search jobs that tenants and individual users (administrators) can run. Jobs include searches and exports. The limits are outlined in Table 18-2.

<b>eDiscovery Premium</b>	<b>Tenant Limit</b>	<b>User limit</b>
Concurrent job limit (any jobs)	100	50
Concurrent job limit (tenant wide job)	50	25
<b>eDiscovery Standard</b>		

Concurrent job limit (any jobs)	50	25
Concurrent job limit (tenant-wide jobs)	5	5
Maximum number of jobs per day	500	
Maximum data size limit per day	2 TB	

Table 18-2: Search job limits

Microsoft says that the limits should not interfere with normal tenant operations and exist to ensure service stability and reliability. Searches can be resource-intensive operations, especially if asked to search through a large number of locations, such as all sites or all mailboxes. The limits exist to stop some organizations running too many jobs at the same time. From a marketing perspective, they also help to emphasize the difference between eDiscovery Standard and Premium.

## Limiting Search Sources for Better Results

A content search can scan tens of thousands of target repositories and return multiple terabytes of results. To ensure that searches execute in a reasonable time, it is important to limit the number of locations where the search looks for relevant information by either restricting the number of targets or by refining the search query so that it excludes irrelevant items from the search results. As the foundation for successful searches is based on precision, the latter approach is always the best one to pursue.

As a very large multi-national company, Microsoft follows many legal and regulatory requirements around the world. Its litigation team uses eDiscovery to manage investigations and has written a [white paper to explain how it handles eDiscovery cases](#). In 2017, Microsoft said that its use of Microsoft 365 compliance technology contributed to an annual cost saving of some \$4.5 million. Your mileage might vary.

## Content Search Scalability

The speed and scalability available to Microsoft 365 content searches originate from the way that the searches use the Microsoft 365 infrastructure. A limit exists for the on-premises searches performed by either Exchange or SharePoint based on the load that a single server can manage. For instance, the older In-place Hold and eDiscovery tool used in Exchange on-premises servers ([retired from Exchange Online in mid-2020](#)) uses a single server to control searches and uses synchronous connections to all the mailbox servers that host mailboxes within the scope of the search. The search eventually collates the results returned by individual servers to form the set of found items.

Content searches use the Microsoft 365 server fabric to process searches and split work across multiple servers. Asynchronous messages pass between the servers doing the work to keep them updated about search progress. This implementation reduces the potential for failure and parallelizes the workload to scale up to deal with far higher volumes of data (such as over 700,000 mailboxes in a single operation). Microsoft gathers statistics about the time needed to complete content searches and cites the guidelines shown in Table 18-3.

<b>Number of mailboxes</b>	<b>Average search time</b>
100	30 seconds
1,000	45 seconds
10,000	4 minutes
25,000	10 minutes

50,000	20 minutes
100,000	25 minutes

Table 18-3: Average time for content searches (source: [Microsoft](#))

Searches need some time to spin up, so the time cited by Microsoft is indicative rather than precise. Once the search starts, it rapidly processes target locations to find items based on the keywords and conditions in the search criteria. The number of mailboxes included in a search is the biggest single factor influencing how long the search will take. Apart from user-created items in the mailbox, the Microsoft 365 substrate stores many other items in user and group mailboxes, including “digital twin” copies of data from SharePoint Online and OneDrive for Business and compliance records for workloads like Teams and Yammer. Online mailboxes are usually larger than their on-premises counterparts and can have large recoverable items and archive components. All contribute to the number of items stored in mailboxes and increase the time needed to search mailboxes. The bottom line is that if you want fast records, be precise about the mailboxes included in a search.

Content searches also include some retry logic to handle the situation where a required mailbox or site is offline for some reason. Usually, a retry is enough to complete a search. See this page to understand more about [the limits applying to content searches](#).

## SharePoint and Exchange Support for Sensitive Information Types

You can combine keywords to build a search query that covers multiple conditions. SharePoint Online and Exchange Online have different abilities to use keywords. Some of the keywords are specific to documents and some to mailbox items. For example, you cannot include the *ViewableByExternalUsers* keyword in a search that scans Exchange Online mailboxes because this kind of sharing concept does not exist for Exchange. In addition, although both SharePoint Online and Exchange Online support [sensitive information types](#) (like credit card numbers, passport numbers, and national identification numbers) used by Data Loss Prevention and retention policies, only searches of SharePoint content support these keywords. If you include unsupported keywords in a search, the search ignores them when it builds its results. This [page gives guidance about the keywords](#) that you can use for Exchange and SharePoint locations.

## Creating and Running a Content Search

You must be a member of the Compliance eDiscovery Manager or Organization Management role groups to be able to create and execute a content search. In addition, an account used to conduct compliance searches should have a functional Exchange Online mailbox. Although no strict licensing requirement exists for a mailbox, the preview function for search results does not work if the account used for searching has no mailbox.

To access content searches, open the Microsoft Purview Compliance portal and go to the **Content Search** section. You then see a list of the existing searches for the tenant. You can select a search and continue working with it to amend its search criteria and run new queries or create a new search. In the example explained below, we create a new search from scratch. After naming the search and adding a description, we define the locations, keywords, and conditions to frame the search.

### Search Locations

After assigning a suitable name to the search, the next step is to define where the search should scan for matching items. The three basic locations (Figure 18-1) are:

- **Exchange mailboxes:** All Exchange Online mailboxes, including inactive mailboxes. Searching Exchange mailboxes also covers Teams and Yammer messages because the Microsoft 365 substrate captures compliance records for Teams and Yammer conversations in Exchange personal mailboxes (for personal and group chats and Yammer private messages, including the call records for people who participate in meetings or calls) and group mailboxes (for Teams channel and Yammer community conversations).
- **SharePoint sites:** Includes all SharePoint sites and OneDrive for Business accounts. Searching SharePoint includes the sites created to store Teams files, including the wiki (.mht file) created for each channel.
- **Exchange public folders:** If your organization uses public folders, you can search for information stored in any public folder. You cannot search a subset of public folders.

Checking the *Add App content for On-Premises users* option instructs the search to include the cloud-only mailboxes used to store compliance records captured for Teams and Yammer messages sent by hybrid, guest, and federated users.

## Locations

### Specific locations

Status	Location	Included	Excluded
<input checked="" type="checkbox"/> On	Exchange mailboxes <small>Microsoft 365 Groups Teams Yammer user messages</small>	All <a href="#">Choose users, groups, or teams</a>	None
<input checked="" type="checkbox"/> On	SharePoint sites <small>OneDrive sites Microsoft 365 Groups Teams Yammer user messages</small>	All <a href="#">Choose sites</a>	None
<input type="checkbox"/> Off	Exchange public folders		

Add App Content for On-Premises Users. [Learn more](#)

Back

Next

Figure 18-1: Specifying target locations for a content search

## Adding Exchange Online Mailboxes

To refine the set of Exchange mailboxes, click the **Choose users, groups, or teams** link and then use the search box to find the target mailboxes, Microsoft 365 groups, and distribution lists. You can't add dynamic distribution lists to a search. When the set of mailboxes is input, the search expands the current membership of distribution lists and adds the individual mailboxes to the set of target mailboxes.

**Disconnected mailboxes and Content Searches:** Although content searches can scan inactive mailboxes, you can't include disconnected mailboxes in a content search. These are mailboxes belonging to Microsoft 365 accounts where an administrator removes the Exchange Online license from the account. Disconnected mailboxes stay in a soft-deleted state for 30 days following the removal of the license before Exchange Online permanently deletes them from the system. If disconnected mailboxes are in the set of mailboxes used for existing searches, the search ignores them when it generates results.

Continue to add mailboxes (remember to select mailboxes from the list and then click **Choose**) until you have added the desired target selections to the search.

## Adding SharePoint Online and OneDrive for Business Sites

Click **Choose sites** to add new locations to the set of SharePoint and OneDrive sites. For each site, enter its URL or part of its display name into the search box (Figure 18-2). When you've found the full set of sites for the search, click **Add** to add them to the search criteria.

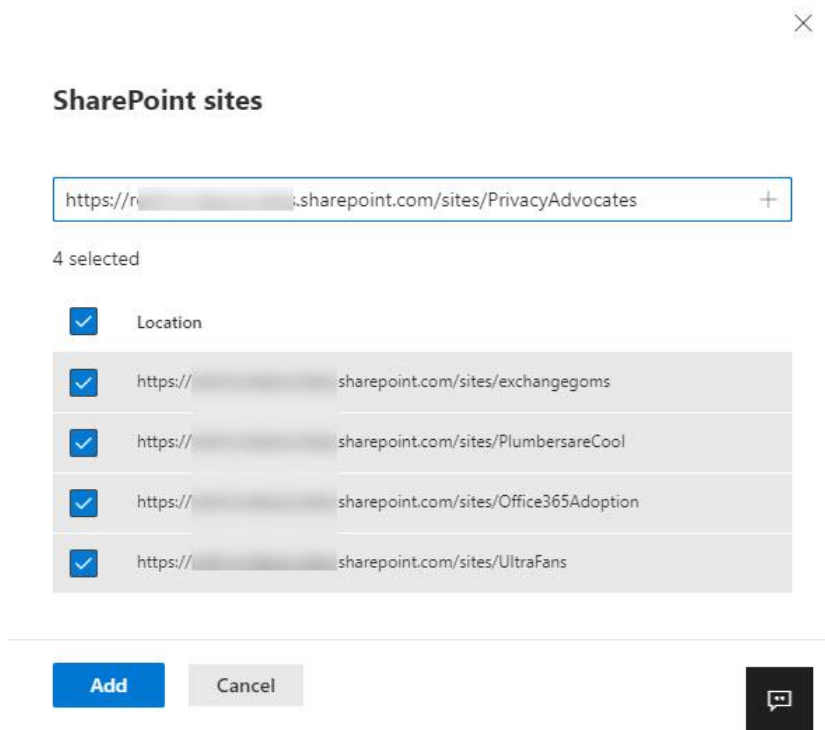


Figure 18-2: Specifying SharePoint Online sites for a content search

To add specific SharePoint Online and OneDrive for Business sites to a search, you must know the complete URL for the target sites. You can find the URL by opening the site in a browser and copying the URL, or you find the site URL with PowerShell. The `Get-SPOSite` cmdlet returns all the sites in a tenant, including their URLs, and is the obvious way to find a site URL, including for OneDrive for Business accounts. To find a list of all OneDrive for Business sites in the tenant, connect to SharePoint Online with PowerShell and run the command below. You could also export the list to a CSV file to keep as a handy reference:

```
[PS] C:\> Get-SPOSite -IncludePersonalSite $True -Limit All -Filter "Url -like '-my.sharepoint.com/personal/'" | Select Owner, Url | Sort Owner
```

If a site belongs to a Microsoft 365 Group or Team, you can find the URL with the command:

```
[PS] C:\> Get-UnifiedGroup -Identity MyGroup | Select SharePointSiteURL

SharePointSiteUrl
-----
https://mytenant.sharepoint.com/sites/mygroup
```

## Adding SharePoint Sites for Teams Private and Shared Channels

Teams private and shared channels keep their conversations and files separate from the other parts of the team to ensure that only channel members can access this content. Each private or shared channel has a dedicated SharePoint site, which is a child of the site owned by the team. Searching the parent team

automatically includes the conversations in a private channel (using the compliance records in the team's group mailbox). To search for conversations in a shared channel, you must add one or more members of the channel to search the compliance records stored in the personal mailboxes. If you want to add the SharePoint site belonging to a private or shared channel to a content search, you need to know its URL. The easiest way is to open the channel in Teams and use the "Open in SharePoint" option in the Files tab to open the channel site in the SharePoint browser interface. Then you can copy the URL shown in the browser navigation bar.

On an admin level, because the settings of these sites come from the parent team site, the SharePoint admin center does not manage the sites used by private or shared channels, but PowerShell can manage the sites. For example, when run by a SharePoint administrator, this PowerShell code returns all the sites belonging to private and shared channels and uses the group identifier to retrieve the team name from the Microsoft 365 group before writing out the display name and site URL.

```
[PS] C:\> [array]$Sites = Get-SPOSite -Template "TeamChannel#1"
ForEach ($Site in $Sites) {
    $SPOSite = Get-SPOSite -Identity $Site.url -Detailed
    $Group = Get-UnifiedGroup -Identity $SPOSite.RelatedGroupID.Guid
    Write-Host "Team" $Group.DisplayName "owns channel site" $Site.URL }
```

Notice that the SharePoint template used for the sites used by both private and shared channels is "TeamChannel#1." This is the current template used by Microsoft when it creates private and shared channel sites. Older private channel sites might have the TeamChannel#0 template. To return sites with both templates, use a command like:

```
[PS] C:\> [array]$Sites = Get-SPOSite -Limit All | ? {$_.Template -eq "TeamChannel#0" -or
$_.Template -eq "TeamChannel#1"}
```

When you examine the URLs created for sites used by private or shared channels, you'll see the following naming convention:

*<https://tenant.sharepoint.com/sites/TeamName-PrivateSharedChannelName>*

For example, the site for the "Legal Discussions" private (or shared) channel owned by the "Corporate Acquisition Planning 2020" team is:

*<https://tenant.sharepoint.com/sites/CorporateAcquisitionPlanning2020-LegalDiscussions>*

## Search Keywords and Conditions

Content searches run queries composed in the Keyword Query Language (KQL) against content indexes to find information. You can use either the condition card builder or KQL editor to compose a query. The KQL editor is also available to create content queries for Premium eDiscovery searches.

The condition card builder allows you to add keywords and conditions in a GUI while the KQL editor supports auto-completion for searchable properties and conditions along with checking for acceptable input, such as the available operators for a property. New search administrators usually find it easier to start with the condition card builder while those who are more experienced use the KQL editor. You can toggle between the condition card builder and the KQL editor as needed while working on a query. This is useful because you can compose a query with the condition card builder that has an error. When you toggle to the KQL editor, the editor highlights any syntax problems it detects to allow you to fix them. The KQL editor is there to make it easier for people to compose queries that don't contain obvious syntax errors. However, there's no guarantee that a query passed by the KQL editor will work or, if it does, will find the information you expect. For this reason, it's important to understand the fundamentals of KQL queries.

Queries are formed from keywords and conditions. The keywords are free-text expressions, meaning that you can use individual words or complete multi-word phrases. If you include phrases, you must enclose them in

double quotation marks. You can leave the keywords box empty if you want to search for everything in the locations that you choose.

Keywords support prefix matching, which means that you can include the wildcard operator (\*) with the beginning or end of a word (for instance, "Office\*"). You can combine free-text expressions with KQL operators such as OR, AND, NOT, and NEAR. For example, this query finds items in SharePoint Online or OneDrive for Business that include the keywords NDA, "Non-disclosure agreement," or "non disclosure" in Word or PDF formats. The reason for specifying non-disclosure and non disclosure is to pick up documents that use either variant:

```
(NDA OR "Non-disclosure agreement" OR "non disclosure") AND (filetype:doc* OR filetype:pdf)
```

This query uses a wildcard operator to find spreadsheets containing "Microsoft 365 Groups" and a special relative date comparison to specify that we're looking for spreadsheets modified in the current year:


```
("Microsoft 3*" AND "Groups") AND (filetype:xls*) AND (LastModifiedTime = "this year")
```

Other special date comparisons include today, yesterday, this week, this month, last month, and last year. We can also find items using date ranges. This query looks for any document or message created in October 2020 which includes the word *shareholder* within 10 terms of the word *agreement* where the author is Tony Redmond.


```
("shareholder" NEAR(10) "agreement") AND (Author:"Tony Redmond") AND (LastModifiedTime >= 2020-10-01 AND LastModifiedTime <= 2020-10-31)
```

To preserve the order in which words appear, use ONEAR instead of NEAR.

### Define your search conditions

Query language-country/region: None 


Condition card builder  
 KQL editor

**Keywords** 



crisis NEAR(10) pandemic


Show keyword list

**AND**

**Date** 

Between

2019-03-01  2022-01-24 

[+ Add condition](#) 

Back
Next
Cancel

Figure 18-3: Entering keywords and contents for a content search

If you include multiple free-text expressions in a query, KQL combines the expressions in the search using the AND operator. If you enter an operator in lowercase (like "and"), the search will offer to uppercase the operators when it checks the query. You should always accept this offer as uppercasing the operators makes




their purpose obvious (and lowercase operators don't work). Figure 18-3 shows an example of a search using some keywords to tell the search that we're interested in items where the words crisis and finance or found within 10 characters of each other together with a date condition to limit the scope of the search.

The Show keyword list checkbox allows you to input a keyword in each row of a list. The reason why you would want to do this instead of typing all the keywords into the keyword box is that the search generates statistics for each keyword. When the search runs, you can review the statistics to understand which keyword is most effective in terms of search results.

If you toggle to use the KQL editor for the query shown in Figure 18-3, you'll see the query in KQL format:

```
crisis NEAR(10) pandemic(c:c)(date=2019-03-01..2022-01-24)
```

The KQL editor supports cut and paste, so you can copy a complex query from one search, paste it into another, and modify the keywords and conditions as needed.

**Keywords with non-English characters:** By default, content searches are language-neutral. However, if you find that a search does not return the expected results, the cause might be that some of the keywords use non-English characters (such as Chinese). You can force a search to be language-sensitive by clicking the **Query-language-country/region** icon  on the top of the search query box and then selecting a language country culture code value for the search (for example, Chinese – Hong Kong SAR).

## Search Conditions

Conditions narrow searches by filtering the results generated by the keywords. When you specify a condition, the search adds a clause (shown using the (c:c) operator) to the search query used to find results. For example, you might only be interested in items created by a certain person, in which case you use the author condition to specify the person. To apply a further filter, you could add the created (date) condition to focus on items created in a certain period. When the search query runs, items must satisfy the keyword query and one or more conditions before the search includes them in the results.

To make it easy for you to include conditions in searches, the **Add conditions** button in the condition card builder exposes a dialog listing the most common conditions. [Some conditions](#), like Sender/Author, are valid conditions for searches against both Exchange Online and SharePoint Online items. Others are specific to a workload. If you attempt to add a condition that won't work for the chosen target locations, the Microsoft Purview Compliance portal flags an error to tell you which condition has a problem. For instance, if you try to use the message type condition to search SharePoint Online, you'll see an error because SharePoint doesn't support that condition. To proceed with the search, you must remove the condition.

To complete a condition and add it to the search criteria, you must define what the condition checks for. For instance, if you add a retention label condition to the search, you must say what compliance tag (classification label) or tags you want to look for. If you add a date to a search, you must input the date range for the search and say whether you want to look for items between two dates, before a date, or after a date.

Remember that people can change their names for good reasons, such as marriage or divorce, which means that you might need to include addresses for people other than their current user principal name or primary SMTP address. Queries can handle this situation by including variants of names in the sender/author condition.

**Hybrid SharePoint searches:** If a hybrid connection is in place, it is possible to use hybrid SharePoint searches to find content stored within on-premises SharePoint farms. However, when you run content searches, the search filters out the results from the on-premises sources because no method exists to allow

Microsoft 365 to apply holds to on-premises SharePoint data or to export data from an on-premises SharePoint location.

## Run the Search

After defining the search criteria, the Microsoft Purview Compliance portal launches a preview for the search automatically. Before this happens, the portal checks the query for any syntax errors and identifies any potential improvements. The search then runs the query against the target locations. A preview search is not a full search. Instead, the preview uses content indexes to retrieve a sample of items that a full search will find to allow the search author to decide if the criteria are accurate enough to find the desired information. The preview also generates estimates for the number of items a full search will find. The actual number of items that a full search will find will probably differ from the preview, but the difference is usually of minor importance. For instance, preview searches against SharePoint sites often include ASPX files (used by SharePoint) that are of no interest to an investigator. It's also possible that users create or delete items matching the search criteria between the time of a preview search and when the actual search happens, or indeed that a lag in indexing stops a file from appearing in a preview when a subsequent search finds it.

A full search might have to process thousands of mailboxes and sites, so it is better to get quick results back from a preview rather than having to wait for the full search to complete. Equipped with the preview information, you can assess the effectiveness of the query and then tweak the search parameters. The process of tweaking search criteria might need multiple iterations before you are happy with the results delivered by the search and proceed to perform the full search and export its results.

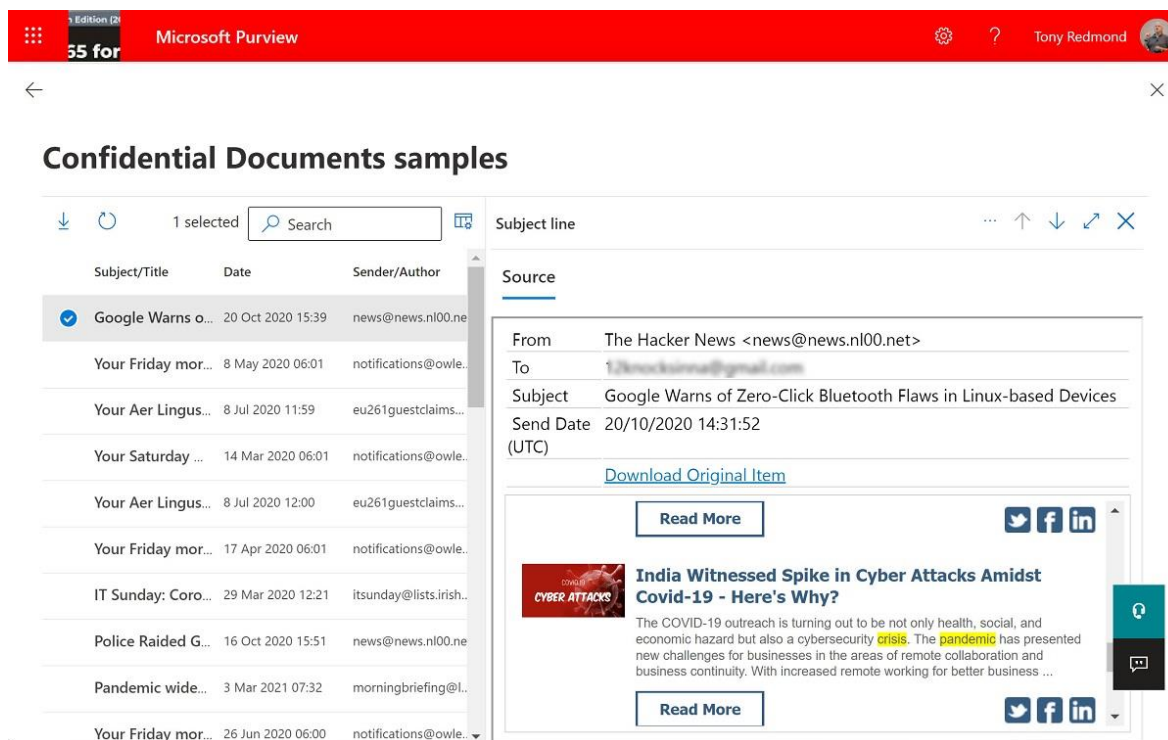


Figure 18-4: Reviewing sample items retrieved by a content search

When the preview search is ready, the search selects samples from the top locations (those that hold the most matches). You can view the items retrieved by the preview search by selecting the **Review sample** option from the search summary. As shown in Figure 18-4, sample items include email messages, Teams compliance records, and files (list items are not available for preview). If an item has an attachment, it also shows up as viewable if the attachment is an Office document. Preview can decrypt protected files and messages too. Naturally, the fewer the number of results, the less time it will take for the search to show the preview results.

If the preview of the source for a found item displayed in the viewing pane isn't sufficient to fully understand the content of an item, you can use the *Download Original Item* link in the item preview pane to download a copy of the item for review.

The Microsoft Purview Compliance portal [limits the number of items selected for preview](#). This is acceptable because the intention behind the preview sample is to help investigators understand the effectiveness of the search query in terms of finding the required data. The preview sample is not a tool to browse through the complete set of items uncovered by a search. Table 18-4 lists the limits for the preview screen in terms of its ability to display items found by searches.

<b>Limit</b>	<b>Number</b>
Maximum number of items per user mailbox displayed	100
Maximum number of items from all mailboxes displayed (the newest items are shown)	1,000
Maximum number of items found in SharePoint Online and OneDrive for Business sites displayed	200
Maximum number of SharePoint Online and OneDrive for Business sites that can be previewed for search results	200
Maximum number of items per public folder mailbox displayed	100
Maximum number of items from all public folder mailboxes displayed	200
Maximum number of public folder mailboxes that can be previewed	500

Table 18-4: Display limits for search preview

The search summary screen includes search statistics for the estimated items that a full search will find (Figure 18-5). The three sets of statistics are:

- **Search content:** Displays the estimated items for the locations scanned (SharePoint and Exchange), the number of locations with hits (matching results), the number of matching items found, and the size of those items.
- **Condition report:** Displays the location type (SharePoint or Exchange), the condition (search query used), the number of locations with hits, the number of matching items found, and the size of those items. If you use a keyword list, you can see how effective each keyword is in terms of locating items (a result recorded as Primary means the complete search query; Keyword means the results from a specific keyword). You can also see whether any unindexed items matched the query. Unindexed items are often graphic files like a bitmap or JPEG file, or the MP4 files used for Teams meeting recordings.
- **Top locations:** The sites and mailboxes where most matching items were found, including the number of those items and their size.

## Confidential Documents

Summary Search statistics

Search content

Estimated items by location

**33 items**

Estimated items by location

Exchange (33)

Estimated locations with hits

**3 location(s)**

Estimated locations with hits

Exchange (3)

Data volume by location (MB)

**14.5 MB**

Data volume by location

Exchange (14.5 MB)

Condition report

Download your search condition report.

Location type	Part	Condition	Locations with hits	Items	Size (MB)
Exchange	Primary	((("crisis") NEAR(10) ...	3	33	14.51
Exchange	Unindexed	Unindexed Items	3	36838	2836.48

Top locations

Actions

Review sample

Close

Figure 18-5: Reviewing statistics for a preview search

## Expired Searches

If you look at the details of a search and see that *"the search has expired,"* you know that it is more than seven days since the last run of the search. As such, Microsoft 365 considers the result results to be unreliable because it doesn't include potentially important information added to the search locations in the last week. To proceed to use other search features such as exporting results, you must first rerun the search to make it current.

## Exporting Search Results

Eventually, you will be happy that a search finds the right information. At this point, you might want to export the files and messages found by the search to perform a more detailed examination of individual items, give the data to external investigators, or make it available to someone who does not have the necessary permission to run searches in the Microsoft Purview Compliance portal. The export function allows you to extract search results from the source locations, copy them to a secure holding point in Azure, and then copy the data from Azure to PSTs (for messages), individual files, or ZIP files. You can also export a report of the search results, meaning that instead of exporting the actual data, the search creates and exports reports listing the files and messages that it would export for a search.

A content search export can process items found in up to 100,000 mailboxes and will fail if a search covers more than this number. If you need to export information from larger location sets, you should use Premium eDiscovery.

## Selecting Export Options

To begin, make sure that the search is current as you will not be able to export results if they are more than seven days old. It is a good idea to refresh search results before beginning an export as this ensures that the exported data will be completely up to date. To start the export job, select **Export results** from the **Actions** menu. Figure 18-6 shows the screen used to gather information for the export process. While the search always exports SharePoint and OneDrive documents as individual files, you can export items extracted from Exchange mailboxes to:

- **A single PST for all Exchange content:** This option is convenient for investigators and is usually the best option when dealing with small amounts of information. Exchange Online does not include the *New-MailboxExportRequest* cmdlet available in the on-premises version to export content from a mailbox to a PST. However, you can mimic the functionality by creating a content search for all content in a single mailbox and exporting the results (the entire mailbox) to a PST.
- **To multiple PSTs** (one per mailbox): This choice allows the people who must process the results to split the workload across multiple individuals. If an export processes more than 10 GB of data, the job will automatically split into multiple PSTs, each of which is up to 10 GB (it is possible to change this limit with [a registry setting](#)).
- **To individual MSG files:** This choice exists because some third-party investigation tools do not support importing data from PST files. In addition, if you want to export messages encrypted with rights management, you must export them as individual files. The export job writes a copy of each message uncovered by the search to the target destination in the file system. The administrator can then move the copied files from there as needed. The export job organizes the individual MSG files into a folder structure. The folder for a mailbox is named after the user principal name of its owner. Underneath this root, the found items are divided into:
  - **Recoverable items:** Items extracted from folders in the Recoverable Items structure are here. For example, if an item was purged by the user but kept due to a hold placed on the mailbox, it is in the *Recoverable Items\Purges* folder in that user's mailbox.
  - **Top of Information Store:** The export job places items extracted from folders visible in the user's mailbox here. The export job creates a separate file system folder for each folder in the mailbox where the search found a matching item. Given the complex folder structure that exists within some mailboxes, the 260-character maximum path to a Windows folder may be reached. When this happens, the export job truncates the folder names to stay within the limit.
- **To a compressed (zipped) folder:** This option includes both Exchange and SharePoint content. Exchange items are in separate MSG files (with attachments) while items found in SharePoint and OneDrive are exported as individual files. Exporting to a compressed folder avoids the issue that sometimes arises when the file path to a SharePoint item exceeds 260 characters. Like PST files, if an export exceeds 10 GB, the search splits the export into multiple ZIP files (you can use the same registry setting mentioned above to control PST sizes to change the maximum size for a ZIP). Note that files in a ZIP folder only have the modified date for files, not the created dates. However, the created date is stored in the XML manifest for the export job.

**Export results**

When you start this export, we'll begin getting these search results ready for download. This may take a while depending on the size of your search results. [Learn more](#)

**Population**  
Searchable Files: Confidential Documents

**Output options**

All items, excluding ones that have unrecognized format, are encrypted, or weren't indexed for other reasons

All items, including ones that have unrecognized format, are encrypted, or weren't indexed for other reasons

Only items that have an unrecognized format, are encrypted, or weren't indexed for other reasons

**Export Exchange content as**

One PST file for each mailbox

One PST file containing all messages

One PST file containing all messages in a single folder

Individual messages

Enable de-duplication for Exchange content

Include versions for SharePoint files

Export files in a compressed (zipped) folder. Includes only individual messages and SharePoint documents.

**Estimation**

After starting the export, a new export object with name "Confidential Documents\_Export" will be created in the Export table. To see status and download results, select the "Export" menu option.

**Export** **Cancel**

Figure 18-6: Defining settings to export items found by a content search

If the search includes SharePoint or OneDrive for Business sites, you have the choice to include versions for documents found by the search. Versions are only available for documents if versioning is enabled for document libraries. Versions come from the preservation hold library of the sites where the search finds documents. You'll recognize these as versions because their name is composed of the original file name plus a date-based identifier created by SharePoint to make the file name unique.

Due to the way that Exchange Online delivers individual copies of messages to mailboxes, a search will likely find multiple copies of a message. To export just a single copy, enable the de-duplication option for Exchange content. Each message has a unique identifier, so when de-duplication is enabled, the export processes only the first copy of a message returned by the search. If it's important to recover a copy from specific mailboxes (perhaps for legal reasons), you can always run another export to recover copies for each mailbox where a message exists.

System data found by a search in a user mailbox, such as the information used by Viva Insights, doesn't appear in a preview, but the search exports this data to a folder called *Other Office 365 Data*. This data also includes Microsoft Forms owned by the user. Search exports metadata and Forms Q&A in JSON format. These items don't appear in search previews.

## Decrypting Search Results

The export process decrypts email messages protected with sensitivity labels including any attachments, but only if you export messages as individual MSG files. Protected items exported to a PST remain encrypted. Search marks items decrypted by the export process with “Decoded” in the Decode Status column of the Results.csv file in the export destination.

The Standard eDiscovery and content search export process does not decrypt protected documents exported from SharePoint Online and OneDrive for Business sites. If necessary, you can decrypt these files by running PowerShell cmdlets from an account that holds rights management super-user privilege. See the section on this topic in the Information Protection chapter. The Premium eDiscovery export process can decrypt protected documents, so if you have the necessary licenses, you can use this facility to export decrypted copies of protected files.

## Handling Partially or Unindexed Items

Tenants are likely to have a certain percentage of unindexed or partially indexed items in Exchange Online, SharePoint Online, or OneDrive for Business sites. More will be in Exchange than the other workloads simply because users generate more messages than documents. Partially indexed means that the index contains an item’s metadata (like subject, author, and creation date) but some or all of the item’s content might not be. Content searches can find partially indexed items using metadata.

The reasons why these items exist are varied. Some attachments or documents might be in an unsupported format; some items are too large (greater than 150 MB); some messages have more attachments than the supported maximum (250); and some formats have specific limits, like Excel’s 4 MB limit. The limits for content searches [are available online](#) and the statistics reported for a content search tell you how many unindexed items were found. To help organizations understand the number of these items in their tenant, Microsoft wrote a PowerShell script to analyze how many partially indexed Exchange Online items exist and report the reasons why indexing failed. You can read [the article online](#) and fetch a modified version of [the script from GitHub](#).

Partially indexed items can be of interest to investigators, and you should include these items in exports if you want to be sure that an investigation can consider every possible issue. After all, a human might make sense of a file included in an export where a search cannot. Three options are available:

- **All items, excluding unindexed items.** This is the default and means that the search only exports items meeting the search criteria. Unindexed (partially indexed) items are excluded.
- **All items, including unindexed items:** The export includes unindexed items, but only if the search also finds items matching the search criteria in a site.
- **All unindexed items:** The export includes all unindexed items from all sites in the search, even if the search does not find matching items.

The usual process is to exclude unindexed items from exports and then decide whether a deeper examination is necessary incorporating these items.

## Downloading Exported Results

After selecting the options for the export job, click **Export**. A background process starts to extract the information from source locations and copy it to the holding area in Azure. To follow the progress of the export job, click the **Export** tab in the menu bar to see a list of all the export jobs processed for searches. Select the export job for your content search (it has the same name as the content search with a suffix of “\_Export”), to display the status for the export job. Any search which uncovers tens of thousands of items spanning gigabytes of data will need some time to export the matching data. In this case (Figure 18-7), the number of search results is small, and we can see that the export is complete.

Before going ahead to download the exported results from Azure, we must copy the export key. The export key is a [shared access signature](#) to grant access to the secure storage area holding the export results in Azure and is of the form:

```
?sv=2014-02-14&sr=c&si=eDiscoveryBlobPolicy9%7C0&sig=PpqtiOKpBMzZPtA1ksuY8iciP6jsYYI2VHePjDXY45w%3D
```

The export key becomes the credentials to authorize the download of the exported data using the Microsoft 365 eDiscovery Export Tool. Essentially, the report key is a token that Azure Data Services recognizes when offered by a process that wishes to access some data belonging to another entity. You cannot use the key to access data held in Exchange Online or SharePoint Online with a browser or another client. After copying the export key to the clipboard, you can paste it into Notepad or another editor to make sure that it is always at hand. However, if you are ready to go ahead and download the results, click **Download Results**. It doesn't matter if the search has not finished preparing data for export yet because the export tool checks with the search when it downloads information and will pause to let the export complete.

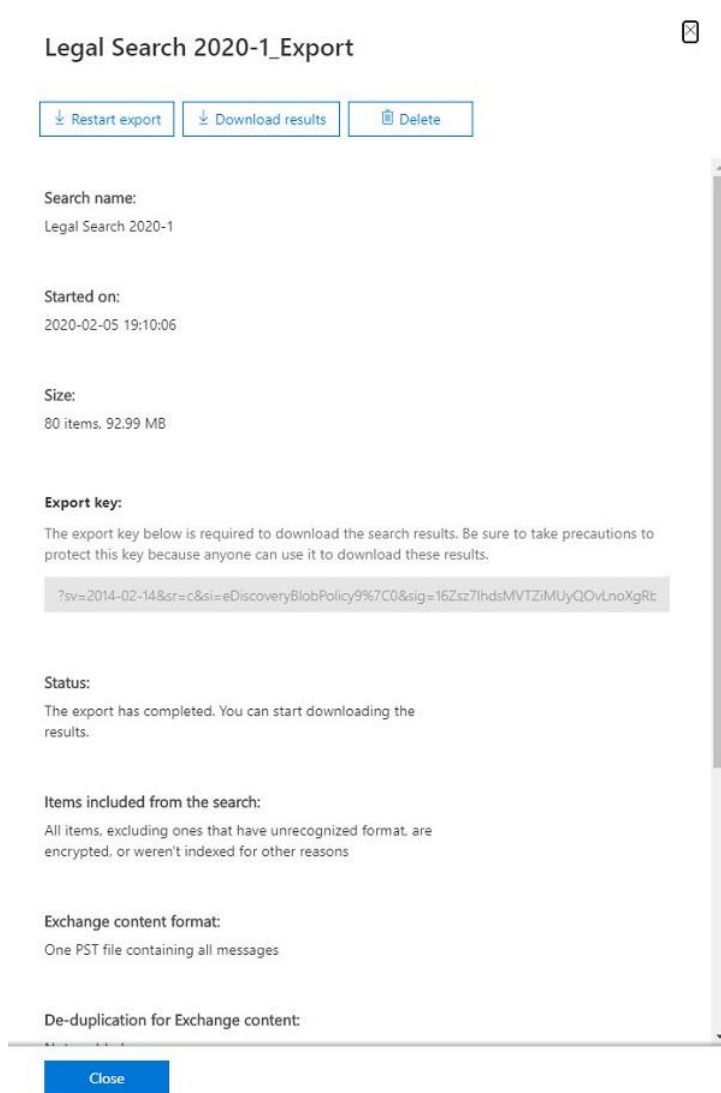


Figure 18-7: Preparing to download the exported results from a search

You must use the Microsoft Edge browser to download and install the export tool. You can use any other browser supported by Microsoft 365 to perform all other search functions. The first thing to do is configure Edge to enable *ClickOnce* support (the ability to [install and run an application created with the ClickOnce technology](#)). To do this, open a tab in edge and go to `edge://flags/#edge-click-once` and make sure that the value in the drop-down list is Enabled. If you don't do this, the export tool won't be able to run.



**After Edge downloads the export tool, it installs the tool on the workstation if necessary and then starts the tool. When export results are available in Azure, you can proceed** with the export by entering (paste) the export key and the target location for the tool to copy the exported data (Figure 18-8). The export tool then authenticates its access with Azure and downloads the exported data to the selected destination, including metadata for SharePoint and OneDrive documents. If you don't choose to use the advanced option to name the PST to hold exported Exchange content, Purview exports the results to a file called Exchange.PST. If necessary, you can use **Restart Export** to recommence the export process. This action removes any earlier search results stored in Azure and then recopies the search results from the search locations to the holding location in Azure.

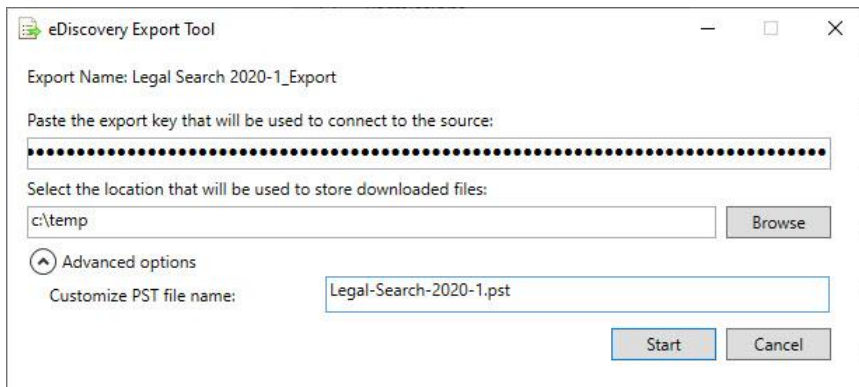


Figure 18-8: The eDiscovery Export tool begins to export information from Azure

The export job creates several sub-folders in the destination to hold the export data. The folder named after the date and time of the export includes an Electronic Discovery Reference Model manifest listing all the exported items. The Exchange folder holds the PST files used to export the data. Figure 18-9 shows the content exported from a SharePoint Online document library. In this case, the export job copied the files to a ZIP file. Once exported to disk, it is easy to give the information found by a search to an external company, such as specialized legal investigators. The external company can then apply whatever tools they choose to analyze the search results.

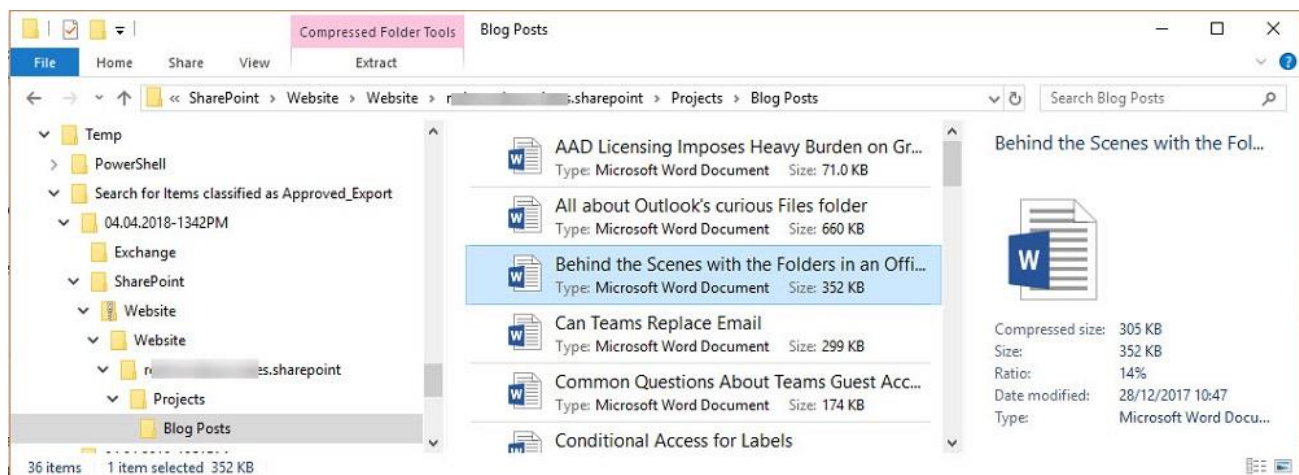


Figure 18-9: SharePoint documents exported from a search

You can export up to 2 TB of data for a single search from up to 100,000 mailboxes (you can [use PowerShell to export results for a larger number of mailboxes](#)). If a search finds more than this amount of data, you will have to split the search to get under the 2 TB limit. One way of doing this is to create several searches, each of which uses the same keyword query but different date ranges. Reflecting that it is a multitenant environment, Microsoft 365 limits a tenant to exporting 2 TB of data in a single day. A tenant can run up to ten export jobs concurrently, but a single user can only run three of those exports.

## Protected Documents and Searches

If support for sensitivity labels is enabled for SharePoint Online, the content of documents protected with rights management (Microsoft Information Protection) are indexed and can be found by a content search. If not, a content search cannot find protected documents based on their content and must rely instead on metadata such as the document title, comments, or tags.

Protected documents found by a search are exported like any other item, but they continue to be encrypted and cannot be examined by an investigator unless an arrangement is made to decrypt the files.

The Information Protection chapter contains some advice about how accounts assigned with super-user permission for information protection can run PowerShell code to find and decrypt protected files exported for a search. This is especially important for GDPR data subject requests (and possibly for data erasure requests) because someone must check the content of documents to understand whether they relate to a data subject.

## Export Search Report

Being able to review search results gives investigators insight into the kind of information uncovered by a search. However, it would be unreasonable to try to preview every item found by a search, especially when large numbers are involved, in which case an investigator might want to make a broader check of the data uncovered by a search before going ahead to export the data. To export the reports for a search, select the search and then **Export results** from the **Actions** menu. This downloads the same reports (the XML manifest, summary log, and results log) holding details of the information included in the full set of downloaded search results. When you choose to export the report, an export job generates the reports and stores them in Azure. You receive an export key to access the data and can use the same download tool to export the reports to the nominated destination. After checking that the correct results are available, you can then continue to perform a full export.

**Export for third-party review:** The eDiscovery industry spans many companies that specialize in the analysis of information recovered from IT systems like Microsoft 365. Microsoft has a program to support the export of information recovered by searches in a form that third-party applications can interpret and analyze. The export data exists in a location in Azure Data Services controlled by the third-party. Contact Microsoft to get an up-to-date list of certified partners if you are interested in this capability.

## Reopening a Search

You can go back and open a search at any time. Select **Content search** in the Microsoft Purview Compliance portal to list the available searches and select the one you want to work with, and then **Edit search** from the **Actions** menu. If the last search is less than seven days old, you can preview results with the **Review sample** option to see how accurate the search criteria are before making changes. If the search is older, you need to rerun the search before you can preview the sample results.

## Compliance Boundaries

When you have the permission to create and run content searches, you can look for anything across the set of supported locations in a tenant, including the ability to pry into sensitive mailboxes or hunt for interesting documents stored in document libraries that you would not otherwise be able to see. Compliance Security Filters allow tenants to impose control over the data visible to investigators by establishing boundaries for searches. Large companies often divide administrative and other responsibilities along geographic or divisional lines. When the time comes to conduct content searches, they might not want those who run the search to be able to include search locations outside their business, country, or region. This makes a lot of

sense: someone running a content search to respond to a discovery action in France does not necessarily need to look at German mailboxes. Apart from respecting user privacy, Compliance Security Filters also mean that content searches return a smaller amount of data for investigators to review.

A Compliance Security Filter creates a restrictive view of mailboxes or SharePoint and OneDrive sites within a tenant. When users that conduct searches come within the scope of a filter, they cannot see any data returned by searches except that given by the restrictive view. Therefore, we can set things up that U.S.-based eDiscovery administrators only can see results from mailboxes located in U.S. data centers or that only certain eDiscovery administrators can search particularly sensitive SharePoint sites.

Before starting to plan your filter strategy, you should read the Microsoft [support article about Compliance Security Filters](#) and [this blog post](#). The key point to note is that you can only create and manage filters through PowerShell. To guide you in creating a filter, answer the following questions:

- **Who will the filter apply to?** You can specify individual users or use the name of a Compliance role group, including a role group created specifically for this purpose. You cannot use a distribution list, Microsoft 365 Group, or security group to define a set of users.
- **What can the users do?** You can restrict users to individual compliance actions (Export, Preview, Purge, Search) or "All". You cannot specify two or three actions. In most cases, you will want to use Search or All.
- **What can the users see?** You can combine mailbox and SharePoint locations into a single filter that works across multiple workloads. In both cases, you can have filters that look for specific objects (mailboxes or sites) or content (based on KQL queries).

With these points in mind, here is a simple filter that applies to a single named user to allow them to perform all compliance actions while restricting them to searching mailboxes with a specific value in their *CustomAttribute6* property. To create the Compliance Security Filter, connect to the compliance endpoint and run this command:

```
[PS] C:\> New-ComplianceSecurityFilter -FilterName VikFraudSearch -Users "Marc Vigneau" -Filters "Mailbox_CustomAttribute6 -eq 'POI'" -Action All
```

To add a mailbox to the set that searches performed by the users identified in the filter can find, we update the *CustomAttribute6* property as follows:

```
[PS] C:\> Set-Mailbox -Identity Kim.Akers -CustomAttribute6 POI
```

After updating the mailboxes, you can test the filter to check the set returned with *Get-Recipient*:

```
[PS] C:\> Get-Recipient -RecipientType UserMailbox -RecipientPreviewFilter {CustomAttribute6 -eq 'POI'}
```

Now that we know that the filter is valid and returns a set of mailboxes, we can test it with a search. Log into the Microsoft Purview Compliance portal using the account of one of the users restricted by the filter. The user must have the necessary permission to run searches. Create a new search for all mailboxes with a query that you know will find some information. Launch the search and wait for it to complete. The results should reflect the filter in terms of the number of locations scanned and the amount of information found.

Now log into the Microsoft Purview Compliance portal as another administrator who is not restricted by the filter and run the same search again. This time the results should be very different. Figure 18-10 shows an example of a content search run by restricted (left) and unrestricted (right) eDiscovery administrators. The restricted search only scans 2 mailboxes while the unrestricted search looks through 468. The number of found items is also different, as you'd expect.

## Secret Investigation

Summary Search statistics

### Description

Fraud

### Last run on

2021-05-22T15:39:32.813Z

### Searched by

Marc Vilas

### Search conditions

Fraud

### Status

The search is completed

298 items(s) (4.21 GB)

28 unindexed items, 76.51 MB

2 mailbox(es)

All sites

Actions

Review sample

## Secret Investigation

Summary Search statistics

### Description

Fraud

### Last run on

2021-05-22T15:37:48.12Z

### Searched by

Tony Redmond

### Search conditions

Fraud

### Status

The search is completed

2,124 items(s) (4.71 GB)

4,256 unindexed items, 2.67 GB

468 mailbox(es)

All sites

Actions

Review sample

Figure 18-10: The result of applying a search filter

Microsoft suggests many other examples of Compliance security filters [in its documentation](#), including how to filter mailboxes based on the ISO 3166-1 code (for example, 124 is the three-digit code for Canada while 372 is the code for Ireland). One example of obvious interest is a filter that restricts access to confidential or sensitive mailboxes. The code to create such a filter first finds the distinguished name of a distribution list called the "Senior Leadership Team". The members of the list are the mailboxes that we want to restrict. We include the distinguished name in the filter to stop anyone who runs a search against these mailboxes from being able to preview items found by the search.

```
[PS] C:\> $DG = (Get-DistributionGroup -Identity SLTDL).DistinguishedName
New-ComplianceSecurityFilter -FilterName NoSLTPreview -Users All -Filters "Mailbox_MemberOfGroup -ne '$($DG)'" -Action Preview
```

The *Get-ComplianceSecurityFilter* cmdlet reveals the details of the filter:

```
[PS] C:\> Get-ComplianceSecurityFilter -FilterName NoSLTPreview
FilterName : NOSLTPREVIEW
Description :
Action     : Preview
Users      : {all}
Filters    : {Mailbox_MemberOfGroup -ne 'CN=DL Senior Leadership
Team,OU=tenant.onmicrosoft.com,OU=Microsoft Exchange Hosted
Organizations,DC=EURPR04A002,DC=prod,DC=outlook,DC=com'}
```

This kind of filter effectively stops casual browsing of preview samples from sensitive mailboxes by people who should know better. It does not stop searches from finding items, nor does it stop eDiscovery administrators from being able to export items found by the searches. The downside of using a filter is that it applies to all content searches, including those executed by people that might legitimately have reason to preview content found in the sensitive mailboxes.

Users cannot preview items found by a content search in the designated mailboxes, but they can preview documents and other items (like PDFs) found in sites covered by searches. To ensure full confidentiality for the Senior Leadership Team, you need to define a site filter to protect these locations. This filter restricts access to documents found in the document library. In this context, "All" means members of the eDiscovery manager role group rather than all users.

```
[PS] C:\> New-ComplianceSecurityFilter -FilterName NoSLTPreviewDocs -Users "All" -Filters "Site_Site -ne 'https://tenantname.sharepoint.com/sites/SLTGroup'" -Action Preview
```

You can create a filter with clauses for multiple workloads (a filter list). Here is a filter with two clauses: one for U.S.-based mailboxes and the other for a specific SharePoint site:

```
[PS] C:\> New-ComplianceSecurityFilter -FilterName CountryFilter -Users annb@contoso.com -Filters "Mailbox_CountryCode -eq '840'", "Site_Site -eq 'https://tenant.sharepoint.com/sites/Confdocs'/" -Action All
```

## The Effect of Filters on Content Searches

In its documentation, Microsoft explains that:

*"The permissions filter is added to the search query when a Content Search is run. The permissions filter is essentially joined to the search query by the **AND** Boolean operator."*

and:

*"In a Content Search query, multiple permissions filters are combined by **OR** Boolean operators. So results will be returned if any of the filters are true. In a Content Search, all filters (combined by **OR** operators) are then combined with the search query by the **AND** operator."*

If you use multiple filters, the filters are joined with the query (AND) and then combined (OR). This allows the search to find all the content per the query and then apply each filter to arrive at a combined set of results. The exact results that any action (search, preview, or export) produces depend on the actions specified in each filter. The ability to combine filters with content searches creates a great deal of flexibility in what you can do to control searches, even if it might take some time and effort to arrive at the filters needed to generate the desired result.

Note that you cannot exclude specific public folders using a search filter. The filters only work for user and group mailboxes and SharePoint and OneDrive sites.

# Auditing of Search Activities

Administrators and eDiscovery managers access and view user content through content searches. Being able to access messages, documents, and other potentially sensitive user content is necessary as otherwise, you'd never be able to perform a compliance search and export its results. Because searches use privileged access to user data, the Microsoft Purview Compliance portal captures audit records when investigators create and run content searches and eDiscovery cases, and for search actions against results like preview, export, and purge. Having audit information to hand enables organizations to ensure that they meet their requirements to protect confidential user information under regulations such as GDPR.

You can review the audit records for search activities through the audit log search in the Microsoft Purview Compliance portal. This is a reasonable approach when you know the time and date when an activity occurred, and you only need to review a small number of activities.

If you want to process a lot of audit records, perhaps to verify that administrators are not abusing their access to user information, it is better to use PowerShell to interrogate the audit log and extract the records you might be interested in analyzing. This example (see Chapter 21 for a more extensive discussion of using audit

log data) extracts records captured when users view, export, or preview the results of content searches, and outputs a CSV file. The CSV file can be opened and reviewed in Excel or imported into Power BI.

```
[PS] C:\> $StartDate = (Get-Date).AddDays(-7) ; $EndDate = Get-Date
$Records = (Search-UnifiedAuditLog -StartDate $StartDate -EndDate $EndDate -Operations
"SearchExportDownloaded", "SearchViewed", "ViewedSearchPreviewed" -ResultSize 1000)
If (!$Records) {
    Write-Host "No audit records for content search activities found." }
Else {
    Write-Host "Processing" $Records.Count "audit records..."
    $Report = [System.Collections.Generic.List[Object]]::new()
    ForEach ($Rec in $Records) {
        $AuditData = ConvertFrom-Json $Rec.Auditdata
        $ReportLine = [PSCustomObject]@{
            TimeStamp = Get-Date ($Rec.CreationDate) -format g
            User       = $AuditData.UserId
            Action     = $AuditData.Operation
            Exchange   = $AuditData.ExchangeLocations
            SharePoint = $AuditData.SharePointLocations
            Query      = $AuditData.Query }
        $Report.Add($ReportLine)
    }
}
$Report | Export-Csv c:\temp\SearchAuditRecords.csv -NoTypeInformation
```

Microsoft Purview captures audit records for many other content search and eDiscovery activities, and it is easy to modify the PowerShell code shown above to include multiple other activities and export those records for later analysis.

## Data Subject Requests

According to [Microsoft's Data Subject Guide](#), 90% of an organization's data stored in Microsoft 365 is in Word documents, Excel spreadsheets, PowerPoint presentations, OneNote files, and Outlook messages. The data are indexed and searchable. Some data cannot be found by content searches, so the information you can find through a DSR case is not necessarily all the personal information belonging to a data subject existing in a tenant. For instance, videos featuring the user stored in Stream cannot be found through eDiscovery, so extra effort is needed to review and retrieve this information if necessary.

Article 15 of the GDPR grants a data subject (a person) the right to have a data controller (the organization owning a tenant) provide them with a copy of their data. To help tenants respond to these requests, the Microsoft Purview Compliance portal supports the creation and management of special eDiscovery cases, called [Data Subject Request \(DSR\) cases](#). These cases are created under User data search in the eDiscovery section. This capability is free as the cases use a customized version of regular content searches.

A DSR includes a content search for the personal data belonging to a defined individual. An eDiscovery case is created including a search of the person's mailbox (if it exists), all public folders, Teams chats and channel conversations, and SharePoint Online and OneDrive for Business sites for any information relating to the data subject. If the tenant's Yammer network is configured in native Microsoft 365 mode, its messages can be included in the search.

Things to remember about DSRs include:

- The search locations are not dynamic. If new mailboxes or sites are added to the tenant after you create the DSR, the search will not look in those locations.
- A DSR can only search cloud locations. If you run a hybrid organization, you must run separate searches to find relevant information in the on-premises locations. Exchange 2010 and later and SharePoint 2010 and later include the ability to run eDiscovery searches that can be tailored as DSRs.
- GDPR says that DSRs must be responded to within a month of the request. It is therefore important to keep an eye on the progress of DSRs and highlight instances when DSRs are delayed.

- DSRs do not address the need to remove information about a data subject (the right to be forgotten defined in article 17). However, the reports generated for a DSR will tell you where data matches are found and allow you to then check the individual items to decide whether items should be removed. Remember, not all data found for a data subject needs to be removed as it is legally permissible to keep data under certain circumstances, such as to comply with a legal obligation.

When you are happy that the searches return the information necessary to satisfy the data subject request, you can export the information as normal and prepare it to give to the data subject. A further check is always needed to exclude any information in the exported data that is unrelated to the data subject. For instance, if the name of the data subject is common (like John Smith), some of the matches returned by the searches probably do not refer to the data subject. It is also possible that some of the information is commercially sensitive and needs review by the business and potentially by legal advisors before it can be released.

## Priva Subject Rights Requests

The Priva Subject Rights Requests solution is part of the [Microsoft Priva privacy management suite](#). It is a more developed and comprehensive version of DSR processing which Microsoft licenses on a per-request basis. Organizations can run a free trial of Priva Subject Rights Requests for up to 90 days or 50 requests (whichever limit is reached first). Priva Subject Rights Requests include:

- Customized content searches to find information relating to a data subject.
- Retrieval of information to an Azure blob storage container for review (annotation, tagging, redaction, etc.). During retrieval, Priva attempts to identify high-priority items for reviewers to consider.
- Generation of reports.
- Integration with Teams for collaboration (by default, each request creates a team for those working on the request to use with the person who creates the request being the team owner).
- Workflow based on Power Automate (customizable by the organization).
- Processes requests in line with compliance regulations including GDPR, the California Consumer Privacy Act (CCPA), UK Data Protection Act (UDPA), U.S. State Breach Notification Laws (USBNL), and U.S. Patriot Act (USPA).

# Microsoft Purview eDiscovery

Content searches are useful in a variety of situations. Sometimes it is to find information that a user has “lost,” and sometimes it is to find information for more serious reasons, such as looking for evidence of corporate or personal malfeasance needed to prove a point for internal or external purposes. If the information exists in an indexed location, content searches will find it.

Good as content searches are at finding information, searching is only part of the eDiscovery lifecycle. And while the results of an individual search might be critical to proving or disproving a point, many eDiscovery projects involve teams of investigators using multiple searches looking for different sets of information. eDiscovery cases deliver the structure needed by investigators to organize their work, including searches, holds, and exports.


Microsoft Purview eDiscovery functionality divides into standard and premium. Standard eDiscovery is available to tenants with E3 licenses; you need E5 licenses or the Microsoft 365 E5 Compliance license to use Premium eDiscovery.

## Standard eDiscovery

To access the cases created by standard eDiscovery, go to the **eDiscovery** section in the Microsoft Purview Compliance portal and select Standard. This reveals the set of eDiscovery cases available to the logged-in user

(Figure 18-11). Some of the cases are "Active," meaning that investigators might still be working on these cases. Closed cases are still available, and investigators can reopen them if necessary.

## eDiscovery (Standard)

 Remove from navigation

After creating an eDiscovery case and choosing who has access to it, use the case to search for email, documents, Skype for Business conversations, Teams data, and other content in your organization. You can then preserve the content and export the search results for further analysis. [Learn more](#)

+ Create a case   ↓ Download list   ↻ Refresh   14 items      Group   Filter   Customize columns

Name	Status	Created date ↓	Last modified	Last modified by
Project Clifden investigation	Active	3 Jun 2021 18:13	3 Jun 2021 18:13	Tony Redmond
Project Cleggan Investigation	Active	3 Jun 2021 18:00	3 Jun 2021 18:00	Tony Redmond
Bing Stuff	Active	8 Dec 2020 20:25	8 Dec 2020 20:25	Tony Redmond
Improper management transactions LD-20...	Active	14 Feb 2020 11:42	14 Feb 2020 11:42	Tony Redmond
Bad Executives	Active	3 Feb 2020 15:12	3 Feb 2020 15:12	Tony Redmond
Test for EAC Migration	Active	8 Jan 2020 16:32	8 Jan 2020 16:32	Global Tenant Administrator
Inactive Tenant	Closed	4 Jul 2018 11:44	29 Apr 2019 14:27	Tony Redmond
Improper management transactions (LD-2...	Closed	4 Apr 2018 19:03	29 Apr 2019 14:26	Tony Redmond
Personal Harassment Investigation LD-201...	Closed	4 Apr 2018 18:31	29 Apr 2019 14:27	Tony Redmond
Contoso Office 365 Investigation	Closed	28 Apr 2017 20:14	19 Apr 2018 00:29	Tony Redmond
Project Alpha Foxtrot	Closed	28 Apr 2017 20:14	29 Apr 2019 14:26	Tony Redmond
Stock Trading Case Reference 2017-001	Closed	28 Apr 2017 20:14	29 Apr 2019 14:26	Tony Redmond
Board Minutes	Closed	28 Apr 2017 20:14	19 Apr 2018 00:29	Tony Redmond
Inactive Exchange Mailboxes	Active	28 Apr 2017 20:14	21 Jun 2019 17:03	Tony Redmond

Figure 18-11: A set of Standard eDiscovery cases

## eDiscovery Case Components

If you open an eDiscovery case, you see that its major parts are:

- **In-place holds to ensure that users cannot remove information.** Holds ensure that Microsoft 365 workloads retain information needed for an eDiscovery case in their original location. A hold covers a set of locations (Exchange mailboxes, SharePoint and OneDrive for Business sites, and Exchange public folders) and criteria to specify the items in those locations that come under the scope of the hold. A hold can be as broad as to include all items in all available locations within a tenant or as specific as to hold just one or two items found with a highly-specific phrase in a selected mailbox.

All items that fall under the scope of the hold stay in the specified sites and mailboxes (including those used by Microsoft 365 Groups). Users can try to remove items that come under the scope of the hold, but if this happens, the workloads keep a copy of the item until the hold expires. In-place holds on Exchange mailboxes include both the primary and archive mailboxes. Standard content searches do not include the ability to place holds on sources. Note that a user account must have at least an Exchange E3 license before you can place it on hold. The license must stay assigned to the account for the duration of the case. An eDiscovery case can include multiple holds, each of which has its own target set of locations and hold criteria.

- **Searches are used to gather** information of interest to the case. The searches in an eDiscovery case behave very much like content searches. An eDiscovery case might only have one search, but it might equally deploy multiple searches, each of which focuses on different material. Each search covers a different aspect of the case and might look for content based on different queries, various locations,



or even content created in different languages. The intention is that multiple searches allow investigators great flexibility in how they approach looking for information because they can create a series of searches to interrogate information with different criteria. The searches created for eDiscovery cases do not appear along with other content searches because they belong to an eDiscovery case. Both types of searches use the same technology to find information using the content indexes populated from Exchange Online and SharePoint Online data. The names assigned to searches belonging to eDiscovery cases must be unique and not clash with standalone content searches.

One difference between regular content searches and those executed within an eDiscovery case is that in a case, you can specify the special *"Locations on hold"* target for a search, meaning that you want to search the held content in all the locations included in a case. For example, assume two holds exist for the case. The first covers two mailboxes and two sites and the second span a further twelve mailboxes. In this instance, *"Locations on hold"* means searching the held content in the fourteen mailboxes (including their archive mailboxes if these exist) and two sites. If the scope of any of the holds belonging to a case changes (for example, the addition of new locations), the scope of *"Locations on hold"* changes to match the holds when you refresh the search results. Searching against on-hold content is a current workflow scenario in eDiscovery situations where investigators place locations such as mailboxes or sites on hold before starting to search those locations.

- **Exports of information gathered by the searches.** An eDiscovery case might use several searches to find different sets of information. You can export results from a single search using the same steps as for exporting results for a normal content search, or you can combine the results of multiple searches belonging to a case into a single set of data to give to external investigators.

When you combine results from multiple searches, the search combines the queries using the OR operator to form an overall query and runs the query against the locations to generate the search results. In other words, the search results created for export come from a single search instead of manually combining the sets generated for the individual content searches. The results of the searches are deduplicated so that the export includes only a single copy of an item found in a location. Because of the way that a combined search brings multiple search queries into one, the overall keyword limit for a query (500) might be met. If so, you will see an error and the search will end. To achieve the desired result, you must combine fewer searches or simplify the queries.

## Case Members

Each case has one or more members, each of whom must be a member of the eDiscovery Manager role group. The members are the only people who can access the results of the searches associated with a case. The eDiscovery Manager role group divides into two sub-groups:

- **eDiscovery Managers** deal with specific cases. They only have access to the content belonging to those cases. When a user accesses the Microsoft Purview Compliance portal, they can only see the cases where they are the eDiscovery Manager. They cannot see the cases belonging to other eDiscovery Managers.
- **eDiscovery Administrators** have oversight over all eDiscovery cases and can view and edit any case within the tenant regardless of who is the eDiscovery Manager for the case. To access a case, an eDiscovery administrator adds themselves as a member of the case.

Logically, a very small set of users within a tenant should be eDiscovery Administrators.

## Case Management

When you open the eDiscovery section, the Microsoft Purview Compliance portal shows you all the cases that your account manages or has access to. You can then select an existing case or **Create a case** to initiate a new case. When you open an existing case, you can use the **Settings** tab to access the properties of that case, including the members working on the case and its status. Figure 18-12 shows that the selected case has three members plus anyone who is a member of the Senior Investigators role group.

### Access & permissions

Update successful

#### Members (3)

+ Add  Remove

Name	Email
Ben Owens (Business Director)	Ben.Owens@office365itpros.com
Kim Akers	Kim.Akers@office365itpros.com
Tony Redmond	Tony.Redmond@redmondassociates.org

#### Role Groups (1)

+ Add  Remove

Name	Description
Senior Investigators	Senior Investigators

 Exit

Figure 18-12: Viewing members for a Standard eDiscovery case

You can use the same form to manage the status of the case by closing or removing the case. Due to the nature of corporate investigations, eDiscovery cases are often long projects that span several years. It is not unknown for cases to last five years or more.

## Creating a New eDiscovery Case

From the eDiscovery section of the portal, click **Create a case** to begin. You can then input the name of the new case and some details to describe the need for its creation. Figure 18-13 shows a typical situation. The name of the case clearly shows its intent while the description tells anyone reviewing the case what its purpose is and who authorized the investigation. Click **Save** to continue and be returned to the portal.



### New case

Enter a name and description

Give this case a friendly name so you can easily find it again later.

Case name \*

Project Clifden investigation

Case description

Investigation of potentially improper transactions relating to Project Clifden. Law department reference LD-2021-18764

Save

Cancel



Figure 18-13: Creating a new Standard eDiscovery case

Creating a new case creates the eDiscovery case container for searches, holds, and other activities. Nothing much exists in the case at this point, so we should select and open the case to begin adding these elements. The case screen then displays to allow us to access the different components.

## eDiscovery Case Holds

The next step is to create one or more holds to preserve the content needed for the case. You can use multiple holds for an eDiscovery case but in most cases, one hold that applies to Exchange mailboxes and SharePoint and OneDrive for Business sites is enough. Click **Holds** to access that section of the case and then **Create (+)** The steps to create a hold are:

- **Name the hold.** Assign a unique name for the hold along with an optional description. Ideally, the description should tell an eDiscovery manager the intention behind the hold and link the hold to the eDiscovery case. For instance, if the case name is "Investigation 2020-003," then you might call the holds in the case "Investigation 2020-003 #1," "Investigation 2020-003 #2," and so on, following whatever naming convention makes sense.
- **Define the locations** the hold will cover. The locations divide into:
  - Exchange email: User, shared, and group mailboxes. Group mailboxes include group conversations and calendars for Microsoft 365 Groups and the compliance records created for Teams and Yammer conversations. Compliance records for Teams personal chats and private channel conversations come from user mailboxes. If you select a distribution list, Exchange expands the membership and adds valid mailboxes to the hold.
  - SharePoint Online sites (including those used by Teams and Microsoft 365 Groups) and OneDrive for Business accounts.
  - Exchange public folders. You can include or exclude all public folders, but you cannot select specific public folders.
- **Define the search criteria** to find items to hold. The criteria are like those used for content searches and include keywords and conditions.

A hold does not have to include a query, and if it does not, it means that you want to apply a hold to every item in the selected locations. You do this when you want to preserve complete mailboxes or sites because you are unsure of the material that you want to hold. Figure 18-14 shows the summary for the creation of a new hold.

## New Hold

**Review your settings**

Name

**Name** Finance Trading Hold LD-2021 17-15

**Description** Hold on executive mailboxes

Edit name

Choose locations

**Exchange email** Jessica.Chen@office365itpros.com, Chris.Bishop@office365itpros.com

**SharePoint sites**

**Exchange public folders**

Edit locations

Query

**Query** "Trading Profits"

Edit query

Back Submit Cancel

Figure 18-14: Creating a hold for a Standard eDiscovery case

After creating the hold, the Microsoft Purview Compliance portal publishes its details to the workloads for the selected locations. It can take some time to synchronize across all workloads to make the hold effective everywhere, but it should certainly be in place within a few hours.

To view information about a hold, select it to reveal a details pane. Here you find the current hold statistics provided by the relevant workloads to show how many items come under the scope of the hold, the date and time of the last modification for the hold, and some statistics about the items which come within the scope of the hold. If a change is needed for the hold parameters, you can edit the hold settings.

**Applying Holds to Teams Private Channels:** Two steps are necessary to apply a hold to information belonging to a Teams private channel. First, you assign the hold to the SharePoint site for the private channel by specifying the URL of the site. Second, you add at least one mailbox of a member of the private channel to the hold. Compliance records for messages posted to private channels are copied to the mailboxes of all members, so messages for the private channel come within the scope of the hold when you add a member's mailbox. The hold might lapse if that member's account is removed from Microsoft 365, so it's best to add at least two member mailboxes.

## eDiscovery Case Searches

After applying a hold to make sure that users can no longer remove any content that we might need for our eDiscovery case, we move on to searching. The simplest eDiscovery case has just one search, but more complex cases might use multiple searches, each of which takes a different approach to look for material of interest. The idea is that an investigator can use different searches to home in on the information they need. Apart from the ability to use different queries to focus on different material in each search, the searches in a case might target different locations. One search might focus on a set of mailboxes and look for a specific item. A second search might look for some documents in a SharePoint Online site and a third might concentrate on a single user and retrieve a much wider range of content from their mailbox. Like other

searches, you can reiterate several times to refine the results retrieved by search criteria until you find the desired content.

Click **Searches** to go to the search page and then **New search** to create a new search. The same steps to create a regular content search occur when creating a search in a Standard eDiscovery case. These are:

- Name the search.
- Identify the target locations from Exchange mailboxes, SharePoint sites, and Exchange public folders. Select all locations or choose individual mailboxes or sites. If you want the search to apply only to locations on hold for the case, select **Locations on hold**.
- Enter the keywords and conditions for the search to find items.
- Review the search parameters and submit them for processing.

The Microsoft Purview Compliance portal goes ahead and launches a preview search. When the preview search finishes, you can review sample items to understand if the search finds the right items.

## eDiscovery Case Exports

After investigators run searches to find the information they need, they can then export the search results. This process works like exports for content searches, with the notable exception that you can select to export the results for multiple searches in the same operation. The same general approach is used.

- Access the search whose results you want to export.
- Click the **Actions** button and select **Export Results**.
- Set the export characteristics:
  - Decide what items the search finds to include in the export (all items, all items except those that are in an unrecognized format, encrypted, or indexed, or just the items in an unrecognized format, encrypted, or unindexed).
  - Decide what PST structure to use for mailbox items (one PST per mailbox, one PST for everything, one PST with all messages in a single folder) or to use individual MSG files.
  - Decide whether to deduplicate the search output.
  - Decide whether to include versions (if available) for SharePoint and OneDrive for Business documents.
- A background job named "Search Name\_Export" then exports the search data to a secure Azure location.
- When the search data is available, you download it to a workstation using a secure key generated by Microsoft 365 as credentials to access the data.
- Investigators use the downloaded PSTs, MSG files, and documents to review the information and assess its content.

## Closing eDiscovery Cases

Eventually, when the investigation for an eDiscovery case winds down, the case manager can close the case by selecting the **Close case** option. When you close a case, the underlying workloads release the holds for the case. It can take some time for workloads to process the hold release commands. After the workloads respond to confirm that they have removed the holds, the case status is set to closed, and the time and date of closure and the user who invoked the closure are captured in the case properties. Later, if the case records are not needed, you can remove the case after first releasing any holds belonging to the case. These holds are inactive, but they exist in case someone wants to reopen the case and reestablish the holds, so they must be removed before you can remove the case.

**Warning!** Depending on the age of the data involved, closing a case could result in background maintenance processes removing the items that are no longer on hold. Exchange Online keeps mailbox

items until the deleted items retention period expires while SharePoint Online and OneDrive for Business keep items for 30 days after the hold terminates to avoid any inadvertent data loss. Even so, closing a case is not something to do on a whim. If you make a mistake and need to reactivate a case, the case manager can reopen the case. However, reopening a case does not reestablish the holds that were previously in place and they will have to be recreated by going to the Holds section of the case, selecting each hold, and then taking the **Turn It On option** in the action pane. The gap in time between removing the original holds and reapplying new holds creates the potential that data will be removed from the sources during this period. The exact amount of data that might be lost is unpredictable because it depends on whether the background processing to remove data from sources has run and removed data.

## Premium eDiscovery

The cost of large-scale eDiscovery actions can be staggering. It is expensive enough to retrieve all the items necessary to satisfy a discovery order handed down by a judge. It can be extraordinarily expensive to individually process each piece of information delivered in response to a discovery order. The number of individual messages or documents can easily mount into the low millions and, when very large companies are involved, quickly grow into tens or hundreds of millions of items. The problem then becomes how to locate the proverbial needle in the haystack.

In the early days of eDiscovery, it was common to print off copies of emails and documents for lawyers to review. This process was tedious, paper-bound, and expensive. Lawyers are paid by the hour and lots of IT effort was necessary to generate the material. However, the technique worked reasonably well then because a small volume of email or electronic documents was involved, and the major focus was on paper documents and communications such as faxes and telexes. Today, the situation has changed because paper files are no longer the focus of business documentation and the volume of items stored and available to be discovered has grown massively. If a discovery order turns up 50,000 items, you don't want to incur the cost of having a professional check each item.

### Basics of Premium eDiscovery

Although the standard approach to eDiscovery (hold, search, and export) is suitable for investigations that generate tens of thousands of items, human examination of every item retrieved by a search is a time-consuming and expensive exercise. To address the problem, Premium eDiscovery applies algorithms to refine very large sets of search data retrieved to make it easier for investigators to find what they are looking for. Premium eDiscovery is currently capable of ingesting data sets spanning up to several million items. Microsoft expects that they will be able to lift the limit and at that point, Premium eDiscovery will be able to process as much data as any organization might need.

Consider a situation where a search uncovers a hundred thousand items. The choices are then to:

1. Refine the criteria for content searches to return the most precise set of discovered information and then review all the found items to decide which are useful and which are not. This is an acceptable tactic for relatively targeted searches where the desired material can be accurately described in terms of search criteria. For instance, looking for evidence that a specific phrase was discussed in an email sent between four known individuals.
2. Ingest the output of content searches, use analysis to parse the set of discovered files, and home in on the material that needs close examination, followed by a review of sample items by expert investigators to find items that are of relevance and those that are not. The items deemed to be of relevance are then fed through machine learning algorithms to construct filters (think of them as very complex search queries) that are used to interrogate the complete set of files to locate the desired information.

The first choice described above is roughly what's possible using the Standard eDiscovery features. The second is what becomes feasible with Premium eDiscovery.

The current implementation of Premium eDiscovery uses a workflow described in the [Electronic Discovery Reference Model](#) (EDRM). The stages in the workflow are:

- **Identification:** Knowing whose data should be searched is essential to an investigation. People who might own information needed by the investigation are called **data custodians** or **people of interest**. These are Microsoft 365 accounts that you might want to add to a Premium eDiscovery case and search for data (Figure 18-15). The mailbox and OneDrive for Business site (custodial locations) for each custodian can be included in the case along with other locations they access, such as Teams, shared mailboxes, and SharePoint sites.
- **Preservation:** After you add custodians to a case, you can create **in-place holds** to preserve selected data needed for the case. You can also add holds that cover people who aren't custodians. Premium eDiscovery includes the ability to **send notifications** to custodians by email to tell them that their accounts are under hold and then track responses from the custodians.
- **Collection:** After identifying relevant data sources, you collect information from those sources by using a **special form of content search**. Information remains in the data sources and users can continue to work with it as usual.
- **Processing:** Once searches have located relevant data for the case, you process the data. Like when you export results from a content search, the eDiscovery case **copies data found by searches** to an Azure storage location called a review set. The review set is a static view of the case data that can be analyzed and reviewed for relevance.
- **Review:** Investigators can now look at specific items in the review set to decide if they have the right information or need to query the data to reduce the set to what is most relevant to the case. Investigators can annotate and tag items during this phase. If necessary, case managers can [load data from outside Microsoft 365](#) for inclusion in the review.
- **Analysis:** Premium eDiscovery includes a set of tools to help reduce the data from the review set down to the most relevant information. Reducing the amount of data for investigators (and lawyers) to consider limits the costs of eDiscovery.
- **Production and Presentation:** When the final set of most relevant items is identified, you can export them from the review set. Export can be in the native format of documents or an EDRM-specified format suitable for ingestion into third-party review applications used by external experts.

Premium eDiscovery can process exports involving up to 5 million documents or 500 GB (whichever is smaller). To deal with such large amounts of data, the export splits processing across multiple ZIP files, which investigators can later combine in a single location to reform the complete set. One thing to be aware of is that Premium eDiscovery processing is slower than many people expect. It takes time to perform the complex background processing performed by many actions. Patience is certainly an advantage when dealing with these cases.

Premium eDiscovery is a specialist activity and apart from assigning the necessary permissions to accounts used by investigators, it is unlikely that most tenant administrators will need to become involved in these cases. Instead, it will probably be specialized compliance managers that create and run the cases.

Advanced eDiscovery &gt; LD-12910 Fraud Investigation

Overview Data sources Collections Review sets Communications Hold Processing Exports Jobs SettingsStart your case by setting up people as custodians to quickly identify and preserve data sources with which they are associated. [Learn more](#)

Add data source ▼ 🔄 Refresh 8 items









Name	Source type	Status	Hold	Indexing job status
 Chris Bishop	Custodian	Active	true	Submitted
 James Gangley (Inact...	Custodian	Active	true	Successful
 James Abrahams	Custodian	Active	true	Successful
 James Joyce	Custodian	Active	true	Successful
 Boris Johnson	Custodian	Active	true	Successful
 Brian Weakliam (Ope...	Custodian	Active	true	Successful
 Oisín Johnston	Custodian	Active	true	Successful
 Kim Akers	Custodian	Active	true	Successful

Figure 18-15: Custodians for a Premium eDiscovery case

**Premium eDiscovery Licensing:** While the accounts of users who create the input set for a Premium eDiscovery search only need Office E3 licenses, every user included in the scope of a Premium eDiscovery analysis needs an Office E5 license or Microsoft 365 E5 compliance license. This might sound like Premium eDiscovery is an expensive proposition, but the purpose of the content search is to refine the set of content for Premium eDiscovery to analyze. Refining the input set should mean a reduction in the number of users covered by the set too, so the number of Office E5 licenses you need is probably less than you think. From April 16, 2021, users need an appropriate license to create new Premium eDiscovery cases. Cases created before that date remain accessible to administrators without a license.

## Using PowerShell with Content Searches

Many compliance operations are quite complex, and it is usually best to execute eDiscovery and search operations through the portal. The information about running searches with PowerShell offered here gives some insight into what happens in the background so that you have a starting point for further investigation, if necessary. Remember that to run these cmdlets, you must connect a PowerShell session to the compliance endpoint. In addition, the account used must have the required Compliance permissions to perform whatever action you wish to take. If the account does not have the right permissions, not all cmdlet parameters are available. For example, if the account does not hold the eDiscovery Manager role, you cannot add the export action to a search.

### Creating and Running a New Content Search

To create a new content search, run the *New-ComplianceSearch* cmdlet, followed by the *Start-ComplianceSearch* cmdlet to start the search with the specified query to find items. For example, let's assume that we are interested in finding out whether any items exist in user mailboxes that have a specific phrase in their subject delivered in a certain period. Perhaps these items belong to a potential phishing attack. This command creates a simple content search to look for items based on some subject text delivered in a certain date range.



```
[PS] C:\> New-ComplianceSearch -Name "Look for Phishing Items" -Description "A search to locate suspicious phishing items" -PublicFolderLocation All -ExchangeLocation All -ContentMatchQuery '(Subject: "Phishing") AND (Received:06/01/2016..04/05/2021)' -AllowNotFoundExchangeLocationsEnabled $True
```

Setting *AllowNotFoundExchangeLocationsEnabled* to *\$True* means that the search will check the cloud-only mailboxes used to hold data for hybrid (on-premises) and guest accounts.

After creating the search, we use the *Start-ComplianceSearch* cmdlet to execute the query. Over the lifetime of a search, any time it has been longer than seven days since the search was run, you should run *Start-ComplianceSearch* to scan the target locations and refresh the results.

```
[PS] C:\> Start-ComplianceSearch -Identity 'Look for Phishing Items'
```

The *Get-ComplianceSearch* cmdlet fetches information about the search. For instance, this command checks the status of the search:

```
[PS] C:\> (Get-ComplianceSearch -Identity 'Look for Phishing Items').Status
```

The search status begins with "Starting" as the search engine initializes and evaluates the search parameters. It moves to "InProgress" when the search processes items in the target location. The search status reaches "Completed" when search results are ready for review. The number of items returned by the search is always interesting, so we can see this data with:

```
[PS] C:\> (Get-ComplianceSearch -Identity 'Look for Phishing Items').Items
```

Although content searches are much faster than their on-premises counterparts, a search across many locations can take some time to complete. Inside a script, we might want to have a loop to check the search status and then continue to some other processing once the search completes. Here is a simple loop that checks the search status every five seconds.

```
[PS] C:\> $Search = "Contoso Review of Patent Information for Project Alpha"
Start-ComplianceSearch -Identity $Search
do
{
    Write-Host "Searching..."
    Start-Sleep -s 5
    $Test = (Get-ComplianceSearch -Identity $Search).Status
}
While ($Test -ne 'Completed')

Write-Host "All done. Items found:" (Get-ComplianceSearch -Identity $Search).Items
```

The full output of the *Get-ComplianceSearch* cmdlet includes more detail than the search status and the number of found items. For example, the information returned about this content search shows the number of items found by the search in different mailboxes:

```
[PS] C:\> Get-ComplianceSearch -Identity 'Search for Phishing Items' | Format-List Items, Size, SearchStatistics
```

The information delivered in the *SearchStatistics* property is in JSON format. To see the essential statistics, do the following:

```
[PS] C:\> $Stats = Get-ComplianceSearch -Identity "Contoso Review of Patent Information for Project Alpha" | Select -ExpandProperty SearchStatistics
$Data = $Stats | Convertfrom-Json
$Data.ExchangeBinding.Search

Name           :
Sources        : 12
SourcesRaw     : 12
ContentItems   : 684
ContentSize    : 154.32 MB
```

```
ContentSizeRaw : 161813705
HasFaults      : False
```

To see the search results for each location, examine the *SuccessResults* property:

```
[PS] C:\> Get-ComplianceSearch -Identity 'Contoso Review of Patent Information for Project Alpha' |
Select -ExpandProperty SuccessResults
```

```
Location: Kim.Akers@office365itpros.com, Item count: 34, Total size: 24165076
```

There can be many lines of information detailing locations where a search finds items, including the special mailboxes used to store compliance records generated by Teams and Yammer by guest and hybrid accounts. To report the set of locations where a search finds information, we can use [subexpressions and the special \\$Matches hashtable](#) to capture results. For example:

```
[PS] C:\> [array]$Results = Get-ComplianceSearch -Identity "Project Derrigimlagh" | Select-Object
-ExpandProperty SuccessResults
$RecipientList = [System.Collections.Generic.List[Object]]::new()
$Data = $Results -Split '[\r\n]+'
ForEach ($Item in $Data) {
    If ($Item -match 'Location: (\S+),.+Item count: (\d+)' -and $Matches[2] -gt 0) {
        $RecipientDetails = [PSCustomObject][Ordered]@{
            "Recipient"      = $Matches[1]
            "Items found"    = $Matches[2] }
        $RecipientList.Add($RecipientDetails)
    } # End if
} #End ForEach Data
```

```
$RecipientList
```

Recipient	Items found
-----	-----
Tony.Redmond@office365itpros.com	489
Deirdre.Smith@office65itpros.com	160
Ben.Owens@office365itpros.com	11
Customer.Services@office365itpros.com	9

A content search with a broad search query is likely to return hundreds or even thousands of items. The easiest way to validate that the search finds items of interest is by reviewing the sample items retrieved by the search through the GUI (Figure 18-16). Apart from anything else, this allows you to examine the items carefully to ensure that the search is working as expected. The next step is then to export the items for closer examination or to refine the search keywords and conditions to focus on a more refined set.

## Look for Phishing Items samples

The screenshot displays an email search results interface. On the left, a list of search results is shown with columns for Subject/Title, Date, and Sender/Author. The selected item is an email from Patrick Peterson (Agari) dated Mar 5, 2020, with the subject '[Webinar] KnowBe4 & Agari: Transforming Phishing Protection March 18'. On the right, the detailed view of this email is shown, including the 'From', 'To', 'Subject', and 'Send Date' fields. The body of the email contains a message about a new partnership with KnowBe4, aimed at combating phishing and BEC attacks. A red arrow points to the word 'phishing' in the subject line of the selected email in both the list and the detailed view.

Figure 18-16: Viewing sample items found by a content search

## Searching SharePoint and OneDrive for Business

The locations searched by the example query used above are user mailboxes and public folders, including the cloud-only mailboxes holding compliance records for non-tenant users. We can add OneDrive for Business and SharePoint sites to the set of search locations, but only if the target resources support the query. In this case, if we try to add non-Exchange locations to the search by running the *Set-ComplianceSearch* cmdlet, we will see a warning message because the search query is based on email-specific properties (Subject and Received date) that OneDrive for Business and SharePoint Online don't support.

Here is how to create a content search that specifically targets SharePoint and OneDrive for Business sites. In this case, we search for any item stored in any site (the *SharePointLocation* parameter is set to "All") that includes some credit card information (the search uses the sensitive information type to find these files – see Chapter 19 for more information). We can further refine the search query by only looking for items created on or after 1 January 2015:

```
[PS] C:\> New-ComplianceSearch -Name "SPO and OD4B search for credit card data"
-Description "A search of SharePoint and OneDrive to locate files containing credit card data"
-SharePointLocation "All" -ContentMatchQuery '(SensitiveType:"Credit Card Number:2..") AND
(created>=01/01/2015)'
```

Once again, after creating a new content search, we start it with the *Start-ComplianceSearch* cmdlet to retrieve some results.

The example query searches all sites. If we try to refine the set of sites that we want to search, we can do so by specifying the URL for each site as shown below. In this instance, we also add a couple of Exchange Online mailboxes to the search locations. This will cause an error because the search properties we are trying to use are not available for all locations.

```
[PS] C:\> Set-ComplianceSearch -Identity "SPO and OD4B search for credit card data" -
SharePointLocation "https://office365itpros.sharepoint.com/sites/O365ITPro/",
"https://office365itpros.sharepoint.com/Projects/" -ExchangeLocation "Tony Redmond", "Paul
Cunningham"
```

PowerShell allows you to go ahead with the content search even though the problem with search properties is flagged. It is a mistake to do this as the net effect is usually to find everything in the search location that does

not support the properties used in the search query. In this example, because we use search properties specific to SharePoint and OneDrive for Business but have included some Exchange mailboxes, the search returns every item in those mailboxes, which is probably not what you intended.

See this [blog post](#) and this [support article](#) for tips about how to format KQL queries to interrogate SharePoint and OneDrive for Business sources.

## Content Search Actions

Behind the scenes, many of the actions you take after a search runs use the *New-ComplianceSearchAction* cmdlet to add a new action to the search. For example, when you export data from a search, the cmdlet runs to add an export action to the search. The following parameters are the most used:

- *SearchName*: The name of the search to add an action.
- *NotifyEmail*: The email address(es) of users to receive a notification when an action is complete.
- *NotifyEmailCC*: The email address(es) of CC recipients for notification messages.
- *Scope*: Specifies what kind of items (indexed, unindexed, both indexed and unindexed) the search will process.
- *Preview*: Add an action to preview items found by the search.
- *Export*: Add an action to export items found by the search to a PST or folder.
- *EnableDeDupe*: Removes duplicate messages during exports of search items.
- *IncludeSharePointDocumentVersions*: Input *\$True* to include all versions of discovered SharePoint documents or *\$False* to include only the most recent version. Remember that SharePoint Online captures many versions of Office documents during edit sessions due to the autosave feature.
- *RetryOnError*: Retry the search if an error occurs. Do not specify this parameter when using a content search action to remove items from mailboxes as it will stop the action.
- *Purge*: Add a search action to remove items found by the content search. Exchange Online is the only workload to support purging items. Because the search depends on the content indexes, it cannot remove [unindexed items](#) (like encrypted messages that aren't protected by sensitivity labels) from mailboxes. Users who don't hold the compliance Organization Management role cannot use the purge action.
- *PurgeType*: Specify how the content search should remove items. Items can be soft-deleted by passing the *SoftDelete* value. Users can recover soft-deleted messages during the deleted items retention period (default 14 days). You can also specify *HardDelete* to force hard deletion, which makes the items irrecoverable by the user.

You cannot add a new action to a search if the search is in progress or the search results are stale. For export operations, a search is stale if it is more than seven days old. For destroy or purge operations, a search stays usable until it is more than ten days old.

## Using a Content Search to Purge Mailbox Items

Administrators often need to purge items from mailboxes to eliminate:

- Phishing messages before users have the chance to read them and activate the harmful links.
- Messages that hold malicious attachments (viruses or other code).
- Messages sent in error and hold information that the organization would like to withdraw (insofar as is possible). You cannot retrieve messages that users send outside the organization or when delivered by Outlook to a PST.

In these circumstances, the usual approach is to create a search targeting the mailboxes holding the problem messages. The easiest way to do this is to create the search in the Microsoft Purview Compliance portal and refine it by tweaking search criteria until you are happy that the search finds the right messages. Because you

are going to remove data from user mailboxes, the search must be as precise as possible. Think of keyhole surgery rather than a massive incision. To help focus the search, you should capture the essential characteristics of the targeted message and use these as the basis of the search. For instance:

- An exact word or phrase that occurs in the message subject. For example: "Great Opportunity to Purchase." You can use this phrase for the Subject property in the search query.
- The sent date for the message, or even a limited date range. Use this date as the Received property in the search query.
- The SMTP address of the sender.

Combining several properties will generate a more precise search than when you rely on a single property, such as the message subject. You can use the search preview facility to examine the details of the items found by the search and verify that the search criteria are correct and working as expected.

You can create a purge action for a search against a maximum of 50,000 mailboxes. If you need to purge items from more mailboxes, you must split processing across multiple searches. In addition, a purge action can only remove ten items at a time from a mailbox. You cannot create a purge action for a search covering data held in SharePoint Online or OneDrive for Business.

Another thing to consider is the effect of litigation and in-place holds. If holds are in place for mailboxes, it might not be possible for a content search hard delete purge to remove items. The search can find the items, but any attempt to remove the items fails because they are on hold. Soft delete purges will work because the items remain in the mailbox. The golden rule is that if you want to hard-delete mailbox data, make sure to first remove any holds covering the mailboxes.

**Look before you leap:** No one wants to remove items from user mailboxes in error. Before committing a search to remove items, it is best to check that the search completes successfully and finds the right items. Use the Preview function to verify that the items you expect to find are present and that no unanticipated items have turned up. You might need to refine the search query several times before the search does exactly what you want it to do and is ready to go ahead and remove the items.

## Using the Purge Action

When you are happy the search finds the correct results, add the *Purge* action to the search. Note that you can only do this if you are a member of the *Compliance Organization Management* role group. For example, this command invokes the search called "Look for Phishing Items" and instructs it to soft delete any items matching the search criteria. Before the cmdlet starts the purge, you must confirm that the action should proceed.

```
[PS] C:\> New-ComplianceSearchAction -SearchName "Look for Phishing Items" -Purge -PurgeType SoftDelete
```

After a few minutes, check the status of the purge operation by running the *Get-ComplianceSearchAction* cmdlet. The purge reports its progression in percentage terms from zero to 100. The name of the action is formed by combining the search name and a "\_Purge" suffix. Given that the name of the search referenced above is "Look for Phishing Items," the name of the action is "Look for Phishing Items\_Purge." The following command retrieves the progress of the purge action.

```
[PS] C:\> Get-ComplianceSearchAction -Identity "Look for Phishing Items_Purge" | Format-Table SearchName, JobStartTime, JobProgress, Status -AutoSize
```

SearchName	JobStartTime	JobProgress	Status
Look for Phishing Items_Purge	18/08/2019 19:04:38	100	Completed

To see how many items were purged, look at the Results property of the purge action.

```
[PS] C:\> (Get-ComplianceSearchAction -Identity "Look for Phishing Items_Purge").Results
Purge Type: SoftDelete; Item count: 4; Total size 104080; Details: {Location: ; Item count: 2; Total
size: 20443; Failed count: 0; ,
Location: ; Item count: 1; Total size: 54554; Failed count: 0; ,
Location: ; Item count: 1; Total size: 29083; Failed count: 0; }
```

If you need to remove more than 10 items, let the purge job complete, remove it, and then resubmit. Do this until all the items are removed:

```
[PS] C:\> Remove-ComplianceSearchAction -Identity " Look for Phishing Items_Purge " -Confirm:$False
[PS] C:\> New-ComplianceSearchAction -SearchName "Look for Phishing Items" -Purge -PurgeType
SoftDelete
```

You can purge items with content search actions in two ways:

- The *HardDelete* action moves items into the *Purges* subfolder of Recoverable Items. Users can't access items in the Purges folder. The items remain in the mailbox until they are removed by the Managed Folder Assistant (see below).
- The *SoftDelete* action moves items into the *Deletions* subfolder of Recoverable Items. Users can recover items from Deletions using the Recover Deleted Items feature in Outlook and OWA.

The Managed Folder Assistant processes items held in the Purges or Deletions folders and removes the items after all retention mechanisms expire. These include:

- The mailbox's deleted items retention period (items are kept between 14 and 30 days) together with Single Item Recovery to ensure that items can't be removed until the deleted item retention period passes.
- A personal retention tag is applied to the item.
- Retention policies (Microsoft 365 or Exchange Online) are applied to the mailbox.

Items remain in place to be available for eDiscovery. It's important to remember this because a search will continue to find items in mailboxes even when they have been removed from user view by a purge. In other words, even after you run a content search and purge its results, the messages remain until anything holding them back is removed. However, mailbox owners can no longer access or recover any hard-deleted items.

Of course, software being software, it is wise to verify results when a risk exists that a malicious email might wreak havoc on mailboxes. After running any command to remove items, you can review the properties of the search to see the mailboxes where items were removed, or (if you have permission), sign into a mailbox that you know held a problem item to check that the items are gone.

You can download an example script to create and run a search to find mailbox items before using a search action to purge the discovered items [from GitHub](#). Note that the *Search-Mailbox* cmdlet remains available at the time of writing and can be used to remove items from mailboxes.

## Be Sure About Removal

The removal of items from a user mailbox through a search generates a *SearchResultsPurged* audit record in the audit log, which tells you what query was used. Neither this audit record nor the *Get-ComplianceSearchAction* cmdlet reports the locations from which items were purged, but (using the *Search-UnifiedAuditLog* cmdlet) you should also find a matching *HardDelete* or *SoftDelete* record in the audit log that contains details of the purged items. The lack of obvious feedback on what the search action did is a good reason why you need to make sure that the content search returns the items you want to purge. To validate the results of the content search, use the preview function and then, if necessary, export the items to a PST (you might need to export the items anyway to keep evidence of what is removed from user mailboxes).

**Limit any possible damage:** Removing data from user mailboxes is an operation that is fraught with error. The potential of making a mistake and removing something that you should not delete is always there. To

reduce the potential for harm, purging based on a content search will only remove a maximum of 10 items from a mailbox per run (if the search covers multiple mailboxes, more than ten items can be removed in a single run). The purge function is an “incident response” feature to remove small amounts of problematic data from user mailboxes. Limiting removal to ten items should not be a problem unless malware floods mailboxes through multiple attacks over a short period. Purges should use a precise, focused search to find data. The ideal situation is to find and remove a single item. In addition, by limiting deletion to 10 items per mailbox, search and destroy operations finish much faster (Microsoft reckons on being able to process 100,000 mailboxes in 30 minutes). Using soft deletes rather than permanent removals allows users to recover items if the search removes them in error. If you need to purge more than 10 items from a mailbox, you must run a purge action multiple times, removing the purge action and recreating it each time.

## Exporting Search Data

When a content search is tuned to find the right data, the next step is usually to export the data to allow expert examination and review. For email, the export can be to PSTs or as individual messages while exports from SharePoint or OneDrive are as individual documents, including earlier versions if necessary.

You can use PowerShell to automate the export to Azure part of an export operation. After that, things become trickier to call the tool used to download results from Azure and pass the necessary access token to the tool (see [this discussion for a suggested solution](#)). The example below shows how to start the export process for search results. In this case, we want to export messages in PSTs (the format is “*FxStream*”) with messages for each mailbox stored in a separate PST. A loop tracks the progress of the job. When complete, the data is ready in Azure to be downloaded.

```
[PS] C:\> $Search = "Investigation LD-17166"
New-ComplianceSearchAction -SearchName $Search -EnableDedupe $True -Export -Format FxStream
-ArchiveFormat PerUserPST -Scope BothIndexedAndUnIndexedItems
$ExportJob = $Search+"_Export"
Write-Host "Export started at" (Get-Date)
do
{
    Start-Sleep -s 3
    $ExportStatus = (Get-ComplianceSearchAction -Identity $ExportJob -Details)
    Write-Host "Current status:" $ExportStatus.Status "% progress:" $ExportStatus.JobProgress
}
While ($ExportStatus.Status -ne 'Completed')
Write-Host ""
Write-Host "Export ended at" (Get-ComplianceSearchAction -Identity $ExportJob).JobEndTime
```

To download the exported results, select the search in the Microsoft Purview Compliance portal, followed by **Download exported results**. The export key is visible. Copy the export key to the clipboard. Later, you will paste the key into the download tool to authorize it to copy data from Azure to the export target destinations (such as PSTs or individual files).

If you want to restart an export, run the *Remove-ComplianceSearchAction* cmdlet to remove the exported data from Azure. Remember to append “\_Export” to the end of the content search name to create the name of the export action.

```
[PS] C:\> Remove-ComplianceSearchAction -Identity "Investigation_#2_Export"
```

# Using PowerShell to Manage eDiscovery Cases

You can use PowerShell to create and manage eDiscovery cases. To create a case, use the *New-ComplianceCase* cmdlet.

```
[PS] C:\> New-ComplianceCase -Name "Stock Trading Case Reference 2017-001"
```

The *Get-ComplianceCase* cmdlet returns details of an eDiscovery case. However, because an eDiscovery case is a wrapper around content searches and holds, it does not return much information as no members or sources have yet been added.

```
[PS] C:\> Get-ComplianceCase -Identity "Stock Trading Case Reference 2017-001"
```

```
RunspaceId      : 883d2fb0-ef1b-484f-918d-bedf1930ef11
TenantId        : b662313f-14fc-43a2-9a7a-d2e27f4f3478
Identity        : 2fd9411c-3fd4-4293-a6f1-8fca699f51ed
Name            : Stock Trading Case Reference 2017-001
Description     :
CaseType        : eDiscovery
Status          : Active
ClosingStatus   : Unknown
CreatedDateTime : 11/01/2017 18:13:32
LastModifiedDateTime : 11/01/2017 18:13:32
ClosedDateTime  :
LastModifiedBy  : Tony Redmond
ClosedBy        :
ObjectState     : New
```

Note: a Premium eDiscovery case is a special form of eDiscovery case. You can retrieve these cases by specifying the *CaseType* parameter:

```
[PS] C:\> Get-ComplianceCase -CaseType Adv
```

After finding the case to work with, we can update its details. For instance, we can add some members to the case with the *Update-ComplianceCaseMember* cmdlet, specifying the user principal name or display name for each member in a comma-separated string. For example:

```
[PS] C:\> Update-ComplianceCaseMember -Case "Stock Trading Case Reference 2017-001" -Member "Tony Redmond", "Kim.Akers@Office365ITPros.com"
```

The *Update-ComplianceCaseMember* cmdlet replaces the existing member list for a case. If the user who runs the cmdlet does not specify their name, the cmdlet includes it automatically. If you want to add a specific user to an existing member list, use the *Add-ComplianceCaseMember* cmdlet.

To make sure that no one can remove information needed by the case, we must set up a hold to retain information. A hold includes a rule and a policy. We create the rule with the *New-CaseHoldRule* cmdlet. The simplest form of rule has no query, in which case all content in the search sources is covered. For instance:

```
[PS] C:\> New-CaseHoldRule -Policy "Stock Trading Case Reference 2017-001" -Name "Stock Trading Case Reference 2017-001"
```

It's more usual to include a query, defined in KQL syntax. This example sets up a condition to look for any item where the words "stock" and "trading" occur near each other.

```
[PS] C:\> New-CaseHoldRule -Policy "Stock Trading Case Reference 2017-001" -Name "Stock Trading Case Reference 2017-001" -ContentMatchQuery "Stock NEAR(10) Trading"
```

We now use the *New-CaseHoldPolicy* cmdlet to add the hold policy for the case. The hold policy defines the locations that the case covers. Sources are



- Exchange mailboxes.
- Public folders.
- SharePoint and OneDrive sites.

You can include either individual mailboxes or distribution lists. The membership of groups is expanded into individual mailboxes and added to the policy. You can also include public folders by passing the "All" value to the *PublicFolderLocation* parameter. SharePoint sites are specified with the site URL. For example:

```
[PS] C:\> New-CaseHoldPolicy -Case "Stock Trading Case Reference 2017-001" -Name "Stock Trading Case Reference 2017-001" -ExchangeLocation "Nancy Anderson", "Sanjay Patel" -SharePointLocation https://office365itpros-my.sharepoint.com/personal/ben_owens_office365itpros_com -PublicFolderLocation All -Enabled $True -Comment "Search sources added programmatically"
```

To check that the hold exists, we can run this command:

```
[PS] C:\> Get-CaseHoldPolicy "Stock Trading Case Reference 2017-001"
```

Or, to see all holds in place for all eDiscovery cases:

```
[PS] C:\> Get-ComplianceCase | % {Get-CaseHoldPolicy -Case $_.Name}
```

After you create an eDiscovery case, you probably want to create one or more content searches and link the searches to the case. We discussed the creation of content searches earlier. This example is like those reviewed then with the exception that we use the *Case* parameter to associate the search with the case.

```
[PS] C:\> New-ComplianceSearch -Name "Stock Trading Case Reference 2017-001 - Search 1" -Description "Search associated with Case Reference 2017-001" -ContentMatchQuery "Stock NEAR(10) Trading" -Case "Stock Trading Case Reference 2017-001" -ExchangeLocation "Nancy Anderson", "Sanjay Patel"
```

Once the search is available, you can start it in the normal manner by calling the *Start-ComplianceSearch* cmdlet.

The queries and search locations used in the content searches for eDiscovery cases do not have to match the hold rule or policy. This is because a case might span multiple searches, each of which looks for different information from different sources. The combination of all those results constitutes the information for case investigators to review.

## Removing Cases

You can remove a case by running the *Remove-ComplianceCase* cmdlet, but only after removing any holds associated with the case. For example, let's assume that we want to remove the Contoso Alpha Case July 2015 case. To remove the holds, we issue the command:

```
[PS] C:\> Get-CaseHoldPolicy -Case "Contoso Alpha Case July 2015" | Remove-CaseHoldPolicy
```

The command to remove the holds is published to the workloads. It takes some time for the workloads to process the removals. When this process finishes, and no holds show up when you run the *Get-CaseHoldPolicy* cmdlet for the case, you can run the *Remove-ComplianceCase* cmdlet:

```
[PS] C:\> Remove-ComplianceCase -Identity "Contoso Alpha Case July 2015"
```

## PowerShell and Data Subject Requests

GDPR data subject requests are special forms of eDiscovery cases. The same is true for Subject Rights Requests created in the Privacy management solution. To view details of these cases, you must specify the case type. Regular cases have a case type of "eDiscovery," which the *Get-ComplianceCase* cmdlet assumes if no other type is present. For example, to see DSR cases, specify the DSR case type.

```
[PS] C:\> Get-ComplianceCase -CaseType DSR
```

Name	Status	CreatedDateTime
James Ryan DSR	Active	30/04/2018 21:08:19
Tony DSR	Active	01/05/2018 15:10:56

To list Privacy management cases, specify the *PrivacyManagementDSR* type. To access details of an individual case, pass the name of the case:

```
[PS] C:\> Get-ComplianceCase -CaseType DSR -Identity "James Ryan DSR" | Format-List
```

```
Identity           : 16a90128-1e5b-4506-a6ca-160d0992fbaf
Name               : James Ryan DSR
Description        : James Ryan DSR
CaseType           : DSR
Status             : Active
ClosingStatus      : Unknown
CreatedDateTime    : 30/04/2018 21:08:19
LastModifiedDateTime : 30/04/2018 21:08:19
ClosedDateTime     :
LastAccessTime     : 01/01/0001 00:00:00
LastModifiedBy    : Tony Redmond
```

The other cmdlets used with eDiscovery cases work as documented with DSRs.

One specific aspect of GDPR DSRs is that the regulations say that they should be processed within 30 days. We can check the remaining time for active cases with some code like that shown below, which highlights any case that has seven days or less left to respond:

```
[PS] C:\> $Cases = (Get-ComplianceCase -CaseType DSR | ? {$_.Status -eq "Active"})
ForEach ($Case in $Cases) {
    $Days = (New-TimeSpan -Start $Case.CreatedDateTime -End (Get-Date)).Days
    If ($Days -gt 23) {
        $DaysLeft = (30 - $Days)
        Write-Host "Warning!" $Case.Name "has only" $DaysLeft "days remaining to complete and
respond..." }}
```

## Adding eDiscovery Managers

Although it is usually best to manage the membership of sensitive role groups through the portal, you can add or remove users from the eDiscovery role groups through PowerShell. To do this, connect to the compliance endpoint and run the *Add-RoleGroupMember* (to add an eDiscovery Manager) or *Add-eDiscoveryCaseAdmin* cmdlet (to add an eDiscovery administrator). For example, these cmdlets add a user to the eDiscovery Managers role group and then check that the user is in the group:

```
[PS] C:\> Add-RoleGroupMember -Identity eDiscoveryManager -Member "James Abrahams"
[PS] C:\> Get-RoleGroupMember eDiscoveryManager
```

Only an eDiscovery administrator can add another user to the eDiscovery Administrators role group. If you hold this status, you can run the cmdlets to add a user and check that the add worked as follows:

```
[PS] C:\> Add-eDiscoveryCaseAdmin -User "James Abrahams"
[PS] C:\> Get-eDiscoveryCaseAdmin
```

## Reporting Holds for eDiscovery Cases

The [Office 365 for IT Pros GitHub](#) repository contains a modified version of [a Microsoft script](#) to report the holds that exist for eDiscovery cases. You can use it as an example of how to navigate through eDiscovery cases to unpick and report on components.

The expected output is something like shown below, with cases organized by status and the holds described for active cases. The result of the analysis is held in the *\$Report* variable, so it is very possible to create different reports from the data without changing the script.

EDiscovery Cases found: 9  
 Active Cases: 7  
 Closed Cases: 2  
 Active Holds: 11

Case	Status	Created	HoldCreated
Board Minutes	Closed	28/04/2017 19:14	
Contoso Investigation	Closed	28/04/2017 19:14	
Improper management transactions (LD-2018-002)	Open	04/04/2018 18:03	04/04/2018 18:19:55
Inactive Exchange Mailboxes	Open	28/04/2017 19:14	28/04/2017 19:12:27
Personal Harassment Investigation LD-2018-001	Open	04/04/2018 17:31	04/04/2018 17:38:33
Personal Harassment Investigation LD-2018-001	Open	04/04/2018 17:31	04/04/2018 17:44:42
Project Alpha Foxtrot	Open	28/04/2017 19:14	28/04/2017 19:12:27
Stock Trading Case Reference 2017-001	Open	28/04/2017 19:14	28/04/2017 19:12:27
Stock Trading Case Reference 2017-001	Open	28/04/2017 19:14	15/02/2018 17:16:56

## In-place Holds and Litigation Holds

An eDiscovery case can include one or more in-place holds to ensure that workloads retain information even if someone tries to remove or edit it. Each hold has a search query to find the information, which stays in the source location until something happens, like a user trying to remove or edit the content. At this point, Exchange Online captures a copy of the original content. Holds applied by Exchange and SharePoint or by retention policies and eDiscovery cases all keep content in place. It is an efficient mechanism that avoids the need to duplicate information unnecessarily.

← ×

### Manage litigation hold

When a mailbox is put on litigation hold, users can purge items from their mailbox, but those items will still be retained by Office 365.  
[Learn more](#)

Turn on litigation hold

Hold duration (days)

Note visible to the user

Web page with more information for the user

[Save changes](#)

Figure 18-17: Placing a mailbox on litigation hold

Litigation or legal hold is a somewhat cruder but very effective mechanism available for Exchange Online mailboxes. A litigation hold places the entire mailbox on hold for a set period or indefinitely. Again, all items stay in place, but every deletion means that Exchange keeps a copy of the deleted item. As you can imagine,

many of the items contained in mailboxes fall into the banal category and are of no interest whatsoever to discovery actions. However, there are instances where it is necessary to keep everything so that there is no chance of missing anything which might remotely be of interest. A litigation hold keeps everything in a mailbox, as will an in-place hold that has no search criteria. You can use litigation holds alongside in-place holds. All mailbox content is held when a mailbox is subject to both types of hold.

Placing a mailbox on litigation hold is easy. First, select the mailbox owner's account in the Microsoft 365 Admin Center, go to the **Mail** tab, and select **Manage litigation hold**. If a hold is already in effect on the mailbox, you'll see when it was set and by whom. To enable litigation hold for a mailbox, check *Turn on litigation hold* and enter values for the optional properties as shown in Figure 18-17.

- **Hold duration:** How long the hold will last in days. Leave blank to set an indefinite hold.
- **Note visible to user:** Text entered here to explain the reason for the hold is visible to the mailbox owner in the Account Settings section of Outlook's "backstage."
- **Web page with more information:** A URL to a web page that should contain more information for the user to know why the organization puts mailboxes on hold and what it means to them. It might also include relevant citations of local laws governing user privacy and the steps taken by the organization to ensure that user privacy is respected. If present, Outlook displays a *More information* link in Account Settings. No check is made to ensure that the URL is reachable.

Click **Save changes** to make any changes to hold settings effective.

To make a litigation hold effective, Exchange Online updates several mailbox properties:

- **LitigationHoldEnabled:** Set to True.
- **LitigationHoldDate:** Set to the time and date when an administrator applies the hold to the mailbox. For example, 2-Feb-2022 19:45:50.
- **LitigationHoldOwner:** Set to the account that placed the mailbox on hold.
- **LitigationHoldDuration:** Set to the retention period. For example, 90.00:00:00, or 90 days.
- **RetentionComment:** The free-text note to inform the user that their mailbox is on hold. You don't have to enter this comment if you do not want to.
- **RetentionUrl:** The URL pointing to a page holding additional information.

You can also set litigation hold on by setting mailbox properties in EAC or with PowerShell. Here is an example of putting a mailbox on litigation hold using PowerShell is:

```
[PS] C:\> Set-Mailbox -Identity "Ben Owens" -LitigationHoldEnabled $True  
-LitigationHoldDate "1-May-2018 19:45:00" -LitigationHoldOwner "Administrator"  
-LitigationHoldDuration 90.00:00:00 -RetentionComment "Your mailbox is on hold"  
-RetentionUrl "http://www.contoso.com/compliance.htm"
```

A hold can be indefinite. If you want to pass a duration in days, Microsoft's documentation states that the limit is 2555 days (7 years), but the cmdlet is happy to accept 100,000 days (>273 years). I doubt anyone now alive will be worrying if such a hold expires before its due date.

# Chapter 19: Managing Data Loss Prevention

## ***Tony Redmond***

It would be nice if users handled confidential or sensitive information correctly all the time, but that is not how things happen in real life. It is human nature to err, and a common mistake is including confidential or sensitive information in emails or via a document shared with external people. Data Loss Prevention (DLP) is a technology designed to prevent users from sharing sensitive information inappropriately. Policies encapsulating rules to dictate the sharing of sensitive information are how organizations deploy DLP. First introduced in Office 365 for Exchange Online, DLP covers Exchange Online, SharePoint Online, Teams, and OneDrive for Business, with plans in place to extend coverage further to protect data drawn from across Microsoft 365.

This chapter describes:

- Data Loss Prevention concepts.
- How Microsoft Purview implements DLP.
- How to build suitable DLP policies to help applications prevent the loss of sensitive information.
- How to know whether your DLP policy is effective.

Data Loss Prevention is not a magic bullet and its capabilities do not extend to every method that someone can use to share data with another person. For instance, someone could post sensitive information like a credit card number in a Facebook conversation. You can handle that issue through third-party monitoring software, but user education might be equally effective. DLP is just another part to fit into the compliance puzzle that contributes to the overall security of information within companies.

## Leak Prevention with Software

The need to prevent the loss (or leakage) of sensitive information such as credit card numbers, personal information like passport numbers, and so on is well understood, possibly because there have been many instances in which company or personal information has been compromised by being made public deliberately or accidentally. Graphic examples exist where large public companies have lost millions of records. The result can be brand erosion, legal expenses, customer loss, and an almost guaranteed public relations disaster. It is easy for a user to attach a document and send it to someone else; it is also easy to make a mistake that ends up with information leaking outside the organization. For example, you might not notice that Outlook has auto-completed an address with the wrong recipient and end up sending a confidential attachment outside the company. The speed of modern email systems makes any attempt to recall messages futile, so if someone realizes that they have made a mistake, all they can do is call the recipient and ask them to remove the content from their system.

Companies devote enormous effort to protecting IT systems. Much of the focus in the past has been on erecting traditional security barriers to stop attackers from penetrating internal systems. This approach might keep out hackers and scammers, but it does nothing to prevent employees from causing data loss, usually by accident. DLP is not new, and vendors have offered solutions for the last decade or so. Microsoft entered the space with Exchange 2013. Although a late entrant to the DLP market, Microsoft enjoys certain advantages over third-party solutions because of its ability to deeply integrate DLP checks into clients and servers:

- Integration of DLP functionality to protect email into Outlook, which is a critical desktop application in large enterprises. The integration analyzes text using a mixture of algorithms including pattern matching and [regular expressions](#) as users type into a message body to detect possible sensitive information considered to be of concern to the company. OWA also includes the ability to perform DLP checking. Both clients display policy tips to users when they detect sensitive information so that users become more aware of the kind of data that they are dealing with before they send messages with that content.
- Integration of DLP functionality into the desktop and online versions of other Office applications such as Word, Excel, and PowerPoint. The inclusion of DLP checks throughout the Office suite means that users have less chance of making mistakes.
- The ability to implement checking at points where data is guaranteed to pass through. For Exchange Online, checking happens in the transport system as all messages must pass through it before leaving the system or are delivered to internal recipients. Sensitive information can therefore be intercepted in messages by integrating DLP checks into the transport flow. For documents stored in SharePoint Online and OneDrive for Business, a crawler to detect sensitive information checks files in document libraries as users add or share content. The crawler is already indexing all content stored in SharePoint and OneDrive for Business sites, so adding a check for sensitive content there creates the desired oversight. When DLP policies are deployed to protect Teams, the examination happens in the chat service.

The need to protect against the disclosure of confidential data does not exist in a vacuum. To be effective, it is important to incorporate DLP into an overall corporate compliance strategy alongside other tools to protect data both inside and outside the organization. For example, the strategy should consider how to use sensitivity labels to encrypt and protect the most confidential content so that even when this material circulates outside the organization, it continues to be protected and cannot be accessed except by authorized people. It's also important to consider how to protect data created before the introduction of the compliance strategy as old documents and messages can contain highly confidential information.

Clients display policy tips when a DLP violation is detected to emphasize that user education and policy communication is very important. People must know why the policy tips appear and what they should do next. It is also important to stress that DLP is not a guaranteed block against users sharing sensitive data. If someone wants to send information externally, they will be able to find another way to do so.

DLP policies to process Exchange Online, SharePoint Online, and OneDrive for Business items require an Office 365 E3 license (or, for Exchange Online only, an Exchange Online E2 license). However, DLP transport rules will block messages in violation of policy no matter what license a user has. DLP for Teams requires Office 365 E5 or Microsoft 365 E5 Compliance licenses.

The importance of DLP and its growing use is illustrated by statements made by Microsoft at the Ignite 2017 conference, when they reported that DLP use within Office 365 tenants had a year-over-year increase of 600% while the number of users protected by DLP grew by 750% in the year to September 2017. That growth has tapered off a little, but there's still lots of interest in using DLP, especially in enterprise tenants.

## Microsoft Purview Data Loss Prevention

Microsoft's goal is to implement a common DLP capability across multiple workloads based on checks evaluated against common conditions and data types (including retention and sensitivity labels) existing within files and messages. DLP currently covers:

- Exchange Online.
- SharePoint Online and OneDrive for Business.
- Teams chats and channel conversations.

- Windows 10 devices (endpoint DLP).

Microsoft Purview DLP policies are managed through the Microsoft Purview compliance portal and are the focus of future development. Some organizations use DLP policies implemented through Exchange transport rules (ETRs) to ensure that the same policy coverage can be implemented for on-premises servers. For more information on ETRs, see the Exchange Online chapter in the companion volume.

The biggest advantage of Microsoft Purview DLP policies is their cross-workload nature. As of March 2021, [the range of conditions, actions, and exceptions available to process email](#) through Microsoft Purview DLP policies are more capable than those available in ETRs, which means that there is no reason to deploy new ETRs.

In October 2021, Microsoft announced that they will [remove the ability to maintain ETRs in the Exchange admin center sometime in the April-June 2022 period](#) (it'll still be possible to make changes with PowerShell).

This is a strong signal that Microsoft is close to the point when they will remove DLP ETRs from Exchange Online. Organizations using DLP ETRs should begin the process of transitioning DLP ETRs to Microsoft Purview DLP policies, taking the following into account:

- The rules implemented in ETRs always take precedence over Microsoft Purview DLP policies when processing email. This is natural because ETRs are explicitly designed to process email. It also means that messages that are blocked by an ETR will never be checked by a Microsoft Purview DLP policy.
- Messages that pass through ETRs are processed by Microsoft Purview DLP policies and could be blocked at this point if the rules detect a condition that calls for this action. If a “stop processing” action is invoked by an ETR, it will not affect the processing of Microsoft Purview DLP policies.

Conversion from ETRs to Microsoft Purview DLP policies is not automatic. An effort is needed to plan the conversion, create and test the new policies, and ensure that the combination of old and new policies function together during the conversion period. To ease the transition, Microsoft has developed [a playbook to help tenants move from ETRs to Microsoft Purview DLP](#). With the playbook in hand and feature parity attained between the two sets, most organizations have no good reason to remain using ETRs. The transition requires careful planning and attention to detail and the switchover should happen at a time of low user demand such as a holiday period.

## Sensitive Information Types

The basis for a DLP policy is understanding what kind of sensitive data you want to protect and how you want to protect the data. Several common data types come to mind when you consider what kind of information you prefer users to not mishandle. Credit card numbers or passport numbers are obvious examples of data usually regarded as sensitive. Although credit cards use a worldwide format, personally identifiable information (PII) and other sensitive information types differ from country to country.

It would not make much sense if Microsoft required tenants to analyze the essential characteristics of data types like credit card numbers and passport numbers to create custom definitions to match these types. Mistakes would happen, and inconsistencies in definitions would abound. Fortunately, Microsoft Purview includes [over two hundred and sixty definitions of sensitive information types](#) for use in DLP policies and other solutions. Some of the data types are like other types (for example, the definition for passport numbers or driving licenses often does not vary much between countries). Among the set of standard sensitive information types are:

- ABA routing numbers.
- Credit card numbers.
- U.S. Social Security numbers.
- Canada bank account numbers.
- European Union debit card numbers.

- Australian passport numbers.
- German driver's license numbers.
- European Union tax identification number.

Over time, Microsoft has expanded the inventory of sensitive information types to increase coverage for data loss prevention and other capabilities. Many of the sensitive information types added recently are to detect country-level data like passports, identity cards, and bank numbers. In addition to new types, Microsoft also tweaks existing definitions to improve their accuracy.

You can see the current set of sensitive information types available in Microsoft Purview by running the `Get-DlpSensitiveInformationType` cmdlet (after connecting to the compliance endpoint):

```
[PS] C:\> Get-DlpSensitiveInformationType | Format-Table Name, Description
```

Name	Description
Canada Driver's License Number	Detects Canadian driver's license number.
EU Debit Card Number	Detects European Union debit card number.
Israel National ID	Detects Israeli national identification number.
Credit Card Number	Detects credit card numbers for American Express, Diner...
U.S. Social Security Number (SSN)	Detects formatted and unformatted US social security nu...
German Passport Number	Detects German passport numbers or Reisepass.

Sensitive information types break down into bundled and unbundled entities. An unbundled entity is a sensitive information type that stands on its own. It can be used in a DLP policy to detect specific information, or it can be used as part of a bundled entity. Each sensitive information type is defined by describing its characteristics in the form of some recognizable pattern, often captured in a regular (regex) expression. For instance, a social security number is a nine-digit number usually formed in three groups of three, two, and four digits separated by hyphens. Many passports have some alpha characters followed by seven digits, and so on.

A bundled entity is simply a collection of sensitive information types managed as a single type. For example, if you use the *All Medical Terms and Conditions* type to detect content in a DLP policy, DLP finds any medical term or condition found in SharePoint, Exchange, and Teams content to which the policy applies.

The use of sensitive information types is not restricted to DLP. You can use sensitive information types in:

- DLP policies.
- Communications compliance policies.
- Auto-labeling policies (retention labels by background processes).
- Auto-apply retention labels (by applications).
- Auto-apply sensitivity labels (by applications).

Later we cover how to create custom sensitive information types. You can use custom sensitive information types in the same way as standard sensitive information types.

**A simple number is never enough:** A sixteen-digit number is not by itself enough evidence to prove that it belongs to a credit card. DLP applies several tests to ensure that data contained in message bodies or attachments are sensitive. In the case of credit cards, the first check compares the number with the Luhn algorithm to ensure that it matches the rules set down for credit card numbers. DLP looks for further evidence such as the presence of a card expiry date (like 11/26), a name, or a CVC/CVV (card verification value). If these elements exist near the credit card number, it lends weight to the argument that the data matches the policy, and enough confidence exists to invoke the policy actions. The same is true of nine-digit U.S. social security numbers that do not have a checksum. If DLP applied a test of just looking for nine-digit numbers, the likely result is that many false positives would be signaled. In this case, the extra evidence is the presence of the word "SSN" in conjunction with nine-digit numbers. Different validation rules are employed to process the various sensitive information types so the thing to remember is that DLP



must be reasonably confident that a violation exists before it will signal a policy tip to the user or block a message.

## Bypassing DLP Checks

It's important to recognize that scanning based on sensitive information types will only detect instances where data matches the patterns defining the data type. If someone wants to get around DLP checking, they can by disguising text to avoid the pattern being detected. For instance, if you use a DLP policy to block email going outside the organization if messages contain credit card numbers, people can get around the block by spelling out the numbers (for example, use "six" instead of 6) or by including a picture of a credit card number. Messages with these variants of credit card numbers will be passed by the policy because the detection routines don't examine graphic data or look for blocks of text that might be a credit card number. In other words, DLP policies can catch most occurrences when people misuse sensitive data, but DLP checks will not detect disguised data.

## Microsoft Purview DLP Policies

To understand how Microsoft Purview DLP policies work across multiple workloads, we'll go through the process of setting up a new policy and see what happens when the policy is operational. First, some broad concepts:

- A DLP policy is composed of rules, locations, and actions. The locations are where the DLP service checks for policy violations. These can be Exchange Online mailboxes, SharePoint Online sites, OneDrive for Business accounts, and Teams chats and conversations.
- Each rule defines conditions that the DLP service applies against data to decide if any violations exist. DLP policies don't need many rules to be effective. For example, the [DLP policy to prevent users from sharing Teams meeting recordings outside the organization](#) uses one simple rule.
- The rules tell the DLP service what actions to take if it detects a violation. For instance, a rule violation could stop users from sharing a document.
- A rule can allow a user to override an action.
- Rules can display policy tips to users to help them understand why it is a problem if they share some information.
- Rules can also generate incident reports to inform compliance administrators when violations occur.

Now that we understand the basics, let's discuss some differences which exist when the DLP service processes different forms of content.

## Checking Documents for Sensitive Data

Despite the obvious differences that exist between processing email and documents, the same need exists to find a single point in the system to apply policies reliably. Using transport rules for Exchange DLP policies guarantees that the transport system will apply policies to all messages as they pass through the transport pipeline. Microsoft Purview DLP policies do not use Exchange transport rules. Instead, DLP policy checking occurs during the indexing of Exchange messages. This happens quickly enough to ensure that violations are detected and stopped before messages leave the organization.

A similar natural chokepoint does not exist for SharePoint Online, so Microsoft needed to take a different approach to offer protection for SharePoint and OneDrive for Business libraries. To detect DLP violations in documents, SharePoint Online uses a mixture of real-time policy evaluation together with a special "crawler" process and a background timer job to find documents containing sensitive data. Microsoft refers to this as a mixture of synchronous and asynchronous checking that applies equally well to new documents as users upload files to document libraries and to existing documents as users update their content.

Crawlers create content indexes by scanning all the document libraries in sites to look for new or changed material to index. The DLP crawler looks for sensitive information types in new or changed documents and can detect them soon after an update occurs (as fast as indexing occurs). However, background processes are subject to throttling and checking might not happen if servers are under a heavy processing load. When system load reduces, the crawler detects any lingering violations as it processes items for content indexing. Excel posed a difficulty in terms of how it stored data in spreadsheets and Microsoft had to create a special file handler to ensure that DLP could detect sensitive information types in these files.

Synchronous or real-time evaluation does not mean the kind of in-place checking that occurs in Outlook when the body of messages is examined for matches against applicable DLP policies as the user updates the message body. Nor is it like the OWA implementation where message content is sent to the server for examination periodically. Real-time checking for SharePoint and OneDrive documents means that the content is examined when documents are added, changed, or shared.

When SharePoint detects a violation, it applies the rules specified in the DLP policy. If the policy rules block access to items containing restricted data outside the organization, SharePoint and OneDrive for Business incorporate the check into the file-sharing dialog so that users cannot commit DLP violations (Figure 19-1). The idea is that prevention at the outset is better than fixing a problem after a violation occurs. If the policy calls for the generation of an incident report, this is when the DLP service creates and sends it.

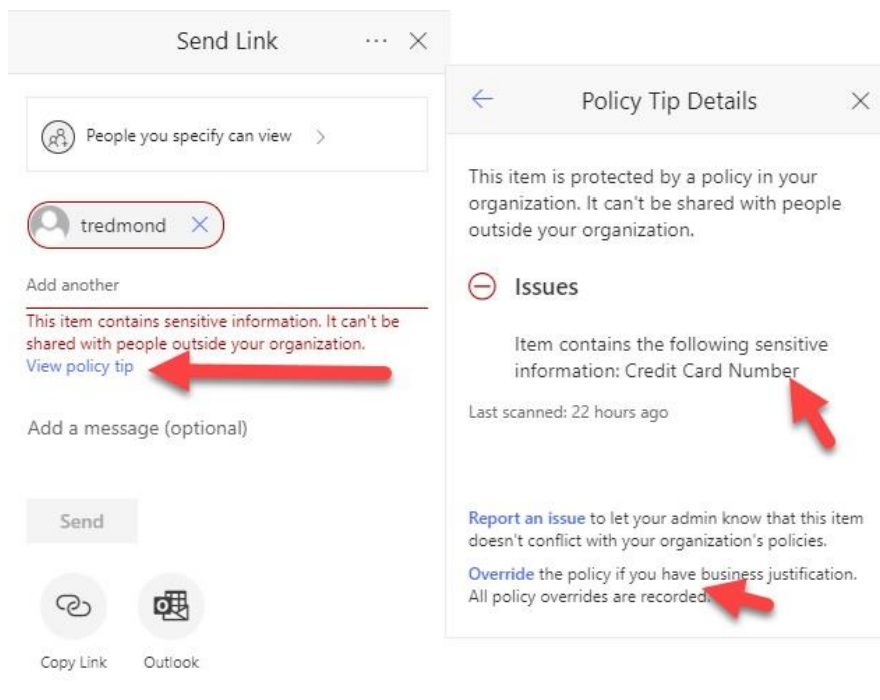


Figure 19-1: DLP blocks a user from sharing a file

Detection of a violation can also cause the display of a policy tip defined in the policy, but only in applications that support the necessary user interface, like SharePoint's browser interface. Policies might allow users to override the blocks imposed on documents holding sensitive information. If so, users can do this through the browser by opening the document properties, viewing the policy tip, and clicking **Override**.

SharePoint does not generate and send incident reports when the crawler detects a violation in content that existed before a rule became active. If SharePoint generated incident reports when it scanned content after an administrator adds a new rule or updates an existing rule, it is easy to imagine how the results might be a mail storm of thousands of incident reports. In a practical sense, it would be impossible to go through all the incident reports and resolve all the newly-detected violations. DLP still enforces any blocking actions contained in the rules; suppression occurs only for the incident reports.

## Sensitive by Default

Because DLP policies are not processed immediately for new files uploaded to SharePoint Online and OneDrive for Business libraries, a short period exists when users can share sensitive data outside the tenant and cause data leakage. If this is a problem, you can configure SharePoint Online to treat all new files as sensitive by default until they have been processed by a DLP policy. To do this, you:

- Run the *Set-SPOTenant* cmdlet in the SharePoint Online PowerShell module to block access to new files. It can take up to 15 minutes before the change is effective. The block applies to all sites in the tenant and you can't exclude sites from its effect.

```
[PS] C:\> Set-SPOTenant -MarkNewFilesSensitiveByDefault BlockExternalSharing
```

- Implement at least one DLP policy to scan all the SharePoint sites in the tenant.

With the block in place, users can still share documents with external people (if not blocked by the sharing settings for the tenant), but those people will be unable to access the content until the document is processed by DLP. The result of the processing will either pass the document for external access (because DLP doesn't detect a policy violation) or block it (because DLP detects some content that violates the policy if shared externally).

To revert the block, run *Set-SPOTenant* to allow sharing without waiting for DLP processing:

```
[PS] C:\> Set-SPOTenant -MarkNewFilesSensitiveByDefault AllowExternalSharing
```

Sensitive by default is an effective way to stop external sharing from SharePoint Online until DLP processing is done (Microsoft is working on its implementation for OneDrive for Business). However, it's a broad-brush policy that covers all sites in a tenant. Applying sensitivity labels to restrict access to documents containing important information might be a better approach, especially when auto-label policies are used to find and apply labels to documents at rest.

## Policy Tips

Overriding a policy tip through the **Override the policy** link allows a user to go ahead and share a document when a DLP policy would otherwise block this action. SharePoint records user overrides in the audit log in a "DLPRuleUndo" audit entry. The *ExceptionInfo* section of the audit entry captures the override text:

```
{
  "Reason": "Override",
  "Rules": [
    "413957f5-c3ab-4765-9322-33d2983dabfe"
  ],
  "Justification": "This document is authorized for sharing!"
}
```

**Mobile Policy Tips:** The OneDrive for Business mobile apps for Android and iOS support DLP policy tips. The clients download the DLP policy from Microsoft Purview when it connects, and the application is then able to scan text input into documents for sensitive information types and, if detected, display the correct prompt. The Teams mobile clients also support DLP policies in that they report violations in the activity feed, show when messages are blocked, and allow authors to override a block if allowed by policy.

## Teams DLP Policy Processing

Teams clients don't scan for DLP violations as users enter text or when items are indexed. Instead, Teams checks messages against DLP policies after users send messages and clients submit those messages to the Chat service. Both personal chats and channel conversations (including private channels) are monitored. Teams DLP policies only work for users with an Exchange Online mailbox.

Teams DLP policies can be scoped to cover all users or to include or exclude up to 1,000 selected accounts. Instead of specifying individual user accounts, distribution lists or mail-enabled security groups can be used to define to whom Teams applies a policy. Using distribution lists or security groups is an excellent way to set the scope for policies that apply on a country-wide or department-level basis. For example, a policy can be applied to Japanese users to prevent the sharing of Japanese identity card numbers. A background process checks for membership changes to relevant distribution lists and security groups to ensure that policies include new members and stop applying to accounts removed from membership. It can take up to 30 minutes after the creation of a new distribution list or security group before it can be used in a policy.

**Note:** Because all Teams messages pass through a message aggregator, DLP policies applied to the host team also apply to messages posted to conversations in Teams private and shared channels and the documents stored in the SharePoint sites belonging to the private channels. However, because private channels don't have an identity within the tenant, you can't add a private or shared channel as an explicit inclusion or exclusion in a Teams DLP policy.

Teams messages remain inside the company, so the division between internal and external users is harder to make than it is for SharePoint or OneDrive, where it's obvious when a document is shared with an internal or external account. When you create a DLP policy for Teams, remember that guest users and federated external users (including those who join meetings) are considered external to the organization; tenant users are inside. Because of this separation, it might be easiest to create separate DLP policies for Teams for external and internal users. In this respect, Teams is different from Exchange and SharePoint. DLP policies applied to email and documents usually focus on external sharing rather than internal communications.

## Checking Teams Messages

As messages flow through the chat service, Teams compares message content against the criteria defined in DLP policies covering Teams locations. The normal matching procedure is used to identify sensitive data specified in the policies by comparing values like credit card numbers together with accompanying information. If a violation is detected, Teams generates a block signal to a "Rethook" (a mechanism used by Teams to flag actions to clients) to instruct clients to block the message. After they receive the signal, the client used by the person who sent the blocked message tells them that the block is in place (Figure 19-2) and, if the policy allows, offers them the chance to override the block (through the *What can I do* link). Users who received a blocked message might see its content momentarily but once the block is in place, the message is hidden, and they can't see it any longer.

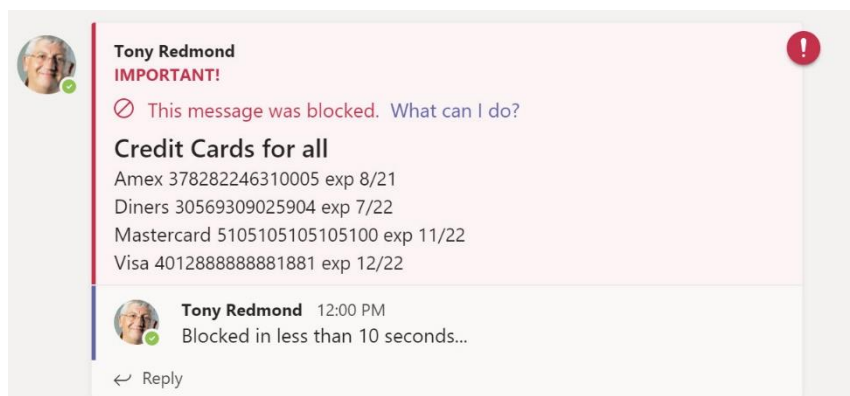


Figure 19-2: Teams blocks a message containing credit card numbers

Users receive a notification in their Activity Feed when Teams blocks one of their messages. In addition, if the user is logged into Teams with a mobile device, they'll receive a notification on the device that their message is blocked.

**Teams DLP Policy Recommendation:** Organizations that use Teams but don't have any Teams DLP policies see a Microsoft recommendation in the Microsoft Purview compliance portal that a policy should

exist to protect Teams messaging. Accepting the recommendation creates a pre-packaged Teams DLP policy to protect against the sharing of financial data, passport and social security numbers, and credit card numbers. You don't have to accept the policy settings as defined and can edit them before making the policy effective. Remember that you might need extra licenses to use Teams DLP policies.

## DLP Audit Records

When a DLP policy blocks an item, it records that fact in an audit entry in the audit log. You can see these entries in the audit log (search for the *DLPRuleMatch* operation). Alternatively, you can view DLP events in the report viewer. Click the *DLP policy matches* widget in the dashboard to be brought to the report viewer and then click *View details* table to see information about the DLP matches for all services covered by DLP policies

DLP events are also accessible via the [Office 365 Management Activity API](#), a REST-based API covering audit information gathered from Exchange Online, SharePoint Online, OneDrive for Business, and Azure AD that's intended for ISVs to develop analysis and reporting products. Two types of DLP events are logged. One is a non-sensitive event and contains data such as the document or email that triggered a violation, the user, the policy rule, actions taken, and the type of sensitive data involved. The other is a sensitive event and returns all the data for non-sensitive events plus the value of the data, such as a credit card number.

## Default DLP Policy

It is always good to have a building block to start with when approaching the introduction of new technology. To help tenants start with DLP, Microsoft includes a default DLP policy to protect credit card information for tenants where DLP is not already in use. Credit cards are the most common data protected by active DLP policies, so they are a natural choice for companies to start their DLP journey. If your tenant has a default DLP policy, you must edit and enable it to make the policy active.

## Creating Microsoft Purview DLP Policies

After becoming used to DLP policy structure and operation, the next step is to create a policy tailored for your organization. You can edit the default DLP policy or create a new one. To create a new DLP policy, go to the Data loss prevention section of the Microsoft Purview Compliance portal and click **Create policy**. Figure 19-3 shows a set of policies for a tenant, one of which is disabled. Note the offer to update the policies to extend coverage to Teams. Do not accept this choice unless you have the necessary Office 365 E5 licenses.

**Data loss prevention** Show in navigation

Overview Policies Alerts Endpoint DLP settings Activity explorer

Use data loss prevention (DLP) policies to help identify and protect your organization's sensitive info. For example you can set up policies to help make sure information in email and docs isn't shared with the wrong people. [Learn more about DLP](#)

+ Create policy ↓ Export Refresh 8 items Search

Name	Order	Last modified	Status	Update options
Teams DLP Credit and Debit Cards	0	Jan 24, 2021 6:46 PM	Enabled	
Block Anyone Links	1	Feb 16, 2021 7:14 PM	Disabled	Extend to Teams
U.S. State Social Security Number Confidentiality Laws	2	Jan 23, 2021 10:01 PM	Enabled	
Block Sharing of Ultra Confidential Files	3	Nov 11, 2020 1:55 PM	Enabled	Extend to Teams
Ireland Personal Info Policy	4	Dec 30, 2020 12:11 PM	Enabled	
Teams Channel Financial Data	5	Jan 24, 2021 7:18 PM	Enabled	
Stop people sending Microsoft 365 passwords	6	Feb 22, 2021 9:38 AM	Enabled	
Exchange Encrypted Email	7	Apr 5, 2021 10:35 PM	Enabled	Extend to Teams

Figure 19-3: DLP policies

Before you go ahead and create a new policy, it is worthwhile to chart out what kind of sensitive data you want to protect and where is that data kept. Every DLP policy follows the same basic structure:

- **Locations:** What workloads to protect. You can choose from Exchange Online, SharePoint Online, OneDrive for Business, and Teams. Yammer is covered if the tenant network operates in Microsoft 365 mode. DLP policies do not currently cover Planner but do cover the SharePoint content associated with this app. You can scope the policy to include or exclude specific mailboxes, sites, or teams. The account pickers for Exchange Online, Teams, and OneDrive for Business support distribution lists and security groups to make it easier to apply policies to large sets of people without having to select each account individually. DLP evaluates distribution lists periodically to detect membership changes.
- **Optional Locations:** DLP is spreading its ability to analyze content to detect violations in different places. For example, you can [use DLP to monitor non-Microsoft apps like DropBox via Microsoft Defender for Cloud Apps](#), [Windows 10 devices via EndPoint DLP](#), or [scan on-premises resources](#). Microsoft deems these to be advanced capabilities, so you'll need appropriate licenses like Microsoft 365 E5 Compliance.
- **Rules:** What kind of protection do you want to achieve. Rules specify what kind of sensitive data you want to protect, conditions that say what must happen before a violation occurs, and actions that the policy takes afterward. Actions include user notifications and the creation of incident reports.

## Basic DLP Policy Settings

You can then start to flesh out the policy by creating the different sections of the policy. The policy creation wizard is divided into these parts:

- **Choose the information to protect:** To make it easier for customers to create DLP policies, Microsoft Purview includes a large number of DLP templates, each of which defines a set of sensitive information types that organizations commonly need to protect in certain fields of activity. For instance, if you want to avoid the loss of personal data, you can select **Privacy** to see templates such as "U.S. Personally Identifiable Information (PII) Data" or "U.K. Privacy and Electronic Communications.". When you select a template as the basis for a policy, the wizard automatically creates the necessary DLP rules to protect the information defined by the sensitive information types defined in the template. If none of the standard sets work for you, select **Custom**, which allows you to construct your own rule set.
- **Name your policy:** If you select a template, the policy name inherits the template name. You can overwrite the name with your own choice. You can also enter a description for the policy. Some organizations used this to refer to internal documentation for the policy, including information about who created the policy.
- **Choose Locations:** The policy wizard displays a set of supported locations with a slider for each location. If you leave the slider on for a location, the policy applies everywhere in the location (for instance, every SharePoint Online site in the tenant). Alternatively, you can select to include or exclude specific:
  - Exchange mailboxes. Specify all mailboxes or select mailboxes to include or exclude. You can use distribution lists to include or exclude sets of mailboxes covered by a policy, but you can also add or exclude individual mailboxes by typing their names into the search field. Dynamic distribution lists can scope DLP policies for mailboxes based on attributes in their Azure AD accounts.
  - SharePoint sites. Specify all sites or enter the URL for the sites that you want to protect (up to a maximum of one hundred sites) or choose to include or exclude specific sites. If you need to include or exclude more than a hundred individually-specified sites in a policy, consider using an org-wide policy.
  - OneDrive for Business accounts. Specify All or select mailboxes using distribution lists or individual mailboxes for inclusion or exclusion. The limit of a hundred accounts also applies.

- Teams. Choose all to cover all teams or select individual accounts by name or using a distribution list to include or exclude in a policy.
- **Policy settings:** Each policy contains at least one rule used to match items with potential violations. If you create a policy from a template, you can accept the default processing settings for the rules copied along with the template. For instance, a rule might be set up to block content when shared outside the organization. If the default rule settings don't do what you want, you can customize the rules or create new rules from scratch. A policy created from a template is likely to have at least two rules, each of which should be checked to ensure that the settings in one rule cannot interfere with settings in other rules. We will discuss rule settings in more detail later.
- **Enable policy:** Define whether the policy is on or off, or if it is in test mode.
- **Review your settings:** Before you save a policy, you can review its settings to ensure that everything is as you expect. Click **Create** to continue.

When you create a new DLP policy or update an existing policy, Microsoft Purview checks the policy to ensure that it is valid. For example, you might have input the URL for a subsite instead of a site. If this happens, the error is flagged, and you must fix the problem before you can save the policy. Once everything is satisfactory, Microsoft Purview publishes the new policy to the various workloads to enact the policy. A certain amount of complexity is involved to publish a policy to all the workloads as multiple servers are involved. Microsoft's SLA is for a new or updated policy to be effective within an hour of publication, but the experience of many is that this SLA is often breached and that you might have to wait up to a day before everything settles down and new DLP rules are active.

A further delay occurs before workloads begin to detect violations. SharePoint relies on the crawler within the content indexing process to detect sensitive data within documents. SharePoint does not index content immediately after an item is updated and (depending on system load) the crawler might take between ten minutes to an hour before it processes an updated document. It can take up to an hour before a new DLP policy for Teams or an update to a DLP policy for Teams is effective. The exact time varies depending on the current load on the Teams infrastructure.

## Resolving Multiple Potential Violations

Large organizations or those with complicated DLP requirements will likely use several DLP policies. This creates the possibility that an item will match several policies. In this situation, the following applies:

- DLP policies have a priority order for evaluation from 0 (most important) downwards. Policies are evaluated in this order.
- Rules within DLP policies also have a priority order.
- DLP rules can be set to stop the processing of other rules when a match is detected.
- In general, DLP applies the most stringent policy on the basis that it is better to be safe rather than sorry.
- The exception is where a rule in a high-priority policy stops the processing of other rules. In this case, the violation flagged by this rule is used and any other potential violation which might be detected by rules in lower-priority policies is ignored.

With this guidance in mind, organizations should:

- Order DLP policies so that the most important policy is assigned the highest priority.
- Order rules within DLP policies so that the most important rule is assessed first.

## Default Rule Settings

If you don't customize the rules in a DLP policy, the rules inherit some default settings, including:

- **Who can receive sensitive content:** DLP can apply the policy inside or outside the organization. People with email addresses belonging to a domain registered for the tenant in the Microsoft 365 admin center are internal; those with addresses belonging to other domains are considered external.
- **How DLP protects content:**
  - **Display policy tips:** It is usual to display policy tips to users to tell them that a problem exists with a document.
  - **How many instances of sensitive data exist in an item before DLP flags a violation:** The default for a rule is that 10 instances of the same sensitive information type (for example, a credit card number) must be present in a file or message before a violation occurs. You can reduce or increase the number of matches to change how the rule works. If you want to make a rule easier to match, reduce the instance count (for example, from 10 to 7). Likewise, to make it harder to match, increase the count.
  - **Who receives the incident report and what is in the incident report:** Typically, incident reports go to a DLP administrator, who checks the incident to assess whether it is valid and if a violation is present, decides what action to take to resolve the violation.
  - **Restrict or encrypt access:** Should sharing the information in a detected item be blocked or encrypted? Encryption is only supported for Exchange Online messages. For SharePoint and OneDrive content, you can restrict access to content after detecting a DLP violation. As shown in Figure 19-4, the normal situation is to block external access to the content while allowing internal users continued access. The options are:
    - Block everyone from accessing the content.
    - Block people outside the organization. Anyone with a tenant account can access the content, including guest users.
    - Block people who can access the content through an Anyone link.

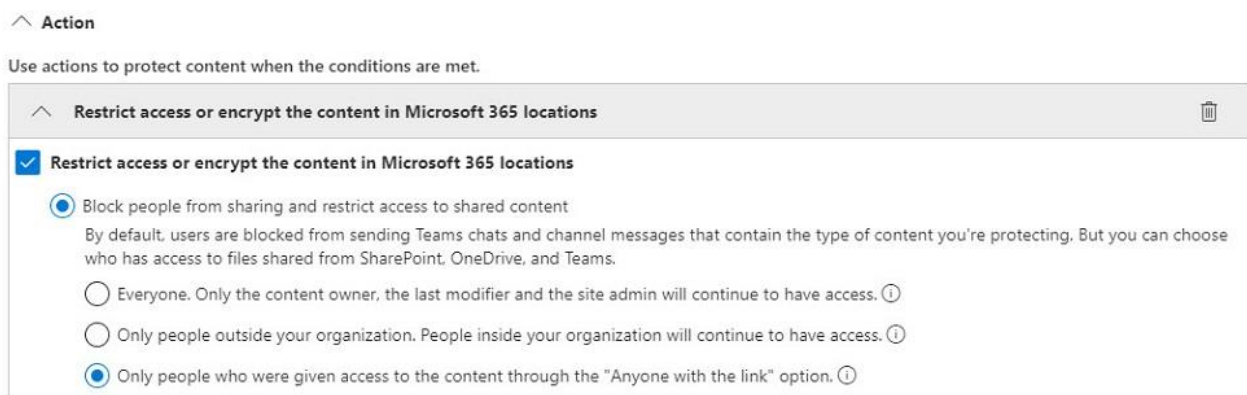


Figure 19-4: Customizing access to content setting for a rule in a DLP policy

When you opt to block external users from seeing sensitive content that violates a DLP policy, the SharePoint and OneDrive for Business sharing dialog includes a warning to users if they try to share sensitive content outside the tenant. The warning appears in real-time and helps to educate users about the need to protect sensitive content.

## Advanced Rule Settings

To have full control over the rule settings, select **Create or customize advanced DLP rules** when creating a new policy. You can then use the full rules editor to edit rules copied from the template (or created from scratch) to meet your exact needs. The settings for rules are:

- **Rule name:** Free-text name assigned to the rule.
- **Description:** A brief description of what the rule does.



- **Conditions:** The type of data present in the content and the context surrounding the data (for example, the number of instances of some sensitive data) to cause the rule to fire. For example, a rule might ignore content where a single credit card number is present but fire when it finds more than three instances of credit card numbers. Figure 19-5 shows where to change the number of detected sensitive data instances for a policy to check (in this case, at least one credit card number must be present in an item to fire the rule).  
DLP conditions can include aspects of documents or messages that must be present for the rule to fire. These conditions might only apply to a single kind of data. For example, a rule might check for email messages sent from a specific IP address or email with attachments that cannot be scanned or can only partially be scanned, or a rule could look for attachments of a specific type or those that are password protected. Another example of a condition is when you use a document property to identify documents that come within the scope of a policy. Any managed property supported by SharePoint search can be used in this manner. Finally, if sensitivity labels exist in the tenant, a DLP rule can include a condition that a certain label must be present for the rule to fire.
- **Exceptions:** Define whether any conditions exist that will cause the rule not to fire. The exceptions include if the content has a specific form of sensitive data (or label), and the sharing of content inside or outside the organization. Exchange Online policies can focus on a characteristic of email (IP address, attachment type, and recipient domain).
- **Actions:** What happens when a violation occurs. The action is to restrict access to the content by removing permissions for everyone except the primary site administrator (or global administrator), site owner, and the person who last changed the item. Microsoft Purview automatically restores the original permissions when the owner or person who last changed the document acts to remove the offending content and bring the document back into compliance. Encryption is available to protect email with sensitive content, but only for policies scoped to apply to Exchange content. See the note below.
- **User notifications:** DLP supports two forms of notifications. First, it can send email notifications to the site owner (or other designated users) and whoever last changed the item to inform them that a problem exists and why (for instance, someone shares the item with external people and includes a debit card number). A link in the message allows the recipient to open the document and fix the problem detected by the policy. The second method is to flag the problem visually with a policy tip. You can customize the text of the policy tip to add organization-specific advice to users.
- **User overrides:** The rule can allow users to override the block because of a false positive. In other words, the user (who knows the content and can put it into business context) says that the rule is incorrect to flag an item for a policy violation. When this happens, DLP removes the block and records the text given by the user to justify the override for auditing purposes.
- **Incident reports:** You can arrange for email notifications to go to one or more people when a policy detects problems. These notifications are known as incident reports. Each report has enough information about the problem to allow an administrator or compliance officer to understand whether a problem exists and if so, whether they need to take further action.

You can also set rule priority to have DLP execute the rules in a certain order and terminate processing after a rule matches conditions and is executed.

**Edit rule**

## ^ Conditions

We'll apply this policy to content that matches these conditions.

^ Content contains
🗑️

Any of these ▾ 🗑️

**Sensitive info types**

Credit Card Number	High confidence ▾ ⓘ	Instance count	1	to	9	🗑️
U.S. Bank Account Number	Medium confidence ▾ ⓘ	Instance count	1	to	9	🗑️
ABA Routing Number	Medium confidence ▾ ⓘ	Instance count	1	to	9	🗑️

Add ▾

Create group

**AND**

^ Content is shared from Microsoft 365
🗑️

Detects when content is sent in email message, Teams chat or channel message, or shared in a SharePoint or OneDrive document.

with people outside my organization ▾

Save
Cancel

Figure 19-5: Defining settings for an individual rule in a DLP policy

When you finish working with the settings for a rule, click **Save** to return to the rule set so that you can edit the next rule. When all the rules are to your liking, click **Next** to return to the policy. Remember that if you make a mistake or want to change something, you can always edit the policy and then update settings for individual rules or the policy.

**Email message type exceptions:** DLP supports exceptions based on the message type. These are used in cases where you might not want to apply rules to read receipts or automatic replies. Care should be taken with the three types which cover encrypted email. Encrypted means S/MIME encrypted; Signed means S/MIME signed; and Permission Controlled means that a sensitivity label or OME template protects the message. The latter is the type usually required to process Exchange Online messages.

## Confidence Levels and Match Accuracy

DLP policies match data found in documents and emails against sensitive information types by looking for patterns and other forms of evidence. The primary element (like a regular expression) in a pattern for a sensitive information type defines the basic match for that type. Secondary elements (like keyword lists) add evidence to confirm that a match truly is an example of a sensitive information type. The more evidence is gathered, the higher the match accuracy and the more confidence exists that the data found is an instance of the sensitive information type. Microsoft refers to the match accuracy as the confidence level. Policy rules use the confidence level to decide whether action is necessary when matches are found.

The higher the percentage confidence level specified by a rule, the more evidence DLP must find before it can match a rule. Three confidence levels are used:

- Low (65% match accuracy).
- Medium (75% match accuracy).
- High (85% match accuracy).

For instance, to meet the bar set for a low confidence match for a credit card number, text must be in the right format (sixteen digits) and pass [Luhn's algorithm](#) to establish that the number could be a credit card. To

increase the confidence that the number is a credit card to 85% confidence, the rule looks for other evidence like a keyword (such as "Visa" or "Amex") or an expiry date in the MM/YY format.

If a policy includes multiple rules that check for the same sensitive information type, the suggestion is to:

- Set the minimum and maximum confidence level for the rule that takes the least aggressive action when a low-confidence match is found.
- Set the confidence level range for the rule that takes more aggressive action when a match is detected with higher confidence.

The idea is that content with lower confidence matches receives an action with less impact on the user (like displaying a policy tip) while content with higher confidence matches gets a more restrictive action (like being blocked).

## Rule Priority

Each rule in a policy receives a priority number in the order in which it is created. When DLP evaluates content against a policy, it processes the rules in priority order; if the content matches multiple rules, DLP applies the most restrictive action. For example, if rule 1 displays a policy tip to users and rule 6 restricts access to the content, DLP applies the settings from rule 6. In addition, the policy tip for the most restrictive rule is displayed to the user.

## Protecting Sensitive Email with Encryption

If DLP detects outbound emails containing sensitive content, you can have Exchange apply encryption to the messages before they leave your tenant. You can choose to apply any of the Microsoft Information Protection templates (not sensitivity labels, because a label might not include encryption) available in the tenant to messages, including the default Encrypt Only and Do Not Forward templates. See Chapter 20 for more information about how to create and use Microsoft Information Protection to protect documents and email.

If you choose to apply encryption to messages in a DLP policy, that policy can only cover Exchange. You can't apply encryption to documents through a DLP policy.

## GDPR Support

Tenants that come within the scope of the European Union's General Data Protection Regulation (GDPR) need to meet specific requirements for the protection and use of personal data. To help, Microsoft Purview includes six common sensitive information types for use in DLP policies or to classify sensitive data. The information types include:

- EU Debit Card Number.
- EU Driver's License Number.
- EU National Identification Number.
- EU Passport Number.
- EU Social Security Number (SSN) or equivalent ID.
- EU Tax Identification Number (TIN).

In addition, Purview includes many country-specific sensitive information types like the Estonia Driver's License Number.

Microsoft Purview includes a General Data Protection Regulation DLP template. If you create a DLP policy using this template, DLP adds the six EU-wide sensitive information types listed above to the policy. You can edit the policy to include other sensitive information types or include the EU data types in other policies.

An example of the difficulties involved in detecting content using sensitive information types happened when emails imported into SharePoint Online caused DLP to flag violations for the EU TIN sensitive information type because eight- and nine-digit numbers in the email headers resembled tax identification numbers. Microsoft

resolved the problem by adjusting the data definition to look for more proof that a number was a TIN before flagging a violation.

## Sample Test Data

When you start working with DLP, the issue of how to generate good test data to use to check policies always arises. The [DLPtest site](#) offers several different sets of sample data, including social security numbers and credit card numbers, that you can use for testing.

## PowerShell support for DLP policies

DLP policies use several cmdlets (included in the compliance module) to work with policies and rules. The cmdlets include:

*New/Get/Set/Remove-DlpCompliancePolicy*: Create, access, manipulate, and delete DLP policies. For example:

```
[PS] C:\> Get-DlpCompliancePolicy -Identity 'Confidential Patent Information DLP Policy'
```

*New/Get/Set/Remove-DlpComplianceRule*: Create, access, manipulate, and remove the rules used in DLP policies. For example, here is how to check the settings for a rule (edited output shown):

```
[PS] C:\> Get-DlpComplianceRule -Policy 'Check for SSN Data' | Format-List

ParentPolicyName           : Credit Card data check
ContentContainsSensitiveInformation: {System.Collections.Hashtable, System.Collections.Hashtable, System.Collections.Hashtable}
AccessScope                 : NotInOrganization
ContentPropertyContainsWords : {}
Workload                   : Exchange, SharePoint, OneDriveForBusiness
BlockAccess                 : True
GenerateIncidentReport     : {DLPIncidents@Office365ITPros.com}
IncidentReportContent      : All
NotifyUser                  : {SiteAdmin, LastModifier, Owner, Tony.Redmond@office365itpros.com}
NotifyAllowOverride        : WithJustification
NotifyEmailCustomText      : You can't keep this kind of information in documents!
NotifyPolicyTipCustomText  : Whoops - bad SSN data found here
ReadOnly                    : False
Priority                     : 0
Comment                     : A rule to detect SSN data in documents
Disabled                    : False
CreatedBy                   : Tony Redmond
```

We can tell that this rule:

- Is associated with the DLP policy "Credit Card data check".
- Looks for three types of sensitive data (the hashtable references in the *ContentContainsSensitiveInformation* property – the names are not shown).
- Checks information shared with users outside the tenant (scope is "NotInOrganization").
- Scans for three supported workloads (Exchange, SharePoint Online, OneDrive for Business).
- Blocks access to content deemed to contain sensitive data.
- Generates a DLP incident report.
- Generates an email notification to the site administrator, the last user who modified the content, its owner, and a nominated other user.
- Allows the policy tip to be overridden with a business justification.
- Is active (*Disabled* is False).

Given that DLP policies can be reasonably complex to set up, it is best to manage DLP policies and rules through the Microsoft Purview Compliance portal and only use PowerShell to check individual objects as needed.

# Endpoint DLP

[Microsoft Purview Endpoint DLP](#) is a solution that uses signals generated by actions performed on Windows 10 workstations to evaluate against DLP policies. Supported actions include copying files to removable media like a USB or to a network share, printing files, uploading to a cloud app, or copying data to the clipboard. The code necessary to detect actions and submit them for evaluation is incorporated into Windows 10 (version 1809 or later) and the Edge or Chrome browsers. No additional agent is necessary to monitor activity on a workstation.

Before you can use Endpoint DLP, you need Microsoft 365 E5 licenses or either the Microsoft 365 E5 information protection and governance or compliance add-ons. Workstations used by licensed accounts can be onboarded (enabled) through the Microsoft Purview compliance portal to start the flow of signals for DLP evaluation, unless administrators have already enrolled the devices for Windows Defender, in which case Endpoint DLP works without any further configuration.

Once a workstation is enabled, actions taken by the user are monitored for potential violations against policy using the same kind of conditions as used to monitor Office 365 activity. For example, attempts to upload documents containing credit card numbers can be detected and stopped. Supported file formats include Office documents, PDF, text, and source code.

Endpoint DLP settings for the organization can be adjusted in the Microsoft Purview compliance portal to reduce the amount of noise in signals by excluding certain folders like the recycle bin, temp folder, or folders used for non-work files. It's also possible to allow uploads to specific cloud services without generating a violation. Policy thresholds can be set to generate alerts when multiple events of the same type happen over a short period. For instance, a policy could alert administrators if someone prints more than twenty documents assigned the Confidential sensitivity label.

When Endpoint DLP is available in a tenant, DLP policies can cover a target location called Devices, just like choosing SharePoint or Exchange as policy locations. The normal approach is to separate device policies from those used with Office 365 workloads, but you can combine them. Device policies have separate settings for restrictions to enforce when policies match conditions (Figure 19-6).

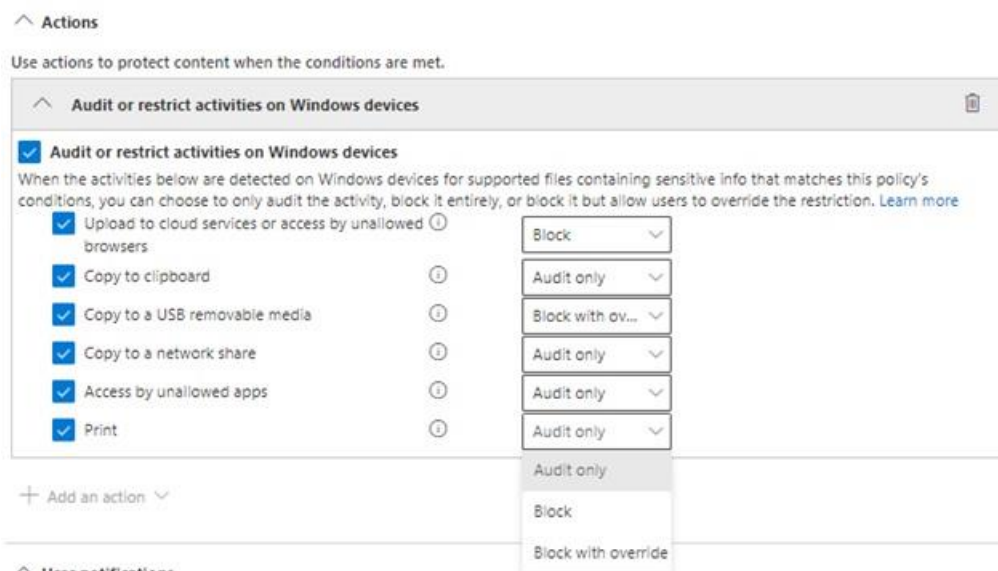


Figure 19-6: Endpoint settings in a DLP policy

Edge is the preferred browser because it understands how to respect endpoint DLP policies, and you can block other browsers from accessing files protected by policies. For instance, you could block Firefox from opening a Word document if a specific retention label is present.

Apart from being used by DLP, the signals generated by devices can be gathered and analyzed in a SIEM. An example using Azure Sentinel is [described in this article](#).

# Creating Custom Sensitive Information Types

Although a large set of standard sensitive information types exists, organizations often have specific ideas about the characteristics of confidential data which are not satisfied by the set of standard sensitive information types. You can create your own custom sensitive information types through the **Data Classification** section of the Microsoft Purview Compliance portal. Custom sensitive information types can be used in DLP policies, retention policies, and auto-label policies. They can be created from scratch or by copying one of the standard types and altering the copy to meet requirements.

The basic need for any sensitive information type is a pattern to detect matches in messages and documents. The simplest patterns have a primary element to instruct how to detect a specific type of sensitive information and a confidence level when a match occurs. A pattern with just a primary element might define its confidence level as medium for any match, while patterns with supporting elements that add more evidence that a match exists might increase the confidence level as supporting elements are found near the matched element. The following methods can be used as a primary element:

- **Regular expression:** A Regex to detect certain patterns of data. Sites like [regex101.com](https://www.regex101.com) are useful to help form regular expressions. For example, a regular expression of `\s[A-Z]{3}-[0-9]{4}\s` can be used to match project numbers like "PRJ-1384" or "PRO-1847".
- **Keyword list or Keyword dictionary:** Both ways to define keywords (like "Finance" or "Funding") that you want to use for matching. Typically, these methods are used as supporting elements rather than for primary matching.
- **Function:** Microsoft publishes a [set of functions](#) for different types of data, such as an Alabama driver's license. You can use these functions to create custom sensitive information types.

Supporting elements add confidence that matched data is what you are looking for and not just a random collection of letters and numbers. For example, as noted above, a social security number is described by a number in the format 999-99-9999 with added evidence coming from the keywords "SSN" or "Social Security" near (in terms of characters) the number. The same is true for credit cards, where the 16-digit number is confirmed by words such as "credit card," "Visa," "MasterCard," "expiry date," and so on.

## The Azure AD Password Sensitive Information Type

In general, it's bad practice to circulate passwords in messages, so let's see how we can stop people from sending Azure AD passwords in email and Teams messages. Some cases exist when sending passwords around is necessary, such as when an administrator resets a password for a user (a tenant can enable [self-service password reset](#) to let users reset their passwords). To detect Azure AD passwords, we need a regular expression to detect passwords matching Azure AD password policies. By default, Azure AD password requirements are:

- A minimum of 8 characters and a maximum of 256 characters.
- Requires three out of four of the following:
  - Lowercase characters.
  - Uppercase characters.
  - Numbers (0-9).
  - Symbols (@ # \$ % ^ & \* - \_ ! + = [ ] { } | \ : ' , . ? / ` ~ " ( ) ; and blank space)

Many suggested regular expressions to validate Azure AD passwords are shared on the internet. However, when you input regular expressions to use with DLP policies, Microsoft validates the code to ensure that it works and will perform well (the [rules are documented here](#)). Most of the suggested expressions do not meet

Microsoft's rules. Fortunately, MVP James Cussen [created an expression](#) that passes the test and matches passwords, which we can use to illustrate how to create a custom sensitive information type. Here's the expression:

```
((?=[\S]*?[A-Z])(?=[\S]*?[a-z])(?=[\S]*?\d)|(?=[\S]*?[A-Z])(?=[\S]*?[a-z])(?=[\S]*?[^\a-zA-Z0-9])|(?=[\S]*?[A-Z])(?=[\S]*?\d)(?=[\S]*?[^\a-zA-Z0-9])|(?=[\S]*?[a-z])(?=[\S]*?\d)(?=[\S]*?[^\a-zA-Z0-9]))[\s]{8,256}
```

The regular expression matches strings in the right format for Azure AD passwords. To provide additional evidence that a string is a password, we can add a supporting element in the form of a keyword list. Typically, when people send passwords in messages, they include some explanatory text. Therefore, if a keyword occurs close to the matched password, the likelihood increases that the term is a password. For instance, someone might send a message saying "Here's your new password: AzurePW123!@." The string AzurePW123!@ matches the regular expression and becomes the anchor for DLP to look for keywords within the window set for keyword proximity checks. The window extends to the left and the right of the match, which means that a match against "password" occurs because the keyword is within a few characters of the anchor. The confidence that a good match exists is therefore higher.

In Figure 19-7, a keyword list of terms that help to confirm passwords is visible. If an organization is multi-lingual, you should include terms from different languages as shown here. It is also a good idea to consider including some misspellings of terms in the keyword list, such as "passwrđ" or "passwrod" as there's no guarantee that people will spell everything correctly when they send messages.

## Edit keyword list

Keyword lists identify the words and phrases you want this info type to detect. For example, the keyword list to identify Netherlands VAT numbers is 'VAT number, vat no, vat number, VAT#'. [Learn how to create keyword lists](#)

☰ Choose from existing keyword lists

ID \*

Password components

Keyword group #1 \* ⓘ

Case insensitive

password, pwd, passcode, passwd, wachtwoord, mot de passe, Passwort, contraseña

Case sensitive

Enter keywords, separated by commas. Each keyword is limited to 50 characters, and exact casing is required to detect matches.

Word match  String match

Figure 19-7: Defining a keyword list

The components used to define the pattern for the new custom sensitive information type are:

- The primary element is a regular expression to detect the pattern of a valid Azure AD password.
- The supporting element is a keyword list of password terms.

- The confidence level is medium to start. This level is set when a match occurs against the primary element. It will increase to high confidence if DLP also detects a keyword from the keyword list in the supporting element.
- The character proximity defines how close one of the password terms must be to the primary element. Purview usually suggests a proximity window of 300 characters, but in this instance, it is better to adjust the proximity window to around 80 characters because the likelihood is that any mention of a supporting term will be very close to the password. If policies using the custom sensitive information type generate too many false positives, you can consider reducing the proximity to 50 or 60 to see if that reduces false positives. The option to check for anywhere in the document is not suitable because it is likely to result in many false positives.

Figure 19-8 shows the components of the new pattern during the creation phase. The primary element is saved with a name for easy reference later as is the keyword list. After defining the pattern, we set the recommended confidence level to look for in policies. The final step is to save the new type.

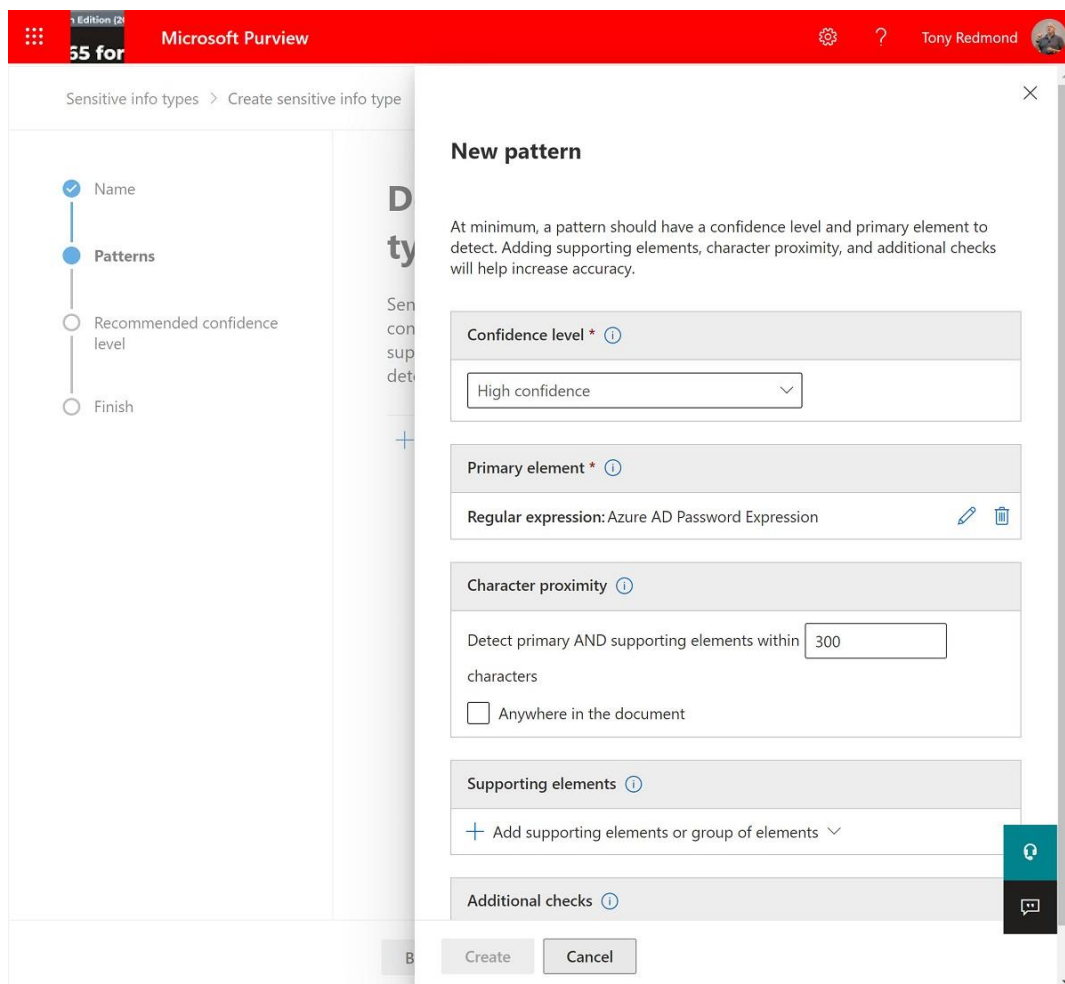


Figure 19-8: Defining a new custom sensitive information type

## Testing a Custom Sensitive Information Type

After creating a new custom sensitive information type, it is wise to test it before attempting to use the new type in a policy. Select the custom sensitive information type from the set available in the tenant. Test is one of the available options.

At least two text files are needed for testing. One should contain a set of values that you expect the sensitive information type to detect; the other contains values that it shouldn't. For a more comprehensive test, you can create a suite of text files to check various combinations against the new type.



Click **Test** and browse to select a test file that you've prepared. Microsoft Purview uploads the file and tests its contents against the custom sensitive information type in the same way that a comparison occurs in a DLP policy. You'll then see the results of the test and the matches (Figure 19-9). We can see that the test made three matches using the primary element for three passwords found in the test file (65% confidence). The level of confidence increased to 75% because some of the keywords listed in the secondary element were detected within the defined proximity window.

**Match results**

We have detected the following in *TestPwd.txt*

- Azure Active Directory password**  
65% Confidence - 3 matches

Matches	Supporting elements
TestStuff124298!	"password", "password", "pwd"
NewThingToTry56@	"pwd", "password"
GatheringDust12	"password", "pwd", "password"

- Azure Active Directory password**  
75% Confidence - 3 matches

Matches	Supporting elements
TestStuff124298!	"password", "password", "pwd"
NewThingToTry56@	"pwd", "password"
GatheringDust12	"password", "pwd", "password"

Back Finish Cancel

Figure 19-9: Testing a custom sensitive information type

After the new custom sensitive information type passes its tests, you can go ahead and deploy it like any of the out-of-the-box types. It's important to test the custom sensitive information type in a variety of circumstances before you put it into production. For example, you should consider if it is necessary to exclude some email addresses from the policy (the easiest way to do this is to create a distribution list and use it to hold the accounts allowed to send passwords). Another thing to consider is whether a policy that works well for Teams (which mostly deals with internal communication) is equally effective for emails when you can exert less control over what external people might include in messages. The policy described here will, for instance, block email invitations for Zoom online meetings because the links to the meetings contain "pwd" and the string identifying the meeting looks like an auto-generated password. Finally, it's a good idea to consider how best to explain the rationale for the policy to end users so that people know what to do if the policy blocks their messages, including how they can override the policy (if policy settings allow an override).

The new custom sensitive information type can be managed in PowerShell. For instance, here's how to list all the custom sensitive information types in the tenant. The example we just defined shows up together with other types created by document fingerprinting (covered later).

```
[PS] C:\> Get-DlpSensitiveInformationType | ? {$_.Publisher -ne "Microsoft Corporation"} | Format-Table Name, Publisher
```

Name	Publisher
U.S. Tax Documents	Office 365 for IT Pros
U.S. Tax Form W-8BEN	Office 365 for IT Pros
U.S. Tax Form W-4 (2015)	Office 365 for IT Pros
U.S. Tax Form 1040 (2014)	Office 365 for IT Pros

## Custom Keyword Dictionaries

In addition to defining custom sensitive information types, you can set up keyword dictionaries. A dictionary is a set of keywords that you want to check for in DLP policies. For instance, you could define a set of inappropriate words like *blast*, *damn*, and *bugger* that you don't want people using in internal email and Teams messages or a set of product code words that you do not want people to share outside the company.

The easiest way to create a keyword dictionary is with a simple editor like Notepad. Put each word on a separate line and save the file in Unicode format. You can then load the file into a PowerShell variable and use the variable to create the dictionary. In this example, we read the contents of a file called *Codewords.txt* and use it to create a dictionary with the *New-DlpKeywordDictionary* cmdlet.

```
[PS] C:\> $FileData = Get-Content c:\temp\Codewords.txt -Encoding Byte -ReadCount 0
[PS] C:\> New-DlpKeywordDictionary -Name CodeWords -Description "Code Words for major projects"
-FileData $FileData

RunspaceId      : 3ab0f9ee-24ff-40c7-901c-c2bd8e43a830
Identity        : c1247f79-8cdd-4137-91a0-9c6dfb79192e
Name            : CodeWords
Description     : Code Words for major projects
KeywordDictionary : bitterball
                  gandalf
                  redsnark
                  torchcraft
                  x1050
                  x1051
                  x1052
                  yellowplane

IsValid        : True
ObjectState    : Unchanged
```

To update a keyword dictionary, use the [Set-DlpKeywordDictionary](#) cmdlet. You can input a new set of terms interactively or load them from a file. After the keyword dictionary is created, you can use it as a primary or secondary element when creating a new custom sensitive information type.

## Document Fingerprinting

Microsoft supports a comprehensive set of standard sensitive information types for use in DLP policies, but they cannot be aware of data that is uniquely sensitive to a company, like forms used in HR hiring processes, capital acquisition requests, employee review forms, and so on. Most companies would not like this information to be circulated externally and probably have security and privacy policies to govern external disclosure. DLP can help to educate users that confidential company documents should not be circulating in emails unless good business reasons exist.

Like the recognizable patterns that can be used to detect instances of data like social security numbers, documents have characteristics in terms of layouts, fields, and text blocks. You can use a process called fingerprinting to capture the characteristics of a document and create a digital signature in the form of a hash value. That signature, or fingerprint, then becomes a valid sensitive information type that can be used in DLP processing.

Microsoft Purview can create a digital fingerprint from a document in any of the formats supported by Microsoft Search. This includes any of the Microsoft Office formats plus Adobe PDF, so most documents used in corporations are candidates to serve as the basis of a digital fingerprint. It is always best to select a blank form rather than one that is filled in so that the resulting fingerprint is not affected by text or other markings

that will not appear in other copies of the form. Sample documents should not be password-protected or be graphic-heavy (text-based documents generate the best results).

Two methods exist to create document fingerprints:

- The first is in the Manage Document Fingerprints section under Data Loss Prevention in the classic **Exchange admin center**. The document fingerprints created here are used with Exchange transport rules. We describe the process to create document fingerprints via the EAC in Chapter 8 of the Companion Volume.
- The second option is to use **PowerShell** (compliance module). [Document fingerprints](#) are used to create new sensitive information types that can be used across Microsoft 365. This is the preferred approach and is described below.

After locating a suitable template file, use the *Get-Content* cmdlet to read the content of the file into a variable. In this case, I'm using the Irish Personal Tax Form 11.

```
[PS] C:\> $IrishTax11 = Get-Content "C:\Temp\Ireland Tax Form 11.pdf" -Encoding byte -ReadCount 0
```

Next, create a digital fingerprint from the variable generated by Get-Content:

```
[PS] C:\> $IrishTaxFingerPrint = New-DlpFingerprint -FileData $IrishTax11 -Description "Ireland Tax Form 11"
```

Finally, create a new sensitive information type from the fingerprint:

```
[PS] C:\> New-DlpSensitiveInformationType -Name "Ireland Tax Form 11" -Fingerprints $IrishTaxFingerPrint -Description "Ireland Personal Tax Form 11"
```

Microsoft 365 responds by listing the properties of the new sensitive information type. You can also check by running a command like:

```
[PS] C:\> Get-DlpSensitiveInformationType -Identity "Ireland Tax Form 11" | Format-Table Name, Type, Description
```

Name	Type	Description
Ireland Tax Form 11	Fingerprint	Ireland Personal Tax Form 11

The new sensitive information type is now usable in Purview DLP policies or Exchange transport rules

# Chapter 20: Managing Information Protection

## **Tony Redmond**

This chapter discusses how to use information protection technology to protect email, documents, and other files. The topics covered include:

- The rights management service and Microsoft Purview Information Protection (MIP).
- Office 365 Message Encryption (OME).
- How to protect documents and messages with sensitivity labels.
- How to protect SharePoint Online document libraries and lists and OneDrive for Business sites.
- How rights management templates work and how sensitivity labels use rights management.
- Using PowerShell with protected content.

In some quarters, rights management enjoys a reputation of being a technology difficult to understand, implement, and manage. It is certainly true that deploying this kind of protection for on-premises infrastructures can take significant effort. Hopefully, we'll prove that although Microsoft Purview Information Protection uses many of the same concepts as found in on-premises deployments, it is easier to implement rights management-based protection in the cloud because Microsoft takes care of maintaining the infrastructure to issue and check user rights to access information.

## The Need to Protect Data

We live in a world where encryption is pervasive for both businesses and consumers. It is unthinkable to consider conducting a banking transaction through a website not protected with encryption. Even so, an awful lot of email traffic still moves over unencrypted SMTP links (Exchange Online uses TLS to secure mail traffic in transit), perhaps because setting up reliable encryption for email has always run into the difficulties of agreeing on a common standard for external transmission and the cost of deploying and supporting an infrastructure to assure secure email internally. The earliest versions of Exchange Server included the Key Management Server (KMS) to manage the storage and distribution of keys for message encryption. At the time, the great hope was that the availability of KMS would encourage many organizations to adopt encryption to protect email. That hope never fully materialized.

Rights management delivers the capability to protect confidential messages or documents to control what recipients can do with the content (enforce usage restrictions). Without protection, an email recipient can forward messages they receive, print the messages, cut and paste message content, and so on. Rights management implemented through sensitivity labels allow an organization to define sets of rights that users apply to messages and documents to restrict what a recipient can do. The functionality protects companies by allowing them to circulate confidential or sensitive information under control and avoid situations such as "information leakage" which happens when, for instance, a disaffected employee forwards messages to journalists or other interested parties. Another example of similar functionality in a different context is the use of DRM to protect music downloaded from various sites. In addition, settings also control how long users can access information, after which the information becomes inaccessible. Microsoft offers several ways to deploy rights management to protect content:

- On-premises customers can deploy a Rights Management (RMS) server to manage the encryption keys used to protect data by both on-premises and hybrid users. [Hold your own key](#) (HYOK), which also uses keys managed by on-premises servers, extends the coverage to cloud applications. Because the complexities involved in using HYOK to protect data differ from customer to customer, we do not cover the topic here.
- Tenants can generate encryption keys and import the keys into Microsoft's data centers, where the keys are under the control of the tenants. This implementation is known as [Bring your own key](#) (BYOK) or HYOK.
- The most common method to implement rights management in the cloud is when Microsoft manages the encryption keys for tenants (a Microsoft managed key, or MMK). As with BYOK and HYOK, the keys issued to tenants protect information including Exchange Online messages and SharePoint Online and OneDrive for Business document libraries and lists.

Although the concept of rights management extends back to the early 1990s, Microsoft began implementing the technology for products like Exchange and Outlook in the middle of the decade. Rights management proved capable of protecting information, but customers did not take to the technology for a few reasons. Part of this was because of the culture shift needed inside companies to take the protection of information seriously, including strong executive leadership to champion the case for deploying the technology. Part of the reason was due to the extra infrastructure plus the time needed to deploy and run Active Directory Rights Management Services (AD RMS), the service underpinning protection within an on-premises environment. Because Microsoft delivers the required infrastructure and integration for cloud services, tenants need to dedicate less effort to secure protection.

The work to expand protection across Microsoft 365 is a journey. Newly-introduced capabilities include auto-application of sensitivity labels for Exchange Online messages in transit. We'll get to the new functionality as the chapter unfolds.

## Rights Management

Microsoft 365 uses the Azure Rights Management protection service as the foundation for rights management and encryption to allow users to protect information through features like:

- **Office 365 Message Encryption (OME)**, a set of features to encrypt Exchange Online email including:
  - Out-of-the-box encryption for the Outlook and OWA clients through special *Encrypt Only* and *Do Not Forward* options. When OME encrypts a message, it places a special wrapper around the message that Exchange Online or Outlook.com recipients can automatically decrypt. Other recipients can access the content through the OME portal. OME's ability to protect email sent to any recipient using the Encrypt Only feature can replace the need to deploy S/MIME or other third-party tools. Outlook mobile clients can create and read messages protected by OME.
  - Recipients of encrypted messages outside the tenant can read the content in the OME portal by authenticating their access to the portal or using a one-time passcode. The same occurs when "unenlightened" clients (ones that don't understand rights-management based encryption) connect to Exchange Online. For instance, if someone uses an IMAP4 client to connect to Exchange, they must go to the OME portal to read any protected messages they receive.
  - Protection applied to outbound email by email transport (mail flow) rules. For example, a rule can ensure that any message sent to a partner domain is encrypted. Transport rules can apply the standard OME *Encrypted* and *Do Not Forward* protection as well as sensitivity labels defined in the tenant.

- Users and policies assign **Sensitivity labels** to messages and documents. While some labels only apply visual markings to content, others can invoke protection with rights management permissions. Clients that apply and process sensitivity labels use the MIP Software Development Kit. MIP is a framework or platform for developers inside and outside Microsoft to add protection to applications.
- **Information Rights Management protection** for documents downloaded from SharePoint document libraries. This form of protection is now obsolete and should be replaced by sensitivity labels.

Some differences might exist in the capabilities depending on the data center region used by your tenant. For example, the sovereign cloud region in [China](#) does not support the full suite of protection capabilities and users within these clouds have limited options to protect email and documents. Check with Microsoft about the current situation.

## Protection and Permissions

Rights management protects Items through a set of rights or permissions granted to a recipient. Logically, the author or originator of content always has full control over the content. You can compare applying a sensitivity label to a message as being analogous to registering a postal letter. The recipient is only able to open and access the content if the rights management service recognizes their access – the recipient gains access by having a known account (authenticated against Azure AD) with access rights defined by the sensitivity label that the application can apply to the item. If protected content ends up in the hands of an unknown user, they will not be able to access the content because it will stay within an encrypted “wrapper” that only intended users can open.

It is important to understand the available rights do not cover every possible circumstance or method that a recipient might use to interact with protected content. For example, if you don't grant the right, recipients can't take a screen capture of the content on Windows devices because Windows respects the denial of the right. However, iOS and Mac devices allow screen captures because those operating systems do not permit applications to restrict screen captures. This difference in behavior between operating systems illustrates the point that rights management can go so far in restricting deliberate attempts to share protected content (taking a photo of a screen with a mobile device is another way around rights management). Rights management is therefore part of an overall solution to manage critical and sensitive information: it is not a complete answer if deployed without supporting user education, policies, and possibly other technology.

### Other Ways of Using Protection

In addition to allowing authors to decide what level of protection to assign to their work by applying a label to a message or document, administrators can create transport rules to apply labels to messages that meet certain criteria as the messages travel through the transport pipeline. This is an excellent way to automatically protect confidential messages if you can create criteria to find those messages. For example, all messages that mention the word “Confidential” in the message subject or all messages sent to the “Corporate Planning” distribution list. Transport rules protect messages without user intervention and users cannot override what they do. Best of all, this approach works for messages sent from any client.

Protecting messages through transport rules allows great control over information circulated through email. It also means that the content of any message sent externally is inaccessible because external recipients cannot retrieve the use licenses necessary for decryption. Fortunately, you can use encryption to protect messages sent to external recipients in a way that they can access the content securely.

While more difficult to manage, it is possible to use [an on-premises Active Directory RMS server](#) to deliver a Rights Management service to users. That setup is outside the boundaries of the discussion presented here.

## Licensing Requirements for Microsoft Information Protection

Microsoft 365 has two Information Protection service plans. The standard version is part of Office 365 E3 and the premium version is bundled with Office 365 E5. These service plans govern the functionality available through sensitivity labels to protect email and documents in Exchange Online, SharePoint Online, and OneDrive for Business. The Office 365 E3 license covers the use of sensitivity labels to manually classify and protect content stored in Exchange Online, OneDrive for Business, and SharePoint Online. The Office 365 E5 license enables functionality like automatic labeling at rest and label analytics. Tenants don't need licenses to apply sensitivity labels to items using Exchange Online transport rules.

All Microsoft 365 licenses include the ability to consume protected documents (in other words, to open protected documents). For instance, an account with an Office 365 E5 license could protect a document with a sensitivity label and circulate the document to users with frontline (F1) licenses. Even if allowed to, those users are not licensed to change the protection on the document, but they can certainly open and view its content. They can also upgrade their capabilities with a premium license to gain the ability to apply protection to documents.

Office 365 Message Encryption is part of the Office 365 E3 and E5 plans (enterprise, academic, and government) plus Microsoft 365 Business Premium. For other plans like Exchange Online Plan 1 or 2, you'll need to purchase the Microsoft Information Protection add-in.

Premium licenses such as Microsoft 365 E5, Microsoft 365 E5 Compliance, and Microsoft 365 E5 Information Protection and Governance cover advanced functionality like automatic labeling. When you see references to premium licenses in this chapter, we usually mean that the feature under discussion needs one of the above licenses. However, given the number of features and plans available in this space, the issue of licensing can be quite complex and it's wise to check exactly what you need. Microsoft [publishes guidance](#) to help tenant administrators and licensing coordinators understand when premium licenses are required. A useful [Microsoft 365 compliance comparison spreadsheet](#) is also available to show which license covers each feature. The spreadsheet also identifies gaps in terms of desirable features not covered by licenses held by a tenant.

In some cases, a feature might not enforce the stated licensing requirement. This could be because the necessary code is not yet available. The code might or not appear soon. In any case, a tenant must have licenses to use functionality. It's a bad place to be in if features the business depends on suddenly stop working because Microsoft updates its license enforcement code.

### Licensing the Protection of Content Stored Outside Microsoft 365

Your organization may want to protect files both inside and outside Microsoft 365. Office 365 E3 and E5 licenses cover the use of sensitivity labels and OME to protect documents and messages stored inside its repositories and to protect emails sent outside the organization. To apply sensitivity labels to files stored in external repositories or to files belonging to applications that don't include native support for Microsoft Information Protection, you must:

- Deploy the unified labeling client to workstations.
- Assign Information Protection licenses to accounts that apply protection to external files.

Licenses are available standalone or bundled in the Enterprise Mobility and Security and Microsoft 365 enterprise plans. The basic rule here is that using the unified labeling client to protect a file (or to change the protection on a file), requires a license. On the other hand, no license is necessary for someone who only opens and accesses the content in protected files.

The unified labeling client does not check licenses. If an organization has a Microsoft Information Protection subscription, all clients within the organization can download the policy and use labels. The use of labels is what invokes the need for a license, so to avoid this happening you can force the workstations used by

unlicensed users to run the client in [protection-only mode](#) to stop clients from downloading policies. See [this page](#) for more information.

## The Flow of Protection

When a tenant activates rights management (the default state), clients use an automatic process to receive certificates from the rights management service to allow users to access protected content even when they work offline. The service keeps a copy of the user's Rights Account Certificate (RAC) so that it can issue it to another workstation if the user connects from there. Once authenticated, users can send and receive protected messages or apply sensitivity labels to protect files. When a user protects a message or document (for example, by applying a sensitivity label), the client embeds a unique key (the content key) in the item's header. To protect an item, the client encrypts it (and any attachments which support protection) using the AES 256-bit symmetric encryption algorithm. The content key persists with the item even if the author edits it to create an updated version. Rights management protects the content key with another key (the tenant key), which is common across all protected content. Either Microsoft or the tenant can manage the tenant key.

When it encrypts information, the client generates a certificate called the publishing license which includes a policy with the usage rights for recipients (individual users or groups) as well as any other restriction, such as an expiry date. Like all the licenses used by rights management, the publishing license is an [XrML certificate](#). The settings stored in the publishing license come from the sensitivity label selected to protect the content. The client signs the publishing certificate with a user certificate. The client also encrypts the publishing certificate and the content key using the tenant key. Finally, the client combines the encrypted content with the signed and encrypted publishing certificate to create the protected item. This step ensures that the publishing certificate always stays with the encrypted content. If necessary, the rights management service can use the information in the publishing certificate to create a use license in cases when the template used by a sensitivity label is inaccessible for some reason.

When a recipient reads protected content, the client extracts the publishing certificate from the protected item and sends it and the user certificates to the rights management service, which decrypts and evaluates the set of rights for the item. The service then extracts the content key from the publishing certificate and uses it to create an encrypted use license holding the set of rights allowed to the user and returns the license to the client. The client decrypts the use license with the user's private RAC key to extract the content key, which it uses to decrypt the content before displaying it to the recipient. The client also handles the enforcement of rights given to the user by policy, including limiting access if the license is valid for a certain period. When a use license for an item expires (usually after 30 days), the user must reauthenticate with Azure AD to get another license.

Mobile clients like Outlook for iOS use a simpler transactional flow. When mobile clients protect content, they send the selected policy to the server and receive a publishing license and symmetric key to protect the item. To consume protected content, the client sends the policy to the service and requests a use license. The service responds with the necessary keys and policy information to allow the client to open and display the content.

## External Users

Protection works on the basis that a recipient can authenticate themselves using an Azure AD account or a Microsoft Service Account (MSA) associated with an email address. Therefore, you can send protected items to people in other tenants and other email domains, providing that the policy used allows those recipients rights over the item. See [this document](#) for more information about how to use protection to secure documents shared with people outside your tenant.



Azure AD federates with some other directories, such as Gmail and Yahoo, to allow users of those services to authenticate by signing into that service. Otherwise, the external user must sign in using an MSA account or a [one-time passcode](#) before they can access the document.

## Contacting the Microsoft Information Protection Team

If you have a technical issue with anything related to Microsoft Purview Information Protection (including the SDK, unified labeling client, or another aspect of the technology) or wish to discuss a requirement your organization might have, you can join the conversation with the [MIP team in Yammer](#).

# Enabling Rights Management for a Tenant

Before you can protect content, the rights management service must be enabled. Microsoft [enables rights management automatically](#) for eligible tenants. In other words, if you have an eligible plan like Office 365 E3 or E5 or buy the Azure Information Protection add-on for other plans, you do not have to enable Rights Management. Users with other Office 365 licenses can consume sensitivity labels applied to content but need to have an upgraded license to apply labels to content. See [this documentation](#) for information about enabling the protection service in a tenant.

## Configuring Rights Management for Exchange Online

Although Microsoft configures Exchange Online so that you can use OME and sensitivity labels without doing anything else, you might wish to change some of the settings for the rights management configuration. This section explains how to make those changes. We'll discuss how to configure other relevant settings for how OME handles messages sent to other email systems and protection for SharePoint Online later.

### Configure Exchange Online

The *Get-IRMConfiguration* cmdlet reports the current rights management configuration for Exchange Online while the *Set-IRMConfiguration* cmdlet is used to update settings. Microsoft configures rights management for all eligible tenants and sets the *AzureRMSLicensingEnabled* property in the IRM configuration to *\$True* to enable Exchange Online to use the Azure Information Protection service. See [this page](#) for more information.

### Test the Configuration

To test the configuration, we run the *Test-IRMConfiguration* cmdlet. The templates listed in the results are the default (customizable) set created for a tenant called *Confidential* and *Confidential View Only* and the special *Encrypt* and *Do Not Forward* options created by OME. Of course, you want to see "PASS" as the overall test result.

```
[PS] C:\> Test-IRMConfiguration -Sender Tony.Redmond@Office365itpros.com

Results : Acquiring RMS Templates ...
          - PASS: RMS Templates acquired.  Templates available: Redmond & Associates -
Confidential View Only, Redmond & Associates - Confidential, Encrypt, Do Not Forward.
          Verifying encryption ...
          - PASS: Encryption verified successfully.
          Verifying decryption ...
          - PASS: Decryption verified successfully.
          Verifying IRM is enabled ...
          - PASS: IRM verified successfully.

OVERALL RESULT: PASS
```

If you do not see similar output to that shown above, it is probably because your tenant configuration uses the earlier IRM stack. Repeat the steps to configure IRM and then retest to confirm that the new configuration is in place.

## Check the IRM configuration

Run the *Get-IRMConfiguration* cmdlet to check the rights management configuration. If any changes are necessary, you can make them with the *Set-IRMConfiguration* cmdlet. For example:

```
[PS] C:\> Set-IRMConfiguration -InternalLicensingEnabled $True  
-ClientAccessServerEnabled $True -EnablePdfEncryption $True
```

In most cases, you do not have to amend the default configuration to use start protecting email. If you want to, you can investigate settings such as:

- **AutomaticServiceUpdateEnabled:** Controls if new Information Protection features are automatically enabled in the tenant. The default is *\$True* and it is best to leave this setting alone unless you have good reason to block the deployment of a new feature into your tenant.
- **AzureRMSLicensingEnabled:** This should always be *True* as it controls the ability of Exchange Online to connect to Azure Rights Management.
- **ClientAccessServerEnabled:** Default is *\$True*. Controls whether OWA and ActiveSync clients (including Outlook for iOS and Android) can use IRM to protect and decrypt messages. When true, Exchange Online can fetch use licenses from the rights management service on behalf of these clients (see below).
- **DecryptAttachmentForEncryptOnly:** Controls if Exchange Online decrypts attachments for messages protected with the *Encrypt Only* feature. The default is *\$False*, meaning that attachments remain encrypted, even when downloaded. If you change this control to *\$True*, Exchange Online decrypts the attachments, and recipients have full control over the files. (Note: the older *DecryptAttachmentFromPortal* setting is deprecated).
- **EDiscoverySuperUserEnabled:** Default is *True*, which allows members of the Discovery Management RBAC management role group to access protected information found through (now deprecated) Exchange Online eDiscovery searches.
- **EnablePdfEncryption:** Set to *\$True* to enable OWA and Outlook Mobile clients to apply sensitivity labels or standard OME templates to protect messages with PDF attachments. These clients include the necessary code to apply protection to PDF attachments, but Outlook desktop clients do not. To extend coverage to messages with PDF attachments sent from Outlook desktop clients, use an Exchange mail flow rule or DLP rules to apply a suitable sensitivity label to messages with PDF attachments. When users download protected PDFs, they can be opened with [an app that supports the ISO standard for PDF encryption](#), including the Edge browser.
- **InternalLicensingEnabled.** Default is *\$True*. Controls whether licenses are automatically granted to internal recipients to allow them to access protected content.
- **JournalReportDecryptionEnabled:** Default is *True*. Controls whether the Exchange Online transport service uses its IRM super-user privilege to decrypt protected messages copied by journal rules to an external journaling system. When this happens, the journal recipient receives a journal report with two attachments. One holds the original message, the other the decrypted copy. Decryption works for both the default (*Encrypt Only* and *Do Not Forward*) options and custom sensitivity labels.
- **SearchEnabled:** Default is *True*. Controls whether OWA can search protected items in a mailbox.
- **SimplifiedClientAccessDoNotForwardDisabled:** Set to *\$False* to make the *Do Not Forward* option available for messages in OWA.
- **SimplifiedClientAccessEnabled:** Controls if the Protect button is available in OWA. Set to *\$False* if you don't want users to apply protection to messages. This control is now obsolete as the Protect button has been replaced by the Sensitivity button.
- **SimplifiedClientAccessEncryptOnlyDisabled:** Controls if the *Encrypt Only* option is available in OWA. Set to *\$False* to disable the option.
- **TransportDecryptionSetting:** Default is *Optional*. Controls the access to protected messages for the transport service as they pass through the transport pipeline. *Optional* means that the transport

service tries to decrypt protected messages so that the content is available for checking by transport and DLP rules or by anti-virus agents. However, transport will continue to process and deliver the message even if decryption is not possible. You can disable any attempt to decrypt messages by setting this value to *Disabled*. Setting it to *Mandatory* means that transport will reject any messages that it cannot decrypt and return a non-delivery report to the sender. The transport service automatically re-encrypts decrypted messages when they reach the end of the transport pipeline.

To make things easier for OWA users, if the *ClientAccessServerEnabled* setting in the IRM configuration is *\$True*, Exchange Online imports keys from the Rights Management Trusted Publishing Domain (TPD) and is thereafter able to decrypt content locally on behalf of clients so that OWA can display protected content inline within message windows. Some Exchange ActiveSync clients use the same approach, which is known as prelicensing.

## Bring Your Own Key (BYOK)

Microsoft 365 allows customers who need the highest possible degree of control over all aspects of security the ability to have full control over their tenant key, the element that serves as the root of trust for protection. The key is “pinned” or imported to a FIPS140-2 hardware security module (HSM), usually kept under tight control at a customer’s premises. In collaboration with Thales E-Security, Microsoft uses a secure process to transfer the key from a customer and import it into the Microsoft data centers into [Azure KeyVault](#) to make the key available to serve as the basis for protection. The process is free but needs a good deal of planning and coordination. Both SharePoint Online and Exchange Online support Azure KeyVault, so once a tenant imports its key to become the tenant key, that key becomes the base for encryption for SharePoint and Exchange content, except when encryption is applied through sensitivity labels as the rights management templates used by these labels are cloud-based. See [this link](#) for more information about BYOK.

# Sensitivity Labels

The original implementation of cloud-based rights management followed the approach taken for on-premises deployments and used rights management templates to protect content. Microsoft Information Protection underpins the rights-management driven encryption side of sensitivity labels. Given the complex nature of anything to do with encryption, the deployment of sensitivity labels can take time to achieve, especially in large, complex companies. The basics of sensitivity labels are:

- Administrators manage sensitivity labels in the Information Protection section of the Microsoft Purview Compliance portal. Administrators must publish sensitivity labels to users through label policies before users can apply labels to content.
- The settings defined in sensitivity labels allow the labels to perform three major functions:
  - Apply visual markings such as headers and footers to Office documents and messages.
  - Protect the content of files with rights-management based encryption (Microsoft Information Protection). Encryption is persistent and remains until someone with the necessary authority removes the label. To gain access to protected content, users must authenticate.
  - Manage containers (Teams, Microsoft 365 Groups, and SharePoint Online team sites).Labels don’t have to perform all functions. The scope set for a label establishes its use. Applications use the defined scope to decide if they should reveal a label to users. For instance, client user interfaces do not reveal labels used solely for container management when displaying the set available to users to protect files.
- Applications use the Microsoft Information Protection SDK to incorporate code to recognize and support the protection (encryption) applied through sensitivity labels. Within Office 365, support is available in Exchange Online, SharePoint Online, and OneDrive for Business. Users apply labels to messages and documents through the Microsoft Office apps. Other applications like Power BI use

sensitivity labels to protect objects and when users export content. Teams uses sensitivity labels to manage team settings but doesn't currently use sensitivity labels to protect conversations and chats. Files protected by sensitivity labels are viewable in Teams. A preview is available where [sensitivity labels can protect Azure Purview objects](#).

- The Office 365 E3 and E5 plans include a license to apply sensitivity labels manually to documents. All Office 365 users can consume (read) files protected by sensitivity labels.
- Document understanding models generated by SharePoint Syntex can apply both retention and sensitivity labels files stored in document libraries.
- Sensitivity labels support manual (explicit assignment) by users or automatic assignment by policy (implicit assignment). A label assigned explicitly by a user cannot be replaced by a label assigned by an auto-label policy.
- The unified labeling client can apply labels to files stored outside Office 365 and to non-Office files such as PDFs. Users need a Microsoft Information Protection license to apply sensitivity labels to files stored outside Office 365.

## Retention and Sensitivity

Some confusion might exist over the functions of retention labels and sensitivity labels; a summary of the differences and similarities is:

- Documents and messages can have a single sensitivity label and a single retention label, but not multiple labels of either type.
- Retention labels are only valid within Microsoft 365. A sensitivity label is persistent and remains with an item until removed by someone with the right to do so.
- Both types of labels are applicable manually (by a user) or automatically by an auto-label policy. It's also possible to use the presence of sensitivity labels as a condition for the auto-application of retention labels.
- Sensitivity labels support container management. Retention labels can be the default for a container like a SharePoint site, but they have no management function.

## Marking Important Content

Applying a sensitivity or protection label (think of a label as an adhesive sticker) to a document or email marks the item as having a certain level of importance to the organization. Usually, the higher the sensitivity of the content, the more stringent the protection invoked by the label. Some labels are purely visual indicators for a certain kind of content, like "Public." Others impose visual markings on documents and email to help users understand that the information is more sensitive; for instance, a label named "Confidential" might insert a footer and watermark in documents when applied. The most interesting type of label uses rights management to protect the most sensitive content. For example, a "Secret" label might use permissions to grant view-only access to anyone in the tenant while blocking access to anyone external.

The people who create and manage information often best understand the sensitivity or confidentiality of the content. The ability to apply labels to items is increasingly pervasive across Microsoft 365, meaning that it is very easy for information generators to apply appropriate labels as they create new content. Once applied, the labels are persistent and remain in place for the lifetime of items. Labels can be removed or replaced with other labels by people with permission to make those changes.

Microsoft Endpoint Manager and Microsoft Defender for Cloud Apps support sensitivity labels. Policies can apply labels to ensure that users don't copy unprotected content to a third-party app like DropBox or removable storage like a USB drive.

## Files, Messages, and Containers

Sensitivity labels support:

- **Files and messages:** The traditional use of sensitivity labels is to apply visual markings and protection to files and email. The definition of files includes Power BI and Microsoft Purview schematized data assets like SQL.
- **Containers:** If configured for the tenant, labels can apply management settings to Teams, Microsoft 365 Groups, and SharePoint Online sites. Applying labels to containers does not protect the items stored in the containers. Microsoft is building out the set of controls settable through labels to make this capability more useful and powerful. Only administrators and container owners can apply labels to containers.

Configuration for container support is a one-time operation involving an update for the *EnableMIPLabels* setting in the Azure Active Directory Policy for Microsoft 365 Groups using PowerShell. When the value of the setting is True, applications that support container management know that they should replace the older text-only classifications with sensitivity labels. Here are the necessary commands:

```
[PS] C:\> Connect-MgGraph
$TenantSettingsId = (Get-MgDirectorySetting | ? {$_.DisplayName -eq "Group.Unified"}).Id
$TemplateId = (Get-MgDirectorySettingTemplate | ? {$_.DisplayName -eq "Group.Unified"}).Id
Update-MgDirectorySetting -TemplateId $TemplateId -DirectorySettingId $TenantSettingsId -Values
(@{'name'='EnableMIPLabels';'value'='true'} | ConvertTo-Json)
```

Before users can apply sensitivity labels to files, messages, or containers, the labels must be published via label policies.

Assigning sensitivity labels to files and messages for visual marking and protection through rights management-based encryption is known as information protection. Applying labels to containers is container management. The settings defined for a label sets its scope, which can be:

- Information protection only (file, email).
- Container management only (site, unifiedgroup).
- Both information protection and container management (file, email, site, unifiedgroup).

When the Microsoft Purview Compliance portal displays a list of labels, it shows the scope for each label as its content type (listed in parenthesis above). As discussed later, it is often easier to manage sensitivity labels when organizations create dedicated and separate sets of labels for the two purposes.

Applications check labels to understand their scope and filter out the labels that they cannot use. For example, the set of labels available in the Office apps includes labels with information protection settings. The set displayed in the Teams client when someone creates a new team is those that have container management settings.

## Native Support in Microsoft 365 Apps

Native support means that an app includes the necessary code (built using the MIP SDK) to fetch policy and label information from the unified store and comply with policy settings, including encryption, as users assign sensitivity labels to items. Microsoft Office is a good example of an application that replaced the need to run a separate client with native support. Among the advantages gained by the integration are better cross-platform support, increased performance, lower memory demand, and avoiding problems that can happen when apps like Word or PowerPoint load add-ins. In addition, native support for sensitivity labels is included in the regular updates issued for the Office apps and you don't have to update a separate component each time an update is available.

The downside is that the apps don't display the information protection toolbar inserted into apps by the old AIP client (because of screen constraints in mobile clients). Instead, interaction with sensitivity labels is through the Sensitivity button to apply, update, or remove sensitivity labels when working with Office files. The user's signed-in account appears at the top of the list of labels revealed by the Sensitivity button. Documents assigned a label display the name of the label in the bottom information bar. A label icon appears when the applied label doesn't include encryption and a padlock icon when it does.

Page 1 of 2    659 words    English (Ireland)    Secret

Native support for sensitivity labels is available in the Office mobile apps, the desktop apps for Windows and Mac, and the Office Online apps. Table 20-1 summarizes the current support for sensitivity labels in different client platforms.

<b>Client</b>	<b>Notes</b>
Microsoft 365 Apps for Windows (Outlook, Word, PowerPoint, Excel)	Office desktop clients for Windows and Mac include native support for sensitivity labels as described in this section. You can access sensitivity labels with older versions of the Microsoft 365 Apps by installing the <a href="#">Unified Labeling client</a> . The presence of this client on a workstation exposes the Sensitivity button in the app menu to allow users to work with sensitivity labels. The MSI versions of Outlook (2016 or 2019) don't include the necessary code for sensitivity labels.
Office online apps	Supported by Word Online, PowerPoint Online, and Excel Online.
Office mobile apps	Available for Word, PowerPoint, Excel, and Outlook mobile on iOS and Android.
Other non-Office files	Use the unified labeling client to apply labels to these files. You can then import the files into SharePoint Online or OneDrive for Business.

Table 20-1: Client support for sensitivity labels

## Unified Labeling Client

If you don't use a version of the Office apps which includes native support for sensitivity labels or want to apply sensitivity labels to files stored outside Office 365, you can deploy the unified labeling client. The unified labeling client:

- Adds the Sensitivity button to the menu bar of the Office desktop applications to allow users to assign sensitivity labels to Office files. The Office applications in Microsoft 365 apps for enterprise include these UI elements. The code to encrypt and decrypt files is in the unified labeling client.
- Integrates with Windows File Explorer to allow users to assign sensitivity labels to files stored outside Microsoft 365 through a Classify and Protect option in the right-click menu.
- Installs a viewer to allow the display of protected content when an application doesn't include the necessary Microsoft Information Protection code to process rights management and encryption.
- Installs a PowerShell module to allow administrators to manage protected files.

The unified labeling client is available only for Windows. The [generally available version of the Unified Labeling version of the client](#) applies sensitivity labels to content stored inside Microsoft 365. The unified client uses the Information Protection SDK, gets its label and policy information from the unified store, and manages labels through the Microsoft Purview Compliance portal. Although designed for use with Microsoft 365 content, you can apply sensitivity labels to files stored in external locations like a Windows drive. See [this page](#) for the release history of the unified labeling client. The unified labeling client entered maintenance mode on January 1, 2022. This means that Microsoft will no longer develop new features for the client as they concentrate on expanding native support for sensitivity labels in applications. However, Microsoft will support the unified labeling client and issue bug fixes for problems reported by customers.

Microsoft developed an older version of the AIP client ([the classic client](#)) to apply labels to files stored outside Microsoft 365 but [deprecated the classic client](#) on March 31, 2021, and no longer supports its use.

A comparison of [the unified labeling client and the labeling incorporated in Office applications](#) (on Windows) is available online. To assist with user adoption of sensitivity labels, [Microsoft has published some information](#) to help educate users about how and when to apply labels to content.

### Third-Party Support for Microsoft Information Protection

An SDK is available to help ISVs develop integrated solutions based on MIP. [Some examples](#) include:

- Symantec DLP integration with MIP: decrypts protected content for DLP scanning.
- McAfee MVISION: applies labels to sensitive content discovered during scans.
- Relativity Trace: examines protected content during communication monitoring.
- VMware Boxer and Workspace ONE: apply labels to protect sensitive content.

### Double Key Encryption

Double key encryption is a form of protection where the customer holds one key and Microsoft holds the other. Both keys are needed to decrypt content and the mechanism is intended to allow customers to protect their most confidential and sensitive information. As Microsoft only has one key, they cannot access the protected information. [Double-key encryption](#) requires Office 365 E5 or Microsoft 365 E5 licenses and supports sensitivity labels.

## Planning Sensitivity Labels

Before creating any sensitivity labels, it is sensible to chart out a plan for protection within the tenant. The plan should include a discussion of points such as:

- **Consider applying a default label:** The policies used to publish sensitivity labels to users can dictate a default label for documents and messages. One school of thought is that people are more likely to understand the use of labels if new documents and messages receive a default label. If you choose to go along this path, consider not using encryption with the default label. The argument against applying default labels is that the application of a sensitivity label should have some purpose, and rights management-based protection is a great way to convey the level of importance, confidentiality, or sensitivity of an item. Users should therefore make considered decisions about the right label to apply to a specific item. The decision as to which course to take is highly dependent on the organizational culture and the goals for the deployment of sensitivity labels. MIP applies default labels to new Office documents and to existing documents when a user modifies those files.
- **Restrict the number of labels.** Although an organization can deploy up to 500 labels, it's debatable if more than a small number of labels meet the needs of all but some corner cases. Using the keep it simple principle, the set of labels available to "the average user" shouldn't include more than ten labels to make it easy for the user to understand what label to apply when. The general set of labels published to all users can be supplemented with extra labels for specific groups if they are needed.
- **Scope labels.** Sensitivity labels can be applied to items such as files and messages. They can also apply settings to containers, such as Teams and SharePoint sites. The settings defined in a label determine its scope. Although you can have "cross-over" or multi-purpose labels, it is usually simpler and therefore better to separate labels into a set used for information protection and a different set used for container management.
- **Names and descriptions associated with company workflow and culture are more understandable.** For example, if "Confidential" has been used in the company to mark highly sensitive documents for years, it's a good choice to use it as a label name. Don't put confidential information (like "Company Merger") in label names – choose a code name instead. Test your label names with real users before committing to the final set of names. Include a description of the kind of information that the label is intended to be used with. For files and emails, the description should tell the user about the level of sensitivity of the information that the label is appropriate to be used with.

For containers, the description should say something about the settings that will be applied, such as privacy, guest access, and the site sharing capability.

- **Choose enduring label names.** Labels can be assigned to items for many years. Choose names that will last rather than those associated with one-time events. Give labels used for container management names different from those used for information protection.
- **Use sublabels sparingly.** Sub-labels are variations of a parent label. If a label defines the sensitivity of an item, a sub-label allows the sensitivity to be varied for a certain scenario. If you have too many sub-labels, it might make it difficult for users to understand which variation applies.
- **Protection and marking are applied by labels.** Some labels are designed to serve as visual indicators of an item's sensitivity without taking any protective action. Other labels will impose actions such as marking and/or encryption. The protection given by a label should match the expected sensitivity set by its name, and the rights assigned in the label must support its use within the organization (or outside, if labeled items are shared with external people). Remember that a label can also apply visual marking to items (watermarks, headers, and footers), so consider if this is needed.
- **Involve more than IT.** It's unlikely that IT understands every business and legal ramification flowing from how information is used within the company, so involve other expertise before settling on a final design.

The goal is to create a practical set of sensitivity labels that make sense to people and meet business requirements without over-complicating matters. It's bad to have labels that people never use, and good when sufficient labels are available to allow users to select granularity in terms of confidentiality. In some cases, it might be possible to have a general-purpose label such as "Company Confidential" that allows full access to any tenant account while blocking external access. Such a label might deliver enough protection for a large percentage of use cases. In other circumstances, specific labels might be necessary to restrict access to certain sets of users or to grant rights to selected groups of users. In other words, the plan should define a set of labels to meet the needs of the business while limiting the number of labels to a manageable set.

The plan can be laid out in a form like that shown in Table 20-2, ordering the set of labels used for information protection in order of sensitivity from least sensitive to most confidential. The Microsoft Purview Compliance portal organizes sensitivity labels in this order. Office applications use the order to know whether a user has replaced a label with one of higher or lower sensitivity.

<b>Label Name</b>	<b>Description</b>	<b>Marking</b>	<b>Protection</b>	<b>Extended Protection</b>
Public	Content approved for sharing outside the company	Footer	None	None
Internal	Content that should remain internal but can be shared	Footer	None	None
Confidential	Content that should remain inside the company	Footer	Yes	None
Secret	Content that should never go outside the company	Footer and watermark	Yes	7-day expiry
Ultra	"Eyes-only" content for restricted circulation	Footer and watermark	Yes, restricted to certain groups	3-day expiry

Table 20-2: Planning Sensitivity Labels for Information Protection

If you like to use colors to highlight different types of labels, you can include the color (hex value) for each label in the plan. Some recommend the use of a traffic light scheme for labels where labels with a green color are for general access, yellow labels mark confidential material, and red is used for labels with the highest level of confidentiality.



In addition to the set of general-purpose sensitivity labels, the plan might include labels for use by certain projects or departments. For instance, you could create a label for the Legal Department to mark documents associated with a patent application or other aspects of intellectual property.

Figure 20-1 shows a set of sensitivity labels displayed in the Microsoft Purview Compliance portal. Remember that the portal lists labels in priority order with the most important labels at the bottom of the list. The labels are in the order described above, with some of the labels created for use by projects and departments coming after the general-purpose set. The names of the labels appear rather than the display names, which is what users see in applications. Usually, the name and display name are the same for a label, but they can differ.

The screenshot shows the Microsoft Purview Information Protection portal. The left sidebar contains navigation options like Home, Compliance Manager, Data classification, and Solutions. The main content area is titled 'Information protection' and shows a list of sensitivity labels. The labels are ordered from lowest to highest sensitivity. A context menu is open over the 'Office 365 Book Source' label, showing options to 'Add sub label', 'Move up', and 'Move down'.

Name	Order	Scope	Created by	Last modified
Public	0 - lowest	File, Email	Tony Redmond	14 Mar 2021 18:06:09
No Encryption	1	File, Email	Tony Redmond	10 Jan 2021 19:18:57
Non-business use	2	Site, UnifiedGroup	Tony Redmond	7 Jun 2021 16:26:44
Partner-Accessible Content	3	File, Email	Tony Redmond	10 Jan 2021 19:24:41
General Access	4	Site, UnifiedGroup	Tony Redmond	23 May 2021 14:35:03
Internal	5	File, Email	Tony Redmond	23 Feb 2022 23:11:22
Office 365 Book Source	6	File, Email	Tony Redmond	10 Jan 2021 20:32:18
Guest Access		Site, UnifiedGroup	Tony Redmond	23 Feb 2022 23:26:45
Confidential		File, Email	Tony Redmond	3 Mar 2021 13:32:39
Secret	9	File, Email	Tony Redmond	2 Aug 2021 17:53:50

Figure 20-1: A set of Sensitivity Labels in the Microsoft Purview Compliance portal

You can also see the choices available in the ellipsis menu to reorder labels or to create a sublabel. When you move a label up in the order, its level of sensitivity decreases. For instance, the Public label, which is at the top of the list, is used to mark less sensitive information than the Secret label further down the list. Conversely, if you move a label down, Office applications regard the label as being more sensitive. You can also use the *Set-Label* cmdlet to assign a priority to a sensitivity label. For example:

```
[PS] C:\> Set-Label -Identity "Eyes Only" -Priority 26
```

Label names should leave users in no doubt as to the relative importance or sensitivity of the information in an item. After all, they have no idea of the position a label has within the list managed by Microsoft 365: all they see is the label name, so it's critical to give the labels meaningful and understandable names. Because label names might appear in mobile applications on devices with limited screen estate, it's best if the names are short rather than long.

**Note:** If you use different labels for container management and information protection, the position of the container management labels in the priority order affects the signaling of label mismatches as documents are uploaded. See the later discussion on this topic.

## Sublabels

Sensitivity labels used for information protection can have sublabels. This means that you have a main or group label that isn't used for labeling. Instead, the main label is a container that serves as a logical collection of other related labels, each of which can have different marking and encryption settings. Although main labels appear in applications, you can't assign them to content. Instead, users can select from one of the sublabels, which is then assigned to the content.

## Labels Outside a Policy

Sensitivity labels become available by publishing the labels in a policy to all or some users in a tenant. The users included in the policy form the target set for the policy. Sometimes users will receive items with a sensitivity label that isn't available to them in the policy published to their account. In this case, because policies include all label definitions, including those unavailable to the user, the client can display the label. However, the user won't be able to apply the label to other items.

## Planning for Container Management

If you use sensitivity labels for container management, a separate implementation plan should describe the set of labels for assignment to Teams, Groups, and Sites. Table 20-3 describes an example of what a container label plan might look like. Once again, the plan describes labels from least sensitive to most confidential.

<b>Label Name</b>	<b>Description</b>	<b>Privacy</b>	<b>Guests Allowed</b>	<b>Sharing</b>	<b>Endpoint</b>
General Access	Container holding non-sensitive, general-purpose information.	Public	Yes	Anyone	Unrestricted
Guest Access	Container holding restricted information.	Private	Yes	Existing guests	Web access
Limited Access	Container holding restricted information.	Private	No	Existing guests	Web access
Secret Access	Container holding highly confidential material.	Private	No	No external	Block Access

Table 20-3: Planning Sensitivity Labels for Container Management

Some organizations like to include a sensitivity label to mark groups and teams dedicated to non-business use, such as sports or other non-work activities. Such a label would have the same settings as the "General Access" label described above.

Given the more limited set of setting permutations available for container management, it's normally the case that fewer container management labels exist than the set of information protection labels. Keeping to a small set is good in any case because it makes it easier for container owners to choose and apply the right label.

## Rescoping Labels for Container Management

Given that sensitivity labels started with a single scope (files and email), it is possible that an organization has labels used for both information protection and container management and now wishes to separate labels into two distinct sets. The suggested approach is:

1. Create a set of new labels specifically for container management. Make sure that these labels do not have information protection (files and email) settings. Give names to labels that reflect their use for container management and make sure that the descriptions used are informative and give guidance to container owners about the proper use of each label.
2. Publish the new labels to make them available to container owners. It's best to use a separate label publication policy for the labels used for container management. Optionally, you can choose to force users to assign a label to new containers and choose a label that will be applied by default to new containers.
3. Unpublish the information protection labels that have container management settings. This must happen before you can remove the container management settings.
4. Remove the container management settings from the set of information protection labels. Within a few hours, these labels should no longer be available to assign to new containers. Existing labels and their settings will remain in place in the containers to which they are applied.
5. Republish the information protection labels to make them available to users again. The labels can now only be applied to files and email.
6. Map the set of new container management labels to the set currently applied to containers. Some of the mappings will be 1:1, other labels might be able to replace multiple labels, and some new labels might not be used. For example, this mapping might be used:

<b><i>New container management label</i></b>	<b><i>Current information protection label</i></b>
General access	Public
Guest access	(None)
Limited access	Internal Employee confidential Market sensitive Financial data
Confidential access	Confidential Secret Ultra-confidential

7. Update the containers with new labels as defined by the mapping exercise. This won't take long to do manually if there are only a few containers to update. If not, it's easy to script it with PowerShell (an [example is available from GitHub](#)). The functioning of the containers won't be affected if the new labels have the same settings as the old labels. In some cases, it will be more appropriate to apply a new label to a container. These cases can be handled on an individual basis after the remapping exercise is complete.

As always, testing is recommended before you make a change like this in a production environment.

## Tracking Label Changes for Containers

Organizations use sensitivity labels to apply policies to containers, but group, site, and team owners can change the label assigned to their containers, which might affect the settings applicable to the container. Microsoft doesn't support locked labels for containers (defined here as labels that cannot be changed by group owners), so organizations need to make arrangements to monitor label changes and respond appropriately. This [article discusses an example of how to approach the problem](#).

## Creating a New Sensitivity Label

Management of sensitivity labels is through the Information Protection section of the Microsoft Purview Compliance portal. Creating a new sensitivity label goes through six steps:

1. **Name:** This is the name that users see in applications. You can also add a tooltip that the apps display when users hover over the label name and an administrative description that can be whatever makes sense for your organization.
2. **Encryption:** Decide if the label invokes rights management. If yes, you need to decide what rights (permissions) to assign to different users and groups receive. A label used purely for marking purposes doesn't need to use encryption.
3. **Marking:** Decide if the label should insert text in the header or footer of documents and messages, or as a watermark (documents only).
4. **Site and Group settings:** A label can control settings for Microsoft 365 Groups, SharePoint Online sites, and Teams. The range of settings includes privacy, guest access, external sharing capability, and access from unmanaged devices.
5. **Auto-labeling:** Office apps can automatically apply a default label to emails and documents if their content matches a sensitive information type (like a credit card number) or classifier (like a resume).
6. **Endpoint data loss prevention:** Microsoft Endpoint Manager app protection policies check files to ensure that users do not move the files to unauthorized devices or storage. The definition of a managed device is one known to Azure AD through Microsoft Endpoint Manager. If a device isn't known to Azure AD, it is unmanaged and therefore liable for restriction if dictated by the label.

After going through the steps, you can review and change settings before finally saving the new label. In the following example, we create a new sensitivity label called "Employee Restricted" that protects (encrypts) files and documents and applies markings to items. To finish up, we create a label policy to publish the new label so that it shows up in applications for assignment to items by users.

### Naming

As shown in Figure 20-2, creation of a new label starts with the population of four label properties:

- **Name:** The name must be unique. Internally, Microsoft Purview also creates an immutable GUID for internal use.
- **Display name:** Clients display this value to users. Unlike the label name, you can update the display name later by editing the label in the Microsoft Purview Compliance portal or with the *Set-Label* PowerShell cmdlet.
- **Description for users:** Clients display this text to users in a tooltip to help them understand how to use the label. Apps like Word display the information when a user hovers over a label in the list shown by the Sensitivity button
- **Description for admins:** Users never see this text. It is for administrative purposes and intended to hold notes about the use and history of the label. Although it is mandatory to enter values for the other properties, you do not have to enter anything here.

## New sensitivity label

- Name & description**
- Scope
- Files & emails
- Groups & sites
- Schematized data assets (preview)
- Finish

### Name and create a tooltip for your label

The protection settings you choose for this label will be immediately enforced on the files, email messages, or content containers to which it's applied. Labeled files will be protected wherever they go, whether they're saved in the cloud or downloaded to a computer.

**Name \*** ⓘ

**Display name \*** ⓘ

**Description for users \*** ⓘ

**Description for admins** ⓘ

**Next** **Cancel**

Figure 20-2: Providing essential details for a new sensitivity label

## Label Scope

## New sensitivity label

- Name & description
- Scope**
- Files & emails
- Groups & sites
- Schematized data assets (previ...
- Finish

### Define the scope for this label

Labels can be applied directly to files, emails, containers like SharePoint sites and Teams, schematized data assets, and more. Let us know where you want this label to be used so you can configure the applicable protection settings. [Learn more about label scopes](#)

- Files & emails**  
Configure encryption and content marking settings to protect labeled emails and Office files. Also define auto-labeling conditions to automatically apply this label to sensitive content in Office, files in Azure, and more.
- Groups & sites**  
Configure privacy, access control, and other settings to protect labeled Teams, Microsoft 365 Groups, and SharePoint sites.
- Schematized data assets (preview)**  
Apply labels to files and schematized data assets in Azure Purview. Schematized data assets include SQL, Azure SQL, Azure Synapse, Azure Cosmos, AWS RDS, and more.  
 ⓘ To apply this label to Azure Purview assets, you must first turn on labeling for Azure Purview. You can do this from the Labels page. [Learn more about labeling for Azure Purview](#)

**Back** **Next** **Cancel**

Figure 20-3: Setting the scope for a new sensitivity label

As explained earlier, sensitivity labels have a scope that defines their use. In this instance, the label is available for both files and email and containers (groups and sites). To make things simple, we'll configure this label to

have a scope of Files and Email. Figure 20-3 shows that both Files and Email and Groups and Sites are configured, so we'll uncheck the second box.

If you [enable the preview feature](#), you can amend the scope of the label to support its application to files and Microsoft Purview schematized data assets, including SQL, Azure SQL, Azure Synapse, Azure Cosmos, and AWS RDS.

Because we've chosen to limit the scope of the label to files and email, the next step is to decide on what kind of protection the label can apply to these items. The choices are:

- **Marking:** Apply custom headers, footers, and watermarks. Be aware that if a label applies a custom header or footer to a document, it overwrites any header or footer information already present. This is a replacement operation, not a merge of some label-specified text with whatever's already there.
- **Encryption:** Use rights management encryption to limit access to content.

You don't need to choose either of these options. If you don't, the label becomes a purely visual indicator that users can assign to files and emails. In most cases, you'll choose to apply either marking or encryption or both.

## Encryption and Permissions

The encryption settings allow you to configure the type of protection to apply to items or to remove encryption from items. Labels that remove encryption can only be applied by people who have sufficient usage rights (Export or Full Control) for the label currently applied to an item or be the owner of the message or document. Super-users (see later) can also remove encryption from items.

When a label applies encryption (Figure 20-4), rights management defines the protection given to items. The settings are:

- **Assign Permissions now or let users decide:** The heart of rights management is the permissions given to recipients of an item. Permissions are granted by the author to individual accounts, groups, or collections of people (such as every authenticated user in a tenant). For sensitivity labels, you can decide to create a set of predefined permissions that will be the same for all labeled items, or you can allow users to assign permissions when they apply a label to messages or documents. If you assign predefined permissions, you must add permissions for at least one user or group.
- **User access to content expires:** If never, the user can continue to access labeled content without hindrance. In some cases, like a draft for a plan, you will define an expiry date. In others, you know that content is only valuable for a certain period, so you can set access to expire a specific number of days after a user applies the label to an item.
- **Allow offline access:** When a label allows offline access, the use license obtained by the user and downloaded along with the content is used to access the content. The use license is a certificate that attests to the user's right to access content and the encryption key used to decrypt the content. The normal validity of a use license is 30 days (you can choose any period between 1 and 100 days), during which the user does not need to reauthenticate to prove their access. However, you can block offline access completely (for very sensitive information) or allow access for limited periods when offline. You can then further limit access by requiring the application to check with the rights management service when online after a set period to ensure that access is still valid and hasn't been revoked. When a check is performed, group membership is re-evaluated to ensure that someone who gains access through group membership is still a member.

## New sensitivity label

**Encryption**

Control who can access files and email messages that have this label applied.  
[Learn more about encryption settings](#)

Remove encryption if the file is encrypted

Configure encryption settings

**Assign permissions now or let users decide?**

Assign permissions now

The encryption settings you choose will be automatically enforced when the label is applied to email and Office files.

**User access to content expires**

A number of days after label is applied

Access expires this many days after the label is applied

60

**Allow offline access**

Always

**Assign permissions to specific users and groups**

[Assign permissions](#)

Users and groups	Permis...

Use Double Key Encryption

Back Next Cancel

Figure 20-4: Configuring the encryption settings for a sensitivity label

If you decide the label should have predefined permissions, assign them in the **Assign permissions to specific groups** section of the form. Click the *Assign permissions* link to display the form to enter permissions (Figure 20-5). It is common practice to grant the Viewer permission to **All users and groups in the organization** as a catch-all to ensure that any account belonging to the tenant can read files. This is especially true for sensitivity labels intended for application to information that is freely available within a tenant such as confidential messages sent to large distribution lists or documents intended for internal consumption that are posted to intranet sites. The other permission types are:

- **Any authenticated users:** Grants access to anyone with an authenticated identity. The item will be accessible to anyone in any other Office 365 (Azure AD) domain or with a Microsoft Services account. Depending on the federation configuration for the tenant, the definition of an authenticated user might be even wider, including users who authenticate using a one-time passcode.
- **Users or groups:** Grants access to specific users or groups in the tenant. Use this option to assign specific rights to people who need it to interact with the content. For example, if you want a set of people to have full control over documents, assign the Co-Author permissions to a distribution list or Microsoft 365 group (this can be a dynamic group) containing these people. It's always easier to use a distribution list or group than to specify permissions for individual users.
- **Specific email address or domains:** Grants access to specific people outside your tenant (by email address) or complete domains. Use this option when you need to collaborate with protected content with people outside your tenant. For example, if you wanted to share protected information with

Microsoft, add Microsoft.com as the domain and assign an appropriate permission (like Viewer). You cannot grant permissions to the guest accounts used by applications such as Teams and SharePoint Online. Instead, you must add explicit assignments using the email addresses of the accounts (or grant access to Microsoft 365 Groups containing the accounts). External users must be able to authenticate before they can access the content. In practice this means:

- Using an Azure AD account (probably belonging to another Office 365 tenant).
- Using a Microsoft Services account, like Outlook.com.
- Using an account from a federated directory (Gmail.com or Yahoo.com). These accounts must create a Microsoft Services account using their email address.
- Another email or identity service. These accounts need a Microsoft Service account using their email address.
- On-premises directories. These accounts need to register for [RMS for Individuals](#) using their corporate email address. RMS for Individuals is a free service for users who need to open files encrypted by Microsoft Information Protection.

✕

## Assign permissions

Only the users or groups you choose will be assigned permissions to use the content that has this label applied. You can choose from existing permissions (such as Co-Owner, Co-Author, and Reviewer) or customize them to meet your needs.

- + Add all users and groups in your organization
- + Add any authenticated users ⓘ
- + Add users or groups
- + Add specific email addresses or domains ⓘ

Permissions assigned to

██████████.onmicrosoft.com



Choose permissions

Viewer  
VIEW.VIEWRIGHTSDATA.OBJMODEL

Save

Cancel



Figure 20-5: Assigning predefined permissions for a sensitivity label

## Use Double Key Encryption

For extra security, organizations can run a Double Key Encryption (DKE) service to manage encryption keys. When a sensitivity label uses double key encryption, two keys are necessary to access a protected item's content. Microsoft Information Protection manages one key; the DKE service manages the other key, which remains under the control of the organization. If the DKE service is inaccessible, the content cannot be decrypted. To configure double key encryption for a label, you must provide the URL of the DKE service. Because of the additional cost and complexity, most organizations don't use double-key encryption. [Double-key encryption](#) requires Office 365 E5 or Microsoft 365 E5 licenses. While deploying a DKE service on Azure is the fastest implementation method, it's also possible to [run the DKE service using an on-premises Windows server](#).

## Permissions and Usage Rights

Sets of individual usage rights form the permissions used in sensitivity labels. The usage rights are:



- View content
- View rights.
- Edit content.
- Save.
- Print.
- Copy and extract content.
- Reply.
- Reply all.
- Forward.
- Edit rights.
- Export content.
- Allow macros.
- Full control.

To make usage rights easier to manage and assign, information protection uses predefined permission sets such as Co-Author, Reviewer, and Viewer. Each permission set is composed of a set of usage rights. Assigning predefined permissions to someone is a convenient way of giving them all the usage rights defined in the permissions. However, if none of the predefined permissions meet your needs, you can choose to define custom permissions made up of appropriate usage rights.

When someone receives protected content, information protection checks their signed-in account against the access granted to the users and groups specified in the label. If the account is not present in the list of permissions, they won't be able to open the content. Whenever possible, it is best to grant rights to groups rather than individuals as this makes permissions much easier to manage.

Information protection uses Azure AD accounts to authenticate access to content protected by encryption imposed by sensitivity labels. This isn't an issue if you share documents with people in another Microsoft 365 organization as the recipients authenticate against their home tenant directory. Likewise, it's not an issue if an external recipient has a Microsoft Services account because these accounts also authenticate against a Microsoft directory.

If you need to share protected content with people who don't have an Azure AD account, you can create guest accounts in your directory and the external people can sign in using those accounts, just like external recipients of SharePoint Online sharing links sign in to access shared documents. The [SharePoint Online and OneDrive for Business integration with Azure AD](#) creates guest accounts automatically for sharing. If you use this integration, guest accounts may well exist for the people with whom you share protected content.

If you assign access to an external user in a label, it's a good idea to check if their email address belongs to a Microsoft 365 organization, and if not, create a guest account. Remember to inform the person that they'll need to use the guest account to sign in to access protected content received from your organization. It's also wise to remind the external user that they must use a suitable application to access protected content, such as Microsoft 365 apps for enterprise or the standalone edition of Office 2019. See [this page](#) for more information.

## User-Defined Permissions

If you don't assign preset permissions to a label, you can allow users to control the permissions assigned when they apply the label to an item. You can enable protection for Outlook or the other Office apps, or both (shown in Figure 20-6):

- Outlook: the permissions on the message can be set to either *Do Not Forward* or *Encrypt Only*. Unprotected Office attachments inherit the same protection from the message. Files that have protection (through another label) before being attached to a message keep the protection assigned

by that label, even if the label assigned to the email is more sensitive than one assigned to an attachment.

- Word, PowerPoint, and Excel: the app prompts authors to define the permissions (like Viewer, Reviewer, or Co-Author) they wish other users to receive.

### Encryption

Control who can access files and email messages that have this label applied. Learn more about encryption settings

- Remove encryption if the file or email is encrypted
- Configure encryption settings

Turning on encryption impacts Office files (Word, PowerPoint, Excel) that have this label applied. Because the files will be encrypted for security reasons, performance will be slow when the files are opened or saved, and some SharePoint and OneDrive features will be limited or unavailable. Learn more

Assign permissions now or let users decide?

Let users assign permissions when they apply the label

The labeling behavior for these settings varies depending on which operating system platform is used to apply the label. Learn more

- In Outlook, enforce one of the following restrictions
- Do Not Forward
  - Encrypt-Only
- In Word, PowerPoint, and Excel, prompt users to specify permissions

Figure 20-6: Defining user assigned permissions for a sensitivity label

If you don't set the Outlook checkbox, the label doesn't appear in the set shown in Outlook clients. The same is true for Word, Excel, and PowerPoint if you don't select this setting.

User-assigned permissions is a premium feature limited to the Office apps on Windows and Mac. Online apps don't currently support user-assigned permissions because SharePoint Online doesn't offer the same support for files protected with these labels as it does for files encrypted by labels with fixed permissions. Microsoft is due to deliver [a preview for labels with user-defined permissions in SharePoint Online](#) in July 2022.

### Encryption Applied by OWA and Outlook Mobile

Outlook desktop clients encrypt messages before submitting emails to the Exchange Online transport service for onward processing. Other clients do not include the necessary code to protect messages. These clients stamp outbound messages with the metadata for a label and rely on the transport service to apply appropriate protection. You know if the transport service applies protection to a message if the copy of the message in the Sent Items folder has a label but is unencrypted.

OWA can apply the OME Do Not Forward and Encrypt Only protection to outbound messages before submission to the Exchange transport service, but OWA cannot encrypt messages for sensitivity labels using either preassigned or user-defined permissions. Outlook Mobile always submits messages to Exchange Online for encryption.

### Document Owners

When someone applies a sensitivity label with encryption to a message or document, they are regarded as the owner or issuer of rights for that content. The owner always has the right to access content, even if the policy sets an expiry date. Likewise, the owner can always access content offline or open it after it has been revoked.

### Content Marking

If you decide that a sensitivity label should apply content marking to files and documents, you can configure settings (Figure 20-7) to control the visual indicators inserted by applications after a user applies a label. Word, Excel, and PowerPoint support headers, footers, and watermarks and insert the markings as soon as a

label is assigned to a file, while Outlook only supports headers and footers and inserts the text when a message is saved.

## Content marking

Add custom headers, footers, and watermarks to content that has this label applied. [Learn more about content marking](#)

 All content marking will be applied to documents but only headers and footers will be applied to email messages.

### Content marking



Add a watermark  
Customize text

Add a header  
Customize text

Add a footer  
Customize text  
Employee Restricted

Back

Next

Cancel

Figure 20-7: Configuring content marking settings for a sensitivity label

Headers and footers can be up to 1024 characters, except for Excel, which limits these markings to 255 characters. Watermarks can be up to 255 characters.

## Performance and Encryption

Microsoft warns that some degradation in performance is possible when accessing protected Word, Excel, and PowerPoint files. This is because users must authenticate with the protection service to establish their right to access the content, after which the app can decrypt the content. These operations are unnecessary for unprotected files. Usually, the slight delay in opening documents is not very noticeable. Large files always take some time to open, so the need to decrypt a large Word document or Excel spreadsheet will add a small amount to the time needed to open these files.

## Client-Side Auto-Labeling with Sensitivity Labels

Client applications like the Office applications and OWA use the MIP SDK to detect sensitive information types or classifiers and apply sensitivity labels automatically when a match is found in a message or document. For example, OWA checks an item's content and applies automatic labeling if matches are found when messages are sent. The check processes the content of the message body (but not header information or the message subject) to decide if a match for the specified sensitive data is present. The other Office applications perform automatic labeling when they save files.

The sensitive information types used for the automatic application of sensitivity labels are those used by other applications such as DLP policies and auto-label policies. Checking for matches uses the same concept of detecting a certain number of occurrences of the data type in content together with meeting a set confidence level that the data type is what it seems before deciding that a match exists (see the DLP chapter for more information).

In Figure 20-8 we see the properties of a sensitivity label with auto-labeling enabled. In this case, the label is applied when a client detects the existence of a single debit card number. Typically, you combine this feature with content marking to apply a header or footer to warn users that sensitive data is present and perhaps

note when the label encrypts content. The *message displayed to the user* property is a policy tip displayed in a banner to communicate with users when their items receive labels automatically.

Several sensitivity labels published in a policy might invoke automatic labeling and match against an item's content. When this happens, the label with the highest priority (as ordered in the sensitivity label policy published to the user) is applied. The automatic application of sensitivity labels is a premium feature.

## Auto-labeling for Office apps

Auto-labeling is supported in Office apps for users who have either Office 365 ProPlus or the Azure Information Protection unified labeling client installed. When we detect sensitive content in email or documents matching the conditions you choose, we can automatically apply this label or show a message to users recommending they apply it themselves. [Learn more about auto-labeling](#)

### Auto-labeling for Office apps



**i** Since encryption is turned on, a large amount of content might be automatically encrypted when this label is applied. Turning on encryption impacts Office files (Word, PowerPoint, Excel) that have this label applied. Because the files will be encrypted for security reasons, performance might be degraded when the files are opened or saved, and some SharePoint and OneDrive features might be limited or unavailable. [Learn more](#)

#### ^ Detect content that contains

^ Content contains
🗑️

All of these
▼
🗑️

**Sensitive info types**

EU Debit Card Number	Accuracy	85	to	100	Instance count	1	to	Any	🗑️
----------------------	----------	----	----	-----	----------------	---	----	-----	----

Add ▼

Create group

+ Add condition ▼

#### When content matches these conditions

Automatically apply the label ▼

#### Message displayed to user

Provide policy tip for the user

Back

Next

Cancel

Figure 20-8: Setting the auto-labeling properties for a sensitivity label

## Configuring Container Management Settings for Sensitivity Labels

Sensitivity labels configured with a scope of "Groups and sites" have additional settings to control the management of containers (teams, groups, and SharePoint Online sites). We suggest that you separate labels used for container management from those used to protect files and email. This approach reduces the number of labels displayed to users and makes it easier for them to select the most appropriate label.

The visual marking aspect of container management via sensitivity labels replaces the previous markings set through text-only classifications defined in the Azure AD policy for Groups (see Chapter 11). Sensitivity labels serve the same purpose as classifications in that users see a visual indicator (like "Secret" or "Confidential") to remind them of the importance or sensitivity of the container. However, sensitivity labels have the extra

benefit of being able to impose access controls on the containers. You can, for instance, assign a label to a team to stop the team owner from being able to invite guests from outside the tenant to join the team membership. To make container management even more useful, Microsoft has said that they plan to increase the range of management settings available through sensitivity labels in the future.

Outlook desktop, OWA, Teams, and SharePoint Online support container management. If you create or edit groups using these apps, you can apply sensitivity labels to the groups. The app will then respect the settings inherited from the applied label and synchronize the label with the other workloads to ensure compliance in those apps.

**Container Labels Don't Affect Content:** It's important to realize that sensitivity labels applied to the container level do not affect the individual messages, conversations, and files stored in these containers. You can certainly assign sensitivity labels to individual files stored in document libraries, but these assignments are independent of anything applied to the container.

Four settings for container controls are available, presented on two screens. You can configure any or all these settings.

**Privacy and external access** (Figure 20-9):

- *Privacy* controls if the group is *Public* (anyone can join) or *Private* (members must be invited to join). The setting can also be set to *None*, which means that the group owner decides which level of access to apply to the group.
- *External user access* controls if the group membership can include guest users. If you block guest users for the group, it assigns an Azure AD policy to the group and sets the value of *AllowedToAddGuests* to *False*. Blocking external access does not remove existing guests from the group membership: it only blocks the addition of new guests.

### Define privacy and external user access settings

Control the level of access that internal and external users will have to labeled teams and Microsoft 365 Groups.

#### Privacy

These options apply to all Microsoft 365 Groups and teams that have this label applied. When applied, these settings will replace any existing privacy settings for the team or group. If the label is removed, users can change it again.

- Public. Anyone in your organization can access the group or team (including content) and add members.
- Private. Only team owners and members can access the group or team, and only owners can add members.
- None. Team and group members can set the privacy settings themselves.

#### External user access

Let Microsoft 365 Group owners add people outside your organization to the group as guests. [Learn about guest access](#)

Back

Next

Cancel

Figure 20-9: Sensitivity label container management settings for privacy and user access

**External sharing and conditional access settings** (Figure 20-10): These settings control access to content in SharePoint Online team sites. Microsoft requires that any account using sensitivity labels to manage settings for group-connected sites has a premium license.

Two settings are available:

- **Control external sharing from labeled SharePoint sites:** The capability to share documents with external users from a site is controlled by a default setting for the tenant which can be overridden at the site level, which is what happens when a label is applied to a site. The available options are listed below (the relevant value used to set external sharing capability for a site by running the *Set-SPOSite* cmdlet is in parenthesis).
  - **Anyone** (ExternalUserAndGuestSharing): Sharing is allowed with all external users, and documents can be shared using anonymous access links (Anyone links).
  - **New and existing guests** (ExternalUserSharingOnly): Sharing is allowed with new external users, who must accept a sharing invitation and go through an authentication process to create a guest account.
  - **Existing guests** (ExistingExternalUserSharingOnly): Sharing is only allowed with the guest users already in an organization's directory.
  - **Only people in your organization** (Disabled): No sharing with external users is allowed.
- **Use Azure AD Conditional Access to protect labeled SharePoint sites:** Limiting access to unmanaged devices depends on Azure AD conditional access policies. Policy evaluation during the authentication process identifies the managed state of a device (by the organization or managed by the user). If the device is unmanaged, SharePoint Online can apply the restriction set in the label. See [this page for directions](#) on how to create the necessary conditional access policy.

## Define external sharing and conditional access settings

Control who can share SharePoint content with people outside your organization and decide whether users can access labeled sites from unmanaged devices.

### Control external sharing from labeled SharePoint sites

When this label is applied to a SharePoint site, these settings will replace existing external sharing settings configured for the site.

#### Content can be shared with

- Anyone ⓘ  
Users can share files and folders using links that don't require sign-in.
- New and existing guests ⓘ  
Guests must sign in or provide a verification code.
- Existing guests ⓘ  
Only guests in your organization's directory.
- Only people in your organization  
No external sharing allowed.

### Use Azure AD Conditional Access to protect labeled SharePoint sites

You can either control the level of access users have from unmanaged devices or select an existing authentication context to enforce restrictions.

- Determine whether users can access SharePoint sites from unmanaged devices (which are devices that aren't hybrid Azure AD joined or enrolled in Intune).
- ⓘ For this setting to work, you must also configure the SharePoint feature that blocks or limits access to SharePoint files from unmanaged devices. [Learn more](#)
- Allow full access from desktop apps, mobile apps, and the web
- Allow limited, web-only access ⓘ
- Block access ⓘ
- Choose an existing authentication context (preview). Each context has an Azure AD Conditional Access policy applied to enforce restrictions. [Learn more about authentication context](#)

Back

Next

Cancel

Figure 20-10: Configuring external sharing and device access settings for container management

Remember that the external sharing capability assigned to a site cannot be less restrictive than allowed by the tenant-level setting. Microsoft Purview does not check the tenant-level setting when adding sensitivity labels,

so someone can create a label that allows a less restrictive external sharing capability than allowed by the tenant. For instance, if the tenant bars *Anyone* sharing links, a label should not select this level. If a label has an invalid sharing capability, SharePoint Online ignores it when it reads settings from the label.

**Label Synchronization with SharePoint Online:** SharePoint caches the container settings in sensitivity labels to improve performance. This means that changes to label settings that affect the operation of a site, such as an update to the external sharing capability will not become effective for at least 24 hours after the update. Settings imposed by new sensitivity labels apply 15 minutes or so after the publication of a new label to SharePoint.

## Conditional Access and Authentication Context

Highly confidential SharePoint sites often need special attention to ensure that connections are secure. An Azure AD authentication context is a way of marking resources like SharePoint sites as needing special processing by Azure AD conditional access policies. Now [in preview](#), a sensitivity label can be associated with an authentication context in its external sharing settings. When such a link exists, Azure AD invokes conditional access policies which include the authentication context whenever someone attempts to access SharePoint sites stamped with the label. For example, if a conditional access policy with an authentication context requires connections to use multi-factor authentication, attempts to connect to sites with labels linked to the authentication context will fail unless authenticated with MFA. Information about the selected authentication context is in the *ProtectionLevel* value in the *LabelActions* setting of the label.

## Coordinating Label Updates Across Apps

A change made to container settings in one app might have unforeseen consequences for another. For example, an Exchange administrator might apply a label that prohibits guest users to a group with OWA. The SharePoint Online and Teams synchronizes the update for the assigned label, and it becomes active for the site and team. While the newly-assigned label won't stop existing guests from accessing the site or team, it will prevent site and team owners from adding new guests, which could come as a surprise to them. For this reason, it's wise to update group owners before applying a more restrictive label to their group.

## Publishing Sensitivity Labels

Before sensitivity labels show up in applications and users can apply the labels to content or containers, a label policy must publish the labels. A label policy consists of:

- One or more sensitivity labels.
- A target audience. The default audience is everyone in the tenant. You can specify an audience by selecting Microsoft 365 Groups, security groups, distribution lists, or individual users (the sole requirement is that the objects must be mail-enabled, which excludes some security groups). Note that sensitivity label policies publish labels to people while retention label policies publish labels to storage locations like sites and mailboxes.
- Settings to define whether one of the labels in the policy is mandatory and applied to emails and Office documents and if users must give a justification if they remove a label or replace a label with a lower classification. A setting is also available to define a custom help page for users to consult to learn about the proper use of sensitivity labels within the organization. Label use settings don't apply to labels created solely for container management.

Clients that support sensitivity labels learn about new labels or changes to existing labels from the underlying workload. The time required for a client to acquire details about sensitivity labels vary, but you should anticipate that several hours are necessary. It all depends on when the client refreshes its cache of label information.

A tenant can have multiple sensitivity label publishing policies, each of which has different sensitivity labels and target audiences. For instance, you might create a general-purpose policy to publish a default set of sensitivity labels to everyone in the tenant and then have a set of specific policies to publish certain labels to specific groups. Another policy might publish a set of container management labels to users allowed to create groups, teams, and sites. Multiple sensitivity label publishing policies might cover the same user. If this is the case, clients combine the labels from all applicable policies to create a single set of labels to display to the user. In addition, if more than one policy requires mandatory labeling, Office applies the sensitivity label specified in the highest priority policy.

## Select Labels and Target Audiences

To create a policy for a single label, select the label from the list and then click **Publish label**. The policy publication wizard starts with the selected label already added. You can add other labels to the policy with the **Edit** link. Otherwise, you can create a new labels publication policy and add the set of labels you want to publish (Figure 20-11).

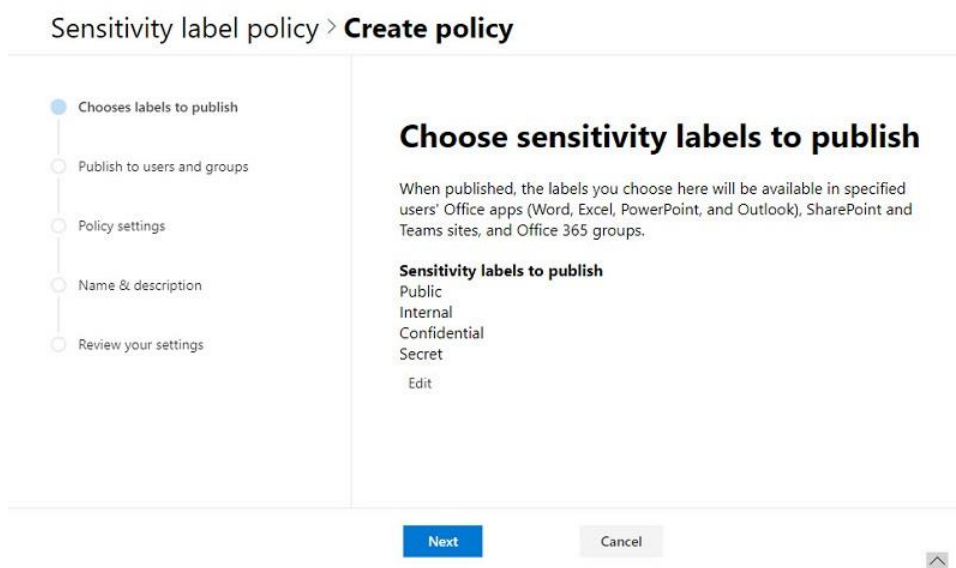


Figure 20-11: Beginning the publication process for a set of labels

The next step is to select the target audience who will be able to use the labels published in the policy. In Figure 20-12, all users and groups receive the policy, which is the norm for labels in general use across the tenant. If you want to limit publication to specific users or groups, click **Choose users or groups**. You can select individual users, Microsoft 365 Groups, or distribution lists. Members of the groups keep access to the labels defined in the policy for as long as they are members. New members added to groups benefit from permissions granted in labels when they join.



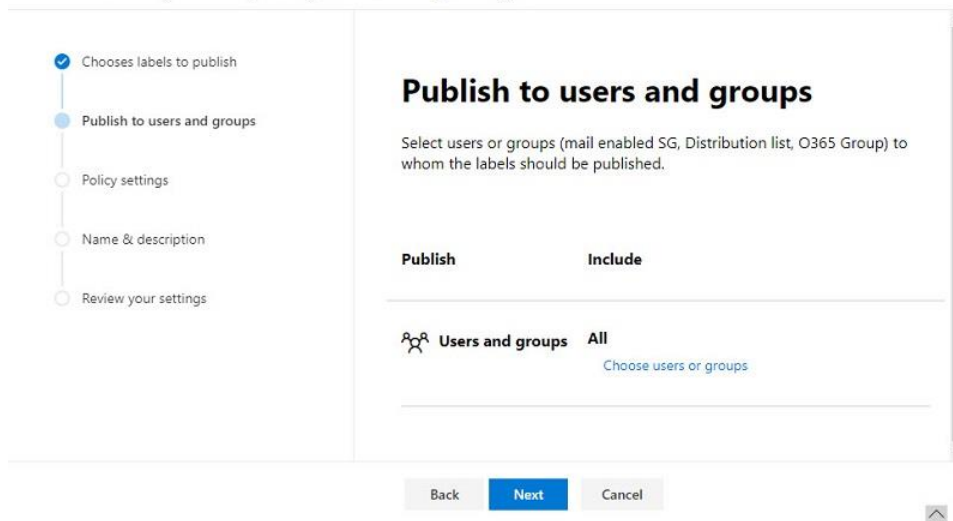
Sensitivity label policy > **Create policy**

Figure 20-12: Specifying the target audience for the sensitivity policy label

When you define a target audience for a policy, make sure that the users in that audience have the necessary rights to access any encrypted content protected by the labels included in the policy. For policies used to publish container management labels, the target users should be the accounts allowed to create new groups. In some tenants, everyone can create new groups; other organizations restrict group creation by policy to the members of a specific group. See the section about the Groups creation policy in Chapter 11 for more information.

### Excluding Mailboxes in a Sensitivity Label Policy

You might notice that the GUI to create sensitivity label policies allows administrators to select All or specific groups as the target (Figure 20-12). All means every user mailbox in the tenant. However, the Microsoft Purview Compliance portal doesn't support the exclusion of specific users when a policy uses the special *All* destination. Fortunately, this is possible with PowerShell. For example, this command excludes the mailboxes of Terry Hegarty and Kim Akers from receiving the labels published in the specified policy:

```
[PS] C:\> Set-LabelPolicy -Identity "General Sensitivity Policy" -AddExchangeLocationException "Terry.Hegarty@Office365itpros.com", "Kim.Akers@office365itpros.com"
```

```
Get-LabelPolicy -Identity "General Sensitivity Policy" | Select ExchangeLocationException
ExchangeLocationException
-----
{Kim Akers, Terry Hegarty}
```

Adding a mailbox to a label publishing policy in this manner does not overwrite the set of excluded mailboxes. To remove an excluded mailbox, run *Set-LabelPolicy* and pass the mailbox name in the *RemoveExchangeLocationException* parameter.

```
[PS] C:\> Set-LabelPolicy -Identity "General Sensitivity Policy" -RemoveExchangeLocationException Kim.Akers
```

Running the cmdlet to add more than a few mailboxes can become tiresome. In these circumstances, it's better to find the set of mailboxes using *Get-ExoMailbox* or another method and pipe the set of mailboxes to *Set-LabelPolicy*. For example, this code extracts the mailbox members of a distribution list into an array and uses the array to add policy exclusions:

```
[PS] C:\> [array]$Members = Get-DistributionGroupMember -Identity "Planner Gurus" | ?
{$_ .RecipientTypeDetails -eq "UserMailbox"} | Select -ExpandProperty PrimarySmtpAddress
```

## Publication Policy Settings

Settings in label publication policies allow the selection of a default label for assignment to new messages and documents. In Figure 20-13, we require users to apply a label to their emails and documents. Default labels only work with clients that support information protection during content creation, such as the Office apps. By comparison, if someone sends an email when connected to Exchange Online via IMAP4 with the Thunderbird client, nothing will happen.

Depending on what types of sensitivity labels a label publishing policy includes, settings are available to control if users must apply labels to:

- Documents.
- Emails (can be the same as for Documents).
- Containers (groups, teams, and sites – if the policy includes these labels).
- Power BI content.

For each content type, you can define a default label to apply. To make it easier to manage, it's often best to have separate label publishing policies for different content types.

You can control if users must provide a free-text justification if they change the assigned label to another label with a lower classification. The justification is logged in the Information Protection logs rather than the Office 365 audit log.

The policy settings for default label assignment are not retrospective, so unlabeled files and emails remain in this state until users or auto-label policies assign them labels. It's easier to deal with unlabeled groups because it's possible to search for unlabeled groups with PowerShell and then assign a suitable label to those groups.

The policy can define a different label for Outlook to apply to new emails or for Power BI to apply to its content. To aid in driving user awareness about information protection, the *Provide users with a link to a custom help page* setting allows organizations to define a web page for users to view if they want added information about how to use sensitivity labels to mark content.

**Policy settings**

Configure settings for the labels included in this policy.

- Users must provide a justification to remove a label or lower its classification**  
Users will need to provide a justification before removing a label or replacing it with a one that has a lower-order number. You can use activity explorer to review label changes and justification text.
- Require users to apply a label to their emails and documents**  
Users will be required to apply labels before they can save documents, send emails, and create groups or sites (only if these items don't already have a label applied).  
Support and behavior for this setting varies across apps and platforms. [Learn more](#)
- Require users to apply a label to their Power BI content**  
Users will be required to apply labels to unlabeled content they create or edit in Power BI. [Learn more about mandatory labeling in Power BI](#)
- Provide users with a link to a custom help page**  
If you created a website dedicated to helping users understand how to use labels in your org, enter the URL here. [Learn more about this help page](#)

Figure 20-13: Defining settings for the sensitivity label policy

If you configure a default label for messages, Outlook creates new messages with the selected label in place. If the user then chooses to send using a different mailbox (group, user, or shared mailbox), Outlook prompts them to justify the removal of the label. This happens because the default label is associated with the original

account, and when the user selects a different address to send from, Outlook removes the default label. It's logical because the account used to send messages might not have permission to use the original label.

The ultimate steps in the publication process are to create a name and description for the policy (you only need to enter a name), review the policy, and then publish it. You can't change the name of a label policy, but you can update its description at any time.

**Audit Who Changes Sensitivity Label Policies:** You might want to know who makes changes to sensitivity label policies. The easiest way to find out is to search the Office 365 audit log for the *Set-LabelPolicy* event. The *AuditData* property in the audit event contains details of the changes made while the user who updated the policy is in the *UserIds* property.

## Some Delay Before Policy Changes Become Effective

Don't expect people to be able to use new or updated sensitivity labels immediately after making changes. The publishing process takes time to enact policy changes by performing tasks like creating or updating labels. Also, multiple services must be updated with the new policy settings. Applications pick up policy updates when they refresh their cache of tenant policies, usually within a couple of hours of an update. At this point, the applications combine the labels published through all the policies that cover the user and present the labels in a single list. Users cannot see details of the policies through which they access labels.

## Multiple Sensitivity Label Policies

You can assign multiple sensitivity label policies to an account. If you do, the account has access to the combined set of sensitivity labels from all assigned policies, and the policy settings which apply are determined by the order the policies are listed in the Microsoft Purview Compliance portal with the [lowest priority policy shown at the top](#) and the highest at the bottom.

## Apply Sensitivity Labels Automatically

For licensing purposes, Microsoft differentiates between manual and automatic application of sensitivity labels to documents and email. The basic rules are:

- Office 365 E3 covers manual application. For instance, a user opens a Word document and selects a sensitivity label.
- Users must have Office 365 E5 or Microsoft 365 E5 compliance licenses to cover the automatic application of sensitivity labels to content. For example, default label assignment to email and documents is automatic labeling, Automatic processing also includes auto-label policies where clients or servers apply labels based on the content of messages or documents.

People might not consider default label assignment as automatic, but it is. A more advanced example is when you create an auto-label policy to assign sensitivity labels based on the content of a document or message. Auto-label policies make it easier for users to do the right thing to protect information but only for new content.

To address the problem, tenants can use server-side processing to auto-label email and documents, sometimes referred to as labeling content at rest (SharePoint Online and OneDrive for Business) or in transit (Exchange Online). Server-side processing happens at scale and isn't dependent on clients, so it's appropriate when tenants have large quantities of information to protect. Organizations can define server-side automatic label policies in the Information protection section of the Microsoft Purview Compliance portal.

## Points to Consider About Server-Side Sensitivity Label Processing

The following conditions apply to server-side sensitivity label processing:

- A tenant can deploy up to 100 auto-label policies to apply sensitivity labels. Each policy can cover all SharePoint Online sites and OneDrive for Business accounts in the tenant (until mid-August 2021, a policy could cover a maximum of ten locations).
- The auto-label process can handle a maximum of 25,000 items per day per tenant.
- Server-side auto-labeling only works for Office documents in Open XML format (older formats like .doc Word documents don't support sensitivity labels). When an auto-label policy applies a sensitivity label to documents, it does not update their modification dates.
- Server-side auto-labeling can replace labels assigned by client-side auto-labeling, but only if the label applied by the client has a lower priority.
- Auto-label policies for SharePoint Online and Exchange Online use different rules to locate items for processing. You can define rules with the same effect for both locations. The rules for Exchange Online processing have additional conditions to match against message properties.
- Policies for Exchange Online can apply to all mailboxes or selected mailboxes.
- When an auto-label policy scans Exchange Online messages, it processes both attachments and message bodies. However, if the policy finds a match, it applies the sensitivity label only to the message. In terms of rights management, the sender of the message is the issuer/owner.
- When a label applies encryption, it encrypts both the message and some attachments. The encrypted attachments are Office documents and PDF files (if *EnablePdfEncryption* is True in the tenant's IRM configuration).
- If an auto-label policy matches messages protected using *Encrypt Only* or *Do Not Forward*, the sensitivity label defined in the policy replaces the protection.
- Encryption applied by sensitivity labels overrides any encryption applied by Exchange mail flow rules or data loss prevention policies. If the sensitivity label doesn't include encryption, then encryption by mail flow rules or data loss prevention policies apply.
- Incoming email passes through the Exchange transport service. It is at this point that auto-label policies evaluate messages and apply labels to matching messages. However, as the sender of these messages doesn't come from your organization, external messages encrypted by an auto-label policy don't have a valid owner. This won't matter in the normal course as recipients will have the necessary rights to read the messages.

Due to the possibility that an administrator mistake in policy configuration might auto-apply incorrect sensitivity labels to large numbers of files, the process of setting up an auto-label policy involves a simulation period where Microsoft 365 reports results of auto-application without assigning labels. This phase allows those responsible for creating auto-label policies to evaluate the effectiveness of the policy settings and adjust where necessary before releasing the policy to assign labels.

SharePoint Online uses a background process to evaluate auto-label policies. The process scans files to find instances of sensitive information types that match the rules set in policies. The scan for Exchange happens when messages pass through the transport pipeline. When a policy detects a match, the process applies the sensitivity label unless a user-applied label exists (explicit assignment always beats auto-assignment). Like labels applied by users, labels applied to documents by auto-label policies remain in place for the lifetime of documents, even if they move out of the target sites processed by policies.

Auto-labels have a simulation or test mode, meaning that you can discover what items match the rules set in a policy. The idea is to allow administrators to tune policy rules by seeing the effect changes to rule conditions have in terms of their ability to detect matching items in the target locations.

## Using Machine Learning to Apply or Recommend Sensitivity Labels

A variant of automatic label assignment uses machine learning to identify patterns in Office documents to decide to apply a sensitivity label automatically. Trainable Classifiers created by Microsoft or custom classifiers created within the tenant describe patterns to identify document types such as a resume (describing a

document with job details) or code (computer code), and some to help identify problematic documents like those containing offensive, threatening, or profane content. Label policy settings then control the automatic application of sensitivity labels. See the section about trainable classifiers in the compliance chapter.

## Removing Sensitivity Labels

It's not a good idea to delete a sensitivity label after users have applied the label to items. When an administrator removes a sensitivity label from the label publishing policies it is in, the label becomes unavailable to clients and invisible to users. Because the label metadata remains in place, any encryption applied by the label is intact. Some apps, like Office, can continue displaying the name of the label. However, users cannot apply unpublished labels to new content.

You can't delete a sensitivity label until after removing it from label publishing policies. If you attempt to delete a label without unpublishing it, Purview flags an error and shows which label policies include the label. To proceed and remove the label from the tenant, you must remove it first from the publishing policies.

After the deletion of a label, the underlying protection template enters an archived state. This ensures that protection remains for labeled content. Users won't be aware that files have labels because applications cannot resolve the label name against an active template. It's not possible to recreate deleted labels through the Microsoft Purview Compliance portal. Instead, you can recreate the label with PowerShell by running the *New-Label* cmdlet and (critically) passing the identifier (GUID) of the removed label in the Identity parameter. This will reconnect the new label to the protection template. After a period to refresh client caches, the reconstituted label should reappear in applications.

Because of the issues involved in deleting sensitivity labels, it is always better to remove unwanted sensitivity labels from all label policies to unpublish them from clients instead of deleting the labels.

The removal of sensitivity labels used for container management is slightly easier. First, create a list of the containers with the sensitivity label you wish to remove. Use the *Get-Label* cmdlet to find the label GUID:

```
[PS] C:\> $LabelGuid = (Get-Label -Identity "Confidential Access").ImmutableId
```

Now find the containers that have this sensitivity label. This example uses the *Get-UnifiedGroup* cmdlet to find the Microsoft 365 groups with the label.

```
[PS] C:\> [array]$Containers = Get-UnifiedGroup -ResultSize Unlimited | ? {$_.SensitivityLabel -eq $LabelGuid.Guid}
Write-Host ("{0} groups found with the sensitivity label {1}" -f $Containers.count, $LabelGuid.Guid)
```

If the check finds some groups with the sensitivity label, you should replace it with the GUID of another label to maintain the settings applied to the containers. A loop like this does the job.

```
[PS] C:\> $NewLabelGuid = "d6cfd185-f31c-4508-ae40-229ff18a9919"
ForEach ($Container in $Containers) {
    Write-Host ("Updating group {0} with new label {1}" -f $Container.DisplayName, $NewLabelGuid)
    Set-UnifiedGroup -Identity $Container.ExternalDirectoryObjectId -SensitivityLabel $NewLabelGuid }
```

Exchange Online will then synchronize the update with SharePoint Online and Teams to make sure that these workloads know that the containers have a different sensitivity label. This process may take a few days to complete. When it is complete, and you see that the containers show the new label in SharePoint Online and Teams, you can remove the sensitivity label that you replaced.

## Remove Locations Rather Than Remove Policies

Along the same line, it's better to make a sensitivity label policy unavailable instead of deleting it as this allows for reinstatement of the policy if needed later. To do this, create a new blank group (or one with just an owner) and edit the policy so that the publication target is just that group. The policy will be withdrawn from

the users and groups it was previously published to and published to the empty group, which has the effect of nullifying the policy and putting it into a state where it cannot be used by anyone. If required, the policy is easily reinstated by editing it to publish the policy to a new set of target users and groups.

## Using Sensitivity Labels with Auto-Signature Products

Many ISV products insert autosignature text into outbound messages. The autosignatures usually include personal details about the sender plus some organizational information and a company logo. If you apply a sensitivity label that protects content with encryption to an email, autosignature products might not be able to process messages because Outlook or OWA encrypts the content when sending the messages. Some products have client-side plug-ins that work by inserting text during message creation. These usually work because the client inserts the autosignature before encryption. [An Outlook API](#) makes it easier for ISVs to apply signatures to protected messages across all Outlook clients.

If you want to send encrypted emails with autosignatures, you should test the available products to find one that supports sensitivity labels. Alternatively, you can use a transport rule to insert an autosignature as messages pass through the transport pipeline. A transport rule can insert autosignature text even when messages are encrypted because Exchange Online uses super-user privilege to decrypt the content, apply the autosignature, and then encrypt the message again.

## Sensitivity Labels and Power BI

If enabled through the [tenant settings section of the Power BI admin portal](#), Power BI users with Pro (not free) licenses can apply sensitivity labels to reports, dashboards, datasets, embedded reports, and dataflows. Sensitivity labels are inherited by Excel spreadsheets generated through Power BI's PivotTable connection and when the Analyze in Excel feature is used. Label inheritance also occurs when new reports or dashboards are created from a dataset. By default, any licensed user can apply sensitivity labels to Power BI objects, but you can restrict access by allowing or excluding specific security groups.

Inside Power BI, sensitivity labels are used as visual markers of the relative sensitivity of the information in items. Labeled objects are not encrypted by labels. However, when objects are exported from Power BI to Excel, PowerPoint, or PDF, encryption is applied to the output file if imposed by the label. Users who export protected information from Power BI can edit the content, but they cannot change the label applied to the file. This is because the owner logged for the file is a service rather than a person, as we can see by running the *Get-AipFileStatus* cmdlet from the AIP module:

```
[PS] C:\> Get-AIPFileStatus "basic data.pptx"
File           : Basic data.pptx
IsLabeled      : True
LabelId        : 81955691-b8e8-4a81-b7b4-ab32b130bff5
Label          : Secret
Method         : Privileged
Date           : 6/16/2020 11:17 AM
RMSGuid        : c7fc2174-097c-4123-9cad-15f1a32cb145
RMSTemplate    : Secret
Owner          : 00000009-0000-0000-c000-000000000000@04dac1c9-6661-42f4-b974-c68551262cff.rms.eu.a
                adrm.com
```

Sensitivity labels appear in the workspace list as a visual reminder to users of sensitive data. Labels are unsupported for template apps, and you can't apply the Do Not Forward label, labels with user-defined permissions, or labels based on HYOK.

# Protecting SharePoint Online and OneDrive for Business

Microsoft 365 includes two methods to protect content stored in SharePoint and OneDrive for Business sites.

- Rights management can protect items downloaded from a library or list. This is the older form of protection also available in SharePoint on-premises designed to limit the set of actions users can take after they download files. It does not protect individual items held in the library or list.
- Sensitivity labels can protect individual documents, folders, and lists in a library. The label and associated protection stay with their assigned items even if the documents move from their original location. This is the preferred mechanism to protect documents stored in SharePoint Online and OneDrive for Business sites. Protection by sensitivity labels is also available for [some components used by SharePoint Syntex](#).

Apart from being an old-fashioned method to protect SharePoint content, the first approach is less preferable because protection depends on files being in a certain library and protection occurs when users download documents from the library. Assigning protection through labels applied to individual files is better because the protection persists no matter where the file travels, including outside the organization.

Optionally, file policies created by Microsoft Defender for Cloud Apps can inspect SharePoint Online documents and if sensitive data is in a document, apply a sensitivity label. [This functionality](#) is outside the scope of this book, but it's worth investigating if your organization uses Microsoft Defender for Cloud Apps.

## Comparing Sensitivity Labels and Traditional SharePoint Protection

Sensitivity labels are the preferred way to protect content stored in SharePoint Online and OneDrive for Business. They are more flexible and powerful than the traditional approach of protecting SharePoint libraries with IRM. The advantages of sensitivity labels include:

- Support for labels in a wide range of clients including desktop, browser, and mobile apps.
- Labels can apply visual markings to content in addition to protection.
- Because rights management underpins labels, granular control is available to control who can do what with a file.
- Labels become part of the metadata of files and messages and protection travels with content as it moves between libraries or outside a tenant.
- Documents can automatically (by label policy, DLP policies, or transport rule) or manually (by users).
- Labels can be used to assign sensitivity markings and some group settings to Microsoft 365 Groups, Teams, and SharePoint containers.
- Documents protected by sensitivity labels support advanced features like co-authoring and auto-save.
- SharePoint Online populates a sensitivity column to show the label applied to files (the column is not available in OneDrive for Business). If you hover over the label, you can see if it was applied manually (by the user) or automatically (by policy).
- Microsoft Search can index documents protected by sensitivity labels. This means that protected content can be found by content searches, available for eDiscovery, and accessible to any feature which depends on the content indexes, like DLP policies.

The benefit of traditional SharePoint “protection on download” is that encryption is applied automatically when files are downloaded from a library. Only people with access to the library can access the files. Users don't have to worry about applying a sensitivity label to protect confidential information.

The long-term strategy for any tenant should be to phase out the traditional SharePoint IRM-based protection and replace it with sensitivity labels as soon as business requirements and user training allow.

## Enabling Sensitivity Labels for SharePoint Online and OneDrive for Business

When users or policies assign sensitivity labels with encryption to Office files stored in SharePoint Online or OneDrive for Business, some functionality cannot process the encrypted content in the same way as it can handle unprotected content. Until you enable support for sensitivity labels in SharePoint Online as described below, several issues exist. For more information, see [this support article](#). If you've invested in applications that update SharePoint Online with elements like custom type schemas, it's important to understand the effect that some [current limitations](#) might have on your deployment.

Exchange transport rules do not have the same issue because Exchange uses super-user permission to examine protected email (and attachments) as messages pass through the transport pipeline.

**Searching for Labeled Documents:** The SharePoint Online search schema includes a managed property called *InformationProtectionLabelId*, which holds the GUID (identifier) for the sensitivity label assigned to a document. You can use this property to search for documents with a specific sensitivity label in SharePoint search or content searches by using the form *InformationProtectionLabelId:GUID*. For example, *InformationProtectionLabelId:2fe7f66d-096a-469e-835f-595532b63560*. The search results are trimmed and only display documents whoever performs the search can access. In the future, you might be able to search using the display name of a label instead of its GUID.

### Supporting Sensitivity Labels in SharePoint Online and OneDrive for Business

To enable support for sensitivity labels in SharePoint Online and OneDrive for Business, follow the directions on [this page](#). You can then enable the opt-in from the Microsoft Purview Compliance portal or by running the following PowerShell cmdlet from the SharePoint Online module:

```
[PS] C:\> Set-SPOTenant -EnableAIPIntegration $True
```

To revert, run the command again and set the switch to *\$False*. This action does not remove any sensitivity labels assigned to documents. If your tenant uses multi-geo capabilities, you must run the command in each data center region (geo-location) used by the tenant.

The following capabilities are enabled by support for sensitivity labels in SharePoint Online:

- Only users (including guest accounts) with access granted by the assigned sensitivity label can open a protected document (the View content right is the minimum needed).
- Apps can use the Microsoft Information Protection SDK to incorporate the necessary functionality to apply and respect sensitivity labels. Office Online and Microsoft 365 apps for enterprise (Windows and Mac) support co-authoring for Excel, Word, and PowerPoint files if everyone involved has the appropriate access to protected files (the feature is in preview for Office mobile). If using the desktop apps, people must [use a version that supports co-authoring of protected files](#). Guest users can't apply sensitivity labels to documents.
- Users can upload protected documents if they have at least viewer permission for the documents. If the uploader doesn't possess sufficient rights, SharePoint will upload the file, but it cannot process the contents.
- SharePoint Online ensures downloaded files retain sensitivity labels. However, if users export protected Office files to a non-Office format (such as PDF), encryption does not protect the exported files.



- A sensitivity column is available in SharePoint Online document views to display labels assigned to documents. In addition, if you hover over a label in the sensitivity label column, SharePoint Online tells you if a user (manual) or policy (automatic) applied the label.
- SharePoint Online and OneDrive for Business decrypt files protected with sensitivity labels to store and index their content. Decryption happens when users upload protected documents to a site for the first time or after editing. Being able to index protected content means that content searches and DLP policies can check both the metadata and content of protected files (SharePoint does not encrypt metadata such as document title and name for protected documents). In addition, Microsoft Search “trims” searches for protected content to ensure that only users who have access to that content see it in the results.

Some limitations currently exist. SharePoint Online cannot process protected documents with labels using:

- **User-defined permissions:** The permissions for these labels are set when a user assigns the label to a Word, Excel, or PowerPoint file using [a supported client](#).
- **Expiring access:** The label settings contain an expiration period after which access to the content is unavailable.
- **Double-key encryption or HYOK:** SharePoint Online doesn't have access to the key managed by the tenant.

The lack of support means that although you can store files protected with these labels in SharePoint Online or OneDrive for Business, SharePoint can't deal with the encryption and doesn't index the content (which affects eDiscovery). Features like preview and auto-save don't work, and labels with these attributes don't show up in Office Online apps. Put another way, the only supported labels are those which apply permissions set by the administrator, do not expire access, and use a single cloud-based encryption key. Some additional limitations exist which Microsoft is working through to resolve. Details of [currently-known limitations of sensitivity label support in SharePoint Online](#) are updated regularly.

When documents assigned supported sensitivity labels with encryption are uploaded into SharePoint Online or OneDrive for Windows, SharePoint Online decrypts their content before storing the document in Azure SQL. This step allows features like indexing and eDiscovery to work. The [metadata describing the sensitivity label](#) remains in place in the file and the data is protected in the store through the encryption at rest applied by Azure SQL. Any time afterward the file is requested, SharePoint checks if it needs to apply encryption and uses the metadata to ensure that the file is correctly protected. For example, encryption is reapplied when someone downloads a document with a label requiring encryption.

## Default Sensitivity Label for Document Libraries

Site owners can define a default sensitivity label for a document library by selecting an appropriate label in the library settings. Any new Office document added to the library will receive the assigned sensitivity label unless it already has a sensitivity label applied by a user or by an auto-label policy. In the case of policy assignment, the default label will replace the assigned label if the assigned label has a lower priority (determined by the priority number for labels set in the Microsoft Purview Compliance portal). Stamping occurs using an asynchronous thread, so the sensitivity label does not appear (or a replacement occur) on the document for one or two minutes after its upload. Existing unlabeled documents in the library remain unlabeled until the next time someone edits the file; at which time the document receives the default label.

The default application of sensitivity labels to new documents is a premium feature that requires Office 365 E5 or another suitable license for all members of the site hosting the document library.

## Sensitivity Label Mismatches

A label mismatch occurs when someone creates or uploads a document to a site that is assigned a higher-level sensitivity label than the site as dictated by the order of sensitivity labels defined in the organization. The order in which labels appear in the list is intended to reflect their relative sensitivity or importance with the

first label (order number 0) being the least sensitive and the last being the most sensitive. As an example of a mismatch, let's assume that a site is assigned the *Confidential* label, at position 4 in the set, and a user uploads a document assigned the *Super Confidential* label, at position 5 in the set. The mismatch occurs because the document label has a higher priority than the site label. Note that the comparison is performed based on the position of parent labels (not sub-labels) within the set of labels in the organization.

When it detects a label mismatch, SharePoint Online sends an *Incompatible sensitivity label detected* email notification to the user who uploaded the document (Figure 20-14) to inform them that the label assigned to the document might be inappropriate for the site. No automatic action happens to rectify the issue and it's left to the user who uploaded the document (who should be in the best position to understand its true sensitivity) to decide what to do to resolve the mismatch.

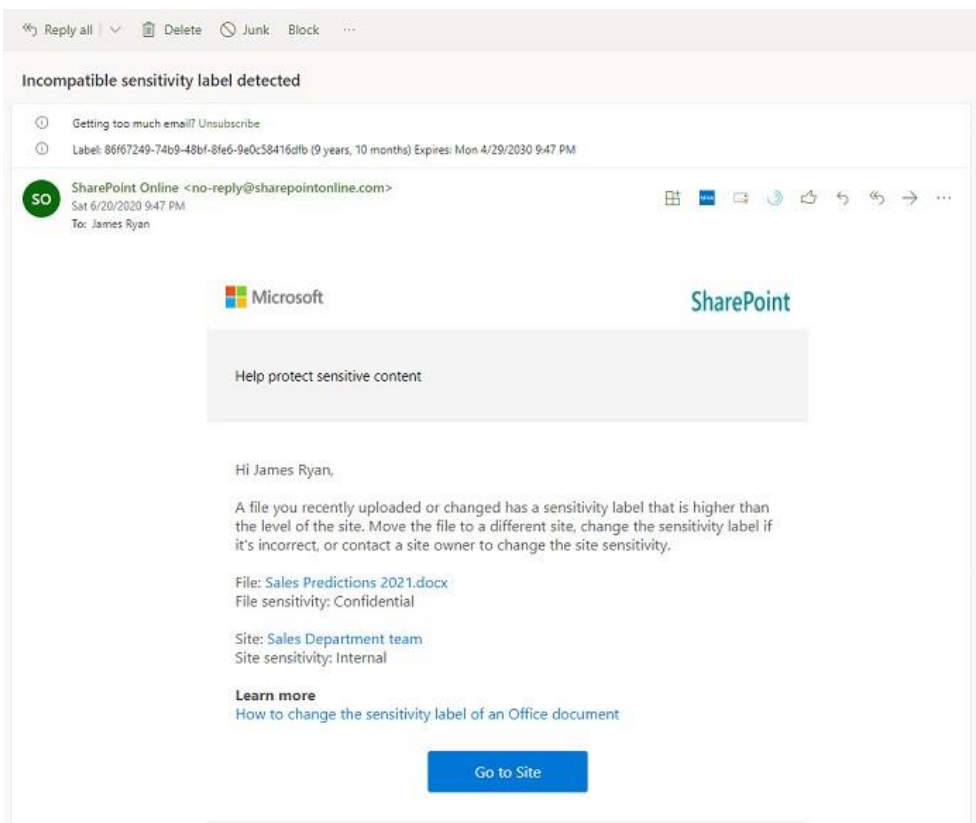


Figure 20-14: Email warning about a sensitivity mismatch for an uploaded document

The message helps to educate users about the different levels of sensitivity labels, but it doesn't explain why this is important. In the example shown, a document labeled as *Confidential* is uploaded to a site labeled as *Internal*. Anyone who is a member of the site can see the document metadata even if they can't open the document. People can learn a lot about a document from its metadata and it's possible that some confidential information might leak because someone can see a document's title. For example, external people might be guest members of the site. The rights assigned in a highly sensitive label might not give some or all guests the ability to open the document to view its content, but the title might tell them what the document is about.

Another reason why mismatches matter is that a site marked with low sensitivity might allow access to people using unmanaged devices. When documents of a higher sensitivity are stored by the site, the risk is created that those documents might be accessed on unmanaged devices, which is not what you might want to happen.

### Label Order Matters for Mismatch Detection

As we know, each sensitivity label has a priority order in the set of labels and the priority order is used to detect mismatches between site labels and document labels. If you separate the labels used for information

protection from those used for container management, you might give the container management labels a higher priority than those used for information protection. In this scenario, you'll never see a mismatch condition occur because the label placed on the container will always have a higher priority than the labels assigned to documents. Conversely, if you place the container management labels at the less sensitive end of the label set, every document labeled will generate a mismatch condition.

If you want label mismatch checking to operate as planned, make sure that each container management label comes **after** all the information protection labels used for documents stored by the site. For instance, a container management label called *Confidential Access* is used to label sensitive sites. An information protection label called *Confidential* is assigned to documents stored in sites labeled with *Confidential Access*. Another information protection label called *Super Confidential* exists which users should not apply to documents stored in sites assigned the *Confidential Access* label. Everything works properly if the priority order used for the labels is as shown below. In this scenario, a mismatch is detected if a document labeled *Super Confidential* is placed in a site labeled *Confidential Access*:

1. Information protection label: *Confidential*.
2. Container management label: *Confidential Access*.
3. Information protection label: *Super Confidential*.

SharePoint won't flag mismatches if the *Confidential Access* label has a higher priority than the *Confidential* and *Super Confidential* labels.

## Searching for SharePoint Files with Sensitivity Labels

Office documents, spreadsheets, and presentations store label information in their file attributes. When you store a labeled file in a SharePoint or OneDrive for Business document library, Microsoft Search captures the label data in crawled properties (something extracted from a file when the SharePoint crawler processes a file). The properties include *Sensitivity*, which stores the local language version of the label applied to the stored document (for example, "Confidential" or "Public"). You can remap the *Sensitivity* crawled property to one of the *RedefinableString* preconfigured managed properties to make the label searchable. The SharePoint schema includes a set of 200 *RedefinableString* managed properties (from *RedefinableString00* to *RedefinableString199*) to allow tenants to customize search behavior.

To make label data searchable, use the Search section of the SharePoint Admin Center to [remap the Sensitivity crawled property](#) to one of the *RedefinableString* properties (for example, *RedefinableString01*). It is also a good idea to give the property an alias (for instance, *SensitivityLabel*) to make it clear what the data is. Because the crawler must process a file before the label data is searchable, you should reindex the sites that hold classified information to force the crawler to find and index the label data. After the label data is added to the index, you can search for it using terms such as "*SensitivityLabel:Confidential*."

## OneDrive for Business Permissions

OneDrive for Business sites are personal and therefore need some individual attention before a site's owner can protect documents. You can read how to apply protection using steps for an individual user or how to use scripts to enable protection for multiple sites [at this link](#). These instructions apply to the older IRM approach to protection. You don't need to do anything special to allow users to apply sensitivity labels to items stored in their OneDrive for Business accounts.

## Protecting Individual Office Documents

Protection is applied to a Word, Excel, or PowerPoint file stored outside Office 365 by selecting a sensitivity label through the **Sensitivity** icon in the Office menu bar. If the sensitivity label includes encryption, the app

applies protection to the file the next time the user (or app) saves the file. Users must sign into a Microsoft 365 account to be able to access the set of sensitivity labels in a tenant.

If the unified labeling client is present on a PC, you can also protect files using the Classify and protect option in File Explorer. Finally, you can run the *Set-AIPFileLabel* cmdlet to protect individual files.

Applications know that files are protected because of metadata inserted in the files. Protection remains when labeled documents are uploaded to SharePoint Online or OneDrive for Business using the browser interface or if they are synchronized to a library with the OneDrive sync client. Protection for files downloaded in SharePoint Online or OneDrive for business also remains unaffected because downloading does not remove the sensitivity label.

## Recognizing Labels Cross-Tenant

Sensitivity labels are specific to and owned by the tenant where they are generated and unknown outside that tenant. Sometimes (for example, during a corporate merger) you might want a tenant to recognize labels owned by other tenants. The owner of the document or message will continue to control the rights assigned by the label, but you can [make some changes to the tenant configuration](#) to recognize and display another tenant's labels.

## Protecting Email

Several methods are available to protect email:

- Users can apply the two default OME options (*Do Not Forward* and *Encrypt Only*) using Outlook desktop (click to run), Outlook for Mac, or OWA. Outlook mobile does not support these options. *Encrypt Only* and *Do Not Forward* are available to tenants after configuring IRM. Users in tenants with Office 365 E3 and E5 can automatically read encrypted messages from other tenants; the IRM configuration must be in place to allow users to initiate new encrypted conversations.
- If configured for the tenant, users can apply sensitivity labels (with encryption) to protect messages. Recipients can only access the protected content if the rights assigned by the label allows them to interact with the content. Outlook Mobile supports sensitivity labels.
- Tenants can apply protection for messages sent to specific domains or sets of recipients by configuring transport rules.

Content searches can scan for and find information in protected messages. However, unless you enable support for sensitivity labels as described later, SharePoint Online and OneDrive for Business only index the metadata of protected documents, which means that searches cannot check attachments for protected messages. To allow transport rules to process protected messages, Exchange Online invokes super-user capability to decrypt. The same occurs to index email content. In addition, Exchange Online Protection can scan the content of encrypted messages to detect malware.

All successful implementations of protection for common Office content such as email and documents involve communication with the user community to explain the need for the technology, how to protect information (including some examples of best practices based on common scenarios), and the consequences (as laid down in HR policies) that might ensue if someone ignores the protection applied to data. No need exists to hit users across the head with a 2-by-4 to enforce policies; instead, all that anyone needs is a common-sense approach from all concerned.

## Enlightened and Unenlightened Email Clients

The Office 365 email clients support inline viewing of protected content. This means that clients decrypt and display encrypted information in the same way as they do for unencrypted messages. At the top of an

encrypted message, clients display an information banner to tell the reader that the message is protected. These applications include the necessary software to use the rights management APIs to retrieve and consume rights policies and licenses. To enable offline access to protected messages, Exchange fetches pre-licenses for messages. Outlook for Windows and Mac (version 16.23.19021400 or later) use these licenses to decrypt the messages without needing further authentication with the server. Browser-based applications, like OWA or OneDrive for Business, run online and connect to the rights management service to obtain use licenses to handle protection for documents.

Table 20-4 lists enlightened email clients on different platforms. Unless they incorporate code from the Microsoft Information Protection SDK (the Samsung Email app for Android is an example of a client which leverages the SDK), other clients know nothing about rights management or its implementation within Microsoft 365. These clients can't obtain the necessary licenses to unwrap the protection around protected messages to display the content. In Microsoft terminology, these clients are called "unenlightened." Users of unenlightened clients must go to the OME portal to open encrypted messages.

<b>Platform</b>	<b>Supported email clients</b>
Windows	Microsoft 365 apps for enterprise (Outlook click to run) OWA Windows 10 Mail app Outlook 2013 SP1 or later
macOS	Microsoft 365 apps for enterprise (Outlook for Mac) OWA
iOS	Outlook for iOS OWA
Android	Outlook for Android Samsung Email app OWA

Table 20-4: Enlightened email clients

Microsoft consumer accounts can use Microsoft 365 apps for enterprise, Outlook mobile, or the OWA browser interface (outlook.com) to access protected email. Accounts belonging to other services can access protected email through the OME portal. See here for more information about [applications that support rights management](#).

## Handling Unenlightened Clients

The Apple mail app for iOS is a good example of an unenlightened application that is popular with users. If an organization makes heavy use of rights management, forcing iOS users to go to the portal to read messages can be a sub-optimal experience. If you want these users to be able to read protected email on their devices, you can configure server-side decryption. The downside of this approach is that decrypted copies of the message then exist on user devices. To enable server-side decryption, run this command:

```
[PS] C:\> Set-ActiveSyncOrganizationSettings -AllowRMSSupportForUnenlightenedApps $True
```

When people use the mail app to view protected messages decrypted on the server, they see a header to tell them that the sender applied protection to the message. Unenlightened apps do not understand or apply the rights assigned to recipients, so a user of the iOS mail app can copy or print the message. However, Exchange Online knows that the original message is protected and, if the user attempts to do something without permission that involves the server, like forwarding the message, the action is blocked when the server processes that message.

In most cases, if you want to protect email within a tenant, it is a better idea to encourage people to use Outlook clients rather than disabling encryption.

## How Rights Management Protects Email

Users can protect email by assigning a sensitivity label (with encryption) or by using the default OME *Do Not Forward* and *Encrypt Only* features. Alternatively, if an organization has invested in S/MIME, it can use this method of protection.

### Protection for Non-User Recipients

The recipient list for a protected message can include a mixture of internal and external recipients, including those who do not use Exchange Online. Subject to the scoping defined for the sensitivity label used to protect a message, recipients will be able to open and access the content. Some restrictions exist when sending protected messages to recipients other than user mailboxes. For example:

- **Protected messages (and attachments) sent to a Microsoft 365 group** can be read by any member of the group, including guest accounts, because they authenticate their access through membership of the group. However, if a group member does not come within the scope of the label used to protect a message sent to the group, they can see that the conversation exists, who contributed to the conversation, and the title of the conversation, but they cannot see the content of the messages that make up the conversation because their credentials do not match the permissions assigned in the template. A banner informs the user that messages can't be displayed. If they click the banner, they see a link to the OME portal. This doesn't help either because the portal won't open protected messages when someone doesn't have the right permissions.
- **Users with full access to a shared mailbox (delegates) can read protected messages delivered to the mailbox** if the client supports this access. This includes access to protected Office attachments. See the section about Outlook access to protected content later.
- **Messages protected with Office 365 Message Encryption** can be sent to a dynamic distribution list. However, the members of the list cannot read the protected content because their email addresses are not in the recipient list.
- **Protected messages sent to the email address of a Teams channel** are rejected by Exchange Online because the transport service cannot re-encrypt the message for delivery to the phantom mailbox used to route messages to Teams. The sender receives a 5.7.1. Delivery Service Notification (DSN). Exchange Online decrypts protected messages to allow transport rules to process their content as the messages pass through the transport pipeline.
- Much the same happens if you try to send **protected messages to a Yammer group**. Again, the transport service cannot re-encrypt the message for onward delivery to Yammer, so it issues a 5.7.1. DSN.

Protected messages keep their status for their entire lifetime. Any replies to messages inherit the same protection to ensure that only the intended recipients can access the entire conversation.

## Office 365 Message Encryption

Office 365 Message Encryption (OME) is a capability available in Office 365 E3 and E5. OME makes two special rights management templates called *Encrypt Only* and *Do Not Forward* available to protect messages. *Encrypt Only* protects messages while allowing the recipient full control over the content. *Do Not Forward* blocks recipients from being able to forward a protected message. Clients encrypt messages protected by these options before they submit the messages to Exchange for onward processing. The only time Exchange applies encryption to email in transit is when a transport rule includes the application of a sensitivity label as an action.

Users can apply *Do Not Forward* or *Encrypt Only* to protect messages addressed to any user of any email system. Outlook and OWA clients connected to Exchange Online and Outlook.com accounts know how to

process OME-protected messages, meaning that the clients can display message contents inline as normal. Recipients using other mail systems must go to the OME portal to read the messages.

*Do Not Forward* and *Encrypt Only* grant viewer permission to all recipients. By contrast, if you apply a sensitivity label to protect messages sent to other domains, the recipients cannot read the content unless the label settings include their email address or their domain in the list of users and groups allowed to access the content. Organizations cannot change these templates to change the permissions assigned to recipients.

Unlike a sensitivity label, which can restrict what a recipient can do after receiving a message, when you use *Encrypt Only* or *Do Not Forward* to protect a message, recipients have full rights over a message. Protection through encryption is like the set of rights defined in sensitivity labels. The sender grants the recipient the right to decrypt and view the content. Put another way, the sender implicitly trusts the recipient to do the right thing with the content when they receive it. This isn't the case with other sensitivity labels, where the assumption is that rights need to be removed from some recipients so that they don't do the wrong thing. The important thing is that both approaches assure the sender that only authorized recipients can access a message and its attachments.

The idea behind *Encrypt Only* is to encourage users to consider the protection of confidential messages as a normal thing to do. Unlike third-party solutions such as S/MIME or PGP, users do not have to configure and manage certificates or install plug-ins. To make the idea even more appealing, encryption works for messages sent to any email address. If they use an enlightened client, a recipient can read encrypted content inline, while users of other email systems can read encrypted messages through the OME portal. These users receive a message with a link to allow them to log into the portal and read the message content. The link lasts for sixty days.

People who receive encrypted messages forwarded to them (by another person, rules, or when a forwarding address is set for a mailbox) cannot read the content because their addresses are not in the original recipient list. However, if a user forwards an encrypted message with OWA or Outlook, the recipient's address is added to the message header, and they can read the encrypted content.

## S/MIME and Microsoft Information Protection

[S/MIME \(Secure/Multipurpose Internet Mail Extensions\)](#) is a protocol that uses X.509 digital certificates to digitally sign, encrypt, or sign and encrypt messages. Like any protection scheme, S/MIME helps people who receive emails to be sure that messages are not tampered with in transit. S/MIME also helps to authenticate the sender; it proves the message is not spoofed. Before sending messages, the sender chooses to sign and/or encrypt the message, which is then processed like any other outbound message. The recipient uses the PKI infrastructure to verify the certificate and confirm the signature of the message. When a message is encrypted by S/MIME, the certificate of the recipient is used. For the sender to encrypt a message, he must have access to the public key of the recipient's S/MIME certificate. This requires the sender to import the recipient's certificate information before encrypting the email. For messages sent within the organization, an administrator can automatically publish that information to the Global Address List.

Outlook desktop, Outlook Mobile, and OWA support S/MIME for signing and encrypting messages. You can use S/MIME signing for both internal and external communications. When the S/MIME setting is enabled, Outlook for iOS and Android automatically disable the **Organize by Thread** setting. This is because S/MIME encryption becomes more complex as a conversation thread grows. By removing the threaded conversation view, Outlook for iOS and Android reduces the opportunity for issues with certificates across recipients during signing and encryption. As this is an app-level setting, this change affects all accounts added to the app.

It is technically possible to combine encryption applied by S/MIME with Microsoft Information Protection (and you can configure a sensitivity label to use S/MIME). However, applying multiple levels of protection only works if you use S/MIME to encrypt a message and then protect it afterward. This can only be done using Outlook desktop (for Windows) because this client is intelligent enough to resolve the inherent conflict

between the two encryption schemes. Of course, just because something is technically possible does not mean that it is feasible in practice. Combining S/MIME and Microsoft protection might result in a message that the recipient cannot process. A more practical approach is to select one scheme and use it everywhere. An organization should only consider using S/MIME if it already has a heavy investment in the technology needed for key management and clients deployed which can consume S/MIME-encrypted messages or if the organization decides to use third-party technology for protection. Outside these circumstances, we recommend that tenants avoid the complications (and extra costs) to deploy, manage, and maintain third-party encryption mechanisms and use the standard features built into Office 365 to protect email and documents.

**Real-world:** When configuring S/MIME, you can use an internal or external (third-party) certificate authority. In both scenarios, the entire certificate chain must be uploaded to Microsoft 365. To do this, you can export the chain from a local machine into an SST file and upload it to Microsoft 365, as described in [this procedure](#). After uploading the certificate, allow at least 30 minutes for the changes to replicate across Exchange Online. If you do not upload the certificate chain, you will see a message that says “An error occurred while sending this S/MIME message. The certificate used to sign this message isn’t trusted by your organization.”

## Encrypt or Protect

Given the choice to encrypt or protect messages, what should you do? Here’s a simple rule of thumb.

- **Encrypt** messages to protect confidential or sensitive data sent to recipients outside your organization. This includes messages sent to Microsoft 365 Groups (which can decrypt and show encrypted content in conversations to group members). The situation for Teams is different because Teams uses “special” email addresses for channels that don’t belong to a regular tenant. When you send an encrypted message to a channel, it is rejected because the Exchange transport service cannot decrypt and re-encrypt the message for the channel email address.
- **Protect** messages with confidential or sensitive data sent to internal recipients. If you define rights in labels for other domains, you can send protected messages to recipients in those domains. If someone receives a message protected by a label that doesn’t assign them the rights to access the content, they won’t be able to open the message.

This rule of thumb is based on the simple fact that OME works for messages sent to any email address, so it is the catch-all solution when a need exists to protect content sent outside the company. Not every destination might be able to understand the limitations imposed by sensitivity labels, but if a label is configured to support recipients in an external domain, it is an excellent way to protect information for the lifetime of the content.

## Protecting Email with Outlook

Outlook desktop clients automatically download the set of sensitivity labels available to users and refresh this information by connecting to the Information Protection service every four hours. Once the client downloads label information, Outlook has two options to protect email:

- The **Sensitivity** button in the menu bar reveals the set of sensitivity labels defined in the tenant and published to your account. Figure 20-15 shows how Outlook desktop (Windows) lists the available sensitivity labels. An Outlook profile can configure accounts from multiple domains, but sensitivity labels are specific to a tenant, so you’ll only see the labels belonging to the tenant configured in the default profile.
- The **Encrypt** button in the **Options** tab applies the OME *Encrypt Only* or *Do Not Forward* templates. The default option is to apply the *Encrypt Only* template. Click the arrow under the button to apply *Do Not Forward* protection to the message.



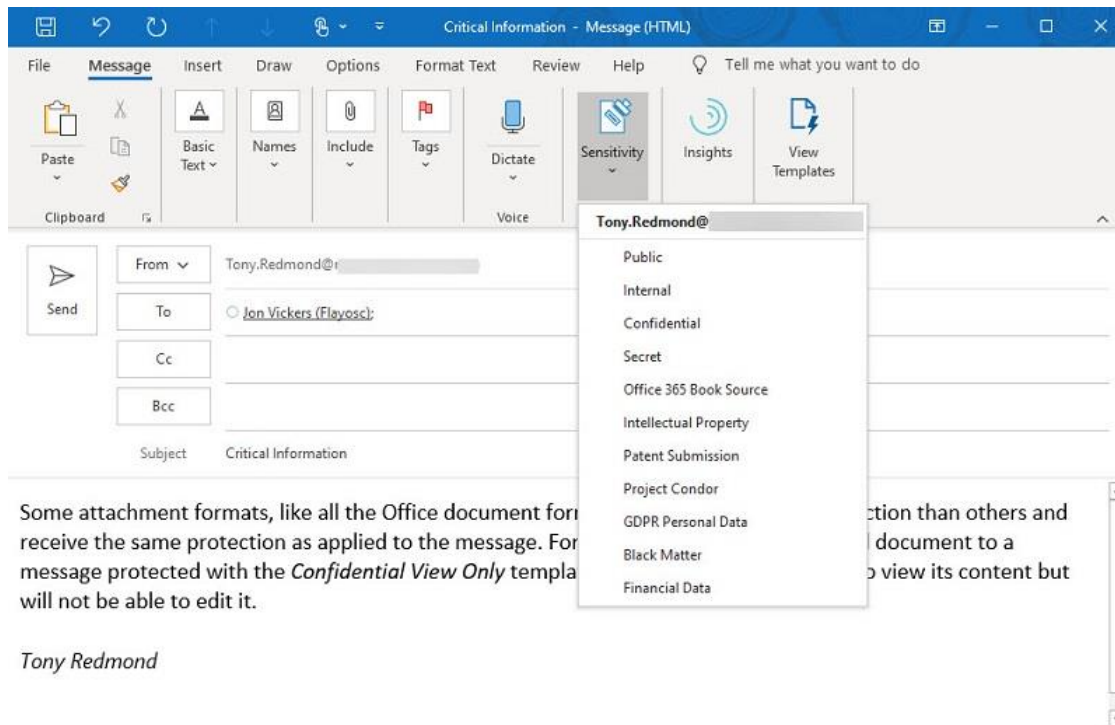


Figure 20-15: Selecting protection for a message with Outlook for Windows

Any replies and forwards (if allowed) created in response to a protected message inherit the same level of protection. Because the author always has full control over the content, they can decide to forward the message to another person, reply to the original message, or apply a different label. Outlook does not support the recall of a protected message. This is not as big a problem as it might seem because the speed of email servers today means that it is almost impossible for a message recall to work, especially when a message travels to another domain.

## Outlook Delegate Access

Delegates are users granted access rights to another user's mailbox or to a shared mailbox. The ability of a delegate with full access to a mailbox to read encrypted messages depends on the client used and the type of mailbox. Only Outlook clients support the ability of delegates to read encrypted messages (and attachments) subject to the following:

- Outlook for Windows clients do not support delegate access to encrypted messages sent to user mailboxes. Delegates can only read encrypted messages if the sender includes the delegate as a TO or CC recipient. In this scenario, the delegate's ability to read the message depends on the rights granted to them as a recipient.
- Outlook for Windows clients support delegate access to encrypted messages sent to shared mailboxes if the delegate has full access and auto-mapping is specified when the delegate receives permission to the mailbox. Auto-mapping is the default used by Exchange Online when delegates receive full access permission through the Microsoft 365 admin center or Exchange admin center. It forces Outlook for Windows to open the shared mailbox as part of the resources available to the delegate.
- The other Outlook clients (OWA, Outlook for Mac, Outlook Mobile, and the Windows Mail app) support delegated access to encrypted messages in both user and shared mailboxes if the delegate has full access to the mailbox.

Microsoft documents [some restrictions that apply to delegate access for encrypted messages](#).

If you wish to prevent users with full access to a user or shared mailbox from being able to view encrypted messages using clients other than Outlook for Windows, you can block their access by running the *Set-MailboxIRMAccess* cmdlet. For example, this command blocks the ability of Kim Akers to read any encrypted messages delivered to the Customer Services mailbox:

```
[PS] C:\> Set-MailboxIRMAccess -Identity Customer.Services@Office365itpros.com -User Kim.Akers@Office365itpros.com -AccessLevel Block
```

To make sure that a block is in place, use the *Get-MailboxIRMAccess* cmdlet:

```
[PS] C:\> Get-MailboxIRMAccess -Identity Customer.Services@Office365itpros.com -User Kim.Akers@Office365itpros.com
```

Identity	User	AccessLevel
-----	----	-----
Customer Services	Kim.Akers@office365itpros.com	Block

A block on delegate access remains in place until an administrator removes it and only affects the ability of a delegate to read encrypted messages using clients that support the block. For instance, the block will stop a delegate reading encrypted messages in a shared mailbox using OWA or Outlook for iOS, but they can switch to Outlook for Windows to see the message content. In addition, blocking access does not hide message subjects, which can contain sensitive information, nor does it prevent a delegate from deleting or moving encrypted messages. The block exists for reading, and only works for clients that support the block.

To remove the block and restore the ability to read encrypted messages to a delegate, run the *Remove-MailboxIRMAccess* cmdlet:

```
[PS] C:\> Remove-MailboxIRMAccess -Identity Customer.Services@Office365itpros.com -User Kim.Akers@Office365itpros.com
```

**Remove Encrypt Only:** Some organizations consider that users should only use sensitivity labels to protect messages and want to remove the standard OME *Encrypt Only* template. The instructions to do this are [available online](#).

## Protecting Email Attachments

Unprotected Office documents (and PDF files if enabled: see the notes about the *EnablePdfEncryption* setting in the tenant IRM configuration earlier) attached to protected messages inherit the same protection as assigned to the message. For example, if you attach a Word document to a message protected with a sensitivity label that only allows edit access for the author, the recipients will be able to view its content but will not be able to edit it. If a higher level of protection is applied to a file before it is attached to a message, that protection is preserved for the attachment.

Figure 20-16 shows a Word document open on an iPhone. This is an attachment to a protected message opened by Outlook for iOS. The user can discover what rights they have over the content by clicking

### Permissions.

Recipients can open and edit attachments in file formats that do not support sensitivity labels, like text files or bitmap images, unless you protect the files with the unified labeling client (covered later) before attaching them to the message. If a user applies a sensitivity label without encryption to a document and attaches the file to a message, Outlook assigns the same label as applied to the message to the document to ensure consistency of protection.

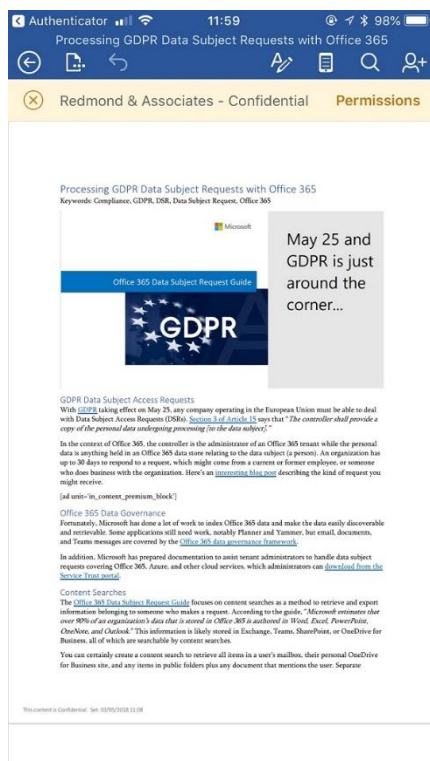


Figure 20-16: A protected Word attachment viewed with Word for iOS

## Automatic Decryption of Encrypt Only Attachments

Dealing with protected attachments is not an issue for Exchange Online clients because they can obtain the necessary use licenses to decrypt the attachments when opening messages. However, if a message using Encrypt Only goes to recipients of other email systems like Gmail, recipients can read the messages after authentication through the OME portal, but they cannot access the contents of a downloaded attachment because the downloaded copy of the attachment remains protected. If you want external recipients to be able to download decrypted copies of attachments, you must update the IRM configuration to instruct Exchange Online to decrypt attachments and remove protection whenever they are accessed or downloaded by an authenticated recipient. The effect is to give recipients full control over the downloaded files. To change the configuration, run this command to update the tenant IRM configuration:

```
[PS] C:\> Set-IRMConfiguration -DecryptAttachmentForEncryptOnly $True
```

This setting only applies to messages protected with the Encrypt Only feature.

## Viewing Rights

When you open a protected item, you see the name of the template and some information about its intended use. Outlook users can click a message header to gain more insight into what they can and cannot do with a protected item (Figure 20-17). It is obvious from the list of rights supported for email that you cannot block every conceivable action that a recipient can take with a message. For instance, although you can remove the Copy right to stop someone from copying text from a message or attachment, including blocking screen captures on Windows devices, you cannot stop someone from taking a screenshot with a smartphone and circulating that image to others. The point here is that protection cannot prevent every kind of unacceptable user interaction with content. Users must accept some responsibility for their actions. Applying a sensitivity label to an item makes users aware that the author protected the information for a reason and that they should therefore deal with that information appropriately. What they do is entirely up to the user.



Figure 20-17: Outlook reveals the rights available for a message

## Protecting Email with OWA

Because OWA runs in online mode, it always uses the current set of sensitivity labels published for a user. In saying this, we should realize that some client-side caching occurs for performance and that a small delay is likely before a new label or a change to an existing label published by the Microsoft Purview Compliance portal becomes available in OWA.

The Sensitivity button is available as an option in the OWA new message window. After a label is set on a message, the label name appears in the banner above the message recipients. In Figure 20-18, we can see that the sensitivity label selected for the message invokes encryption because of the padlock icon beside the label name. A label that applies some markings to items without encryption or simply acts as a visual indicator uses a plain label icon. OWA also displays these icons for labeled items in the read message window. Like Outlook, the protection applied to a message also applies to any of its attachments.

Sensitivity labels can also be applied to replies to messages that aren't previously labeled. In this case, the **Sensitivity** option to apply a label is in the [...] menu of the reply message window. When you assign a sensitivity label to a reply, it does not apply to the previous messages in the thread. However, Exchange automatically assigns the same label to future messages in the thread.

OWA can also use the OME *Encrypt Only* and *Do Not Forward* features to protect messages. Click the [...] menu and you'll find **Encrypt** in the list of menu choices. Using these templates for protection does not assign a sensitivity label to the protected messages.

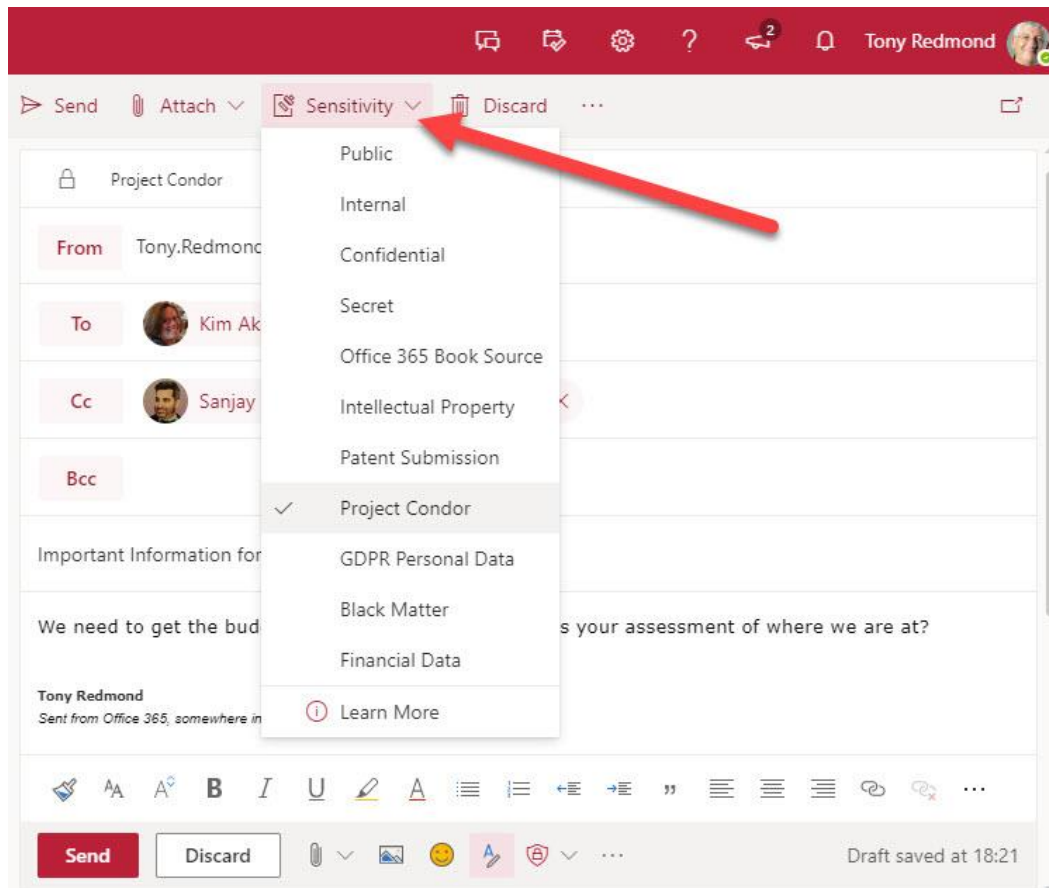


Figure 20-18: Applying protection with OWA

## Microsoft Rights Management Connector

The Microsoft Rights Management (RMS) Connector lets Exchange and SharePoint on-premises servers deployed in a hybrid configuration access the capabilities of cloud-based Information Protection. The connector runs on an on-premises Windows server or a virtual machine. Its role is to act as a communications relay between on-premises servers and the cloud rights management service. The on-premises servers are then able to use the cloud service to protect messages and documents and process protected email in the transport pipeline. In addition, Outlook clients connected to on-premises mailboxes can use the cloud service to access protected email.

For more information about how to install and configure the Rights Management connector, see [this page](#).

## Managing Office 365 Message Encryption

Office 365 Message Encryption (OME) is an online service built on top of rights management to allow users to protect emails sent to recipients in any email system. OME includes:

- The default Encrypt Only and Do Not Forward templates.
- The ability to encrypt messages sent to specific destinations through transport rules. For example, you might decide to encrypt all messages sent to a partner domain. Any sensitivity label with encryption can protect content using a transport rule.
- The message encryption report (available to Office 365 E3 and E5 tenants), is available in the [Reports section of the Microsoft Purview Compliance portal](#). This report details how messages receive protection (by users or by policy) and whether encryption occurs through a default template or sensitivity label. The data is useful in terms of understanding how many people protect email the

organization. Because Exchange Online generates the report data daily, the information is not real-time and can be up to a day behind.

OME-protected messages do not support protecting messages sent to dynamic distribution lists. Exchange Online bifurcates messages to deliver copies to the members of the dynamic distribution list when they pass through the transport pipeline. However, the list members are not present as message recipients, so they never gain the right to open the message. A sensitivity label protects messages sent to dynamic distribution lists if the permissions assigned in the label grant access to everyone who receives a copy.

It's worth emphasizing that OME is all about protecting email without needing the recipient to install any special software. Sensitivity labels can also protect email with encryption, but the big difference is that sensitivity labels also define the actions a recipient can take after they open that content. Organizations can also include sensitivity labels as an action for transport rules to execute on outbound messages.

## Customizing the OME Configuration

When an external recipient (not belonging to another Office 365 tenant or Outlook.com) receives a protected message, they receive a notification message to tell them what they must do to access the content. The notification directs the recipient to sign into the OME portal using an account from one of the federated identity providers. If they cannot sign in or prefer not to, the recipient can get a one-time code to read the content. After signing in, the user sees a modified version of the OWA read message window, which enforces the permissions they have over the content.

You can customize the OME configuration (otherwise known as a branding template) to add custom text used in the notifications and the OME portal. To start, use the *Get-OMEConfiguration* cmdlet from the Exchange Online module to view the details of the standard OME configuration. In this case, the configuration has several customized settings, which we can update further with the *Set-OMEConfiguration* cmdlet.

```
[PS] C:\> Get-OMEConfiguration -Identity "OME Configuration"
TemplateName           : OME Configuration
Image                  : {137, 80, 78, 71...}
ImageUrl              :
EmailText              : Thank you for communicating with our company. To protect our
confidentiality, this message has been encrypted.
PortalText             : Our Great Encrypted Email Portal
IntroductionText      :
DisclaimerText        : Office 365 for ITPros takes no responsibility for the operation of this
portal
ReadButtonText        : Protected Email
OTPEnabled            : True
SocialIdSignIn        : True
ExternalMailExpiryInterval : 00:00:00
PrivacyStatementUrl   :
Identity              : OME Configuration
```

Settings of note are:

- **BackgroundColor:** You can set the background color for OME messages and the portal by passing the hex code color. Usually, people who know about corporate branding will know [the right code](#) to use.
- **DisclaimerText:** A string of up to 1024 characters placed at the bottom of the encrypted message intended as a disclaimer but can convey a different message. If you don't provide a value, OME uses: *"This email message and its attachments are for the sole use of the intended recipient or recipients and may contain confidential information. If you have received this email in error, please notify the sender and delete this message."*
- **EmailText:** A string of up to 1024 characters to tell the recipient that they have received an encrypted message. If you don't provide a value, the default OME text is *"You've received an encrypted message"*

from" plus the SMTP address of the sender. If you change the default message, OME does not include the sender's SMTP address (but it still appears in the message header).

- **Image:** An image file (.png, .jpg, .bmp, or .tiff) of less than 40 KB to show that messages originate from the sender's company. Ideally, you might use a 170x170 pixel version of the company logo. By default, OME does not display a logo.
- **OTPEnabled:** If true (the default), recipients can opt to get a one-time code to access the portal. Set this to False if you want to force users to authenticate using an account from one of the supported providers.
- **PortalText:** A string of up to 128 characters that the OME portal displays when users connect to access a decrypted message.
- **PrivacyStatementURL:** A link to a web page for the company's privacy policy. This link is invoked when a user clicks the Privacy Statement link in an OME notification.
- **ReadButtonText:** A string to display on the Read button in the notification. To make sure that it fits on the button, this text should be no more than 16 characters.
- **SocialIdSignIn:** Defines if a recipient can authenticate to read a message using an identifier from a recognized social network such as Google or Yahoo! If False, the recipient can authenticate using another method such as a one-time password.

Fwd: Looking after our customers during COVID-19

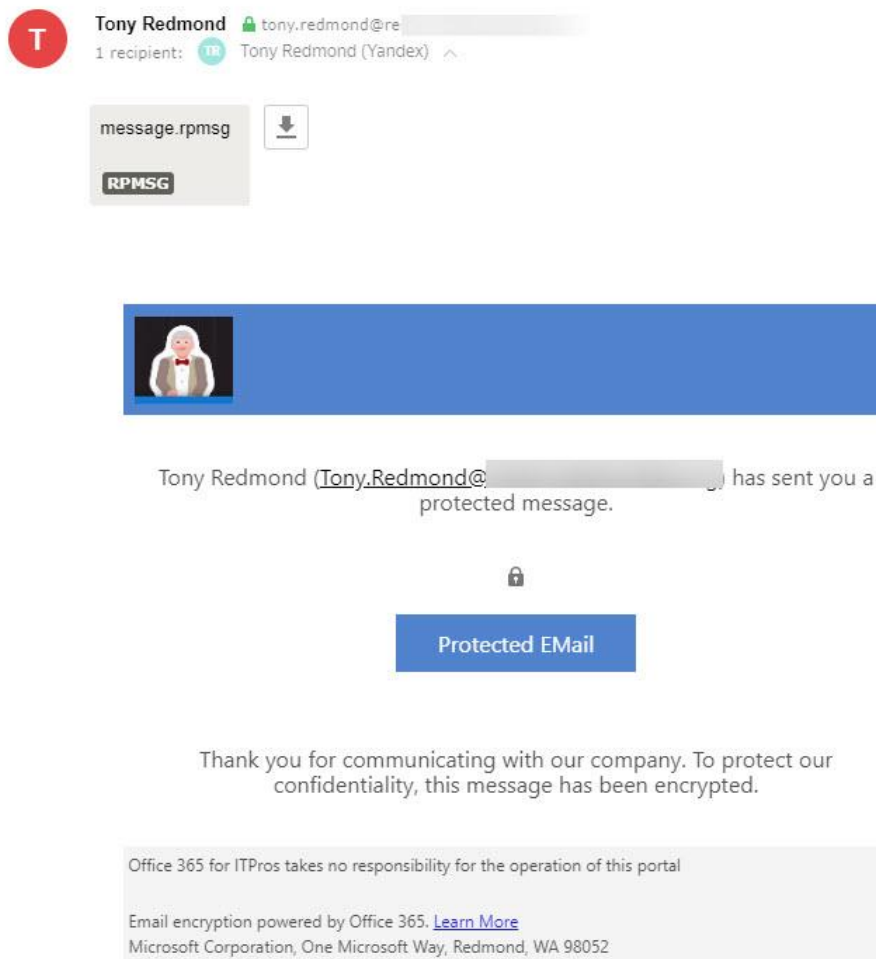


Figure 20-19: The OME notification received to let someone know that they should read some protected email  
For example, this PowerShell command updates several properties in the OME configuration:

```
[PS] C:\> Set-OMEConfiguration -Identity "OME Configuration" -BackgroundColor "#5183cd" -DisclaimerText "We take zero responsibility for anything that happens here." -EmailText "Yippe!"
```

```
You've got some encrypted super-secret email" -PrivacyStatementUrl  
"https://office365itpros.com/privacy.html"
```

Here's how to add a graphic file for the logo displayed in notification messages and the OME portal:

```
[PS] C:\> Set-OMEConfiguration -Identity "OME Configuration" -Image (Get-Content  
"C:\Temp\CompanyLogo.jpg" -Encoding Byte)
```

When you update the OME configuration, Microsoft warns that custom content used in configurations is the tenant's responsibility. In other words, make sure that any graphics and text used are not protected by trademarks or other restrictions.

Figure 20-19 shows the effect of the updated OME configuration. The logo appears at the top of the message, the email text appears under the link, and the customized disclaimer text is in place. The restricted permission attachment (rmsg) is the encrypted content shown when the user accesses the message through the OME portal. Internal recipients do not see notifications unless they use a non-Outlook client that can't obtain the necessary use licenses to decrypt and display the protected content automatically. For example, if you connect to your mailbox using an IMAP4 client, you'll see notifications for protected messages sent by other people in the tenant.

## Advanced Office 365 Message Encryption

The standard OME supports a single branding template, applied by Exchange Online to all protected messages delivered to recipients outside the tenant. Advanced OME adds support for:

- Multiple branding templates and the application of branding templates in transport rules. For instance, you could use a transport rule to protect all messages sent to a partner domain and customize the template that recipients in that domain see. To do this, select the **Apply Custom Branding to OME messages** action in a transport rule along with the custom template to use and the criteria to apply the branding (for example, specific domains). Custom branding templates are updated in the same manner as the general template used by the standard version of OME (see above).
- Expiry of protected email. The custom templates support the ability to expire messages sent to external recipients after a set period of anything from one to 730 days.
- [Revoke access for external recipients to protected email](#). This feature is available in OWA when senders have an Advanced OME license, and the recipient accesses the encrypted content through a link to the OME portal. Alternatively, an administrator can revoke access to a message on behalf of a user using PowerShell (see below).

Advanced OME is an Office 365 E5 feature also licensed using the Microsoft 365 E5 Compliance SKU.

Custom branding is an optional action for transport rules. If you apply custom branding in its own rule, make sure that the rule to apply branding runs (has a lower priority) after any rule which applies encryption. In addition, make sure that any rule which applies encryption or custom branding does not terminate rule processing and let other rules with lower priority run afterward.

## Creating a Custom Branding Template for OME

Custom branding templates are created with PowerShell. This example shows how to create a new branding template and then populate the properties of the new template.

```
[PS] C:\> New-OMEConfiguration -Identity "Office 365 IT Pros Branding"  
  
Set-OMEConfiguration -Identity "Office 365 IT Pros Branding" -DisclaimerText "Office 365 for IT Pros  
takes no responsibility for this portal." -PortalText "Office 365 for IT Pros Secure Messaging"  
-EmailText "Good things happen when you protect email" -ExternalMailExpiryInDays 10  
-IntroductionText "has sent you a secret message" -Image (Get-Content "C:\Temp\SmallBookCover.jpg"  
-Encoding byte)
```



See [this page](#) for more information about how to customize a branding template. Remember that it can take between 15 and 30 minutes for a change made to a transport rule to become effective, so factor this into the time needed to test how custom branding templates work.

To test that branding works, create a transport rule to apply a custom branding template to messages in a specific domain and send a protected message to an account in that domain. When it arrives in the recipient mailbox, check that the custom branding is visible when the recipient opens the message.

## Administrator Revocation of Encrypted Messages

Recipients of protected messages fall into two categories: internal and external recipients. Revocation, which means removing the right of a recipient to view encrypted content, is only possible for link-based messages delivered to individual external recipients because the OME portal controls access to these messages.

Messages to users in other Microsoft 365 tenants or Outlook.com are irrevocable because these users don't need to access the OME portal to read protected messages.

Administrators can revoke a message through PowerShell by finding the identifier of the message to revoke and then running the `Set-OMEMessageRevocation` cmdlet. The easiest method to find a message identifier is to use the [Message Trace feature in the Exchange admin center](#) (or by running a message trace with the `Get-MessageTrace` cmdlet).

Execute a search to find the message, select it in the set of results, and look at its properties. The message identifier is a long string ending in something like "prod.outlook.com." When you have the identifier, run the `Get-OMEMessageRevocation` cmdlet to check that the message is revocable. At this point, Exchange Online knows if the message recipient received a link to the OME portal or could read the message inline. If you see that the message is revocable, run the `Set-OMEMessageRevocation` cmdlet to revoke permission. For example:

```
[PS] C:\> $MessageId =
"AM6PR0402MB3462A8DF67AF7AF9C9F0B20A8BE50@AM6PR0402MB3462.eurprd04.prod.outlook.com"
```

```
Get-OMEMessageStatus -MessageId $MessageId | Select Subject, IsRevocable
```

Subject	IsRevocable
-----	-----
Important Information	True

```
Set-OMEMessageRevocation -Revoke $True -MessageId $MessageId
```

The encrypted email with subject "Plans for next year" and Message ID "AM6PR0402MB3462A8DF67AF7AF9C9F0B20A8BE50@AM6PR0402MB3462.eurprd04.prod.outlook.com" was successfully revoked.

To check the status of the message, you can run the `Get-OMEMessageStatus` cmdlet. The results returned by the cmdlet confirm the revocation of the message:

```
[PS] C:\> Get-OMEMessageStatus -MessageId $MessageId
```

```
RunspaceId      : 4039a474-e2ab-498f-a92c-460154537c8f
Identity        : 04e78939-85a9-4bf8-8d65-9f533648bb37
IsValid         : True
ObjectState     : New
Container       : SystemMailbox{D0E409A0-AF9B-4720-92FE-AAC869B0D201}@office365itpros.onmicrosoft.com
Subject         : Tony Encrypt
ReceivedTime    : 7 Oct 2021 14:03:55
Revoked         : True
```

When the recipient goes to the OME portal and tries to read the protected message, they see "the message has been revoked by the sender" and that they should contact the original sender if they want access to the content. If an administrator revokes a message, there's no way for the original sender to cancel the revocation.

If the administrator makes a mistake and revokes a message incorrectly, the original sender will have to resend the message.

## Using Office 365 Message Encryption with Transport Rules

The basic scenario is when a transport rule protects email sent to a specific destination, which could be a group of users or one or more domains. Messages sent by clients connected to Exchange Online mailboxes must pass through the transport pipeline. As messages flow through the transport pipeline, Exchange Online uses the criteria defined in transport rules to examine messages. If an email matches a rule, Exchange Online applies the action or actions defined in the rule. The rule set can span hundreds of rules to account for different circumstances, and Exchange Online examines each message against each rule until it reaches the end of the rule set, or the action set in a rule causes rule processing to halt.

Adding protection to messages by applying a template is one of the actions available for transport rules. This capability is often used to ensure that protection is applied to messages even if they originate from clients that do not support rights management, such as the default mail apps for iOS and Android. Using transport rules to apply encryption also avoids the need for users to select the right template to apply when they compose a message. If many different templates exist, users might not know what template they can use to protect messages sent to a certain domain or user. For this reason, organizations often use transport rules to apply protection to ensure that confidential information cannot pass in an unencrypted form to recipients outside the tenant. A transport rule can apply protection using criteria such as:

- Any message sent to specific domains.
- Any message sent to specific users.
- Any message that contains a specific phrase in the message text or an attachment.
- Any message that has a specific word in its subject.

When a transport rule protects outbound messages, the senders of those messages might not be aware that this processing occurs because the copy of the item in their Sent Items folder is unprotected. This is because transport rules apply sensitivity labels in the transport pipeline to protect the message copies delivered to recipients. The copy of the message in the sender's mailbox has not been through the transport pipeline and is therefore unprotected. If users want to see a protected copy, they must include their name as a recipient.

## Configuring a Transport Rule for Protection

Transport rules are configured through the **Mail flow** section of the EAC (for more information on this topic, see the discussion about transport rules in Chapter 7). In Figure 20-20, we see details of a transport rule to do the following:

- Monitor messages sent by members of the Executive Committee distribution list; and
- Apply the "Confidential" template to these messages. This is an action in the **modify message security** set.

Microsoft publishes details of a transport rule to encrypt outbound messages with certain sensitive information types (like credit card numbers). The prototype rule they suggest is a good start, but you should [carefully review and adjust the rule](#) to ensure that it meets the need of your tenant.

## Protect Executive Committee messages

Name: Protect Executive Committee messages

\*Apply this rule if...  
The sender is a member of... 'Executive Committee'  
add condition

\*Do the following...  
Apply Office 365 Message Encryption  
add action

Except if...  
x The subject includes...  
add exception

Properties of this rule:  
Priority: 8  
 Audit this rule with severity level:  
Not specified

select RMS template  
RMS template: Super Confidential  
OK Cancel

Save Cancel

Figure 20-20: Applying protection with a transport rule

## Exemptions

Often rules include an exception condition to allow users to override the rule in certain circumstances. For example, if the subject of the message has a specific text pattern, the rule will not apply the sensitivity label. When a rule allows a message to pass without intervention due to an override, because the transport captures the fact that it processed the rule in the message header, Exchange Online will not try to apply the rule again to replies and forwards that flow from the original message.

Another common exception is not to apply protection when the recipient is a member of a specific group. The logic here is that the members of the group might need to change or update the content circulated in emails.

Although distribution lists are an excellent way to specify the users that come under the scope of a transport rule, remember that Exchange Online caches group membership to avoid the performance penalty of going back to Azure AD to expand membership each time a group is in a message header. For this reason, do not expect a change made to a group used in a transport rule to be effective at once. It can take between 30 minutes and an hour before the transport service learns about the new group membership.

**Interlocking protection:** You can apply a mixture of DLP checking and sensitivity labels to make sure that sensitive data does not leave the organization unprotected. If the transport service detects sensitive data protected by a DLP policy in a message, a transport rule can apply a label. This level of interlocking protection helps organizations ensure that people do not misuse sensitive data.

## Handling Protection when Transport Rules Process Protected Messages

Users might protect messages using the default Encrypt Only or Do Not Forward templates that are subsequently processed by a transport rule which also attempts to protect the message. When this happens, the level of protection is adjusted if the protection applied by the rule is more restrictive (higher sensitivity) than chosen by the user, which might affect the rights assigned to recipients.

This is logical because the author has already explicitly applied protection to the message with what they believe is the right level of protection. Overriding protection selected by a user with a rule runs the risk of interference with the recipient's ability to process the message in the way intended by the author.

## Applying Protection in Transport Rules with PowerShell

Two methods are available to apply templates to outbound emails. The traditional approach that we've just explored is to use a transport rule; a newer approach is to specify protection as an action in a DLP policy that fires when the DLP service detects sensitive data in a message. You can find information about how to invoke protection in DLP policies in Chapter 19. Because of the complexities involved in many rules and policies, it is often easier to build them through the relevant GUI (EAC for transport rules, Microsoft Purview Compliance portal for DLP policies), but PowerShell can also be used for the task.

The *New-TransportRule* cmdlet creates a new transport rule. In this example, we create a rule that applies the "Sensitive Board Reports" label to messages sent from one distribution list to another distribution list.

```
[PS] C:\> New-TransportRule -Name "Protect Board Meeting Information" -FromMemberOf "Board Reports"
-SentToMemberOf "Board Members" -ApplyRightsProtectionTemplate "Sensitive Board Reports"
-ExceptIfSubjectContainsWords @("Override") -StopRuleProcessing:$False -Mode Enforce
-Comments "Protect Board Information when circulated" -RuleErrorAction Ignore
-SenderAddressLocation Header
```

Another example is where you want to protect messages that contain sensitive information types such as credit card numbers. A wide range of default sensitive information types are available for use in DLP policies and you can define custom sensitive information types if necessary. Any sensitive information type can be used in a transport rule to identify messages for protection by including it in the *MessageContainsDataClassifications* parameter. Here's a simple PowerShell example that looks for six different sensitive information types. If any are found in a message, Exchange Online applies the Encrypt template.

```
[PS] C:\> New-TransportRule -Name "Encrypt external email with PII content" -SentToScope
NotInOrganization -ApplyRightsProtectionTemplate "Encrypt" -MessageContainsDataClassifications
@(@{Name="ABA Routing Number"; minCount="1"},@{Name="Credit Card Number"; minCount="1"},@{Name="U.S.
/ U.K. Passport Number"; minCount="1"},@{Name="U.S. Bank Account Number"; minCount="1"},@{Name="U.S.
Individual Taxpayer Identification Number (ITIN)"; minCount="1"},@{Name="U.S. Social Security Number
(SSN)"; minCount="1"}) -Mode Enforce
```

Alternatively, you can create a DLP policy that applies a template when messages are shared outside the organization. Two steps are needed to do this with PowerShell. The first creates a DLP policy; the second creates the rule to encrypt email with the same set of sensitive information types specified for the transport rule and attaches the rule to the policy.

```
[PS] C:\> New-DlpCompliancePolicy -Name "Encrypt external sensitive mail" -ExchangeLocation "All"
New-DlpComplianceRule -Name "Encrypt external email with PII content" -Policy "Encrypt external
sensitive mail" -AccessScope NotInOrganization -EncryptRMSTemplate "Encrypt" -NotifyUser
>LastModifier" -NotifyPolicyTipCustomText "This email contains sensitive PII information and will be
encrypted when sent." -NotifyEmailCustomText "This email contains sensitive PII information and will
be encrypted when sent." -ContentContainsSensitiveInformation @(@{Name="ABA Routing Number";
minCount="1"},@{Name="Credit Card Number"; minCount="1"},@{Name="U.S. / U.K. Passport Number";
minCount="1"},@{Name="U.S. Bank Account Number"; minCount="1"},@{Name="U.S. Individual Taxpayer
Identification Number (ITIN)"; minCount="1"},@{Name="U.S. Social Security Number (SSN)";
minCount="1"})
```

Several methods are available to apply encryption via policy to outbound email. It's important to choose either transport rules or DLP policies to protect sensitive data as it is easy to confuse matters if protection is applied for the same content using multiple methods.

## Applying Sensitivity Labels with Transport Rules

If you apply a sensitivity label to a message using a transport rule, the message is protected if the label invokes encryption, but none of the visual markings are applied to the item. Exchange recognizes the presence of a sensitivity label for a message through the presence of an x-header called *msip\_labels*. The value of the header points to the GUID for the sensitivity label. GUIDs are used instead of text values to allow clients to show label names in local languages. If the *msip\_labels* x-header is not present in a message header, clients that understand labels cannot display a label. Therefore, if we want to apply a label to a message, we must do so by adding the label with a client or applying the label with a transport rule as the message passes through the transport pipeline. The technique works for outbound messages. It doesn't work for inbound messages.

Figure 20-21 is an example of applying a label in a transport rule. The rule is very simple and does the following:

- If the message subject or body contains the word *#Confidential*, execute the rule. This is an easy way of allowing people who use clients that don't support rights management (like older Mac devices or mobile devices that don't support an application capable of applying labels to items) to have Exchange protect messages. The rule applies to messages sent inside and outside the organization.
- Apply the Encrypt template to protect the message.
- Add *MSIP\_Label\_1b070e6f-4b3c-4534-95c4-08335a5ca610\_Enabled=True*; as the value of the *msip\_labels* x-header in the message. The closing semi-colon is important as it terminates the x-header.

### Apply Sensitivity Label to Email

The screenshot shows the configuration for a transport rule named "Apply Sensitivity Label to Email".

- Name:** Apply Sensitivity Label to Email
- \*Apply this rule if...:** The subject or body includes... *#Confidential'*
- \*Do the following...:**
  - Apply Office 365 Message Encryption and rights protection to the message with... *Encrypt*
  - and
  - Set the message header to this value... Set the message header '*msip\_labels*' to the value '*MSIP\_Label\_1b070e6f-4b3c-4534-95c4-08335a5ca610\_Enabled=True;*'

Buttons: Save, Cancel

Figure 20-21: Setting a transport rule to add the *msip\_labels* x-header to a message

A big question for anyone wanting to use a sensitivity label in a transport rule is how to find its GUID. The Microsoft Purview Compliance portal doesn't display this information for a sensitivity label, but you can retrieve label GUIDs by running the *Get-Label* cmdlet (see the section covering using PowerShell to manage sensitivity labels). You can combine labels with templates. In other words, even if a label includes protection, you can decide that a transport rule should apply an x-header for one label and protect the message with a

different template (that isn't associated with a label). If you want to use a sub-label, make sure you use its GUID and not that of the parent label.

Remember that the senders of messages processed by this rule will not see a label or protection on the copy of the message in their Sent Items folder because protection only applies to messages that go through the transport pipeline. If you want to check that the rule works as expected, include your address as a recipient and examine the copy of the message delivered in your Inbox with a client that supports sensitivity labels. You should find the *msip\_labels* x-header in the message headers and that the message is both labeled and encrypted. It's also true that the label information will mean nothing to a different email system (or different tenant) because those systems won't be able to translate the GUID into a label. However, the protection will still apply even if the label isn't respected, and that is probably the most important thing.

## Detecting and Blocking Protected Messages in Transport Rules

We know that protected messages have x-headers holding information about the label applied to protect the content. We can exploit this fact to create a transport rule to block messages leaving the organization if their headers include a certain label. In this case, the criteria are:

- Apply to outbound messages.
- Check the *msip\_labels* x-header and if the GUID for the label exists, block the message with the action "Reject the message with the explanation." The text for the explanation is up to you but might be something like "You can't send sensitive messages outside the organization."

For example, let's assume that you have a label with a GUID of *ed4411cc-bec4-444a-b279-c404aaad79d6*. The text that the transport rule should look for in the x-header is:

```
MSIP_Label_ed4411cc-bec4-444a-b279-c404aaad79d6_Enabled=true
```

If found, we know that this message (or one of its attachments) is protected with the label, so the rule can go ahead and block the message.

Exchange Online can read the label metadata for messages and Office attachments. However, it cannot read metadata for protected PDF attachments. Another complication is that OWA and Outlook mobile clients apply protection to messages using special hidden rules executed as items pass through the Exchange Online transport pipeline. It is more efficient to encrypt messages after the processing of other transport rules finishes, so encryption is applied to outbound messages after all tenant-specified rules finish. The net effect is that protection (and the x-header) is only present in messages sent by OWA and Outlook mobile after Exchange has processed all the other rules, so the rule described above can only block protected messages sent by Outlook desktop.

## Hybrid Protection

In a hybrid deployment, the on-premises Active Directory RMS servers support on-premises mailboxes, and Office 365 supports cloud mailboxes. Automatic synchronization of configuration data, including templates, does not occur. You must:

- Export the trusted publishing domain (TPD) data from your on-premises [RMS servers](#).
- Allow [external access](#) to your on-premises servers.
- [Import the TPD data](#) into Azure RMS.
- Enable any on-premises templates that are in use and make them available to cloud users.

It is desirable to ensure that the same templates are available to both cloud and on-premises users and that they run in the same manner. Because no automatic synchronization exists, any time an administrator updates the configuration for either the on-premises or cloud platforms, you must replicate the change to the other platform.

# Protecting Windows Files

The Unified Labeling client integrates with File Explorer to allow users to select and apply protection to files in much the same way as they perform other operations like printing. You can protect multiple files and folders in a single operation. To protect a file, you select the file in File Explorer and select **Classify and Protect** from the right-click menu. The options exist to merely classify a file or to apply full protection where you specify a list of users who can interact with the file together with the permissions you want to give them. For convenience, roles like Viewer – View Only and Co-Owner dictate the actions that recipients can take when they receive protected files. You can select recipients from the corporate GAL or input the email addresses for individual recipients, distribution lists, or complete domains. As you can see in Figure 20-22, you can also add an expiration date after which recipients no longer have access to the content.

After you protect a file, the client updates its properties to reflect the new protection. In some cases, like Office documents, this is a matter of adjusting the built-in attributes that control access to the file. In others, like JPEG files, the client converts the file into another format (in this case, protected JPEG). After protecting the file, you can share it with anyone you like using whatever mechanism is best. For instance, you can attach the file to a message and send it to someone. The protection placed on files works no matter how users share files, including uploading files to cloud sharing services like DropBox.

If the recipient is on the list of the users specified to access the content, they will be able to open it and interact with the content based on the rights granted by the author. Anyone else will be unable to access the content and told that they need to contact the owner to receive a version of the file that includes them in the permissions list. If the file is in a format supported by an application that supports rights management, the recipient can open and interact with the file in that application. However, if the file is in a format that is unsupported by enlightened applications or such an application is unavailable, you can download a [lightweight viewer](#) that understands how to interpret the protection placed on the file and open it for viewing.

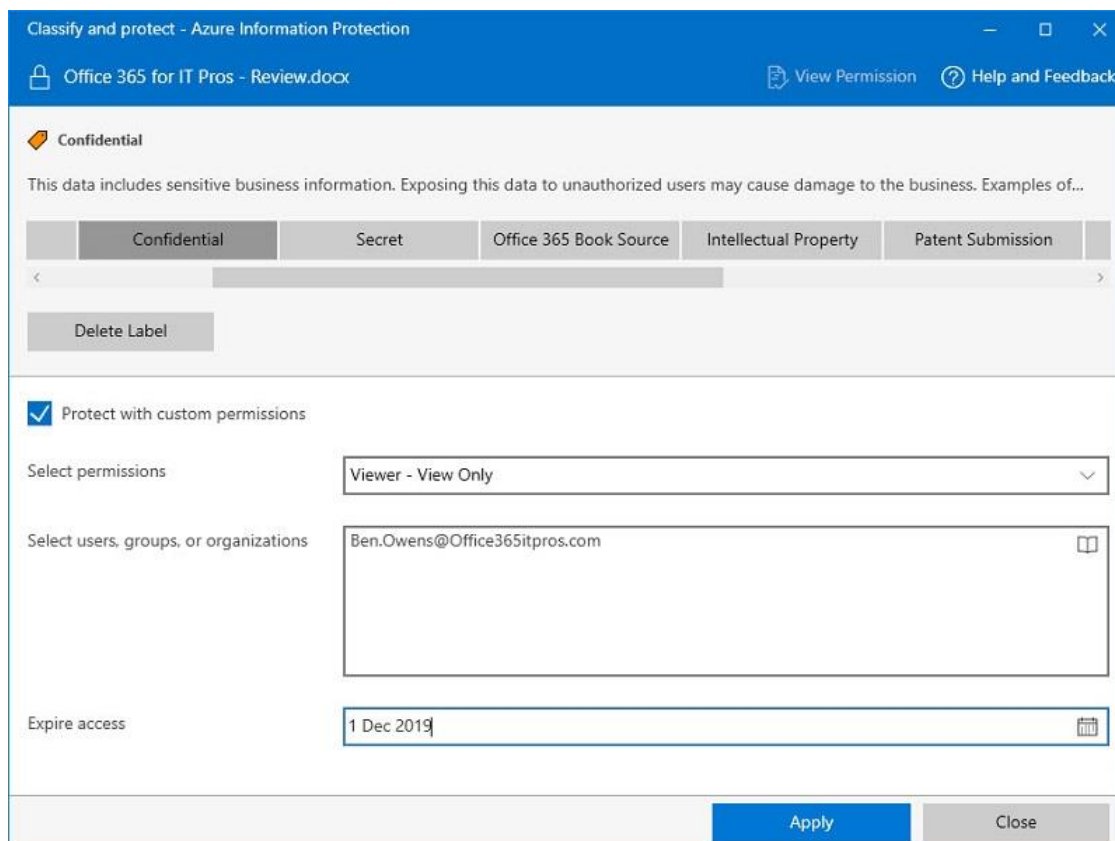


Figure 20-22: Applying a sensitivity label to a Windows file using the Unified Labeling client

As native support for sensitivity labels is available in the Office apps, installation of the unified labeling client on PCs is only necessary when you wish to assign labels to files stored outside Office 365 or to apply a label to a set of files selected in File Explorer (this action can also be performed through PowerShell).

**Disabling Classify and protect:** Some have asked if it is possible to disable the **Classify and protect** option so that it doesn't show up in File Explorer. It is possible by creating a new DWORD value called *LegacyDisable* at *HKEY\_CLASSES\_ROOT\AllFilesystemObjects\shell\Microsoft.Azip.RightClick*. Set the value to 0 (zero) to disable the option. Before doing this, remember that Classify and protect might be the only way that people can apply sensitivity labels to PDFs and other non-Office file types.

## Protected PDFs

[Microsoft and Adobe](#) collaborated to deliver a "native" integration of rights management protection for Adobe PDF documents. Native means that protection applied to PDF files uses the V1.7 of the ISO specification for PDF encryption. Applications that support the standard can open and process the protected content. Protection is unsupported for signed PDFs because this would break the method used to attest to the validity of the signatures.

Unlike earlier third-party implementations of rights-management protected PDFs, which use a PPDF file extension to show the encrypted nature of the files, native-protected files keep their PDF file extension.

To protect PDF files, you need to install the latest version of the [Unified Labeling Client](#). You can then protect PDFs using the **Classify and protect** option in File Explorer to choose a label or to assign custom permissions. You can also apply a label to a PDF using the *Set-AIPFileLabel* cmdlet. Because SharePoint Online can't process protected PDFs, the names of sensitivity labels applied to PDFs outside SharePoint Online do not appear in document views.

To access the content of protected PDFs, you must use a supported reader which understands both the ISO standard and how to display the assigned rights, such as Adobe Acrobat Reader (Microsoft's list of PDF readers that support Information Protection [is available online](#)). To use Acrobat, you need a recent version of Adobe Acrobat Reader or Reader DC and the [Microsoft Information Protection plug-in](#) (the plug-in is bundled with the Adobe installer from June 2022). Only Adobe Reader and not the Adobe Acrobat (paid-for) products support access to protected content. To have the Adobe products display information about the label assigned to protected PDFs, it's necessary to have a DWORD *bShowDMB* with a value of 1 at:

```
HKEY_CURRENT_USER\SOFTWARE\Adobe\Adobe Acrobat\DC\MicrosoftAIP
```

The Edge browser can open PDFs protected with sensitivity labels belonging to any tenant, if the user has the appropriate rights. However, although Edge displays a banner showing that a PDF is protected and tells the user what permission they have, it does not show any information about the assigned label.

Organizations can store protected PDFs in SharePoint Online or OneDrive for Business document libraries, but unless you use the Edge browser, you must download the files and use a supported viewer to view their content as the normal viewer used by SharePoint Online cannot decrypt the files. You can't apply, change, or remove sensitivity labels from PDFs stored in SharePoint Online or OneDrive for Business. Instead, you must download the file and process it with the unified labeling client, and then upload it again.

## PDFs Generated from Office Documents

The production versions of the Microsoft 365 enterprise apps (Word, PowerPoint, and Excel) do not create protected PDFs when they generate PDF files. However, from June 2022, the Office Insider builds (Beta Channel 2206 onward), create protected PDFs for the following actions:

- File – Save as PDF (unsupported for the PDF/A format).
- File – Export as PDF.
- Share – Send a Copy – PDF.



The Print to PDF option removes sensitivity label protection before it can generate a PDF. Because of this, the apps make the option unavailable if mandatory labeling is in force. In addition, password protection is unsupported for protected PDFs.

If the organization has earlier versions of the Office apps in use, protection can be applied to the PDFs generated by Office using the unified labeling client.

## Managing Sensitivity Labels with PowerShell

If you connect a PowerShell session to the compliance endpoint, you can access cmdlets to work with sensitivity labels. The easiest way to connect to the endpoint is to run the *Connect-IPSSession* cmdlet in the Exchange Online Management module. When connected, you can create and manage labels used for information protection and container management.

The *Get-Label* cmdlet returns the set of sensitivity labels defined in a tenant or the properties of an individual label. If encryption is enabled for a sensitivity label, its link to the underlying rights management template is in an object in the *LabelActions* property with the *Type* value set to "encrypt." The *LabelActions* property also holds the GUID for the label in the *TemplateId* field. In this example, we use the *Get-Label* cmdlet to fetch the properties of the Black Matter sensitivity label and then extract the rights definitions from its settings.

```
[PS] C:\> Connect-IPSSession
$Label = Get-Label -Identity "Black Matter"
$Rights = $Null
ForEach ($Action in $Label.LabelActions) {
    $Action = $Action | ConvertFrom-Json
    If ($Action.Type -eq "encrypt") {
        $Rights = $Action.Settings | ? {$_.Value -like "*Identity*"} | Select -ExpandProperty
Value | ConvertFrom-Json
    }
}
If ($Rights) { $Rights | Format-List }

Identity : BoardMembers@Office365itpros.com
Rights   : VIEW,VIEWRIGHTSDATA,OBJMODEL,DOCEDIT,EDIT,PRINT,EXTRACT,REPLY,REPLYALL,FORWARD

Identity : Black.Matter.Team@office365itpros.com
Rights   : VIEW,VIEWRIGHTSDATA,OBJMODEL,DOCEDIT,EDIT,REPLY,REPLYALL,FORWARD

Identity : Sanjay.K.Patel@office365itpros.com
Rights   : VIEW,VIEWRIGHTSDATA,OBJMODEL
```

The rights assigned in this label cover two groups and an individual user. The first set of sets equates to the co-author role. The second equates to the reviewer role and the last to the viewer role.

## Updating Rights in Sensitivity Labels

The *Set-Label* cmdlet updates label properties. For example, here's how to update the display name:

```
[PS] C:\> Set-Label -Identity Ultra -DisplayName "Ultra Confidential"
```

Another example is to use the *Set-Label* cmdlet to update rights assignments for a label. For instance, this command assigns a single set of rights to a domain and a separate set of rights to all authenticated users:

```
[PS] C:\> Set-Label -Identity "Ultra Confidential" -EncryptionRightsDefinitions
"quest.com:VIEW,VIEWRIGHTSDATA,DOCEDIT,PRINT; AuthenticatedUsers:
VIEW,VIEWRIGHTSDATA,OBJMODEL,DOCEDIT,EDIT,REPLY,REPLYALL,FORWARD"
```

The important thing to remember here is that if you run *Set-Label* to update the rights defined for a label, you overwrite the existing rights. It's therefore important that you retrieve the existing rights and include them in the set written back into the label. If you don't, anyone who currently depends on rights to interact with

content protected by the label will lose their access. You don't have to add rights for creators of documents and messages as they always have author rights.

The *New-Label* cmdlet is available to create new labels while the *Remove-Label* cmdlet deletes a label. As explained earlier, it's best not to remove a sensitivity label unless you are sure that the users have not applied the label to protect information. Usually, the better course is to remove the label from all label publishing policies to make it inaccessible to users. The label will continue to work but users cannot the label to new content.

Given the complexity of labels, especially those used with encryption, it is usually best to manage day-to-day updates through the Microsoft Purview Compliance portal and use PowerShell for common operations scripted to remove any chance of an error. For example, let's assume that a label protects content shared with partner organizations. You could:

- Maintain the list of partners in an Excel spreadsheet or other file.
- When a new partner is added, import the list of partner organizations from the file.
- Use the list to create rights definitions and add them to the label.

For more information on how to approach creating such a script, see [this article](#).

Once coded and tested, a process like this is invariably faster, more accurate, and less prone to error than performing a manual update through the GUI (it's also less boring). Any use of label cmdlets, like *New-Label*, *Set-Label*, and *Get-Label*, executed through PowerShell or the Microsoft Purview Compliance portal, is captured in events in the audit log and therefore can be analyzed to understand who is updating labels in the tenant.

## Removing Sensitivity Labels from SharePoint Online and OneDrive for Business Files

Global and SharePoint admins can run the *Unlock-SPOSensitivityLabelEncryptedFile* cmdlet to remove sensitivity labels with encryption from documents stored in SharePoint Online document libraries and OneDrive for Business accounts. The cmdlet is in the SharePoint Online module. It can only remove sensitivity labels belonging to the same tenant; it has no effect when run against documents protected by labels originating outside the tenant.

In effect, the unlock cmdlet delivers the same functionality available to rights management super-users when they use the *Set-AipFileLabel* cmdlet to remove labels from files stored in folders. Because SharePoint Online stores protected files in an unencrypted form (SharePoint applies encryption when users download documents), the cmdlet works by stripping the MIP metadata from documents.

The input parameters to the cmdlet are the full URL for the file and a justification for the removal of the label. These examples remove a label with encryption from files in a SharePoint Online document library and a OneDrive for Business account:

```
[PS] C:\> Unlock-SPOSensitivityLabelEncryptedFile -Justification "Needed to remove label"
-FileUrl https://office365itpros.sharepoint.com/sites/billing/Shared Documents/Invoice Tracking
2020.xlsx
[PS] C:\> Unlock-SPOSensitivityLabelEncryptedFile https://office365itpros-
my.sharepoint.com/personal/kim_akers_office365itpros_com/Documents/Planning.docx -Justification
"Needed to remove label from OneDrive files"
```

Some limitations exist. The only sensitivity labels removed by the cmdlet are those which:

- Include encryption using permissions assigned by an administrator (in other words, an administrator defines the rights for the label when creating or editing the label).
- Do not have user-defined permissions.
- Do not use Double-key encryption (DKE – see earlier section).

The unlock cmdlet doesn't do anything when it processes a file without a label or one with a label that doesn't include encryption.

## Removing Encryption from All Files in a Folder

The need often exists to remove protection from multiple files. Because it's focused on administrative activities, the SharePoint Online PowerShell module does not contain cmdlets to list files in a folder. Suitable cmdlets are available in the SharePoint PnP module. By combining the two we can:

- Connect to a target site and find all the documents in a target folder.
- Use the `Get-FileSensitivityLabelInfo` cmdlet to check each document for the presence of a sensitivity label with encryption.
- If a sensitivity label with encryption protects a document, remove it using the `Unlock-SPOSensitivityLabelEncryptedFile` cmdlet,

```
[PS] C:\> $SiteURL = "https://office365itpros.sharepoint.com/sites/BlogsAndProjects"
$FolderURL= "/Shared Documents/Blog Posts"
$UnLocks++
# Connect to PnP using modern authentication
Connect-PnPOnline -Url $SiteURL -Interactive
# Get all documents in the target folder
$FolderItems = Get-PnPFolderItem -FolderSiteRelativeUrl $FolderURL -ItemType File
Write-Host "Checking" $FolderItems.Count "documents to remove sensitivity labels with encryption"
ForEach ($Item in $FolderItems) {
    $ItemPath = $SiteUrl+$FolderUrl+"/"+$Item.Name
    $ProtectionStatus = Get-FileSensitivityLabelInfo -FileUrl $ItemPath
    If ($ProtectionStatus.ProtectionEnabled -eq $True) {
        Write-Host ("Removing {0} from {1}" -f $ProtectionStatus.DisplayName, $Item.Name )
        $UnLocks++
        Unlock-SPOSensitivityLabelEncryptedFile -FileUrl $ItemPath -JustificationText "Administrator
removed label"
    }
}
Write-Host "All Done." $UnLocks "documents unlocked"
```

When the cmdlet removes a sensitivity label from a file, SharePoint updates the name of the last user to modify a document to *System Account*. If you want to discover the account which ran the cmdlet to remove a sensitivity label, you should check the audit log to examine the *FileSensitivityLabelRemoved* events captured upon the removal of labels. The justification given by the user is captured in the audit event.

An example script [downloadable from GitHub](#) shows how to use a combination of PowerShell and Graph API calls to achieve better performance when the need exists to download large quantities of documents.

## Managing Labels Used for Container Management with PowerShell

Currently, you can't update the label settings applied to a container (Teams, Sites, or Groups) with PowerShell, but you examine the settings in the *LabelActions* property of the label. This property holds the label settings (like encryption, content marketing, and the site and group settings) in JSON format. The site and group settings for a label stored in *LabelActions* are in two sections:

- *protectgroup*: Privacy and external user access. The access type applied by this label is "Public" and group owners cannot invite guest users.
- *protectsite*: Access to SharePoint content from unmanaged devices. Group members can only access documents through online apps using managed devices.

The easiest way to view the label action settings is to extract the settings, convert them from JSON format, and list them like this:

```
[PS] C:\> $LabelActions = (Get-Label "Limited access").LabelActions | Convertfrom-JSON
$LabelActions | Select -ExpandProperty Settings
```

Key	Value
---	----
privacy	public
allowemailfromguestusers	false
allowaccesstoguestusers	false
disabled	false
allowfullaccess	false
allowlimitedaccess	true
blockaccess	false
disabled	false

Sometimes you need to check what labels apply specific controls on containers. This code shows how to build an array of labels with settings to block guest access to groups, teams, and sites. Each entry in the array contains the identifier (GUID) for the label and its display name.

```
[PS] C:\> $Labels = Get-Label
$LabelsBlockingGuests = @{}
ForEach ($Label in $Labels) { # Find out which labels apply blocks to guest users
    $LabelGuestAccess = $True
    $LabelActions = $Label.LabelActions | ConvertFrom-Json
    ForEach ($LabelAction in $LabelActions) {
        If (($LabelAction.Type -eq "protectgroup") -and ($Label.ContentType -Like "*Site*")) {
            $Settings = $LabelAction.Settings
            ForEach ($Setting in $Settings) {
                If ($Setting.Key -eq "allowaccesstoguestusers" -and $Setting.Value -eq "false") {
                    $LabelsBlockingGuests.Add([String]$Label.ImmutableId, $Label.DisplayName)}}
        }
    }
}
Write-Host "The following labels block guest access to Groups, Teams, and Sites"
$LabelsBlockingGuests | Format-Table Value
```

The array is useful in scenarios where it is necessary to check if the label assigned to a group will block guests.

```
[PS] C:\> $LabelDisplayName = $LabelsBlockingGuests.Item($Group.SensitivityLabel.Guid)
If ($LabelDisplayName -ne $Null) { Write-Host "The label assigned to" $Group.DisplayName "doesn't
allow guests"}
```

Another example is to extract the set of labels used for container management. In this case, we build a list with the label identifier, display name, and priority. You could use the list to check if a user has downgraded the sensitivity of a label assigned to a group, team, or site by comparing the priority of the existing label against the label they want to assign.

```
[PS] C:\> $Labels = Get-Label
$ContainerLabels = [System.Collections.Generic.List[Object]]::new()
ForEach ($Label in $Labels) {
    If ($Label.ContentType -Like "*UnifiedGroup*") { # It's a label for container management
        $DataLine = [PSCustomObject] @{
            LabelId      = $Label.ImmutableId
            DisplayName  = $Label.DisplayName
            Priority     = $Label.Priority }
        $ContainerLabels.Add($DataLine) }
}
```

**Label Actions and Advanced Settings:** PowerShell can update both normal and advanced Sensitivity label settings. The normal settings are the set visible when editing labels in the Compliance portal and are in the label's *LabelActions* property. Advanced settings are only accessible with PowerShell. Sometimes, Microsoft moves a setting (like a site's external sharing capability) from the advanced set and makes it available through the Compliance portal. When this happens, the setting moves into the set stored in the *LabelActions* property. The previous advanced setting remains in the label's properties, but the value of the setting in *LabelActions* takes precedence.

## Managing Sensitivity Labels for Containers with PowerShell

You can assign sensitivity labels to containers with PowerShell. Remember to use the GUID of a label instead of its display name. To assign a sensitivity label to a Microsoft 365 group, use the *Set-UnifiedGroup* cmdlet:

```
[PS] C:\> Set-UnifiedGroup -Identity "Banking Team" -SensitivityLabelId "9ec4cb17-1374-4016-a356-25a7de5e411d"
```

Assigning a sensitivity label updates the settings of the group as defined by the label. For instance, if you assign a label that blocks guest access, the group inherits that setting. *Set-UnifiedGroup* won't allow you to assign a sensitivity label to a group unless the group comes within the scope of a label publishing policy containing the label.

The *Set-SPOSite* cmdlet updates the sensitivity label for a SharePoint Online site:

```
[PS] C:\> Set-SPOSite -Identity "https://office365itpros.sharepoint.com/sites/dunesproject" -SensitivityLabel "9ec4cb17-1374-4016-a356-25a7de5e411d"
```

In this example, we search for a set of sites and apply the same sensitivity label to each site in the returned set. If you want to apply the sensitivity label to OneDrive for Business sites at the same time, *Get-SPOSite* will find those sites if you use the *-IncludePersonalSite \$True* parameter in the command.

```
[PS] C:\> $LabelGuid = "27451a5b-5823-4853-bcd4-2204d03ab477"
$Sites = Get-SPOSite -Limit All -Filter "URL -like 'Accounting'"
$Sites | ForEach-Object {Set-SPOSite $_.URL -SensitivityLabel $LabelGuid}
```

Instead of using the *Set-SPOSite* cmdlet to assign a sensitivity label to individual sites and processing multiple sites that way, the *Set-SPOTenant* cmdlet can assign a sensitivity label to multiple SharePoint Online sites. This method is faster when the tenant has more than a thousand sites:

```
[PS] C:\> $Sites | ForEach-Object {Set-SPOTenant $_.url -SensitivityLabel $LabelId }
```

*Set-SPOSite* does not allow assignment of a sensitivity label unless the target site comes within the scope of a label publishing policy containing the label. However, the error is silent, and you won't realize that the assignment failed unless you check the site properties afterward.

To remove labels from one site or a set of sites, use the *Set-SPOSite* cmdlet with the *RemoveLabel* parameter. For example, to remove the labels from the sites in the set used above, we'd run the command:

```
[PS] C:\> $Sites | ForEach-Object {Set-SPOSite $_.URL -RemoveLabel}
```

The *New-UnifiedGroup* and *New-SPOSite* cmdlets support adding a sensitivity label during the creation of a new group or SharePoint site. The *New-Team* and *Set-Team* cmdlets in the Teams PowerShell module have not yet been updated for sensitivity labels. When administrators assign labels to containers from one of the supported workloads, background synchronization processes make sure that the other workloads pick up the change. Yammer-based groups (communities) do not currently support sensitivity labels.

The Microsoft Graph Groups API also supports [assignment of sensitivity labels to Groups](#). However, apart from the MIP SDK, no public API is yet available to assign a sensitivity label to individual documents.

## Fetching Sensitivity Label Information for SharePoint Sites

To respond faster, when used to fetch a set of site objects, the *Get-SPOSite* cmdlet returns a limited set of properties. Sensitivity labels are not one of these properties, so if you want to find out which sites have labels, you must call *Get-SPOSite* to check each site. For example:

```
[PS] C:\> $Sites = Get-SPOSite -Limit All -Template Group#0
Foreach ($Site in $Sites) {
    $SiteDetails = Get-SPOSite -Identity $Site.URL
    If ($SiteDetails.SensitivityLabel.length -ne 0) {
        Write-Host "Site" $SiteDetails.Title "has sensitivity label" $SiteDetails.SensitivityLabel }}

```

## Controlling Default Sharing Links for Sites and Documents

When someone shares a document or folder, SharePoint Online or OneDrive for Business creates a sharing link. The link sets the scope (who can use the link) and permissions (what they can do). The sharing link inherits default settings for the scope and permissions from site settings or the tenant defaults. The default sharing settings can deliver a powerful hint to help people understand the level of confidentiality of information held in a site or individual document by guiding them to a certain mode of sharing. Although users can change a sharing link after SharePoint creates it, they often accept the default settings.

SharePoint administrators can configure the default sharing link settings for sites in the SharePoint admin center or by running the *Set-SPOSite* cmdlet with the *DefaultSharingScope* and *DefaultShareLinkPermission* parameters (which overrides the tenant defaults). This action can also be performed by configuring the advanced settings of a sensitivity label. Even better, the default sharing link settings held in a sensitivity label can apply to:

- SharePoint Online sites: Use the sensitivity label for container management and apply it to a site. The site inherits the settings from the label and overwrites any existing site settings.
- Documents: Apply the sensitivity label to a document. When someone shares the document, the default sharing link settings from the label take precedence over the site setting except when the site settings are more restrictive than the label settings. In that case, SharePoint Online uses the site settings.

Being able to apply default sharing link settings at a document level helps to protect documents stored in a site that should not be shared as widely as permitted by the site settings. For example, you might have a site holding project documents which can be shared with anyone in the organization except for project pricing files which must be restricted to specific people. In this scenario, administrators apply a sensitivity label to the site with a default sharing link settings to allow edit access for anyone in the organization while the pricing documents use a different label whose settings generate sharing links limited to specific people. Being able to control the default sharing link settings through sensitivity labels also helps when users store confidential documents in OneDrive for Business, again because users are likely to accept the default or at least think more about how they share information if they must adjust the sharing link settings before they share a document.

You can only update default sharing link settings for sensitivity labels using the PowerShell *Set-Label* cmdlet. First, run the *Connect-IPPSession* cmdlet from the Exchange Online management module to connect to the compliance endpoint. You can then use the *Set-Label* cmdlet to update sensitivity label settings. These settings are available:

- **DefaultSharingScope** defines the scope of the sharing link. The supported values are *SpecificPeople*, *Organization*, or *Anyone*.
- **DefaultShareLinkPermission** controls what a sharing link recipient can do with a file. The two available options are *Edit* and *View*.
- **DefaultLinkToExistingAccess** is *True* or *False* (default *False*). If set, this overwrites any other sharing link control and sets the link scope to people who already have access.

You can update default sharing link settings separately or together. For example, these commands set the default sharing scope and permission in two steps:

```
[PS] C:\> Set-Label -Identity 'Guest Access' -AdvancedSettings @{DefaultSharingScope = "SpecificPeople"}
[PS] C:\> Set-Label -Identity 'Guest Access' -AdvancedSettings @{DefaultShareLinkPermission = "Edit"}
```

Or set the two values in one command:

```
[PS] C:\> Set-Label -Identity 'Non-Business Use' -AdvancedSettings @{DefaultShareLinkPermission = "Edit"; DefaultSharingScope = "Anyone"}
```

To check the sharing link settings for a sensitivity label, run the *Get-Label* cmdlet. Values only appear for advanced settings after an update:

```
[PS] C:\> Get-Label "Non-Business Use" | Select-Object -ExpandProperty Settings
[contenttype, Site, UnifiedGroup]
[tooltip, Site, team, or group holding information that is confidential to the organization but can be shared with guest users. Private access.]
[displayname, Guest Access]
[defaultsharelinkpermission, Edit]
[defaultsharingscope, Anyone]
```

To enable *DefaultLinkToExistingAccess*, run:

```
[PS] C:\> Set-Label -Identity 'Confidential Access' -AdvancedSettings @{DefaultLinkToExistingAccess = "True"}
```

Like any other changes made to sensitivity labels, it can take up to 24 hours before SharePoint Online respects updates to the default sharing link settings.

## Updating Site Sharing Permissions

Another advanced setting for sensitivity labels controls the sharing permissions for sites. In the SharePoint browser interface, this option is available through *Site Permissions – Site Sharing*. Three settings are available:

- **MemberShareAll:** Site owners and members can share files, folders, and the site. People with edit permissions can share files and folders. This is usually the default setting assigned to new sites.
- **MemberShareFileAndFolder:** Site owners and members, and people with edit permissions, can share files and folders, but only the site owners can share the site.
- **MemberShareNone:** Only site owners can share files, folders, and the site.

To assign a new site sharing permission, run the *Set-Label* cmdlet and update the *MembersCanShare* advanced setting. For example:

```
[PS] C:\> Set-Label -Identity 'General Access' -AdvancedSettings @{MembersCanShare='MemberShareFileAndFolder'}
```

The ability to set sharing permissions for sites through sensitivity labels is currently in preview.

## Advanced Policy Settings

Management of most sensitivity label policy settings occurs through the Microsoft Purview Compliance portal. The unified labeling client supported several [advanced policy settings](#) (or “custom configurations”). Over time, Microsoft is updating sensitivity label policies to support some of the advanced settings in either the GUI or through PowerShell. For example, the label policy created earlier defines the default label to be “Public.” We can create a policy setting to force Outlook to use a different label in the GUI or with PowerShell as follows:

```
[PS] C:\> Set-LabelPolicy -Identity "General Sensitivity Policy" -AdvancedSettings @{OutlookDefaultLabel="2fe7f66d-096a-469e-835f-595532b63560"}
```

To check the normal and advanced settings for a policy, use the *Get-LabelPolicy* cmdlet as shown below. We can see that Outlook uses a different default label to the default label defined for documents, and that container support is enabled to allow super-users to use the *Set-AipFileLabel* cmdlet to remove labels from compressed files.

```
[PS] C:\> (Get-LabelPolicy -Identity "General Sensitivity Policy").Settings
[requireddowngradejustification, true]
[mandatory, true]
[outlookdefaultlabel, 2fe7f66d-096a-469e-835f-595532b63560]
[defaultlabelid, 27451a5b-5823-4853-bcd4-2204d03ab477]
[siteandgroupmandatory, false]
[enablecontainersupport, True]
[disablemandatoryinoutlook, True]
```

Outlook for Windows, Outlook for Mac, Outlook mobile, and OWA read and respect policy settings for:

- **DisableMandatoryInOutlook:** If the sensitivity label policy dictates mandatory labeling, this setting allows Outlook to avoid the need to assign labels to new messages. Set to False if Outlook should apply mandatory labeling, or True to disable mandatory labeling. Even if you disable mandatory labeling for Outlook, it continues to be mandatory for documents created in Word, PowerPoint, and Excel.
- **OutlookDefaultLabel:** If the sensitivity policy dictates mandatory labeling, this setting allows Outlook clients to use a different default label to the one applied to documents (as defined in the *DefaultLabelId* policy setting). The setting contains the GUID (label identifier) for the default label used by Outlook. Note that if a default label is defined for Outlook and mandatory labeling is required (even if disabled for Outlook), Outlook applies its label to all new messages.

Neither PowerShell nor the compliance endpoint validates the name of the advanced setting you update. If you misspell a parameter, PowerShell writes it into the label policy. If you pass an incorrect value, it will end up in the policy too. Always double-check the values you plan to use before updating a policy.

## Locale Settings for Label Names and Tooltips

By default, applications display the name and tooltip defined for a sensitivity label. If you work inside a monolingual organization, this shouldn't cause a problem as the default language is likely the one used by everyone. Inside a multilingual organization, it's a good idea to assign locale-specific values for label names and tooltips. PowerShell is the only supported method to update language values for a sensitivity label. For example, here's how to create a set of locale-dependent display names for the Confidential Access label used for container management. The code does the following:

- Define the label to update.
- Define an array holding the set of supported languages.
- Define an array holding the translated display name for the label in the supported languages.
- Define an array holding the translated tooltips for the label in the supported languages.
- Build JSON structures for the display names and tooltips.
- Run *Set-Label* to update the label settings.

This approach is much easier than the alternative of inputting all the values in one long command.

```
[PS] C:\> $Label = "Confidential Access"
$Languages = @("en-en", "fr-fr","it-it","de-de")
$DisplayNames=@("Confidential Access", "Accès confidentiel","Accesso riservato","Vertraulicher Zugang")
$Tooltips = @("Used for corporate confidential information","Utilisé pour les informations confidentielles de l'entreprise","Utilizzato per informazioni riservate aziendali","Wird für vertrauliche Unternehmensinformationen verwendet")
$DisplayNameLocaleSettings = [PSCustomObject]@{LocaleKey='DisplayName';
Settings=@(
@{key=$Languages[0];Value=$DisplayNames[0];}
```



```
@{key=$Languages[1];Value=$DisplayNames[1];}
@{key=$Languages[2];Value=$DisplayNames[2];}
@{key=$Languages[3];Value=$DisplayNames[3];}}
$TooltipLocaleSettings = [PSCustomObject]@{LocaleKey='Tooltip';
Settings=@(
@{key=$Languages[0];Value=$Tooltips[0];}
@{key=$Languages[1];Value=$Tooltips[1];}
@{key=$Languages[2];Value=$Tooltips[2];}
@{key=$Languages[3];Value=$Tooltips[3];}})
Set-Label -Identity $Label -LocaleSettings (ConvertTo-Json $DisplayNameLocaleSettings -Depth 3 -
Compress),(ConvertTo-Json $TooltipLocaleSettings -Depth 3 -Compress)
```

The key thing is to make sure that the set of language values for the display name matches the set of language values for the tooltips. In the example, both sets have values for default, en-us (U.S. English), French, Italian, and German. If you define a tooltip without a matching display name, the cmdlet will fail to update the label. The length of the tooltips makes them harder to input, so it's wise to compose the command outside PowerShell and paste it into the console.

## Setting Label Colors

The Office click-to-run and online applications support the display of different colors to help users identify labels. Administrators can assign colors to labels through the Microsoft Purview Compliance portal and or PowerShell by running the *Set-Label* cmdlet. For example:

```
[PS] C:\> Set-Label -identity "Confidential" -AdvancedSettings @{color="#008b8b"}
```

The *Set-Label* cmdlet accepts the color for a label in a hex triplet code giving the red, green, and blue combination to build the color. The example above assigns the code for dark cyan to the label. See [this page to interpret hex color codes](#). This code interprets the color settings for labels and translates the hex value using a hash table.

```
[PS] C:\> # Create a hash table to translate color hex values to display values.
$Colors = @{
    "000000" = "Black"
    "0000ff" = "Blue"
    "ff0000" = "Red"
    "ffc0cb" = "Pink"
    "ff8c00" = "Orange"
    "a80000" = "Dark Red"
    "8b0000" = "Darker Red"
    "317100" = "Dark Green"
    "0078d7" = "Dark Blue"
    "8a2be2" = "Bright Violet"
}
$Labels = Get-Label | Select DisplayName, Settings #Fetch label settings
ForEach ($L in $Labels) {
    $ColorFound = ($L.Settings | ? {$_ -match "color"})
    If ($ColorFound) {
        Try {
            $ColorCode = $ColorFound.ToString().Split("#")[1].Split(" ")[0] ; $ColorDisplay =
$Colors[$ColorCode]
            Write-Host "Label" $L.DisplayName "Hex color code is" ("#" + $ColorCode) ("(" +
$ColorDisplay + ")") }
        Catch {
            Write-Host "Error reading configuration for label" $L.DisplayName }
    }
}
```

Applications normally use the Calibri font when listing available sensitivity labels. You can change this by adding an advanced setting to specify which font to use. For example:

```
[PS] C:\> Set-Label -id "Black Matter" -AdvancedSettings @{fontname="Arial"}
```

# PowerShell for Rights Management

If you want to manage the protection service with PowerShell, you must download and install the [AIPService PowerShell Module](#) from the PowerShell gallery (or update the module to make sure that you have the latest release). After installing the module, you can connect to the service on an ad-hoc basis or include commands to connect to the service in your PowerShell profile. When you connect to the service, you can run the `Get-AipService` cmdlet to discover if protection is active (enabled) for the tenant. If the service is not enabled, you won't be able to protect information.

```
[PS] C:\> Import-Module AipService
[PS] C:\> Connect-AipService
A connection to the Azure Information Protection service was opened.
[PS] C:\> Get-AipService
Enabled
```

## Administrator Role for the AIP Service

Your account must have administrative rights for the rights management service before you can run the `Connect-AipService` cmdlet to connect PowerShell to the service. Accounts have administrative rights if they are:

- A global administrator for the tenant.
- A global administrator for the Azure tenant.
- Granted administrator access with the `Add-AipServiceRoleBasedAdministrator` cmdlet. For example, this command grants administrator access to the Kim Akers account. You'll see an error if the account already holds the role.

```
[PS] C:\> Add-AipServiceRoleBasedAdministrator -EmailAddress Kim.Akers@Office365itpros.com
```

To check the current list accounts assigned the AIP Service administrator role, run the `Get-AipServiceRoleBasedAdministrator` cmdlet:

```
[PS] C:\> Get-AipServiceRoleBasedAdministrator | Format-Table DisplayName, Role, EmailAddress
```

DisplayName	Role	EmailAddress
Kim Akers	GlobalAdministrator	SMTP:Kim.Akers@office365itpros.com
Marc Vilas	GlobalAdministrator	SMTP:Marc.Vigneau@office365itpros.com
James Ryan	GlobalAdministrator	smtp:JRyan@Office365itpros.com

The latest version of the AIPService module can [use certificate-based authentication](#) to support scenarios like fetching log information for ingestion into SIEM systems.

## Super-Users

Super-users are accounts that can decrypt any protected content. On-premises deployments often use super-user permissions to decrypt content to allow examination as messages pass through the Exchange transport service or when they review items following retrieval by eDiscovery searches. Assigning highly privileged access like super-user status is something that needs control, restricted to as few people as possible, and audited to track the use of super-user privileges.

However, the cloud is a very different environment. It is under Microsoft's control and tenant administrators do not have access to servers and other basic infrastructure elements that they could use to compromise information. Because Office 365 is a locked-down infrastructure, Microsoft runs a different regime where components decrypt protected content automatically. Transport rules can access protected content, discovery searches uncover protected content with ease, and journaling generates unencrypted reports. You do not

need to nominate any account as a super-user to make any standard functionality in Exchange Online or SharePoint Online work.

A situation might exist where you need to assign super-user status to an account to allow it to access encrypted content. For example, you might recover a set of protected Word documents left by an ex-employee that investigators need to review. Another example is where an eDiscovery search recovers a set of protected documents. Microsoft 365 Search can find protected content, but files remain encrypted when exported in the search results. In both instances, a super-user can decrypt the protected documents. Super-users are even able to decrypt documents if a tenant archives the sensitivity labels used for protection or after a document expires. The only time a super-user is unable to decrypt a document is after the deletion of a template. If the template does not exist, decryption is impossible.

You must enable the super-user feature before you can nominate accounts to be super-users. To enable or disable the super-user feature, first, load the protection cmdlets into a PowerShell session and connect to the rights management service. The set of commands to enable or disable the super-user feature are:

```
[PS] C:\> Connect-AipService
Enable-AipServiceSuperUserFeature
The super user feature is enabled for the Azure Information Protection service.

Disable-AipServiceSuperUserFeature
The super user feature is disabled for the Azure Information Protection service.
```

After enabling the super-user feature, you can add accounts to the super-user list. The email address that you pass must be valid for a user account belonging to the tenant, including those synchronized from on-premises Active Directory. You can add multiple users at one time, separating each email address with a comma. You'll receive an error if you try to add a user who is already on the super user list.

```
[PS] C:\> Add-AipServiceSuperUser -EmailAddress Oisin.Johnston@Office365ITPros.com
Oisin.Johnston@Office365ITPros.com was added to the list of super users for the Azure Information Protection service.
```

The other commands used to manage super-users are straightforward. *Get-AipServiceSuperUser* returns a list of the current accounts that have super-user privilege. Normally this list should be empty and the presence of any account on the list should have justification for the status. To remove an account from the super-user list, run the *Remove-AipServiceSuperUser* command as shown below:

```
[PS] C:\> Remove-AipServiceSuperUser -EmailAddress Oisin.Johnston@Office365ITPros.com
```

Information Protection logs all additions and removals of super users. You can retrieve this information from the AIP Admin log (but not in the Office 365 audit log). This example shows how to download the admin log for the last seven days.

```
[PS] C:\> Get-AipServiceAdminLog -Path "C:\Temp\AipAdminLog.log" -FromTime (Get-Date).AddDays(-7)
```

## Super-User Group

You can also add a super-user group by running the *Set-AipServiceSuperUserGroup* cmdlet. The address specified must point to a Microsoft 365 group or a security group.

```
[PS] C:\> Set-AipServiceSuperUserGroup -GroupEmailAddress MIPSuperUsers@Office365itpros.com
```

If you define a super-user group, the members of the group are all considered super-users. If the group is created with the ability to hold Azure AD roles (something that can only be done when a group is created in the Azure AD portal), it can be used with Azure AD Privileged Identity Management to assign super-user permission for limited periods to individual group members. In other words, instead of having unrestricted super-user access to any protected content in the tenant, the members of the super-user group can only access content when they receive access for a defined period.

## App-Based Super User Permission

The Microsoft Information Protection SDK supports the *Content.SuperUser* permission to allow Graph-based apps access to any protected content in a tenant. Like any other Graph permission, the app must be assigned the permission and consent granted by an administrator before the super user access can be used. See [this page](#) for more information.

## Using PowerShell to Protect Files

A set of PowerShell cmdlets is available to protect and apply labels to one or more files through scripting. These cmdlets are only available when the unified labeling client is present on a Windows workstation. The cmdlets call code in the client to add or remove protection to or from files. Authentication to use the cmdlets is through the account used to sign into the client.

### Retrieving Protection Information

The *Get-AipFileStatus* cmdlet returns protection properties for files, including the classification and templates applied to files.

First, let's check for files in a folder that have a label. Note that the cmdlet cannot return the status for a file if it is open in an application.

```
[PS] C:\> Get-AipFileStatus -Path "C:\Office 365 for IT Pros - Eighth Edition Files" | ?
{$_ .MainLabelId -ne $Null} | Format-Table FileName, MainLabelName, LabelDate
```

Another example shows how to find files in a folder protected by a specific template.

```
PS] C:\> Get-AipFileStatus -Path "C:\Office 365 for IT Pros - Eighth Edition Files" | ?
{$_ .MainLabelName -eq "Viewer - View Only"} | Format-Table FileName, MainLabelName
```

### Applying Protection to Files

To apply a label to files, we need to know the GUID used for the label. The easiest way to get this is to capture the label identifier from a file that we know has the desired classification.

```
[PS] C:\> $LabelId = (Get-AipFileStatus "C:\Temp\Important Staff.XLSX").MainLabelId.Guid
```

You can also find the GUID in the *Label ID* property of a label when viewed through the Microsoft Purview Compliance portal or by running the *Get-Label* cmdlet (after connecting to the compliance endpoint):

```
[PS] C:\> $LabelId = (Get-Label -Identity "Secret").GUID
```

Every use of the *Get-Label* cmdlet creates an audit event in the Office 365 audit log.

We can now use the label identifier to apply the same classification to one or more files. In this case, we look for all Word documents in a folder and its subfolders that do not yet have a label and apply the selected label to those files.

```
[PS] C:\> $TargetLocation = "c:\Temp\"
$TargetFiles = "*.docx"
$Files = (Get-ChildItem ($TargetLocation + $TargetFiles) -File -Recurse)
ForEach ($F in $Files) {
    $FileName = $TargetLocation + $F.Name
    $FileStatus = (Get-AipFileStatus -Path $FileName)
    If ($FileStatus.IsLabeled -eq $False) {
        Set-AIPFileLabel -Path $FileName -Label $LabelId }
}
```

Applying a label to a file updates the file and it can take several seconds to process a file. Be careful when you apply labels to a folder holding many files.

## Reporting Labeled Files

We can use the same technique to generate a report about the labeled and protected files in a folder, including folders holding synchronized files from a SharePoint Online or OneDrive for Business library. This code looks for any Office documents in a folder and then checks each file for the presence of a label. If a label is found, we extract details of the label and any associated rights management template. Note that the *Get-RMSTemplate* cmdlet is available after connecting to Exchange Online.

```
[PS] C:\> $Report = [System.Collections.Generic.List[Object]]::new() # Create output file for report
$Files = (Get-ChildItem "c:\temp\" -Include *.docx, *.xlsx, *.pptx -Recurse)
ForEach ($F in $Files) {
    $FileName = "c:\temp\"+ $F.Name
    $TemplateName = $Null
    $Status = (Get-AipFileStatus -Path $FileName)
    If ($Status.IsLabeled -ne $False) {
        If ($Status.RmsTemplateId -ne $Null) {
            $TemplateId = [GUID]($Status.RMSTemplateId)
            $TemplateName = (Get-RMSTemplate -Identity $TemplateId.Guid -ErrorAction
SilentlyContinue).Name
            If ($TemplateName -eq $Null) {$TemplateName = $Status.MainLabelName }
        }
        $ReportLine = [PSCustomObject]@{
            File           = $F.Name
            IsLabeled      = $Status.IsLabeled
            LabelId       = $Status.MainLabelId
            Label         = $Status.MainLabelName
            Date          = Get-Date($Status.LabelDate) -format g
            RMSGuid       = $Status.RMSTemplateId
            RMSTemplate   = $TemplateName
            Owner         = $Status.RMSOwner }
        $Report.Add($ReportLine)}
}
$Report | Export-CSV -NoTypeInfo c:\Temp\LabeledFiles.csv
```

The output for an individual file that has a sensitivity label with protection is:

```
File       : ABPs and Teams.docx
IsLabeled  : True
LabelId    : 81955691-b8e8-4a81-b7b4-ab32b130bff5
Label      : Secret
Date       : 13 Nov 2018 12:29:42
RMSGuid    : c7fc2174-097c-4123-9cad-15f1a32cb145
RMSTemplate : Secret
Owner      : Tony.Redmond@office365itpros.com
```

After processing the files, the script writes the information collected into a CSV file that can be opened and analyzed with Excel or Power BI.

## Removing Labels

To remove a sensitivity label from a document, run the *Set-AipFileLabel* cmdlet and specify the *RemoveLabel* parameter. You must also give a reason for the removal of the label. This command only works when:

- The logged-in account is an AIP super user.
- The target file is not open in any other application.
- The unified labeling client is on the PC. If you don't install the client, the *AzureInformationProtection* module will not load with the *AIPService* module and you can't run the *Set-AipFileLabel* cmdlet.

For example:

```
[PS] C:\> Set-AipFileLabel "C:\Temp\Important Stuff.docx" -RemoveLabel -JustificationMessage "Label
no longer necessary"
```

Before a super-user can run the *Set-AipFileLabel* cmdlet to remove protection from compressed containers like PST, MSG, ZIP, or RAR files, you must update the advanced settings of the label policy which applies to the super-user account to enable container support. This command is an example of enabling container support for a label policy:

```
[PS] C:\> Set-LabelPolicy -Identity "General Sensitivity Policy" -AdvancedSettings
@{EnableContainerSupport="True"}
```

The *Set-AIPFileLabel* cmdlet only works against items protected with Microsoft Information Protection encryption. It doesn't work against items messages protected with S/MIME, even if a sensitivity label applies S/MIME to an item.

## Processing Protected Documents Found in Content Searches

As an example of how to use the PowerShell cmdlets, let's assume that your organization must respond to a GDPR Data Subject Request (DSR). The Microsoft Purview Compliance portal can create and process DSRs. In this context (see Chapter 18), a DSR is a special form of eDiscovery case that depends on one or more content searches to find information relating to a named individual (the data subject). To satisfy a DSR, we need to find all content relating to a data subject. Protected messages in Exchange Online mailboxes are decrypted when exported by a content search; protected documents stored in SharePoint Online and OneDrive for Business libraries are not.

Before the organization can deliver copies of documents found by searches to the data subject, someone must review the content to make sure that the documents are related to that person. This means that protected documents found by searches must be decrypted to allow the check to proceed. In addition, there's no point in giving someone a set of protected documents that they cannot read.

A content search can export protected documents. When you export the results, you nominate a target folder to receive copies of the found files. Under the target folder, a folder (named after the date and time of the search) holds the search results, including the export summary, manifest, and a folder called SharePoint. Inside the SharePoint folder are folders for each site where the search found items and folders navigating to the point in the site where the search found the content. The path to such a folder might be something like this:

```
C:\Temp\Search for documents_Export\06.05.2018-1106AM\SharePoint\GDPR Planning Mark
II\gdprplanningmarkii\Shared Documents\General
```

The best approach is to gather all the protected documents found by a search in a single folder and process them there. The method used is to examine each file in the target folder and remove the protection if it exists. In this example, after running the *Connect-AIPService* cmdlet to connect to the rights management service with an account granted super user permission, we can form a collection of the files in the target folder and then loop through the set to find any that are protected. We then use the *Set-AipFileLabel* cmdlet to remove the protection.

```
[PS] C:\> $TargetFolder = "C:\Temp\Search for documents_Export\06.05.2018-1106AM\SharePoint\GDPR
Planning Mark II\gdprplanningmarkii\Shared Documents\General"
$Documents = Get-ChildItem -File $TargetFolder
ForEach ($D in $Documents) {
    $ProtectStatus = Get-AipFileStatus -Path $D.FullName
    If ($ProtectStatus.RMSTemplateId -ne $Null) {
        Write-Host $D.Name "is protected with" $ProtectStatus.RMSTemplateName
        $Message = $ProtectStatus.RMSTemplateName + " removed for GDPR DSR"
        Set-AipFileLabel -Path $D.FullName -RemoveLabel -JustificationMessage $Message }
}
```

The output is something like:

```
APC123.docx is protected with Intellectual Property
SPO Protected Content Test.docx is protected with Patent Submission
```

**FileName**

-----

C:\Temp\Search for documents\_Export\06.05.2018-1106AM\SharePoint\GDPR Planning Mark II\gdprplanningmarkii\Shared Documents\General...

C:\Temp\Search for documents\_Export\06.05.2018-1106AM\SharePoint\GDPR Planning Mark II\gdprplanningmarkii\Shared Documents\General...

When all the protected files are unprotected, investigators can review the document content to decide whether to release files as part of the response to the DSR. The same problem of how to deal with protected documents exists for any content search. Remember to remove the account from the list of super users after completing the decryptions.

## Using Microsoft Defender for Cloud Apps to Protect Office 365 Content

[Microsoft Defender for Cloud Apps](#) (MDCA) is a cloud access security broker (CASB) that can ingest and act upon [Office 365 audit information](#). The current set of supported apps includes:

- SharePoint Online.
- OneDrive for Business.
- Exchange Online.
- Teams.
- Dynamics 365.

MDCA is designed to give administrators insight into security-related events for a tenant. Given the number of events that even a small tenant can generate, automation through policies that act when specific criteria are matched is the best way to manage common conditions. For example, what action should happen when someone shares a file outside the tenant or creates a new document in a confidential site. If [Azure Information Protection is integrated with MDCA](#), MDCA retrieves the list of available labels in the tenant hourly, and adding a protection label to Office documents and PDF files is a supported action. Using this capability means that you can automatically apply protection to files matching policy criteria as users interact with them. On the basis that it should not override a decision made by a user, MDCA only applies a label if protection doesn't already exist on a file.

MDCA does not apply protection as users add files to Office 365. Instead, as MDCA ingests events from the Office 365 audit log, it looks for events (like document creation or modification) matching the criteria set in its policies and applies labels as necessary. The elapsed time between something happening a workload and a response occurring in MDCA depends on the ingestion of events from the audit log and the processing of those events in MDCA queues. Depending on the load on the service, the exact time will vary. For example, it might take between ten and twenty minutes before MDCA applies a label to a new file created in a SharePoint document library. If needed to protect an important file overlooked by a policy, an administrator can apply a label to a file from the MDCA dashboard.

The actions taken by MDCA to label files are visible in the Investigate section of its dashboard. You can apply filters to create queries to identify activity for specific applications, users, data ranges, and so on.

MDCA isn't free, but if you are concerned about protecting confidential Office documents stored in SharePoint Online or OneDrive for Business, it's hard to ignore the advantages of using policy-driven application of protection to sensitive content. You could take the manual approach and rely on individual users to do the right thing to protect their documents, but policy-driven automation invariably delivers a more reliable outcome.

# Microsoft Information Protection Auditing

Taking the necessary steps to make protection available to users is one step. Knowing that people use technology to protect important content is another. Microsoft Information Protection captures a wide range of audit data covering [most protection activities](#), such as someone applying a sensitivity label to a document or reading a protected document (but not protected messages). In addition, using the [AIP Scanner](#) generates some discovery actions when it scans sources to find files. Administrators can access the audit data through:

- The Activity Explorer in the Microsoft 365 admin center.
- The Audit feature in the Microsoft 365 admin center.
- The PowerShell `Search-UnifiedAuditLog` cmdlet.

The audit log ingests events for sensitivity labels with the other events generated by workloads. See Chapter 21 for information about how to retrieve and analyze data from the audit log. A [script downloadable from GitHub](#) illustrates how to interpret the content of the events logged for sensitivity label actions.

## Audit Records Captured for Office Applications

When users can assign, remove, or change sensitivity labels to documents using the Office (Word, PowerPoint, and Excel) apps, Microsoft Information Protection captures events for these actions. The events are:

- **SensitivityLabelApplied:** A site owner or administrator applies a sensitivity label to a SharePoint site.
- **FileSensitivityLabelApplied:** A user applies a sensitivity label to an Office document.
- **FileSensitivityLabelChanged:** A user changed a sensitivity label (upgrade or downgrade) for an Office document.
- **FileSensitivityLabelRemoved:** A user removed a label sensitivity from an Office document (in an app or with PowerShell).
- **DocumentSensitivityMismatchDetected:** A mismatch occurs because the sensitivity label applied to a document is higher than the level of sensitivity of the label applied to the site. The `SHAREPOINT\system` account generates the event when it detects the mismatch for a new or edited document.

The desktop versions of Word, Excel, and PowerPoint generate a different set of events:

- **SensitivityLabelApplied:** A user applies a sensitivity label to an Office document. This event differs from that logged for a site because its record type is `SensitivityLabelAction` rather than `SharePoint`.
- **SensitivityLabeledFileOpened:** A user opens a protected Office document on a workstation.
- **SensitivityLabeledFileRenamed:** A user generates a new version of a protected Office document by renaming the file.
- **SensitivityLabelRemoved:** A user removes a sensitivity label from an Office document. Office 365 also captures this event when a user edits a labeled file stored on a local device (not a copy synchronized by OneDrive).

Outlook desktop generates a `MipLabel` event when a user applies a sensitivity label to a message. The exception is when the user protects the message with one of the standard Office 365 Message Encryption templates (Encrypt Only or Do Not Forward). In these cases, Outlook treats the protected messages as if they are documents and logs `SensitivityLabelApplied` events. By contrast, OWA generates `MipLabel` events even when messages use OME. This is because Outlook can apply label-based encryption on the workstation whereas OWA uses the server to encrypt messages.

The Office mobile apps do not generate events when they apply sensitivity labels to messages or documents. In addition, the audit log does not capture events when third-party applications or the unified labeling client apply sensitivity labels to non-Office documents.



**Track that document!** Protected documents usually hold some form of confidential information, and it is good to know who has access to that information and to be able to revoke access if needed. To meet this need, users can track the progress of protected documents by signing into the [document tracking site](#). The ability to track documents is a premium feature (tracking of email messages is unsupported). For more information, see [this page](#). This feature will become more interesting when it is integrated into the Microsoft Purview Compliance portal. Microsoft hasn't said when this will happen.

## Cloud Exit for Encrypted Content

A "cloud exit" is when an organization decides to move some or all its content from a cloud service to some other platform. The other platform could be on-premises or another cloud service. An organization might need to perform a cloud exit after deciding to move all processing back to on-premises servers, to a competing service like Google Workspace, or even temporarily to regain access to content during a major outage. Tenant-to-tenant migrations also need to process protected content before moving data from the source to the target tenant to make sure that users can continue to access the content after they receive new user principal names in that tenant. The bottom line is that organizations that use rights management to protect content move to a new platform, they must do some up-front preparatory work to be able to move protected content to the target platform and maintain access. This often involves the decryption of large numbers of items.

The cloud exit process differs depending on if you use a Microsoft managed key (MMK) or have a company-owned encryption key (HYOK), and is described in [this post by the Information Protection team](#). The majority of Office 365 tenants that use MIP do so with an MMK. To move back on-premises, these organizations must export the keys (TPD) and then import them into an Active Directory rights management server before they can transfer and decrypt the protected content. This is not an operation that happens overnight, so it's necessary to do some up-front planning and preparation. It's also wise to export your tenant keys so that you always have access to this data should a problem arise.

In scenarios where organizations need to decrypt protected content before moving to a new platform, they can use the following approaches to decrypt information stored in Exchange Online, SharePoint Online, and OneDrive for Business:

- The *Set-AipFileLabel* PowerShell cmdlet and a super-user account to remove labels from files.
- The *Unlock-SPOSensitivityLabelEncryptedFile* and a SharePoint administrator account to remove labels from files stored in SharePoint Online or OneDrive for Business. See [this article](#) for an example.
- A content search to find and export protected Exchange Online messages (including protected attachments). Core eDiscovery decrypts the messages during the export process and can generate a separate PST for each mailbox.
- Advanced eDiscovery to find and export protected Office documents stored in SharePoint Online and OneDrive for Business.

None of these methods are fast, especially when tens of thousands of items are involved. They can only deal with content protected by Microsoft encryption technology (sensitivity labels and OME) and not with any other third-party encryption mechanism.

Temporary cloud exits might be necessary to access protected content in the case of an outage. For instance, if you use a third-party backup service to copy Exchange Online mailboxes, then restoring mailboxes to a usable state requires the decryption of any protected messages. Without access to a rights management server and the tenant keys, the mailbox might be restored, but none of the protected content is accessible.

# Chapter 21: Managing Auditing and Reporting

**Tony Redmond**

Microsoft 365 gathers audit information about user and system actions to track the interaction with workloads to add, update, and remove data. This chapter reviews how the audit system gathers audit events from multiple sources into one repository for administrators to query in multiple ways, including Office 365 Cloud App Security and ISV products. We also consider the question of how best to generate reports about user activity and consider the value delivered by the reports included in the standard tools and third-party reporting solutions. Finally, we review communications compliance policies and information barrier policies, two Microsoft Purview solutions that can help Microsoft 365 tenants control the use of information within their organization.

## Auditing Framework

The ability to reliably audit user and administrative operations is an important part of any compliance strategy. Although it is easy to enable auditing for a single workload, creating a common infrastructure to process and store auditing data from many different workloads is more challenging. This is especially true when lines blur between workloads when new applications such as Groups, Teams, and Planner integrate components and data from multiple workloads.

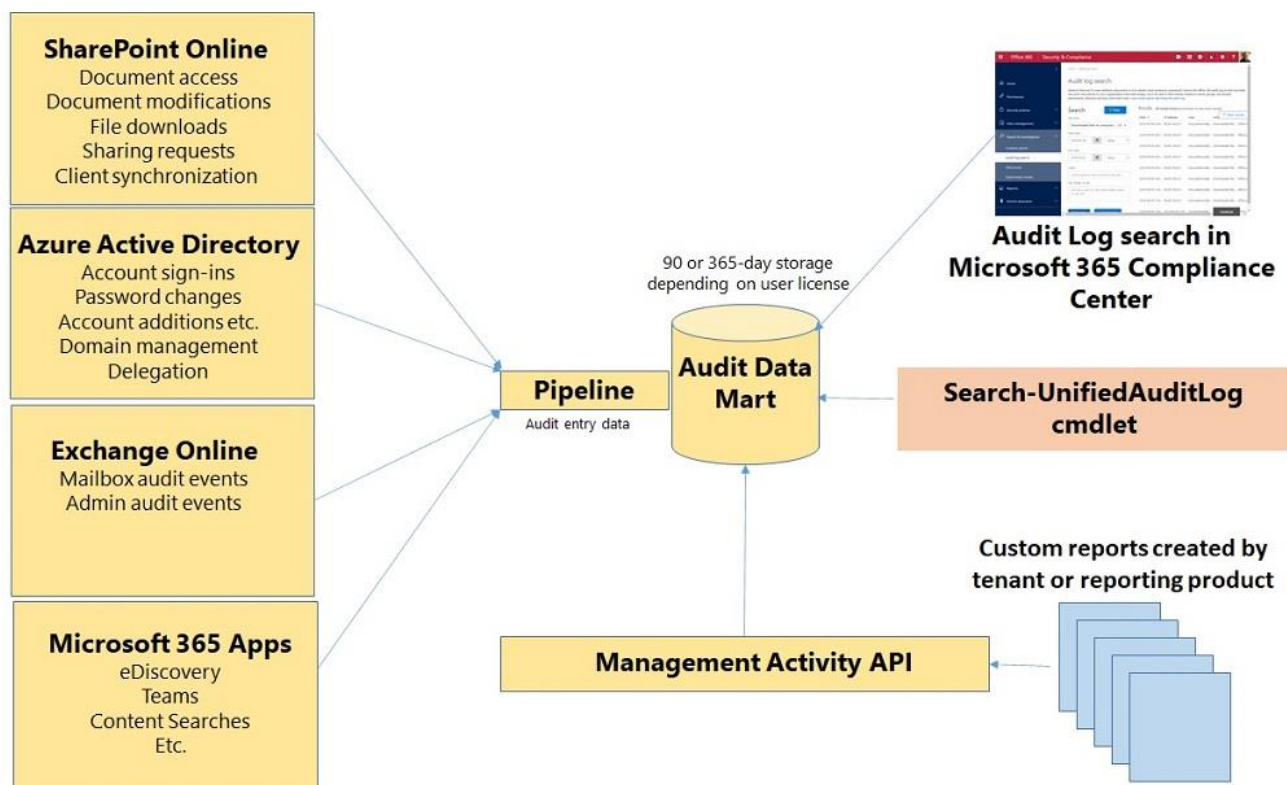


Figure 21-1: Auditing Framework

The background and history of the on-premises Exchange and SharePoint server applications is that each had a different method to enable and control auditing, specific ideas about what audit data to capture, and how to store audit data. Applications also implemented different methods to control access to the audit data. The result is a mess of inconsistencies and no way to ensure audit information remains secure and immutable. Microsoft 365 solves the problem by using a unified auditing framework covering all workloads in the suite. The architecture of the auditing framework in Figure 21-1 includes:

- **Data feeds flow from multiple workloads.** Each workload has varying abilities to capture audit data. It is therefore important to normalize the data retrieved from multiple workloads so that the audit events exist in a consistent format. As the auditing data mart ingests audit events, the ingestion process applies a [common schema](#) to events recorded in different workloads to make sure that all the events include a consistent set of essential information such as a timestamp, the user identity, the client IP address, client type, the action taken, and the object accessed. Events flow into the audit log from many workloads, including Exchange Online, SharePoint Online (including OneDrive for Business), and Azure AD as well as applications like Yammer, Stream, Teams, and Planner. Administrative functions such as eDiscovery operations are also captured. The schema accommodates the need for applications to capture information specific to their activities through product-specific schemas built on top of the common schema. For example, SharePoint Online supports product-specific schemas to describe file operations and sharing events. See [properties of audit log entries](#) for more information.
- **The Audit data mart ingests the data feeds.** The retention period for events stored in the data mart (better known as the audit log) depends on the licenses assigned to user accounts:
  - **Office 365 E3 or Microsoft 365 E3:** 90 days.
  - **Other Office 365 licenses with Microsoft 365 E5 Advanced Compliance add-on:** 365 days (year-long retention of Exchange, SharePoint, and Azure AD audit data is part of [Microsoft Purview Audit \(Premium\)](#).[Microsoft Purview Audit \(Premium\)](#)).
  - **Office 365 E5, Microsoft 365 E5, or Microsoft 365 E5 Compliance add-on:** 365 days.Office 365 Cloud App Security downloads audit data into its store and holds it there for 180 days. If you need to keep audit data for a longer period, third-party products usually extract and store audit data in their repositories for as long as a customer is willing to pay for the storage. The data mart is immutable because administrators cannot remove or change audit records. The data mart must ingest audit entries before records are accessible for reporting purposes. This usually happens within an hour or so but can take longer, depending on the source workload, other activities running within the service, and system events such as software updates. For this reason, the view of audit data is more precise when working with data a day or so after events occur than it is in the short term.
- **Four access methods** are available to consume the audit data:
  - The Audit log search (described below) in the Audit section of the Microsoft Purview Compliance portal. The log search supports ad-hoc queries through a GUI with the ability to export discovered records to a CSV file for later examination.
  - The Exchange Online *Search-UnifiedAuditLog* cmdlet can search the audit data from PowerShell.
  - Cloud App Security for Office 365 consumes audit events. See the later section. Other Office 365 features, such as activity alerts, also depend on audit events.
  - The [Office 365 Management Activity API](#) is available to developers to build third-party tools for audit reporting or analysis (or to extract data for injection into a different audit store – [a GitHub example is available](#)). The Management Activity API is a REST-based web service. The API aggregates actions and events drawn from the source workloads into tenant-specific content blobs, classified by their type and the content they hold (such as Exchange Online, SharePoint Online, or Azure AD). Tenants can use the Management Activity API to build

customized audit analysis tools to handle situations where PowerShell is unviable, as in the case of large-scale tenants where the number of audit records generated daily is too high for PowerShell to process in any reasonable time.

In a multi-geo scenario, some audit data might not be available as expected. For example, if a user has permission to send an email by impersonating another user and the accounts are hosted by different data center regions, Exchange Online captures the *SendAs* audit event in the sender's region but not in the region hosting the mailbox.

**Microsoft 365 workloads generating audit events:** The intention is that audit events should be available and reportable from all workloads. The full set of auditable events is [documented online](#) and includes:

- Exchange Online administrative events.
- Exchange Online mailbox events such as when mailbox delegates send messages (only for user and shared mailboxes, and not for public folder mailboxes or group mailboxes). These events only appear for mailboxes enabled for auditing (the default in Exchange Online).
- Microsoft 365 Groups (created and updated by many different applications). These events are collected by the Azure AD workload.
- SharePoint Online and OneDrive for Business file and folder events.
- SharePoint Online and OneDrive for Business site administration events.
- Azure AD events (like account logins and failed login attempts).
- eDiscovery (events like creating or running a content search), including the use of cmdlets to perform searches and other eDiscovery activities.
- Power BI (see instructions on [how to enable auditing for Power BI](#)).
- Microsoft Information Protection, such as the application of Sensitivity labels to documents.
- Teams, including the creation and removal of teams, channels, connectors, and tabs plus membership management.
- Power Automate (Flow), including the creation and deletion of flows and assignment of permissions.
- Stream, including the creation, removal, and editing of videos, channel activities, uploading and sharing of videos, and even when someone likes a video.
- Threat Intelligence (on the detection of phishing email or malware).
- Office 365 Cloud App Security alerts.

Of the major Office 365 apps, Planner is an outlier when it comes to generating audit events. You can find extra information about Azure AD audit events in the Azure portal. Not all Azure AD events appear in the audit mart. For instance, if you need to investigate failed or suspicious logins, more data is available through the [Azure portal](#).

Some consider that the gathering of audit data from multiple workloads constitutes a Security Incident Event Management (SIEM) or Cloud Access Security Broker (CASB) capability. However, auditing lacks the analysis and investigation features typically found in SIEM or CASB products like Microsoft Sentinel, or available in Office 365 Cloud App Security. The big advantage auditing holds over third-party products is the way that the various workloads integrate auditing into their operations and generate audit events for the audit log.

## Enabling the Audit Log for a Tenant

The audit log collects events generated by user and system activity from workloads on an ongoing basis. Before you can search for events in the audit log, you must enable this capability. If you go to the **Audit** section of the Compliance portal and select **Audit log search** and see a link to **Start recording user and admin activity**, you know that audit searches have not yet been enabled for the tenant. Click the link to turn on the ingestion process to bring audit data in from workloads.

You can also enable the audit log with PowerShell by running the *Set-AdminAuditLogConfig* cmdlet from the Exchange Online module:

```
[PS] C:\> Set-AdminAuditLogConfig -UnifiedAuditLogIngestionEnabled $True
```

The first time you run this command in a tenant, the substrate creates the audit log, and the ingestion of audit events commences. You only need to run the command thereafter if someone pauses the audit log and you wish to reenable it. You cannot pause the audit log through the Compliance portal, but you can do it with PowerShell:

```
[PS] C:\> Set-AdminAuditLogConfig -UnifiedAuditLogIngestionEnabled $False
```

Pausing the audit log allows workloads to capture events, but the events don't show up in the audit log until an administrator releases the pause. When that happens, the flow of events resumes populating the audit log. It might take some time before events flow from all workloads and the ingestion of events into the log stabilizes. During this time, some events might not reach the audit log.

## Searching Audit Events

When a tenant is enabled for auditing, events from the different workloads start flowing into the audit log. Because different methods of auditing are used by the workloads, it should come as no surprise that events flow into the audit log at different intervals after the workloads note details of the underlying actions. Although the exact periods will ebb and flow with demand and volume, the following delays are normal:

- **Exchange Online:** Audit events are available within 30 minutes. Events captured include mailbox audit events configured on a per-mailbox basis and administrative audit events configured for the tenant.
- **Teams:** Audit events for actions such as channel creation show up within 30 minutes.
- **SharePoint Online and OneDrive for Business:** Audit events are often available within 15 minutes but can take up to 30. Typical events include document editing and sharing.
- **Azure Active Directory:** Audit events for user sign-ins are available within 30 minutes and other directory activities within 24 hours.
- **All other workloads:** Audit events are available within 24 hours.

Figure 21-2: Configuring an audit search in the Compliance portal

Selecting **Audit** in the Microsoft Purview Compliance portal displays a form to collect the parameters to create a search of the audit log. Figure 21-2 shows the audit search interface. The process to execute a search is:

- Select the audit events that you want to look for from the drop-down **Activities** list. You can combine events from all the available audit sources in a single search, subject to the 5000-item limit on the

number of entries returned by a search. You might find that some operations are not available in the picker but are available in a PowerShell search. This is because it can take time to update the picker after Microsoft adds new events to the ingestion process.

- Define the date range to use (remembering that audit data are only available for a limited period). The default is to search for the last seven days (not including today). The interface supports searches for events created back as far as 365 days, but audit events for users with Office 365 E3 licenses are available only for 90 days.
- Select the user or users that performed the action you are looking for or leave the **Users** field blank to retrieve data for all users. You can search for specific users, including guest accounts, by entering a couple of characters in the field.
- If you are searching for audit events relating to SharePoint or OneDrive for Business, you can add supplementary information for the search such as the name of a file, folder, or URL. If you want to search for a word in a document title, give the full word rather than a partial substring. For example, if you want to search for events related to a document called "Reporting and Auditing," a search will find it if you specify "Auditing," but will not for "Audit."
- Click **Search** to begin the search.

The usual approach is to start by looking for a specific activity that occurred in a certain time range and then refine the search using tuned criteria to focus on a more precise set of events. You can refine the search by specifying the user whose activity you are interested in or a specific file, folder, or site to which an item belongs. You can combine a wide range of activities drawn from different workloads into a single search. Each event that you include in a search will probably increase the number of entries returned from the audit log and might make it more difficult to find the precise information that you want.

**The GUI Doesn't Show All:** Not every audit record available in the audit log appears in the search results returned in the Compliance portal. This is because many internal events are generated that might confuse or distract a log search. If you want to see the raw information in the audit log, use the `Search-UnifiedAuditLog` cmdlet.

Date	IP Address	User ↓	Activity	Item
Apr 1, 2021 2:10 PM	51.1...	tony.redmond@n...	Deleted file	ESPC21_ONLINE-Presentation-
Mar 22, 2021 1:27 AM	51.1...	tony.redmond@n...	Deleted file	Teams Meetings and Webinars.
Mar 11, 2021 4:57 PM	51.1...	tony.redmond@n...	Deleted file	Decrypting SPO documents - A
Mar 11, 2021 4:57 PM	51.1...	tony.redmond@n...	Deleted file	Decrypting SPO documents - D
Mar 11, 2021 4:57 PM	51.1...	tony.redmond@n...	Deleted file	Decrypting SPO documents.jpg
Mar 11, 2021 4:57 PM	51.1...	tony.redmond@n...	Deleted file	Decrypting SPO documents - S
Mar 26, 2021 8:52 PM	51.1...	tony.redmond@n...	Deleted file	Office 365 Conference Sessions

Figure 21-3: Viewing the results of an audit log search

In this instance, the search results listed in the results pane are for *Deleted file events*, so the search will find records associated with files checked into SharePoint Online document libraries. When the audit search displays its results (Figure 21-3), some details are visible about the found events, including the name of the checked-in file. To speed access to the audit data, the portal fetches and displays the most recent 150 entries. If more matching entries exist, the portal retrieves them in batches of 150 records as you scroll down through

the results. In total, the results pane can display 5,000 events. If more than this amount meets the search criteria, the results pane displays the most recent 5,000 events. Although you can examine the details of large data sets of audit records on-screen, you might lose interest in the search by the time you scan through thousands of lines of audit data.

The current UI doesn't include a way to filter records found by a search. For this reason, it's often better to perform a search using PowerShell and capture the output in an array. You can then run queries against the array to locate the exact audit records you want.

**Auditing and Secure Score:** If you enable auditing for your tenant, your Microsoft Secure Score goes up by 15 points. However, enabling auditing is not enough. It's more important to review what's gathered in the audit log regularly to understand what happens in the tenant and be able to detect when something out-of-the-ordinary occurs.

## User and System Events

It is also important to recognize that both user- and system-initiated operations generate audit events. For example, if offline copies of files from SharePoint Online or OneDrive for Business sites exist on the PC for some reason, a background synchronization process checks the sites periodically to ensure that the local cache holds up-to-date copies of the online files. The synchronization process shows up as a series of "Viewed File", "Accessed File", and "Downloaded File" events and might lead the observer to conclude that the user has accessed many files over a brief period.

Another example of an event that occurs as a by-product of user activity is the Accessed File event logged for the JPEG file for a user's profile photo. It is a fact that someone accessed the file, but it is unlikely that the fact will be important in any sense except in circumstances when you absolutely must prove that someone looked at someone else's profile photo. Examples of system-initiated activity include events recorded for the user "*app@sharepoint*", which relate to background processing of SharePoint and OneDrive sites while those generated by "*NT AUTHORITY\SYSTEM (Microsoft.Exchange.ServiceHost)*" belong to Exchange Online background jobs used to update the organization configuration. For instance, each time a user updates a file in a SharePoint library, you should see a matching *app@sharepoint* event as the crawler re-indexes the updated content. *app@sharepoint* also appears in "Added user or group to SharePoint group" events when users join the membership of a group and that change replicates to the SharePoint group for the team site.

For these reasons, you should take care to restrict the extraction of events to a manageable quantity by using a filter such as a limited time or events for selected users. In addition, you should ask yourself why an event might show up rather than jumping to any conclusions. Microsoft recognizes that the amount of audit data generated by a search can be excessive occasionally and that a filter to separate user-started from system-initiated events might be a sensible future enhancement.

**IP addresses in audit records:** If you look up the IP addresses reported in audit logs, you might find some that seem to be in unexpected places. The IP address might be for a home router connected to an ISP or it might be an IP address managed by Microsoft and assigned to one of their data centers. For this reason, products that use IP-based geolocation sometimes report that a connection comes from places that you know the user has never been. For instance, occasionally an audit record pops up to say that I accessed SharePoint Online from Helsinki. Much as I like the Finnish capital, I have not been there for at least 15 years. However, one of Microsoft's EMEA data centers is outside Helsinki and that is where the IP address originated. [Here's a site](#) that takes an IP address and reports its (assumed) location.

## Examining an Audit Record

To see the full details of an audit record found in a search, select it from the search results. Initially, some key information about the event is displayed to allow the reader to decide if this is an event needing further investigation. The information shown includes:

- **Date:** The timestamp recorded by the originating workload when an event occurred.
- **IP address:** The IP address of the originating client. This can be in IPV4 or IPV6 notation. Some workloads, like the Office Online apps, display the IP address of a trusted application calling into the service instead of the actual address of the device. Admin activities and those performed by system accounts for Azure AD updates don't record IP addresses.
- **User:** The user principal name of the account responsible for the event.
- **Activity:** The registered activity for the event. For example, if someone updates a file in a SharePoint Online document library, the registered activity is "modified file."
- **Item:** A description of the activity.

### Detail

#### Date

2021-05-24 15:25:39

#### IP Address

46.249.81.143

#### Users

vasil.michev@office365itpros.com

#### Activity

Modified file

#### Item

Managing Reporting and Auditing.docx

### Detail

Modified in "Shared Documents/2022 Edition"

#### Id

"e1600c07-b441-49a1-c065-08d91ebfce3e"

#### Object Id

"https://r[REDACTED]sharepoint.com/sites/O365ExchPro/Shared Documents/2022 Edition/Managing Reporting and Auditing.docx"

#### Record Type

6

#### Source Relative Url

"Shared Documents/2022 Edition"

#### User Agent

"Microsoft Office Word/16.0.13801.20442 (Windows/10.0; Desktop WOW64; en-US; Desktop app; Gigabyte Technology Co., Ltd./B450 AORUS PRO)"

#### CreationTime

2021-05-24T14:25:39

#### Id

e1600c07-b441-49a1-c065-08d91ebfce3e

#### Operation

FileModified

#### OrganizationId

Close

Figure 21-4: Viewing details of an audit record



To expose the extended audit information for an event, click on the event. The information revealed here is much more interesting when seeking to understand exactly what happened to an object. Figure 21-4 shows an example of an audit record for a document updated in a SharePoint Online document library. The information exposed here is what's stored in the audit log. We will get to the details of how to parse audit data with PowerShell soon to consider the purpose of each of these fields in the audit record.

## Exporting Audit Data

To capture audit data for further analysis, you can export the audit events to a CSV file. To export the items found in an audit log search, click **Export** and select **Download all results**. The export process downloads all available results from the audit log matching the search criteria.

The download stores the audit data in a CSV file with a name starting with "AuditLog." You can open the file or save it for further processing. The contents of the CSV are raw in that the export process does not expand and format the information for an audit event in the *AuditData* column. The details about the action taken by a user are stored in [JavaScript Object Notation](#) (JSON) format, so further processing is necessary to format the data for easier reading.

The maximum number of audit events that can be exported at one time is 50,000. If you end up downloading 50,000 events to a CSV file, the potential exists that more matching audit events are available that are not in the downloaded set, so it is a good idea to do another search to find events that might be missing (perhaps by using a different date range) and then merge the two sets of results.

## Searching Audit Data with PowerShell

The audit log ingests more than 1,600 different events generated by a range of workloads. The *Search-UnifiedAuditLog* cmdlet searches and returns data from the audit log. As we'll see, the cmdlet has some quirks to understand to execute successful searches, especially when it comes to interpreting the content of the *AuditData* property because different workloads insert different types of information into this property.

By default, *Search-UnifiedAuditLog* returns 100 audit records for any search request unless you specify the number of records to retrieve in the *ResultSize* parameter (up to 5,000). A single search can process a maximum of 50,000 audit records using page retrieval (see below). With the increasing number of workloads generating data for the audit log, a casual search can return hundreds if not thousands of records. It is important to be as specific as possible with search parameters to restrict the number of records returned.

Because the *Search-UnifiedAuditLog* cmdlet is an Exchange Online cmdlet, before you can view data in the audit log, your account must hold the Exchange View-Only Audit Logs or Audit Logs role. These roles are part of the Compliance Management and Organization Management role groups and can be assigned to [other role groups](#) as needed.

### Simple Audit Search

The simplest form of audit search looks for events of a specific type (referred to as an operation or event) for a single user (referred to as a *UserId*) over a short period. For example, to discover who last updated a document, we can search the audit log with a command like the one shown below:

```
[PS] C:\> Search-UnifiedAuditLog -StartDate 1-Nov-2021 -EndDate 1-Dec-2021 -RecordType
SharePointFileOperation -Operations FileModified, FileModifiedExtended -ObjectIds "Important
File.docx" -ResultSize 5 -Formatted | Format-Table UserIds, CreationDate, Operations
```

UserIds	CreationDate	Operations
tony.redmond@office365itpros.com	11/11/2021 14:35:44	FileModified

In this case, we can limit the number of audit records to retrieve by passing precise parameters:

- The **search period**: State the full date and time to delimit the start and end of the search. You can pass a date on its own, in which case the search uses 00:00 as the time to begin (or end) the search. If you use a date that's more than 365 days in the past, *Search-UnifiedAuditLog* returns an error and tells you the first possible date you can use. However, searches only return events as far back as a year for accounts with premium licenses.
- The **object name**: Pass the full name of the document. You don't need to include the URL of the SharePoint Online site or OneDrive for Business account storing the document. However, a document with the same name can be in multiple sites, so you might need to check the information stored in the *AuditData* property of the returned audit records to identify records belonging to a specific document. The same is true if you use a partial document name (such as ".docx") and the audit search returns events for modifications made to several documents. Audit searches don't support wildcard matching against object names.
- The **operation**: The action performed to generate the audit event. SharePoint Online generates *FileModified* and *FileModifiedExtended* events (operation) when someone updates a file in a document library. Every audit event has an operation, and you can filter audit records by specifying one or more operations to find.
- The **record type**: Bypassing *SharePointFileOperation* in the *RecordType* parameter, we tell *Search-UnifiedAuditLog* that we want to retrieve audit records for SharePoint Online file operations like uploading a new file to a document library. The other valid record types [are listed here](#).

Note that the *ResultSize* parameter is set to 5 to tell *Search-UnifiedAuditLog* that it only needs to find five records. Because Microsoft 365 usually returns audit records sorted by date, the first record is the latest update.

If a search finds audit records for multiple documents, a quick way to extract the records for a specific document is to store the found data in an array and then scan the records. For example, this search recovers all document modification events for a week.

```
[PS] C:\> [array]$Records = Search-UnifiedAuditLog -StartDate 18-Nov-2021 -EndDate 24-Nov-2021 -RecordType SharePointFileOperation -Operations FileModified, FileModifiedExtended -ObjectIds ".docx" -ResultSize 1000 -Formatted
```

To scan the records for a specific document, use the *Contains* method to check the contents of the *AuditData* property:

```
[PS] C:\> [array]$DetailedRecords = $Records | Where-Object {$_.AuditData.Contains("Managing Reporting")}
```

You can also use a free text search against audit records. This is slower than finding records of a certain type and filtering the records to find the right ones. However, it might be helpful at times. Here's an example:

```
[PS] C:\> [array]$Records = Search-UnifiedAuditLog -FreeText "*Reporting*" -Formatted -StartDate 1-June-2022 -EndDate 26-June-2022 -ResultSize 1000
```

**Finding the Right Events:** When searching the audit log, it's important to know what you're looking for. This sounds trite but given the number of audit events captured in any reasonably-sized tenant, looking for records for a specific action can be like searching for the proverbial needle in a haystack unless you know the name of the event. A good approach is to take steps to force the workload to generate an audit event, wait for 30 minutes or so, and then run an audit log search for the period to find the events which appear in the audit log. Delaying the search gives Microsoft 365 time to ingest audit records into the log. You can then examine the audit records captured for the period to find the events you want to analyze and use the *Operations* values logged for these events to perform further searches.

## Workload Audit Data

Audit records consist of two parts:

- A set of general properties populated in the same way by all workloads. These properties include the record type, creation date, operation, and user identifier.
- The *AuditData* property contains information specific to the workload generating the audit event. In most cases, you need to interrogate *AuditData* to discover the most important information about an event. Workloads use schemas to describe the properties they insert into audit records. The schemas are [documented here](#). Some trial and error are often necessary to interpret the payload in audit events. The guide to [detailed properties in audit log records](#) is helpful in this respect.

The *AuditData* property consists of a set of multiple attribute-value pairs separated by a comma (in JSON format). The property can extend to several thousand characters. To format the information and make it easier to follow, specify the *Formatted* parameter with the *Search-UnifiedAuditLog* cmdlet.

```
[PS] C:\> Search-UnifiedAuditLog -StartDate '01-May-2018 09:00' -EndDate '20-Jun-2018 17:00'
-ObjectIds 'Ch 20' -Operations FileModified -Formatted

RecordType      : SharePointFileOperation
CreationDate    : 13/06/2018 14:35:44
UserIds         : tony.redmond@office365itpros.com
Operations      : FileModified
AuditData       : {
  "CreationTime": "2018-06-13T14:35:44",
  "Id": "74746ae0-2fee-4399-18ce-08d3ab2af8fb",
  "Operation": "FileModified",
  "OrganizationId": "b662313f-14fc-43a2-9a7a-d2e27f4f3478",
  "RecordType": "SharePointFileOperation",
  "UserKey": "i:0h.f|membership|1003bffd805c87b0@live.com",
  "UserType": "Regular",
  "Version": 1,
  "Workload": "SharePoint",
  "ClientIP": "83.197.88.148",
  "ObjectId": "https://office365itpros.sharepoint.com/sites/0365ExchPro/Shared
Documents/Fourth Edition Files/Ch 20 - Reporting and Auditing.docx",
  "UserId": "tony.redmond@redmondassociates.org",
  "EventSource": "SharePoint",
  "ItemType": "File",
  "ListItemUniqueId": "d9189cf3-8c85-4ae9-9c21-c94dfe76c988",
  "Site": "acfe74d8-edfb-436d-924b-e018666605ee",
  "UserAgent": "Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML,
like Gecko) Chrome/51.0.2704.79 Safari/537.36 Edge/14.14385",
  "WebId": "a2aba197-5f1d-4864-a2c7-4daf0ff6379b",
  "SourceFileExtension": "docx",
  "SiteUrl": "https://office365itpros.sharepoint.com/sites/0365ExchPro/",
  "SourceFileName": "Ch 20 - Audit and Reporting.docx",
  "SourceRelativeUrl": "Shared Documents/Fourth Edition Files"
}
ResultIndex    : 1
ResultCount    : 125
```

The audit data is now more legible, and we can see many items of interest, including:

- **RecordType:** The workload that generated the record. Examples of values include:
  - AzureActiveDirectory: For example, add a member to a group.
  - MicrosoftTeams: For example, a user logs onto Teams.
  - ExchangeAdmin: For example, an update of mailbox properties.
  - SharePointFileOperation: For example, a client downloads a file.
- **CreationTime:** The date and time in UTC format when the user performed the activity. If your time zone isn't UTC, you need to adjust this value to get to a local time.

- **Operation:** In this example, *FileViewed* is a SharePoint Online operation logged when someone accesses an item in a document library. The PowerShell search examples described here demonstrate how to search for and interpret events generated by different operations performed in workloads.
- **OrganizationId:** The unique GUID for the tenant.
- **UserKey:** The identity (in this case, achieved through membership of a group) used to gain access to the item.
- **Workload:** The name of the application workload that logged the event. In this case, it is SharePoint Online. Other values are Exchange Online, OneDrive for Business, and Azure AD.
- **ClientIP:** The IP (V4 or V6) address of the client workstation where the action originated is logged. See earlier note about when the IP address might not be for the workstation.
- **ObjectID:** The full path to the object is logged. We can see that this resolves to a document in a SharePoint Online document library.
- **UserID:** The Azure AD account identifier for the account that caused the action to occur.
- **UserAgent:** The client used to invoke the action. In this instance, a user uploaded a file to the document library.
- **SourceFileName:** The name of the file.
- **UserType:** The type of user that performed the action. "0" indicates a normal user; "1" indicates an action taken by an administrator, while "2" means that the action was taken by a Microsoft data center administrator or data center system account. Because we used the *Formatted* switch for the *Search-UnifiedAuditLog* cmdlet, the numeric value 0 is translated to "Regular."
- **EventSource:** Only SharePoint Online uses this field. It is either SharePoint or ObjectModel.

The *ResultIndex* and *ResultCount* properties are interesting if you need to keep track of a large record set. *ResultIndex* tells us the record number within the set returned. *ResultCount* tells us the total number of records returned. We can therefore say that the record shown is the first of 125 returned in the set. If the search encounters an internal timeout, *ResultIndex* will be -1.

It takes time for administrators to become accustomed to the information contained in audit records and to figure out how best to use this data when confronted with questions such as "who interacted with a document" or "who created new documents in this period." Experience with the PowerShell cmdlets and some trial and error soon shows that the audit log is a surprisingly useful source of information.

## Fetching Large Amounts of Audit Data

In large tenants, or where you need to retrieve information about multiple operations over an extended period, it is likely that a search will return more than 5,000 audit records. In this scenario, to make sure that you retrieve all records, you fetch audit data in pages holding up to 5,000 records at a time until no more data is available. This technique supports the retrieval of up to 50,000 audit records. If more than 50,000 matching records exist, you need to split the work across multiple searches, each of which uses different criteria. You store the results of the searches in an external repository (many organizations use Splunk for this purpose) and run the analysis against that repository.

*Search-UnifiedAuditLog* has two parameters to support the retrieval of large data sets:

- The **SessionId** parameter holds a string value to identify a search session. You can use any value you like from a simple number to a GUID generated with the *New-Guid* cmdlet. The presence of a session identifier tells *Search-UnifiedAuditLog* that it might need to fetch several pages of data.
- The **SessionCommand** parameter tells *Search-UnifiedAuditLog* how to handle large amounts of audit data. The returned data might contain duplicate records. This parameter can be set to:
  - *ReturnLargeSet*: The audit records returned are unsorted. You can fetch up to 50,000 audit records using this method but must remember to sort the data once it is all fetched.

- *ReturnNextPreviewPage: Search-UnifiedAuditLog* returns audit records sorted by date. However, you can fetch only a maximum of 5,000 records using this method. If more matching records exist, attempts to fetch the data will result in an error.

The essential steps to fetching large amounts of audit data are:

1. Create a session identifier.
2. Set up a loop to fetch data in pages.
3. Run *Search-UnifiedAuditLog* several times to fetch all available data.
4. For each run of *Search-UnifiedAuditLog*, store the retrieved data.
5. After fetching all pages, sort the data by date and write it out to a CSV file.

In this example, *Search-UnifiedAuditLog* fetches audit data about SharePoint Online file operations in batches of 4,500 records at a time. The only data stored is in the *AuditData* payload from the audit record. The output of statistics after each page is purely for informational purposes.

```
[PS] C:\> $StartTime = Get-Date; $SessionName = (New-Guid).Guid; $OutData = @()
$OutputFile = "c:\temp\Output.csv"
$EndDate = (Get-Date).AddDays(+1); $StartDate = (Get-Date).AddDays(-90)
$i = 0
Write-Host "Searching Office 365 Audit Log..."
Do {
    $ThisSearchStart = Get-Date; $i++
    [array]$Records = Search-UnifiedAuditLog -StartDate $StartDate -EndDate $EndDate -SessionId
$SessionName -SessionCommand ReturnLargeSet -ResultSize 4500 -RecordType SharePointFileOperation
    If (($Records.Count -gt 0) -and ($Records[0].ResultIndex -eq -1)) { # The audit data returned
is bad
        Write-Host "Error occured fetching audit data. Resetting search and pausing before retrying
..." -foregroundcolor Red
        Start-Sleep -Seconds 240 # Wait for 4 minutes
        $i = 0 ; $SessionName = (New-Guid).Guid; $OutData = @() # Go back to zero
        Continue
    }
    If (($Records.Count -gt 0) -and ($Records[0].ResultIndex -ne -1)) { # Got a good page
        $OutData += $Records | Select-Object -ExpandProperty AuditData | ConvertFrom-Json
        Write-Host ("Completed Page #{1}, returned {2} records in {0} seconds. Total records found so
far {3}" -f [math]::Round((New-TimeSpan -Start $ThisSearchStart).TotalSeconds), $i, $Records.Count,
$OutData.Count) }
} Until ($Records.Count -eq 0) # Until we find no more records

$OutData = $OutData | Sort {$_.CreationDate -as [datetime]}
Write-Host ("All done {0} records found in {1} minutes." -f $OutData.Count, [math]::Round((New-
TimeSpan -Start $StartTime).TotalMinutes,3))
```

The check against the *ResultIndex* property handles the situation where *Search-UnifiedAuditLog* sometimes returns duplicate information due to an internal timeout when fetching data. In this situation, *ResultIndex* is set to -1 (minus one). The search is invalid, and the results are possibly duplicated. The only solution is to zeroize everything and restart the search after a short delay.

Even in smaller tenants, you might want to use *Search-UnifiedAuditLog* to fetch audit data to store in an external repository so that the data can be kept for longer than it is in Office 365. A scheduled job could be run daily to fetch all events for the last day for ingestion into the external repository, which could be another cloud service like Splunk or an ISV reporting product like [Quadrotech NovaQuest Nova](#) (ISV products usually employ a range of Microsoft APIs to fetch audit and other data to log details of workload activity). Many compliance problems are only discovered months after an event occurs and you might not be able to find all the evidence required for an investigation.

## Discovering What Audit Operations Exist

As explained above, you can filter audit records by specifying the type of operations to see. For instance, to see who sends emails on behalf of a shared mailbox, you can look for audit events with the *SendAs* operation.

New operations appear in the audit log on an ongoing basis as workloads enable auditing for new features. The question of how best to discover new events therefore exists. Here's what we do to find if a new feature is captured in an audit event.

- First, use the new feature. Ideally, perform actions several times with different accounts.
- Second, wait for at least an hour to allow the ingestion of audit events from the source workload and appear in the audit log.
- Next, run a search to find all audit events for the current day and group and sort the results by operation. Make sure to specify the user principal name of the account which performed the accounts in the *UserIds* parameter.

```
[PS] C:\> [array]$Records = Search-UnifiedAuditLog -StartDate (Get-Date -format dd-MMM-yyyy) -
EndDate ((Get-Date).AddDays(1)) -ResultSize 2000 -Formatted -UserIds Ken.Bowers@office365itpros.com
$Records | Group Operations | Sort Count -Descending | Format-Table Count, Name
```

You should now be able to browse the sorted list of operations to find unfamiliar actions, such as *Set-LabelPolicy* (logged when someone updates a sensitivity label policy). You can take the same approach with the Audit search feature in the Compliance portal, but not all audit events show up there.

Searching the audit log to find new events also uncovers audit events logged when Microsoft updates tenant settings as part of their normal operations. For instance, Microsoft often updates OWA mailbox policies to introduce a control for a new OWA feature. When this happens, you'll find audit events logged for a user called *NT AUTHORITY\SYSTEM (Microsoft.Exchange.ServiceHost)* for the policy updates.

## Practical Examples of Using Audit Information

Audit records hold a wealth of information in the *AuditData* property. To make the information more approachable, we must break out the important pieces from the JSON-formatted data in this property and use the different elements for whatever analysis is required.

### Deciding What Audit Data is Output

All the examples in this section use a generic PowerShell list called *\$Report* as the output for information about audit data. You can use PowerShell to sort, group, or otherwise process the table to display the audit data however you like. Naturally, because each workload generates both a set of common audit data and some information of its own, you might want or need to adjust the code in the examples to meet your specific requirements by adding new outputs to the object or changing what is written into it.

For instance, somewhat confusingly, audit records have two timestamps to record when an action occurred. The date is the same, but the format is different. The *CreationDate* property outputs the timestamp in the locale defined for the workstation where you run the PowerShell script.

*9 Jan 2019 12:18:47*

The *CreationTime* property in *AuditData* holds the timestamp in ISO 8601 format, as in:

*2019-01-09T12:18:47*

The ISO 8601 format is not dependent on the locale and is what PowerShell uses when it outputs the current date and time using a command like:

```
[PS] C:\> Get-Date -Format s
```

You can use either timestamp as they both store the same data. If you prefer to use a different date format in the reports to match whatever format the consumers of the report like to see, change the line that writes the timestamp into the report. For instance, if you use the *CreationTime* property, you can output it in the general date/time pattern in short form – like *9-Jan-2019 12:18*:

```
[PS] C:\> TimeStamp = Get-Date $AuditData.CreationTime -format g
```

## Using Input Arrays

If you are searching for audit events for multiple operations, it is convenient to populate an array with the operations and use the array as input for the search. For example:

```
[PS] C:\> $Operations = @('FileAccessed', 'FileDownloaded', 'FileModified', 'FileDeleted', 'FileUploaded')
```

An array can also hold user identifiers as input. This example populates an array with user principal names fetched by calling the *Get-ExoMailbox* cmdlet.

```
[PS] C:\> [array]$Users = Get-ExoMailbox -Filter {CustomAttribute1 -eq "Sales"} | Select -ExpandProperty UserPrincipalName
```

Once populated, pass the arrays in the same way as you would pass individual operations or user identifiers:

```
[PS] C:\> Search-UnifiedAuditLog -Operations $Operations -UserIds $Users -StartDate $StartDate -EndDate $EndDate -ResultSize 5000
```

## Examining Retrieved Audit Information

Three options exist to examine the information generated by the search and held in the *\$Report* object.

1. If the number of audit records found is small, you can examine the data on the screen.
2. An alternative is to pipe the information in the table to the *Out-GridView* cmdlet. This makes it much easier to review (and sort) the data.

```
[PS] C:\> $Report | Out-GridView
```

3. Once the number of records grows, it is usually easier to export the data as a CSV file:

```
[PS] C:\> $Report | Export-CSV c:\temp\Report.csv -NoTypeInformation
```

We can then open the CSV file with Excel to format the data as desired or load the CSV file into a tool like Power BI to generate graphs, views, and reports from the audit data.

## Tracking Group Creation

The first example creates a report about who created groups, including the workloads used to create the groups. The relevant event to look for is "Add Group." In this example (code [downloadable here](#)), we search for groups created over the last 90 days. The script processes the records to extract information about who created each group, the group name, and the workload used.

The output for the report should look something like the example below. In this case, we see details of the creation of six groups using a variety of workloads from Yammer to Planner (*ProjectWorkManagement*) to Teams. The *Microsoft.Exchange* workload means that an Exchange client (like OWA) created the group. If you see a workload named with a value like "12128f48-ec9e-42f0-b203-ea49fb6af367," it is the Teams PowerShell module.

TimeStamp	Workload	User	GroupName
9 May 2019 19:44	ProjectWorkManagement	James.Ryan@office365itpros.com	0365Grp-James Plan
8 May 2019 22:11	Microsoft Teams Services	Kim.Akers@office365itpros.com	0365Grp-Acquisitio
8 May 2019 19:01	Microsoft.SharePoint	Joe.Richards@office365itpros.com	Nice Airport Watch
8 May 2019 12:36	Microsoft.Exchange	Brian.Weakliam@office365itpros.com	Brian's Nelson Gro
7 May 2019 17:45	Microsoft.Exchange	Vasil.Michev@office365itpros.com	Sandboxes

If you run this script, you might notice that Microsoft has updated some of the workload names. For instance, the current name for Exchange is "Office 365 Exchange Online" while SharePoint is "Office 365 SharePoint Online."

## User Sign-ins

In another example of how to use the same technique to interpret audit events, here's what you might do to report user sign-ins to different workloads. To make things more complicated, we use two different events, one from Azure AD (to handle most workloads) and the other from Teams. Because we can expect to retrieve many audit records for these operations, we also pass the *ReturnLargeSet* value to the *SessionCommand* parameter and specify that we will accept up to 5,000 records. The code is [available on GitHub](#). Another way to get user sign-in information with PowerShell is via the *Get-MgAuditSignInLogs* cmdlet.

Understanding failed user login events is also interesting because these events might be the result of attempts by attackers to penetrate your tenant. You can find code to search for audit events for failed user sign-ins [on GitHub](#).

## Who Updated That File?

People often want to know who made changes to a document. This search finds audit events created when someone creates or edits a specific file over the last 90 days. The script prompts for a file name to search for and stores its name in the *\$FileName* variable. Depending on the document's history, we could find:

- A single *FileUploaded* event when the document is uploaded to SharePoint Online or OneDrive for Business.
- A *FileAccessed* event when a user opens the document.
- A *FileModified* event when a user updates the document. Events also appear from the background process used by SharePoint Online to maintain files. These appear as the *app@sharepoint* user.

If the AutoSave feature is enabled for the document, multiple *FileModified* events can accumulate over a short period. We can extract information about who did what to the document from the *AuditData* field for the audit events, which contains information about the site (including its URL) where the document is stored.

After analyzing the set of audit records found for the document (use the script [downloadable from GitHub](#)), we can list the results:

```
[PS] C:\> $Report | Select Timestamp, User, Action
```

TimeStamp	User	Action
22 Apr 2020 14:40:41	Jane.Maloney@office365itpros.com	FileModified
21 Apr 2020 15:19:03	Jane.Maloney@office365itpros.com	FileModified
21 Apr 2020 15:02:34	Kim.Akers@office365itpros.com	FileModified
21 Apr 2020 15:01:39	Jane.Maloney@office365itpros.com	FileUploaded

SharePoint Online and OneDrive for Business move deleted files through a two-stage recycle bin. Three different audit events capture each stage:

- **FileDeleted:** The original deletion (by a user, a retention policy, or system process).
- **FileDeletedFirstStageRecycleBin:** A user removes a file from the first stage recycle bin. Any member of a team or group can do this. Normally, files stay in the first stage recycle bin for 30 days.
- **FileDeletedSecondStageRecycleBin:** A site administrator (group or team owner or the owner of a OneDrive for Business account) removes a file from the second stage recycle bin. Usually, files remain in the second stage recycle bin until 93 days after their original deletion. At this point, SharePoint Online removes the file permanently and it is irrecoverable (unless a retention label or policy forces SharePoint Online to keep a copy in the site preservation hold library).

Reporting file deletion events is a straightforward process. An example script to illustrate the process [is available in GitHub](#).



## SharePoint Sharing Events

SharePoint logs audit events when users generate sharing invitations with people inside and outside the tenant. Three separate events are recorded:

- *SharingSet*: Someone shares a document with someone else (inside or outside the tenant).
- *SecureLinkCreated*: SharePoint creates and sends a secure link to the target user. This only happens for external users as users with accounts in the tenant directory can use their accounts to access the shared document.
- *SecureLinkUsed*: The target user uses the secure link to access the document. Again, this only happens when external users access a shared document.

You can search for these audit records with a command like:

```
[PS] C:\> [array]$Records = (Search-UnifiedAuditLog -StartDate (Get-Date).AddDays(-30) -EndDate (Get-Date).AddDays(+1) -Operations SharingSet, SecureLinkUsed, SecureLinkCreated -ResultSize 2000)
```

If you then interpret the audit records, you see this kind of sequence when a sharing invitation is generated and used by an external user.

```
TimeStamp : 2018-04-16T16:21:42
User      : tony.redmond@office365itpros.com
SharedWith : SharingLinks.dd7d0f94-7757-4646-afcd-c1b824f8676d.Flexible.7a5975c0-8ac2-4167-ba4b-70296cd8df9d
SharedType : SPO Sharing Link
Document  : Migration Datasheet.pdf
Action    : SharingSet

TimeStamp : 2018-04-16T16:21:43
User      : tony.redmond@office365itpros.com
SharedType : SPO Sharing Link
Document  : Migration Datasheet.pdf
Action    : SecureLinkCreated
Event     : <Type>View</Type>
Event     : <Permissions granted>Read</Permissions granted>

TimeStamp : 2018-04-16T16:29:45
User      : urn:spo:guest#james.pierrott@yandex.com
SharedType : SPO Sharing Link
Document  : Migration Datasheet.pdf
Action    : SecureLinkUsed
```

Between the records for file accessed and file sharing, you can create a complete picture of what documents are accessed by external users.

## Guest User Access to Documents

Our next example shows how to search the audit log for File Accessed events for guest users. The idea is to understand the files in document libraries guest users have opened. We could expand the search to include File Modified events if you want to know what files guest users update, but this search is enough to prove the point. After finding the data using the [script available from GitHub](#), here's how to sort it to discover the set of documents accessed by each guest.

```
[PS] C:\> $Report | Sort User, Document | Format-Table TimeStamp, User, Document -AutoSize
```

TimeStamp	User	Document
21 Dec 2018 11:39	john_contoso.com#ext#@Tenant.onmicrosoft.com	Summit 2018.one
5 Jan 2019 11:12	mary_contoso.com#ext#@tenant.onmicrosoft.com	Summit 2016.one
6 Jan 2019 09:44	mary_contoso.com#ext#@tenant.onmicrosoft.com	Updates.docx
9 Jan 2019 08:15	terry_contoso.com#ext#@tenant.onmicrosoft.com	Amazon Blurp.docx

We can view the audit data in different ways. For example, to find out how many accesses occurred to each document, you could do this:

```
[PS] C:\> $GroupData = $Report | Group-Object -Property Document
$GroupData | Sort Count -Descending | Select Name, Count
```

Name	Count
Ch 5 - SharePoint Online and OneDrive For Business - Final.docx	17
Project Ambush.docx	11
Interesting data.docx	10
Financial Arrangements 2019.docx	9
Ch 8 - Clients.docx	8
Ch 3 - Basic Workloads.docx	6
Budget 2019.docx	1

## Reporting the Assignment of Retention Labels

The same technique of grouping and sorting data from audit records is useful to report other aspects of Microsoft 365. If you have Office 365 E5 or Microsoft 365 Compliance E5 licenses, you can use the Activity Explorer (see Chapter 17) to see what's happening with retention and sensitivity labels. PowerShell gives you another way to do the job. For instance, to know what the most popular retention label applied to documents is, search for *TagApplied* events, extract the audit data, and put the events into the *\$Report* output. Because audit events often hold different information, this process is slightly complicated by the need to accommodate labels applied as a default for a library and those assigned manually by users. An example script can be [downloaded from GitHub](#). The script generates a list in the *\$Report* variable, and a little grouping and sorting reveals what the most popular retention label is:

```
[PS] C:\> $GroupData = $Report | ? {$_.Type -eq "File"} | Group-Object -Property Label | Sort Count
-Desc | Select @{n="Retention Label"; e={$_.Name}}, Count
```

Retention Label	Count
eBook Content	222
Audit Material	34
Approved	25
Confidential	12
GDPR Personal Data	10
Commercially Sensitive	4

If you see a blank entry in the list, it is for audit events logged when a user removes a retention label from a document. Note that the audit log does not currently capture events for when documents receive retention labels when created in or uploaded to document libraries with a default retention label.

## Who Used the SendAs Permission?

Among the data ingested by the audit log are mailbox audit events from Exchange Online. In the past, you might have used the *Search-MailboxAuditLog* cmdlet to search this information. For example, a common question is who used the *SendAs* permission to send a message impersonating a shared or user mailbox. We can find out by running a command like this:

```
[PS] C:\> Search-MailboxAuditLog -Identity "Customer Services" -LogonTypes Delegate -StartDate
((Get-Date).AddDays(-90)) -EndDate ((Get-Date).AddDays(+1)) -ShowDetails | ? {$_.Operation -eq
"SendAs"} | Select LogonUserDisplayName, LastAccessed
```

LogonUserDisplayName	LastAccessed
James Ryan	2 Nov 2018 12:13:35
James Ryan	2 Nov 2018 11:57:33

You can still use the *Search-MailboxAuditLog* cmdlet to search for mailbox events, but as those events are in the audit log it is better to search there, especially as you can look for events across multiple mailboxes

because the *Search-UnifiedAuditLog* cmdlet is designed to retrieve audit events for multiple users and multiple actions across the entire tenant.

Unlike Exchange on-premises, where all *SendAs* events are generated by delegates sending messages for another mailbox, Exchange Online processes messages sent by other workloads that can show up as *SendAs* actions in audit searches and skew results unless adjustments are applied. Here are some things to consider:

- Audit records with S-1-5-18 captured in the *UserId* property record the generation of a welcome message for a new team.
- Audit records are generated when Teams sends a welcome message.
- Audit records are generated for the group mailbox when a member posts a message to a conversation in an Outlook group using OWA. Records are not generated when messages are posted with other clients or arrive from guest members.
- Audit records are generated for the group mailbox when someone updates a task in Planner.

With these caveats in mind, the script to search for *SendAs* records, process the audit records, and identify events belonging to user or shared mailboxes and those belonging to group mailboxes can be [downloaded from GitHub](#). The former category is normally what people are concerned with because they're looking for instances where someone sent a message from a mailbox rather than posting to an Outlook group or adding a comment with Planner. As you can see in the script, to distinguish between the two categories, we create a hash table of primary email addresses for mailboxes and groups and look up that table for each audit event to decide if it belongs to a user/shared mailbox or a group mailbox.

## Searching for Audit Records Tracking Actions Against an Object

Most auditing attempts to answer "who did what" questions. In other words, you want to know who performed a specific action. Sometimes you need to know what happened to a particular object, like a document or a user. Finding audit events for one or more documents is easy – all you need to do is pass the document names in the *ObjectIds* parameter. In this example, we create an array of document names to search for and then pass the array as the *ObjectIds* parameter for the call to *Search-UnifiedAuditLog*:

```
[PS] C:\> [array]$docs = "New Signature API for Email Signatures.docx", "Controlling default
creation of online meetings with OWA.docx", "Anticipating Microsoft Ignite 2020.docx"
[PS] C:\> [array]$Records = Search-UnifiedAuditLog -ObjectIds $docs -StartDate (Get-Date).AddDays(-
90) -EndDate (Get-Date).AddDays(+1) -ResultSize 500
```

The events found are for all actions performed against the documents, such as being modified or downloaded. The same technique works for users.

```
[PS] C:\> [array]$Users = "Oisin.Johnston@office365itpros.com", "Kim.Akers@office365itpros.com"
[array]$Records = Search-UnifiedAuditLog -ObjectIds $Users -StartDate (Get-Date).AddDays(-90) -
EndDate (Get-Date).AddDays(+1)
```

To emphasize, this search returns events for actions performed for these users (like being added to a group membership) rather than events performed by the users.

Microsoft 365 Groups are not users, so if we want to find the actions performed against a group, we must use the *FreeText* parameter to search audit records for instances of unique values that identify the group we're interested in. Fortunately, the object identifier for a group is a good search term. In this example, we extract the object identifier for a Microsoft 365 group and use it to search for audit events. We then group the audit events to get an overview of the kind of activity performed against our target:

```
[PS] C:\> $ObjectId = Get-UnifiedGroup -Identity "Office 365 for IT Pros" | Select -ExpandProperty
ExternalDirectoryObjectId
[array]$Records = Search-UnifiedAuditLog -FreeText $ObjectId -StartDate (Get-Date).AddDays(-90) -
EndDate (Get-Date) -ResultSize 1000
$Records | Group Operations | Sort Count -Descending | Format-Table Name, Count
```

Name	Count
RecipientChange	17
TabUpdated	10
TabAdded	4
Remove member from group.	3
MemberRemoved	3
Add member to group.	3
Update group.	2
MemberAdded	2
TabRemoved	1
Set-UnifiedGroup	1

The technique also works for finding audit records for security groups (but not for distribution lists). It also works for Azure AD accounts, including guest users, but it's much slower than using the *ObjectIds* parameter. As the name implies, *FreeText* means that a free text search is used to find matching audit events. In a large tenant, a free text search across potentially millions of records won't be fast.

Remember that a single action can result in multiple events. For instance, if you add someone to a group, the *MemberAdded* and *Add member to group* events are captured by different workloads and ingested into the audit log. The duplication is easily detected by comparing the creation date for the events.

## Premium Auditing and Crucial Events

The Microsoft Purview Audit (Premium) solution is available to users with E5 or Microsoft 365 E5 licenses or the Microsoft 365 E5 compliance add-on. The solution covers:

- **Longer retention of audit data.** Instead of purging audit events after 90 days, the retention period for events for Exchange Online, SharePoint Online, and Azure AD is 365 days. The [Microsoft Purview Audit \(Premium\) solution](#) can keep audit data for up to 10 years with an additional per-user license.
- **Audit log retention policies.** Because tenants might not want to hold some audit events for 365 days, they can apply audit log retention policies to remove selected events from the audit log.
- **Crucial or high-value audit events.** These are audit events designed to expose in-depth information for forensic investigation. You don't need to enable the collection of crucial audit events as this happens automatically for all accounts with the appropriate licenses.

The current set of crucial events are:

- **MailItemsAccessed:** Items in a mailbox are accessed (opened) or synchronized.
- **Send:** A message is sent from a mailbox.
- **SearchQueryInitiatedExchange:** A mailbox search is initiated using Outlook desktop or OWA.
- **SearchQueryInitiatedSharePoint:** A search is initiated for a SharePoint site.
- **Other workload events from Forms, Stream, Teams, and Yammer** (see [this link](#)).

Examples of Teams crucial events are:

- **MeetingDetail:** A Teams meeting occurs.
- **MeetingParticipantDetail:** A participant joins a Teams meeting. See [this article](#)

for information about how to find and analyze Teams meetings audit events.

The relevant workloads automatically capture the crucial mailbox events once an account has the necessary license. To capture the search events in Exchange Online, you must update the owner audit configuration for individual mailboxes. For example:

```
[PS] C:\> Set-Mailbox -Identity Kim.Akers -AuditOwner @{Add="SearchQueryInitiated"}
```

## The MailItemsAccessed Event

Exchange Online generates the *MailItemsAccessed* event when licensed users access items in Exchange Online mailboxes using any connectivity protocol from any client. The events are uploaded to the audit log along with other Exchange events. Two kinds of *MailItemsAccessed* events are available: *Sync* (synchronization) and *Bind* (access to a message).

- **Sync events** occur when an Outlook desktop client synchronizes messages from the mailbox to its local cache (the OST for Windows or OLM for Mac). During synchronization, the client downloads copies of all new or changed items from the mailbox and notes any deletions from the server to apply in its local copy. The sync event records that synchronization occurred for a folder. Over a working day, many events record synchronization for folders like Inbox, Sent Items, Deleted Items, and the Calendar, plus any other folder where messages are moved to. In terms of investigation, the assumption is that if a breach occurs and someone can use Outlook to synchronize a folder to a local cache, potentially all items in that folder are copied.
- **Bind events** record access to an individual message. To reduce the number of audit records, Exchange Online generates a single bind event covering access to messages within two minutes. Thus, a bind event might cover access to a single message or ten messages.

To give an idea of the volume of events a user might produce, over two days, 443 *MailItemsAccessed* events were generated for a single mailbox. The majority were Bind events (394), and of these 312 were in the Inbox and 32 in Sent Items. Twenty-five events were captured for the Outbox, which is a transient folder where items exist while awaiting processing by the Exchange transport service. It's also used when Outlook posts to a conversation in an Outlook group, which is why the events for the Outbox folder were captured. The number of events captured by other mailboxes will vary depending on:

- The amount of inbound traffic. For instance, executive mailboxes often have a higher level of traffic. Organizations that use Teams for internal collaboration might find that the level of traffic is lower.
- The actions taken by the mailbox owner (or delegates).
- The number of clients signed into the mailbox.

In the case of very active mailboxes, if more than a thousand events are generated for a mailbox in less than 24 hours, Exchange Online stops generating *MailItemsAccessed* events for that mailbox for 24 hours. Microsoft says that less than 1% of Exchange Online mailboxes are throttled.

## Forensic Investigations of Mailbox Breaches

[Microsoft's documentation](#) lays out the steps that forensic investigators can use to interpret the information captured in *MailItemsAccessed* events if they suspect that an attacker gains access to a mailbox. The basic idea is to:

- Identify when a user's mailbox might have been compromised.
- Discover what folders an attacker might have accessed in the mailbox (Sync events). It is possible that the attacker downloaded (synchronized) the entire mailbox to Outlook.
- Discover the individual messages accessed by the attacker (Bind events). The events store the internet message identifier for messages, which can be used to find the messages in the mailbox.
- Investigate if any messages were sent from the mailbox by the attacker. *Send* events record each message sent. The message subject and its internet message identifier are recorded. Investigators can use this information to discover the recipients of the messages and decide if any information has been sent to an unauthorized address.
- If the account is compromised, the attacker might have performed searches to look for confidential or sensitive information. The *SearchQueryInitiatedExchange* and *SearchQueryInitiatedSharePoint* events record details of searches performed against the mailbox and SharePoint sites. Any evidence of

unusual search activity deserves further investigation to identify if the searches are performed by the account owner or an attacker.

Given the number of audit events that an investigator might have to examine in the steps described above, it's not realistic to use the audit log search GUI. Instead, most investigators use PowerShell to find and analyze audit records to extract relevant and useful information. An example script to parse *MailItemsAccessed* events, including using the *Get-MessageTrace* cmdlet to find the subject of messages (within the last 10 days), can be [downloaded from GitHub](#). Examples of how to parse audit records for the send and search events [are also available](#).

## Audit Log Retention Policies

Microsoft Purview's premium audit functionality includes the ability for tenants to configure audit log retention policies. You define a [retention policy for selected audit events](#) with a set retention period and those items will be purged after that period. A tenant supports up to 50 audit log retention policies.

Management of audit log retention policies is through the [Audit section of the Microsoft Purview Compliance portal](#) or PowerShell (after connecting to the compliance endpoint). This example runs the [New-UnifiedAuditLogRetentionPolicy](#) cmdlet to create an audit retention policy to remove any *SearchQueryPerformed* event executed by the background *app@sharepoint* process after three months instead of the twelve-month retention of audit events if the tenant has E5 licenses.

```
[PS] C:\> New-UnifiedAuditLogRetentionPolicy -Name "90-day Retention SearchQueryPerformed by app@sharepoint" -Description "Remove SearchQueryPerformed events from the app@sharepoint process after 90 days" -RecordTypes SharePoint -Operations SearchQueryPerformed -UserIds "app@sharepoint" -RetentionDuration ThreeMonths -Priority 8
```

### Purging the Audit Log

You can choose to apply retention for any of the events captured in the audit log and keep them for three, six, nine, twelve months, or 10 years. It's a good idea for tenants who either want precise control over the retention of audit data or want to clean up events that don't add much value in terms of investigations. SharePoint is a notoriously "chatty" application when it comes to the capture of audit events, so I can see why tenants might decide to keep important events like *FileUploaded* or *FileAccessed* for as long as possible while removing some of the chatter after 90 days.

## Activity Alerts and Alert Policies

Searching the audit log to find items of interest rapidly becomes boring. It also creates the potential that you might overlook important audit events. Microsoft 365 has two methods to automate checking of user activity within a tenant and alert administrators when something out-of-the-ordinary occurs.

- **Activity alerts** are available to all business tenants. An activity alert checks events recorded in the audit log for specified conditions defined by administrators and fires when those conditions occur.
- **Alert policies** build on the concept of activity alerts and apply extra intelligence to the events recorded in the audit log. Instead of firing when events occur, policies look for patterns of events such as a certain number of file downloads over a brief period. Microsoft includes a set of default alert policies to help tenants understand their use and handle common conditions, including notification of malware attacks or instances when someone gains administrative permissions for Exchange Online. [Several predefined alert policies](#) are available to Office 365 E1 and E3 tenants while the other policies are available to tenants with Office 365 E5, Advanced Threat Protection, or Microsoft 365 E5 Compliance licenses. Tenants can create additional alert policies to meet specific needs or included with apps. For instance, communication compliance policies create alert policies to advise

administrators when a threshold of *SupervisionRuleMatch* events match the conditions monitored to detect possible policy violations occur in a set period.

In both cases, administrators receive notifications via email, and it is then up to the people notified to act to resolve the detected problem. Only accounts that hold the Organization Configuration compliance role can create activity alerts or alert policies.

## Activity Alert and Alert Policies

Experienced administrators know when something is not quite right. At least, their suspicions heighten when they see certain things happening. [Activity alerts](#) help administrators keep up to date with what is happening inside the tenant. Microsoft Purview packages activity alerts into policies to create the mechanism to inform nominated individuals when a certain event occurs. Alert policies try to capture the instinct of experienced administrators to detect patterns of problematic activity, using the ability of software to keep looking for matching patterns repeatedly. When a match occurs between real-time activity as captured in the audit log and the settings defined in a policy, Purview triggers an incident and notifies the recipients defined in the policy. The recipients are then responsible for resolving the incident. Conceptually, you can break down alert policies into four stages:

- **Understand** the normal ebb and flow of user activity to know what you expect to happen within the tenant. This is the activity baseline and can be set manually through your observations and experience or automatically by Microsoft.
- **Define** the characteristics of activity that cause concern and watch for incidents when those characteristics occur. These characteristics are the activity, condition, and threshold checked by an alert policy.
- **Monitor** the events recorded in the audit log against the conditions defined in alert policies.
- **Alert** administrators through email notifications when policy violations occur.

One thing to remember is that Alert policies don't support filtering based on a file or folder name, so you cannot use them to check changes made to a specific item. For instance, if your search is for file check-ins for a document called "Budget" and you use the search to create an alert, Purview generates alerts for all file check-in operations and not just for that document. In addition, activity alerts do not support date ranges, so the alert will fire for any matching activity from when you create it. An alert policy can check for multiple actions across multiple accounts. You can select any of the events logged in the audit log for monitoring. For example, you could have an alert that fires when someone checks a file into a document library or uses the *Send As* permission to send email on behalf of another user.

Purview can trigger alerts for every instance of a certain activity, such as when an administrator grants elevated permissions to another user. It can also trigger alerts based on event aggregation. For instance, users download 50 files from a SharePoint library within 30 minutes. That might be evidence of a hard-working user. On the other hand, it might be a sign that someone is grabbing some valuable intellectual property that they plan to take with them to another job. A more complex form of aggregation is when the threshold for a trigger is set by analyzing up to a week's worth of activities to understand the normal level of activities within the tenant. If something then happens that greatly exceeds the expected norm, the recipients defined in the policy receive notifications.

Purview monitors the stream of audit data flowing from workloads to detect events matching those defined in alert policies. If Purview detects a match, it sends an email notification to the accounts registered for the alert.

Purview creates a set of default alert policies for tenants, including those highlighted in Table 21-1 (see [this page](#) for the most current list and the licensing requirements for some advanced policies).

<b>Alert</b>	<b>Conditions</b>	<b>Severity</b>
<i>Creation of forwarding/redirect rule</i>	Fires when a user creates a rule to redirect or forward their email inside or outside the tenant.	<b>Low</b>
<i>Elevation of Exchange Admin privilege</i>	Fires when an administrator assigns a user account administrative permission for Exchange (for example, adding an account to the Organization Management role group).	<b>Low</b>
<i>Messages have been delayed</i>	Fires when the number of messages waiting to be delivered outside the tenant exceeds 2,000 and the level exists on the queue for an hour.	<b>High</b>
<i>Malware campaign detected after delivery</i>	Fires when an unusually (compared to baseline) number of messages with malware arrive in user mailboxes., Apart from flagging an alert, the problematic messages are removed. [E5/Defender for Office 365]	<b>High</b>
<i>Malware campaign detected and blocked</i>	Fires when Defender detects an attempt by attackers to send an unusual volume of messages with known malware to your tenant. The malware is blocked, and messages do not arrive in user mailboxes. [E5/Defender for Office 365]	<b>Low</b>
<i>Malware campaign detected in SharePoint and OneDrive</i>	Fires when an unusual number of malware detections occur in SharePoint Online and OneDrive for Business sites in the tenant. [E5/Defender for Office 365]	<b>High</b>
<i>Unusual increase in email reported as phish</i>	Fires when there is an increase in the normal volume of email reported by users as phish. [E5/Defender for Office 365]	<b>High</b>
<i>Unusual external user file activity</i>	Fires when external users with access to SharePoint or OneDrive for Business sites in the tenant perform an unusual number of activities like accessing, downloading, and removing files. [E5/Defender for Office 365/AC]	<b>Medium</b>
<i>Unusual volume of external file sharing</i>	Fires when users share an unusual number of files in SharePoint or OneDrive for Business sites with external people outside the tenant. [E5/Defender for Office 365/AC]	<b>Medium</b>
<i>Unusual volume of file deletion</i>	Fires when users remove an unusual number of files from SharePoint or OneDrive for Business sites within a brief period [E5/Defender for Office 365/AC]	<b>Medium</b>

Table 21-1: Default alert policies

You can amend the recipients for alerts generated by the default alert policies, but you cannot change the other settings. The first three policies are available to all enterprise tenants. Access to other policies is through Office 365 E5 (E5), Defender for Office 365, or Microsoft 365 E5 Compliance (AC) licenses. Microsoft introduces new alert policies and updates existing policies as the need arises, including those used for mail flow insights like the alert policy for accepted domains due to expire soon.

The default alert policies cover some generic situations that might or might not apply to your tenant, which is why you can define custom policies. You can only create custom alert policies if your account holds the Manage Alerts role (included in a compliance role group like Organization management). To create a new alert policy, go to the **Policies** section of the Microsoft Purview compliance portal and select **Alert policies**, and then **New alert policy**. You can then add the following information to create the new alert:

- **Name and Description.** These settings are for administrative convenience, and you can enter anything you like. The name appears in dashboards, so it should convey the intent of the alert policy. The description is useful to note who created or last amended the policy and what the policy does.



- **Category:** To help track alerts, you can classify policies into the following categories, which you can use to sort alerts in the **View alerts** page:
  - Information governance. For example, users download more than a certain number of files to their PC over a defined period.
  - Permissions. For example, the elevation of permissions.
  - Mail Flow. For example, Exchange Online Protection detects malware in an email.
  - Threat management. For example, malware outbreaks.
  - Others. This category exists for tenants to use as they wish.
- **Severity:** You can assign alerts detected by the policy to be *Low*, *Medium*, or *High*. The more destructive a condition is, the higher its severity should be.
- **Activity:** The activity that the policy tracks. Alert policies do not yet cover all the activities recorded in the audit log because the intention is that alert policies can deliver near real-time notifications about problems and not all workloads feed events into the audit log that quickly. Microsoft will probably extend coverage across workloads over time. For now, policies can cover most activities in the SharePoint, OneDrive, and Exchange workloads. You can only select a single activity per policy. If activity policies do not support an event you want to check, you can create an activity alert to do the job.
- **Conditions:** For most activities, you can define conditions that must exist before Purview signals an alert. For example, a user downloads a file to a computer with a specific IP address. It is also possible to configure alerts to fire every time a user performs the activity or when users download files from a specific site.
- **Threshold:** How often an activity must occur within a period before a problem condition exists (thresholds aren't available for all types of alerts). The threshold can also be an unusual activity, which is when Purview compares activity for a period against the baseline for the tenant. Setting a threshold too low usually results in many notifications, which can hide real problems (the lowest threshold is 3 activities). Setting the threshold too high means that Purview might not send notifications in situations when administrators need to act. It is easy to tune a threshold after noting how many notifications administrators must process to reach a point where notifications arrive when real problems exist.
- **Email notifications:** You can define that alert notifications go to a list of recipients. Notifications can go to accounts within or outside the tenant. The account that creates a new policy automatically receives notifications while the default policies send notifications to the tenant administrators. You can also set a daily limit for an alert policy so that a policy can only ever generate a certain number of notifications in a single day.

The new policy wizard steps through these sections to create a new alert policy. Figure 21-5 shows how to set the activity and threshold for an alert policy.

Figure 21-5: Defining the conditions for an alert policy

**OneDrive Synchronizations and Alerts:** If you implement alert policies for files downloaded from SharePoint or OneDrive for Business, you should set the threshold for the volume of matches to be higher than the largest number of files in a library synchronized with the OneDrive sync client. If you set a lower number, you will receive alerts each time the OneDrive sync client refreshes its copy of the library.

## Handling Alerts

When Purview detects a pattern of events matching the threshold and conditions set in an alert policy, it triggers an alert and generates an email to the recipient list defined in the policy. The notification tells the recipient:

- The severity level for the alert.
- The policy that triggered the alert.
- The date and time of the alert in UTC.
- The activity that triggered the alert.
- The details of the alert. For example, how many matched activities occurred in what timeframe.

The alert message has a **View incident page** link to open a page with the details of the Defender incident Purview automatically creates for the alert. After viewing the incident, you can manage the alert to decide if the alert is a real incident by recording a status for the alert (Active, In progress, or Resolved) together with any comments to justify the status. You can also classify the incident to indicate if it is expected behavior or a false or true positive. This feedback helps to refine the automatic detection of similar incidents in the future.

Figure 21-6 shows details of an alert generated when a user downloaded a large number of files from a SharePoint Online document library.

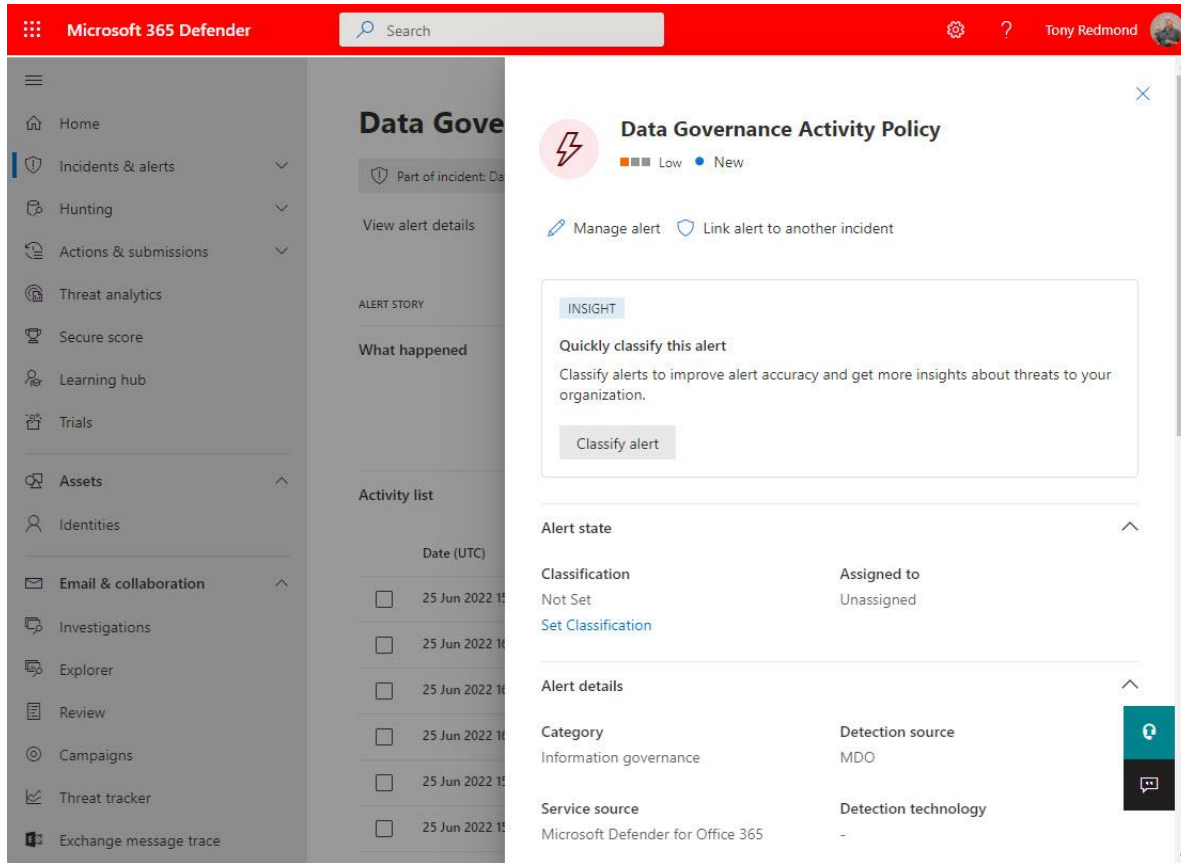


Figure 21-6: Investigating details of an alert

A busy tenant might see many alerts daily. Filters are available to focus on alerts belonging to a certain category or severity level.

## Activity Policies and PowerShell

Because of their complex nature, it is usually best to manage alert policies through the Compliance portal. However, a set of cmdlets are available:

- *Get-ProtectionAlert*: Reports the alert policies defined by the tenant.
- *Set-ProtectionAlert*: Amends an alert policy.
- *New-ProtectionAlert*: Creates an alert policy.
- *Remove-ProtectionAlert*: Removes an alert policy. You cannot remove one of the default alert policies.

For example, to list the set of default alert policies in a tenant, connect to the compliance endpoint with PowerShell and run the command shown below (output edited for space).

```
PS C:\> Get-ProtectionAlert | ? {$_.IsSystemRule -eq $True} | Format-Table Name, Category, AggregationType, AlertScenario
```

Name	Category	AggregationType	AlertScenario
Malware campaign detected after delivery	ThreatManagement	AnomalousAggregation	Protection
Unusual increase in email reported as phish	ThreatManagement	AnomalousAggregation	Activity
Unusual volume of external file sharing	DataGovernance	AnomalousAggregation	Activity
Elevation of Exchange admin privilege	AccessGovernance	None	Activity
Creation of forwarding/redirect rule	ThreatManagement	None	Activity

The *AggregationType* property tells us whether alerts trigger for every occurrence (*None*), based on the volume of the activity in a time window (*SimpleAggregation*), or when the volume of activity greatly exceeds the baseline for the tenant (*AnomalousAggregation*).

## Searching Alerts with PowerShell

Alerts are also recorded in the audit log and can be searched for using the *Search-UnifiedAuditLog* cmdlet. A [script available in GitHub](#) illustrates how to use the technique previously explained to interrogate the audit log for alert records and extract their content. As usual, the contents of the *AuditData* property of the audit records vary from activity to activity. Although this information is useful, the events do not contain details of the status of an alert and interpreting the *AuditData* property can be challenging. Better data is available through the Security/Alerts endpoint in the Graph. An example script showing how to extract current alerts from the Graph and post them to Teams is [available in GitHub](#).

# Office 365 Cloud App Security

Office 365 Cloud App Security (previously called Advanced Security Management) is part of the E5 plan and is also available as an add-on for the other enterprise plans. Unless you use groups to scope the coverage of policies (see the section later), every user in the tenant must have a license for OCAS as it is not possible to exclude the audit data for individual users from the anomaly detection and analysis.

Office 365 Cloud App Security (OCAS) is accessed through the **Manage Advanced Alerts** link in the Policies and Rules section of the [Microsoft 365 Security portal](#). This connects to a special version of [Microsoft Defender for Cloud Apps](#) designed for Office 365. Other [Microsoft 365 licenses](#) are available to extend the coverage and capabilities of Microsoft Defender for Cloud Apps. OCAS keeps up to six months of data about user activities, or twice as much security data as is normally collected for an Office 365 E3 tenant. Data collection begins as soon you enable OCAS for a tenant.

When a tenant opts to use OCAS, Microsoft connects your tenant to an equivalent tenant within the OCAS infrastructure. The link allows the OCAS analytics to extract and analyze the tenant's audit data, which detects suspicious activity and other potential problems. It takes about a week after a tenant is enabled before a satisfactory model exists of its normal activity and builds a baseline to measure suspected anomalies against.

OCAS is part of a long-term plan to give customers much better oversight about what is happening in their tenant based on the data accumulated in the audit log, with the major advantage of the approach being that no need exists to deploy agents or other software to support the gathering and analysis of the data to detect the threats that might lie in the anomalies that are picked up. Analyzing the audit data also reveals how the actions taken by individual users might compromise the security of the organization through suspicious behavior, such as someone downloading all the documents from a library holding confidential information within a brief period. Other indications are considered, such as suspicious IP addresses that might originate from anonymous proxies or known botnets.

OCAS allows administrators to create tenant-specific policies to fire alerts when specific events happen or when a specific pattern of actions occurs. For instance, you could create a policy that will alert administrators by email or SMS whenever certain conditions occur. Microsoft includes many preconfigured anomaly detection policies to get the ball rolling. These policies cover common conditions that should cause suspicion, such as a user logging in from two places widely separated by distance within a brief period. You can add other anomaly detection policies to highlight specific activities that are of concern to the organization. For example, you could create a policy to look for attempted logins from IP addresses outside the corporate IP range. You can tailor policies to turn off or on different risk factors or to increase sensitivity to a risk.

In some respects, apart from the analytics used by OCAS to pick up suspicious activity by correlating events, the technology is not rocket science. You could argue that a skilled administrator who knows what is happening in their tenant is likely to be able to detect and resolve the same kind of issues highlighted by OCAS. However, an application like OCAS scores through its ability to handle massive quantities of

information of the type generated by audit events and to reduce the mass down to what is important. A human can do this too, but will struggle with:

- The volume of data to process (especially as the environment scales).
- The time needed to recognize complex suspicious audit events and to learn the characteristics that mark new threats.
- The need to be consistent in the treatment of events.

It is also likely that the human administrator will forget that some events have happened (or not) in the past, so when something happens, they must consider the event on its merits. Computers are better at remembering things, so OCAS quickly recognizes when an event is rare (and therefore potentially out of the norm) or normal.

In addition, the machine learning that lies behind analytics is much faster at correlating events to detect suspicious activity. Once software learns what it should be looking for, it generally produces more consistent results than a human can, 24 hours a day, 365 days a year, which is why applying technology to automate the collection and validation of information drawn from multiple sources is a good approach to understanding the kind of threat introduced by how individuals behave.

## Alerts

Figure 21-7 shows how alerts appear in the OCAS console. In this case, a set of alerts have been signaled because users have connected to Teams from countries that are not their normal location. Connections are noted from different countries for one account, pointing to someone on a road trip, while another user is flagged from Bulgaria as an infrequent connection. This could be because the user has not connected from this location for a while.

The screenshot displays the OCAS Alerts console. At the top, there's a search bar and navigation icons. The main area is titled 'Alerts' and includes filter options for Status (OPEN, CLOSED), Category, Severity, App, and User name. A table of alerts is shown below, with columns for Alert, Status, and Severity. A dropdown menu is open over the 'App' filter, listing various Microsoft applications. The alert list includes:

Alert	Status	Severity	Time
Impossible travel activity Impossible travel - Chris Bishop	OPEN	High	3/28/21, 6:43 PM
Unusual addition of credentials to an OAuth app (PREVIEW) Unusual addition of credent... - Office 365 - Tony Redmond - MailSendAppDelegat...	OPEN	High	3/25/21, 7:06 AM
Unusual addition of credentials to an OAuth app (PREVIEW) Unusual addition of credent... - Office 365 - Tony Redmond - MailSendApp	OPEN	High	3/24/21, 6:26 AM
System alert: Deprecation of Label Management in the Azure Portal Microsoft Cloud App Security	OPEN	High	3/25/21, 7:06 AM
System alert: Deprecation of Label Management in the Azure Portal Microsoft Cloud App Security	OPEN	High	3/24/21, 6:26 AM

Figure 21-7: Reviewing OCAS alerts

Figure 21-8 shows an alert caused by "impossible travel" for a user. In other words, the IP addresses captured by OCAS for client connections over a certain period originate in multiple countries where it would be impossible for the user to travel between those countries during that time. In this case, the alert flagged

interactions from Ireland and the Netherlands with the 99-minute threshold deemed technically possible to fly between Ireland and the Netherlands. While the problem seems real, it deserves some further examination to figure out if an attacker compromised the user account. The facts are:

- The user signed in from two different IP addresses within a short period.
- The IP addresses indicate connections from Ireland and the Netherlands.
- In both cases, the application was Teams.

The IP addresses were correct, the connections valid, and it looks like a real problem at first glance. Then you realize that:

- Teams uses access tokens to authenticate with different services.
- The tenant is in the EMEA data center region, and the Teams service runs in the region.
- The EMEA data center region includes data centers in Ireland and the Netherlands.

Therefore, the most likely explanation is that the Teams client attempted to use its access token to authenticate. During this process, the original server in The Netherlands data center handled the request to a server in the Dublin data center, which processed the Azure AD secure token login. Azure AD captured details of the connections and sent them to the Office 365 audit log where OCAS picked up the information, analyzed the events, and concluded that evidence existed to point to a potential impossible travel situation. As it happens, I know that this is exactly what transpired, but it's a great example of how tenant administrators need to apply their knowledge about applications and how Microsoft's data center infrastructure operates to assess and resolve a flagged alert.

The screenshot displays the Microsoft Cloud App Security interface. At the top, the title is 'Alerts > Impossible travel activity' with a timestamp of '5/11/21 11:43 PM' and a severity level of 'MEDIUM SEVERITY'. Below the title, there are filters for 'Impossible travel', '2 Services', 'Chris Bishop', '2 IP addresses', and '2 Countries'. A 'Resolution options' dropdown is set to 'Chris Bishop', and a 'Close alert' button is visible. The 'Description' section states: 'The user Chris Bishop (chris.bishop@office365itpros.com) performed an impossible travel activity. The user was active from 51.171.212.129 in Ireland and 84.81.253.164 in Netherlands within 99 minutes. If these are IP addresses that are known and safe, add them in the IP address range page to improve the accuracy of the alerts.' The 'Important information' section lists: 'This user is an administrator in Office 365 (Default)', 'Microsoft Teams (Default) was accessed from Netherlands for the first time in 180 days.', and 'This alert falls under the following MITRE tactic: Initial Access'. The 'Activity log' section shows a table with 10 of 131 activities.

Activity	User	App	IP address	Location	Device	Date
Log on	Chris Bishop	Microsoft...	51.171.212.129	Ireland	Windows	May 11, 2021, ...
Log on	Chris Bishop	Microsoft...	51.171.212.129	Ireland	Windows	May 11, 2021, ...
Log on	Chris Bishop	Microsoft...	51.171.212.129	Ireland	Windows	May 11, 2021, ...

Figure 21-8: Investigating potentially impossible travel

## Resolving Alerts

Cloud App Security assigns alerts a severity of high, medium, or low risk. The risk level is calculated using behavioral analytics to compare normal user interaction against audit data. The analytics are based on

Microsoft's collected knowledge about the threats that exist, and their origin gathered from across Microsoft 365 and other cloud services. Assigning a risk value allows an administrator to filter for high-risk alerts and prioritize their resolution.

Another example of an alert is when an account is detected to have elevated permissions (a "New admin user" alert). Again, if the permissions were assigned purposely, the alert can be resolved, and OCAS knows that it does not have to signal the issue again. However, it could be the case that someone has been assigned permissions in error or that they hold permissions for too long, in which case the resolution is different and might need the account to be suspended or to have its permissions adjusted. User accounts can also be suspended as an action contained in a policy to ensure that action is taken to protect the organization without needing an administrator to do something manually. Suspended users show up as blocked users. If this turns out to be the wrong thing to do, you can reverse the suspension from OCAS or the Microsoft 365 admin center.

An alert may highlight an event that is uninteresting or invalid. In these instances, you can dismiss the alert or mark it as a false positive. These actions are recorded in the Activity Log and the fact that the user's location or their admin status is valid is considered by OCAS when it processes audit events and other data to detect anomalies and suspicious activity in the future.

## Filtered Alerts

OCAS supports activity log filters to focus an investigation on one or more of the Microsoft 365 applications or selected users. The latter filter is valuable when you might be concerned about the activities of a certain individual. You can also search for high, medium, or low severity alerts or for closed alerts. You can also filter by risk category (for example, access control or privileged accounts). The filters can be combined to focus on certain actions, meaning that even a very large volume of alerts can be quickly refined to produce a set of alerts that need to be examined. You can also export alerts to a CSV file if needed.

## Activity Log

The screenshot shows the Microsoft Defender for Cloud Apps Activity Log interface. At the top, there is a search bar and a navigation menu. The main area is titled "Activity log" and includes a "Queries: Select a query" dropdown and a "Save as" button. Below this, there are filters for "App: Microsoft SharePoint Online", "User name: Vasil Michev (Technical Guru) (vasil.mic...)", "Raw IP address: Enter IP address", and "Activity type: FileModified". There is also a "Location: Select countries/regions" dropdown and an "Advanced filters" toggle. The table below shows 1 - 20 of 213 activities. The first activity is "Modify file: file https://..." by "Vasil Michev (Technical Guru)". Below the table, there are options for "SHOW SIMILAR", "General", "User", "IP address", and "Send us feedback...". The table has columns for "Activity", "User", "App", "IP address", "Loc...", "Device", and "Date".

Activity	User	App	IP address	Loc...	Device	Date
Modify file: file https://...	Vasil Michev (Technical Guru)	M...	46.249.81.1...	Bu...	PC Windows 10 10.0	Jun 26...

Figure 21-9: Browsing the OCAS Activity Log to review specific user activity

The activity log contains data from the Office 365 audit log along with other logged items, such as those recorded when an administrator resolves or dismisses an alert. Data is available for the previous 30 days. Many activities generate a large volume of audit log entries, so you'll probably need to use filters to reduce the set to a manageable amount. In the example shown in Figure 21-9, filters extract events relating to file modification activities in SharePoint Online for a certain user. Note the **Save as** option to create a new policy based on the search criteria.

## Policies

The ability to create customized policies to check events and trigger alerts when predetermined conditions occur is one of the most powerful features in OCAS. Using templates or from scratch, you can create various policies to check for various kinds of activity captured in events, including:

- Access policy.
- Activity policy.
- App discovery policy.
- Cloud discovery anomaly detection policy.
- File policy.
- OAuth app policy.
- Session policy.

These policies help administrators to master the vast quantity of events that busy tenants generate. For example, a file policy can check for events when users share documents holding sensitive information with people outside the tenant. You can enable some Data Loss Prevention (DLP) checking to look for specific forms of data, like credit card numbers, and take governance actions, like reporting the problem to a site owner, if the checks discover someone sharing a file when they should not. Administrators can see the results of the policy in the dashboard and opt to receive updates via email or SMS text messages.

It is important to emphasize that OCAS does not replace the DLP policies that you can deploy for workloads. Investigators often review alerts sometime after the fact rather than being actioned at once rather than being brought to the attention of users through visual clues embedded in clients like Outlook or applications like Word. Instead, checking audit events for problems gives the tenant an added layer of protection. The same is true for governance as actions to prevent users from making mistakes are usually better taken through classification and retention policies deeply integrated into individual workloads.

One issue for non-U.S. customers is that OCAS is based on an Azure data store running in a [U.S., U.K, or European data center](#). However, only audit data and information about tenant users and groups are moved to the Azure data store and personal information belonging to tenant users stays within the originating workloads. Microsoft plans to extend OCAS so that its data is stored in other data center regions in the future. When this happens, OCAS data for a tenant will reside in the same region as their other data.

**SIEM integration with OCAS:** You can integrate OCAS with third-party SIEM servers. A SIEM agent runs on the server to pull alerts from OCAS so that you can integrate its alerts alongside alerts generated from other parts of your IT infrastructure. Details of how to perform this integration [are available online](#).

## Scoping Policies to Groups

If you only want to use OCAS to check the activities of a limited set of people, you can tailor policies to limit them to defined groups. These are OCAS groups rather than Azure AD groups. If you want to use Azure AD groups, you must import them into OCAS through the **User groups** choice in the cogwheel menu. Once imported, you can edit OCAS policies and add a group as a filter. Because OCAS then scopes the alerts to just the users in the group, the licensing requirement for OCAS only extends to those users. Make sure that all policies are scoped to groups as otherwise you need to license every user in the tenant.



# Third-Party Auditing Alternatives

Microsoft encourages ISVs to use the [Management Activity API](#) to access audit data through the Microsoft Graph and develop solutions that generate in-depth activity reports, including suspicious or out-of-norm actions, data visualization and analysis to aid planning and oversight for tenants, and to incorporate audit activity in operational dashboards. In short, these solutions will help tenants understand who is accessing their content and whether their compliance framework is working. Quest On Demand Audit (Figure 21-10) is an example of a third-party solution that consumes audit data.

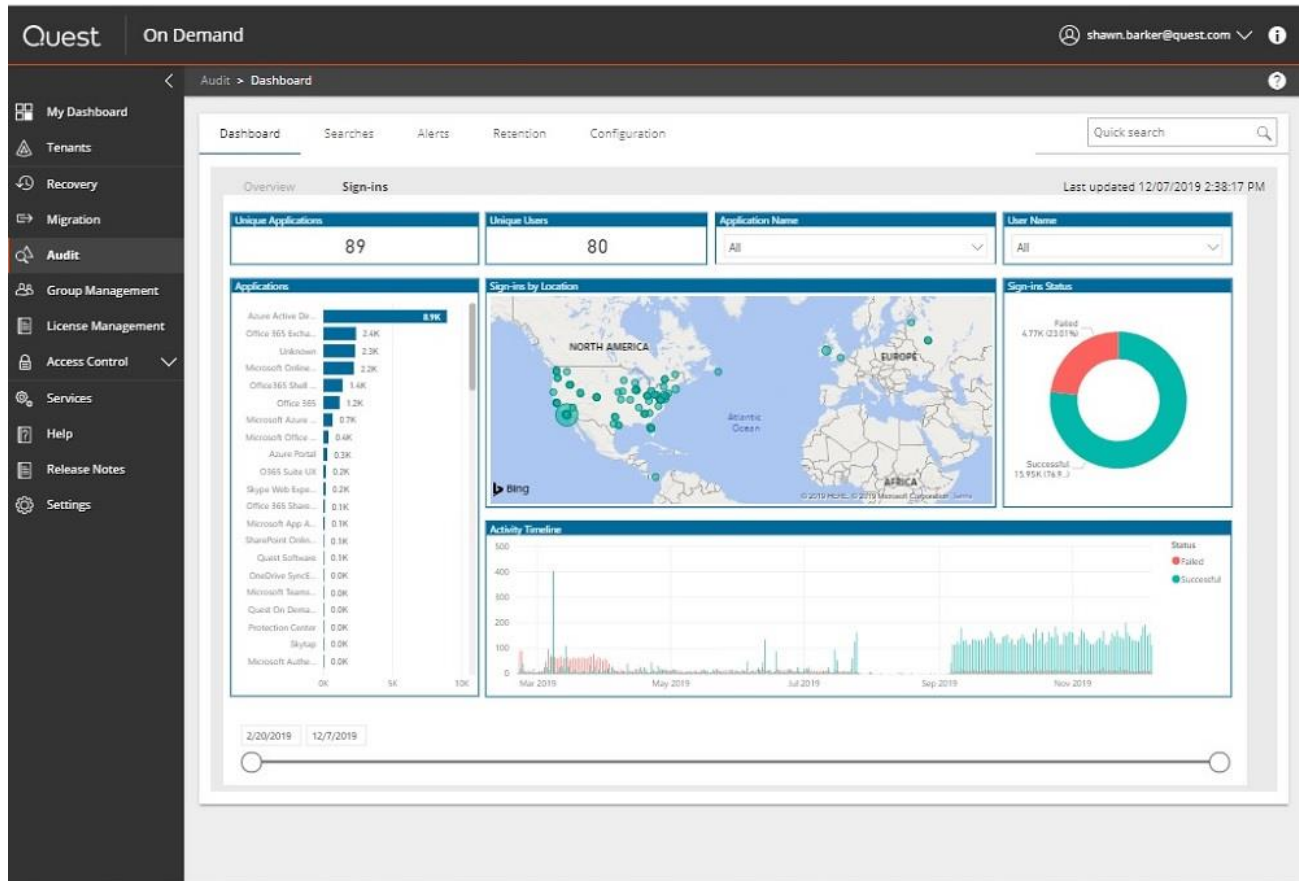


Figure 21-10: Quest On Demand Audit

Apart from taking a different approach to reporting, ISV products often allow tenants to store audit data for longer periods. This is a big advantage for large enterprises that often need to hold audit information for years to meet regulatory requirements. Third-party solutions do not include some of the high-end machine-learning functionality and policy-driven event checking found in OCAS, but they are usually much cheaper and are therefore an interesting choice to investigate if you think a tenant needs more audit reporting and analysis functionality than found in Microsoft 365.

## Exchange Online Administrative Auditing

Exchange administrative auditing is the mechanism used to track the operations performed by administrators when they invoke cmdlets to manage Exchange Online. Another way of putting this is that the audit entries allow the age-old question of “who did that?” to be answered. Administrative auditing is enabled by default for Exchange Online. However, you have little control over what is audited because Microsoft does not allow tenants to change most of the settings in the audit configuration.

Auditing is performed by the Admin Audit Log agent, which is active on every Exchange Online server. The agent evaluates cmdlets as they run against the audit configuration to decide whether the use of the cmdlet needs to be logged. If so, the agent creates an item holding details of the cmdlet and its parameters in the Inbox of the audit mailbox. The audit agent creates separate reports for each object if you execute an action that is performed against several objects. For example, if you use *Get-Mailbox* to fetch a list of mailboxes from a database and then use *Set-Mailbox* to place the mailboxes on litigation hold, the audit agent creates a separate audit event for each mailbox as it is updated.

You can also write custom entries into the audit log with the *Write-AdminAuditLog* cmdlet. This is intended to allow administrators to document actions performed in scripts. Up to 500 characters of text can be inserted in the comment parameter, which is captured in the *CmdletParameters* property of the audit entry:

```
[PS] C:\> Write-AdminAuditLog -Comment "NYC mailboxes placed on litigation hold"
```

## Exchange Online Administrative Audit Configuration

The *Get-AdminAuditLogConfig* cmdlet reveals the administrative auditing configuration used by Exchange Online. An edited version of its output is shown below:

```
[PS] C:\> Get-AdminAuditLogConfig

AdminAuditLogEnabled      : True
LogLevel                  : None
TestCmdletLoggingEnabled  : False
AdminAuditLogCmdlets      : {*}
AdminAuditLogParameters  : {*}
AdminAuditLogExcludedCmdlets : {}
AdminAuditLogAgeLimit     : 90.00:00:00
UnifiedAuditLogIngestionEnabled: True
UnifiedAuditLogFirstOptInDate : 27-Jul-2016 09:00:00
LoadBalancerCount        : 1
RefreshInterval          : 10
PartitionInfo             : {}
```

Apart from *AdminAuditLogEnabled*, which tells us that administrative auditing is in effect, the important settings in the list have the following meanings:

- **LogLevel:** None, meaning that optional information is not captured in audit entries. What is captured is data for *CmdletName* (the name of the cmdlet executed), *ObjectName* (the object that the cmdlet ran against), *CmdletParameters* (parameters passed to the cmdlet), *Caller* (the account used to run the cmdlet), *Succeeded* (whether the cmdlet was successful), and *RunDate* (the date and time when the cmdlet ran).
- **AdminAuditLogCmdlets:** \*, specifying what cmdlets should be audited. In this case, the asterisk means that any cmdlet that creates or changes the properties of an object is recorded. Note that Exchange Online creates far fewer audit entries because it can manage fewer objects. For example, databases and servers are not manageable. Audit entries are never captured for cmdlets that simply retrieve information, like *Get-Mailbox* or *Get-TransportRule*.
- **AdminAuditLogParameters:** \*, meaning that all parameters passed to the audited cmdlets are recorded. For instance, if you use the *Set-Mailbox* cmdlet to modify a setting on a mailbox, the parameter passed to the cmdlet is captured in the audit entry.
- **AdminAuditLogExcludedCmdlets:** Logically, as all cmdlets are recorded, none are excluded.
- **AdminAuditLogAgeLimit:** The default value shown is 90.00:00:00, meaning that audit entries should be kept for 90 days. Administrative events are stored in an Exchange Online arbitration mailbox that is inaccessible to administrators. You cannot change the retention period.

- **UnifiedAuditLogIngestionEnabled:** If True, audit events from workloads are ingested into the audit log. See the earlier section on enabling audit event ingestion for the audit log. The associated *UnifiedAuditLogFirstOptInDate* property records the date when audit log ingestion started.

## Accessing Administrative Audit Log Entries

Several methods are available to access Exchange Online administrative audit events:

- Exchange Online administrative audit log events flow to the Office 365 audit log and can be found in audit log searches performed through the Microsoft Purview Compliance portal or third-party products.
- Some of the reports available in the auditing tab of the compliance management section in the classic EAC, including the “admin audit log report,” use this data. Even more interesting is the “External admin audit log report,” which details the changes made to the tenant’s Exchange Online configuration by Microsoft data center administrators.
- The *Search-AdminAuditLog* and *New-AdminAuditLogSearch* cmdlets can be used to search administrative audit logs through PowerShell. See Chapter 8 in the Companion Volume for some examples of running Exchange audit searches with PowerShell.

In all cases, the preferred method for searching Exchange administrative audit events is to use the Compliance portal or the *Search-UnifiedAuditLog* cmdlet.

## Exchange Online Mailbox Auditing

While administrative auditing helps you understand who changed an administrative setting, mailbox auditing is the way to discover “who did what in that mailbox?” For instance, what account used delegate access to send a message on behalf of a mailbox, or who removed a message from a folder. The audit events captured by mailbox auditing are in three categories of access, each of which has a separate configuration:

- **Owner:** Operations performed by the mailbox owner. Normally there is little point in auditing owner operations as the owner has full control over their mailbox. This was the situation in the on-premises world for many years, but when Microsoft decided to enable mailbox auditing by default, they also decided to enable auditing for owner actions to collect a complete set of audit events for each mailbox. The default configuration for owner actions includes *MoveToDeletedItems*, *SoftDelete*, *HardDelete*, *UpdateFolderPermissions*, *UpdateInboxRules*, and *UpdateCalendarDelegation*.
- **Delegate:** Operations performed by another user who has delegate access to the mailbox via the *SendAs*, *SendOnBehalfOf*, or *FullAccess* permissions. These actions are the usual focus for auditing, especially when multiple delegates have access to a shared mailbox or for mailboxes that hold confidential information, such as the copies of items retrieved by eDiscovery searches that are stored in discovery mailboxes.
- **Administrative:** Operations performed by programs that connect to the mailbox using special administrative access such as a content search.

When auditing is enabled for a mailbox, Exchange captures audit entries for a set of default events for each of the three access categories listed above and stores the data in the Audits sub-folder of Recoverable Items.

Mailbox auditing is only available for user, shared, and group mailboxes. It does not apply to public folder mailboxes or resource mailboxes.

## Mailbox Auditing by Default

Mailbox auditing is enabled by default for all tenants. This means that Exchange Online captures audit records for a default set of actions performed by owners, delegates, and administrative processes in user, shared, and group mailboxes. Exchange Online doesn't support auditing for resource, or public folder mailboxes.

Apart from ensuring the consistent capture of audit records for all user, shared, and group mailboxes (including newly-created mailboxes), the benefit of mailbox auditing by default is that Microsoft manages the auditing configuration to make sure that the audit log ingests events for new mailbox actions as they become available.

To verify that mailbox auditing is enabled for the organization, run the *Get-OrganizationConfig* cmdlet and retrieve the *AuditDisabled* property.

```
[PS] C:\> Get-OrganizationConfig | Select AuditDisabled
```

```
AuditDisabled
-----
False
```

To opt-out of default mailbox auditing and stop Exchange capturing any mailbox audit data for the audit log, change the setting to *\$True*.

```
[PS] C:\> Set-OrganizationConfig -AuditDisabled $True
```

In this case, the organization setting is *False*, so we know that auditing is enabled by default. In this state, you do not need to enable auditing for new mailboxes as Exchange Online takes care of this action. It also means that Exchange Online captures the default set of audit records for all mailboxes, even if the *AuditEnabled* setting for a mailbox is *\$False*. In other words, Exchange ignores the *AuditEnabled* property for a mailbox when default auditing is enabled. If you want Exchange not to capture audit records for a mailbox, you run the *Set-MailboxAuditBypassAssociation* cmdlet as explained later.

### Transmission of Mailbox Events to the Audit Log

When auditing is enabled by default, Exchange Online captures audit events in the *Audits* sub-folder of the *Recoverable Items* folder in the mailbox. You can run the *Search-MailboxAuditLog* cmdlet to interrogate these events. If the license assigned to the mailbox's account is Office 365 E3 or above, Exchange Online can transmit the events to the audit log. Microsoft documentation says that [you must enable mailbox auditing for mailboxes belonging to accounts with Office 365 E3 licenses](#). Our experience is that this isn't necessary and that mailbox events for these mailboxes appear in the audit log along with those from other accounts.

**Mailbox auditing and multi-geo organizations:** Exchange Online does not support cross-geo mailbox auditing. This is particularly obvious in mailbox delegation, such as when a user located in the home region uses their SendAs permission to send messages from a mailbox located in a satellite region. The satellite region does not log the action.

## Audit Actions for User Mailboxes

Table 21-2 lists some of the different actions configurable for capture in each access category. The [complete list of mailbox audit actions](#) is available online.

<b>Action</b>	<b>Description</b>	<b>Owner</b>	<b>Admin</b>	<b>Delegate</b>
<i>ApplyRecord</i>	Mark an item as a record by applying a record or regulatory record retention label.	Yes	Yes	Yes
<i>Create</i>	An item is created in the mailbox.	Yes	Yes	Yes
<i>Copy</i>	An item is copied to another folder.	No	Yes	No

<i>FolderBind</i>	A client opens a mailbox folder, including the original logon to the mailbox. Exchange consolidates folder bind actions and posts a single entry per folder every 24 hours.	No	Yes	Yes
<i>SendAs</i>	A message is sent from the mailbox using the SendAs permission.	No	Yes	Yes
<i>SendOnBehalf</i>	A message is sent from the mailbox using the Send On Behalf Of permission.	No	Yes	Yes
<i>SoftDelete</i>	An item is moved into Recoverable Items.	Yes	Yes	Yes
<i>HardDelete</i>	An item is permanently removed from Recoverable Items.	Yes	Yes	Yes
<i>Update</i>	The properties of an item are updated.	Yes	Yes	Yes
<i>MailItemsAccessed</i>	Messages are synchronized or opened by clients. See earlier section about high-value audit events.	Yes	Yes	Yes
<i>Move</i>	An item is moved into another folder.	Yes	Yes	Yes
<i>MoveToDeletedItems</i>	An item is moved into Deleted Items.	Yes	Yes	Yes
<i>MailboxLogin</i>	The owner logs into the mailbox.	Yes	No	No
<i>RecordDelete</i>	An item marked as a record is soft-deleted.	Yes	Yes	Yes
<i>Send</i>	An email is sent (crucial event). Must be configured before events are gathered.	No	No	No
<i>UpdateCalendarDelegation</i>	Record delegate permissions assigned to the calendar.	Yes	Yes	No
<i>UpdateComplianceTag</i>	A different retention tag is applied to an item.	Yes	Yes	Yes
<i>UpdateFolderPermissions</i>	Record changes to folder permissions, such as allowing a user to view a calendar.	Yes	Yes	Yes
<i>UpdateInboxRules</i>	Record changes to inbox rules.	Yes	Yes	Yes

Table 21-2: Mailbox actions captured for auditing purposes

The *HardDelete* action does not record audit events for items removed using the *Shift+Delete* key combination as this operation only bypasses the Deleted Items folder to move items directly into the Recoverable Items folder. *Shift+Delete* operations are logged as *SoftDelete* actions, which also happens when users remove items from the Deleted Items folder individually or by emptying the entire folder.

**Odd Delegate Audit Records:** You might find some audit records for *SendAs* operations by a delegate for a mailbox that you know has no delegates and wonder what's going on. The answer is usually when a background process impersonates the user to send email. For instance, this happens when Planner generates email notifications when creating new tasks or completing a task (the *ClientInfoString* property is "Client=REST," showing that Planner used the Graph API for this action). The messages that Planner sends are in the Sent Items folder of the mailbox. You can match them with the audit records.

## Updating Mailbox Audit Configurations

To see what actions are captured for a mailbox, run *Get-Mailbox* to fetch details of the *AuditAdmin*, *AuditDelegate*, and *AuditOwner* properties together with the default audit setting for the mailbox:

```
[PS] C:\> Get-Mailbox -Identity "Customer Services" | Format-List Audit*, DefaultAuditSet
```

```
AuditEnabled      : True
AuditLogAgeLimit  : 90.00:00:00
AuditAdmin        : {Update, MoveToDeletedItems, SoftDelete, HardDelete...}
AuditDelegate     : {Update, MoveToDeletedItems, SoftDelete, HardDelete...}
AuditOwner        : {Update, MoveToDeletedItems, SoftDelete, HardDelete...}
DefaultAuditSet   : {Admin, Delegate, Owner}
```

The audit configuration listed above tells us that:

- Auditing is active for the mailbox.
- Exchange keeps audit items for 90 days.
- The set of actions captured for the three categories of access
- Exchange uses the default audit set for each of the three categories.

PowerShell truncates the set of audited events for each category. To see the complete set of audit actions for a category, expand the property holding actions for the chosen category like this:

```
[PS] C:\> Get-Mailbox -Identity "Customer Services" | Select -ExpandProperty AuditDelegate
```

```
Update
MoveToDeletedItems
SoftDelete
HardDelete
SendAs
SendOnBehalf
Create
UpdateFolderPermissions
UpdateInboxRules
```

You can change the set of actions by running the *Set-Mailbox* cmdlet. In this example, we add the *UpdateComplianceTag* action to the set of actions already configured for capture. It can take up to an hour before an updated configuration is effective. Remember too that customizing the set of actions captured for a category stops Exchange from updating the set of default actions should Microsoft introduce new actions in the future.

```
[PS] C:\> Set-Mailbox -Identity James.Ryan -AuditOwner @{Add="UpdateComplianceTag"}
```

To disable auditing for a category, input "None" for the action list. Remember that if mailbox auditing is enabled by default for the organization, Exchange still captures the default set of actions in all three categories.

```
[PS] C:\> Set-Mailbox -Identity "Customer Services" -AuditDelegate None
```

## Managing Default Mailbox Audit Configurations

When mailbox auditing is enabled by default for an organization, Exchange populates the *DefaultAuditSet* property. This property indicates if an admin has changed the default audit configuration set by Microsoft for the Owner, Delegate, and Admin categories for a mailbox. To examine the value run:

```
[PS] C:\> Get-Mailbox -Identity Kim.Akers | Select DefaultAuditSet
```

```
DefaultAuditSet
-----
{Admin, Delegate, Owner}
```

Because the three categories are listed, you know that no change has been made. On the other hand, if you see:

```
DefaultAuditSet
-----
{Admin, Delegate}
```

You know that an administrator has added or removed an action from the Owner set. If a blank value is seen, you know that changes have occurred for all three categories. When the default set is updated, Exchange Online won't update the audit configuration when Microsoft introduces new actions, so it's important to keep to the default set whenever possible and only make changes when necessary. If you want, you can reset

mailboxes by running *Set-Mailbox*. In this example, we scan for mailboxes where the default set is not used and reset those mailboxes.

```
[PS] C:\> $Mbx = Get-Mailbox -RecipientTypeDetails UserMailbox -ResultSize Unlimited | Select Alias,
DisplayName, DefaultAuditSet
$Updates = 0
ForEach ($M in $Mbx) {
    $FullSet = $M.DefaultAuditSet[0] + $M.DefaultAuditSet[1] + $M.DefaultAuditSet[2]
    If ($FullSet -ne "AdminDelegateOwner") {
        $Updates++
        Write-Host "Updating" $M.DisplayName
        Set-Mailbox -Identity $M.Alias -DefaultAuditSet "Admin", "Delegate", "Owner" }}
If ($Updates -eq 0) {Write-Host "No mailboxes updated"} Else {Write-Host $Updates "mailboxes updated
with default audit configuration"}
```

## Mailbox Auditing Bypass

If you want to exclude mailboxes from auditing, you use the *Set-MailboxAuditBypassAssociation* cmdlet to tell Exchange Online not to collect audit events. For example:

```
[PS] C:\> Set-MailboxAuditBypassAssociation -Identity "Kim Akers" -AuditBypassEnabled $True
```

Set the value to `$False` to disable mailbox audit bypass.

Tenant administrators often use the *Set-MailboxAuditBypassAssociation* cmdlet to stop audit events from service mailboxes flooding the audit log. This might have been acceptable in an on-premises environment when the cmdlet first appeared in Exchange 2010. Today, it is not recommended to exclude any Exchange Online mailbox from auditing in this manner. First, Microsoft 365 has much better tools available to search the audit log and a few extra audit events will make no difference. Second, when an account is excluded from mailbox auditing, Exchange does not capture events for any mailbox the account can access. An account could therefore open a shared mailbox and remove items in that mailbox without the recording of any audit data. Last, if you experience something like a Business Email Compromise attack, you need as much audit data as possible to collect to understand how the attack develops and what was its impact.

To see if any mailbox auditing bypasses are in place, run the *Get-MailboxAuditBypassAssociation* cmdlet and filter for accounts with a bypass. It is quite normal to have many mailbox association records returned, including records for guest user accounts and system accounts as these are created when new accounts are created in the tenant. For example, this command finds accounts with mailbox audit bypass enabled.

```
[PS] C:\> Get-MailboxAuditBypassAssociation | ? {$_.AuditBypassEnabled -eq $True} | Format-Table
Name, WhenCreated, AuditBypassEnabled
```

Name	WhenCreated	AuditBypassEnabled
Kim Akers	02/12/2014 15:20:42	True

## Auditing for Group Mailboxes

When mailbox auditing by default is enabled for an organization, audit records are captured for the group mailboxes belonging to Groups. Unlike other mailbox types, you can't change the auditing configuration for group mailboxes. Table 21-3 lists the configuration used to capture audit records for group mailboxes.

Mailbox action	Owner	Delegate	Admin
Create		Y	Y
HardDelete	Y	Y	Y
MovetoDeletedItems	Y	Y	Y
SendAs		Y	Y
SendOnBehalf		Y	Y

<i>SoftDelete</i>	Y	Y	Y
<i>Update</i>	Y	Y	Y

Table 21-3: Audit configuration for group mailboxes

## Reporting Workload Activity

Despite many requests over the years, Microsoft has never delivered good reporting facilities for on-premises applications. This statement was true in the early years of Office 365, but things have improved recently with analytics and usage data appearing in different admin portals. What hasn't changed is the fact that if you want to have truly flexible reporting over extended periods, you should consider ISV products. The Microsoft Graph is increasingly the source of truth for workload activity data, but other sources such as the audit log are also useful.

**Graph Usage Reports:** Microsoft generates a range of workload activity data [accessible through the Microsoft Graph](#). Data for all workloads is not yet available, but you can access information for Exchange Online, SharePoint Online, Yammer, Teams, and OneDrive for Business.

### Using PowerShell to Create Reports

Administrators have used PowerShell for years to generate reports and many examples of scripts to produce nicely-formatted reports are available on the web. You can certainly use PowerShell to create custom versions of some of the standard reports to meet your specific needs. For example, we might want to know the size of user mailboxes. A quick PowerShell query reveals the answer:

```
[PS] C:\> Get-Mailbox -RecipientType UserMailbox | Get-MailboxStatistics | Sort ItemCount
-Descending | Format-Table DisplayName, TotalItemSize, ItemCount -AutoSize
```

DisplayName	TotalItemSize	ItemCount
Tony Redmond	4.116 GB (4,419,110,350 bytes)	36529
Jeff Guillet	1.62 GB (1,739,456,346 bytes)	10411
Vasil Michev (Technical Guru)	1.081 GB (1,161,234,561 bytes)	7019
Deirdre Smith	933 MB (978,334,934 bytes)	6207

Reports like this hold valuable data and are useful on an ad-hoc basis. However, when you set out to use PowerShell, you should realize that:

- PowerShell is an excellent method to query data but is less good when the time comes to format reports for printing. PowerShell can output data in CSV format for later processing in Excel to create charts or to do deeper analysis but generating a report and a few charts in Excel is different from creating a clear and well laid-out report.
- PowerShell has limited ability to analyze data over an extended period (the cmdlets that retrieve data from the reporting data mart sometimes do not give a very granular level of information).
- Exchange Online throttles PowerShell if a script attempts to process details of hundreds or thousands of mailboxes. Throttling ensures that no job soaks available resources and so reduces service to other tenants. It is good if your script can avoid the need to call the *Get-ExoMailbox* cmdlet to create a set of mailbox objects for processing as this is a resource-intensive activity that becomes a candidate for throttling. One way around this is to create lists of mailboxes and store them in CSV files that other scripts can open and process.
- Some components do not support PowerShell or make their data available to PowerShell. Exchange Online has good coverage, but you will not be able to generate reports for Teams or Yammer unless you fetch data from the Graph. An example script showing how to use the Graph APIs with PowerShell to generate usage reports for Exchange Online, SharePoint Online, OneDrive for Business, Teams, and



Yammer workload data and Azure AD user sign-in data is [available in GitHub](#). The data extracted from the Graph is the same as reported in the Microsoft 365 admin center, the Teams admin center, and the Microsoft 365 Usage Analytics for Power BI.

Finally, if you do create custom PowerShell reports, you must be prepared to check and potentially update them regularly to ensure that code does not break when Microsoft introduces updates to cmdlets, new functionality, or new versions of PowerShell. On the upside, some will like the challenge of writing and updating reports, others will find that it is easier and more efficient in the end to buy a reporting package from a company that specializes in this space, and it is always great to be able to reuse the many contributions of PowerShell scripts and code snippets that people make to the community to form the basis of your solution.

## Standard Usage Reports

The Microsoft 365 admin center includes a Reports section (Figure 21-11) offering a set of reports covering various workloads for fixed 7, 30, 90, and 180-day periods. Users do not need full administrative access to access the standard usage reports. Anyone assigned a workload administrator role like the Teams Service administrator, or the Global Reader or Report Reader administrative roles can access the usage reports (see Chapter 4).

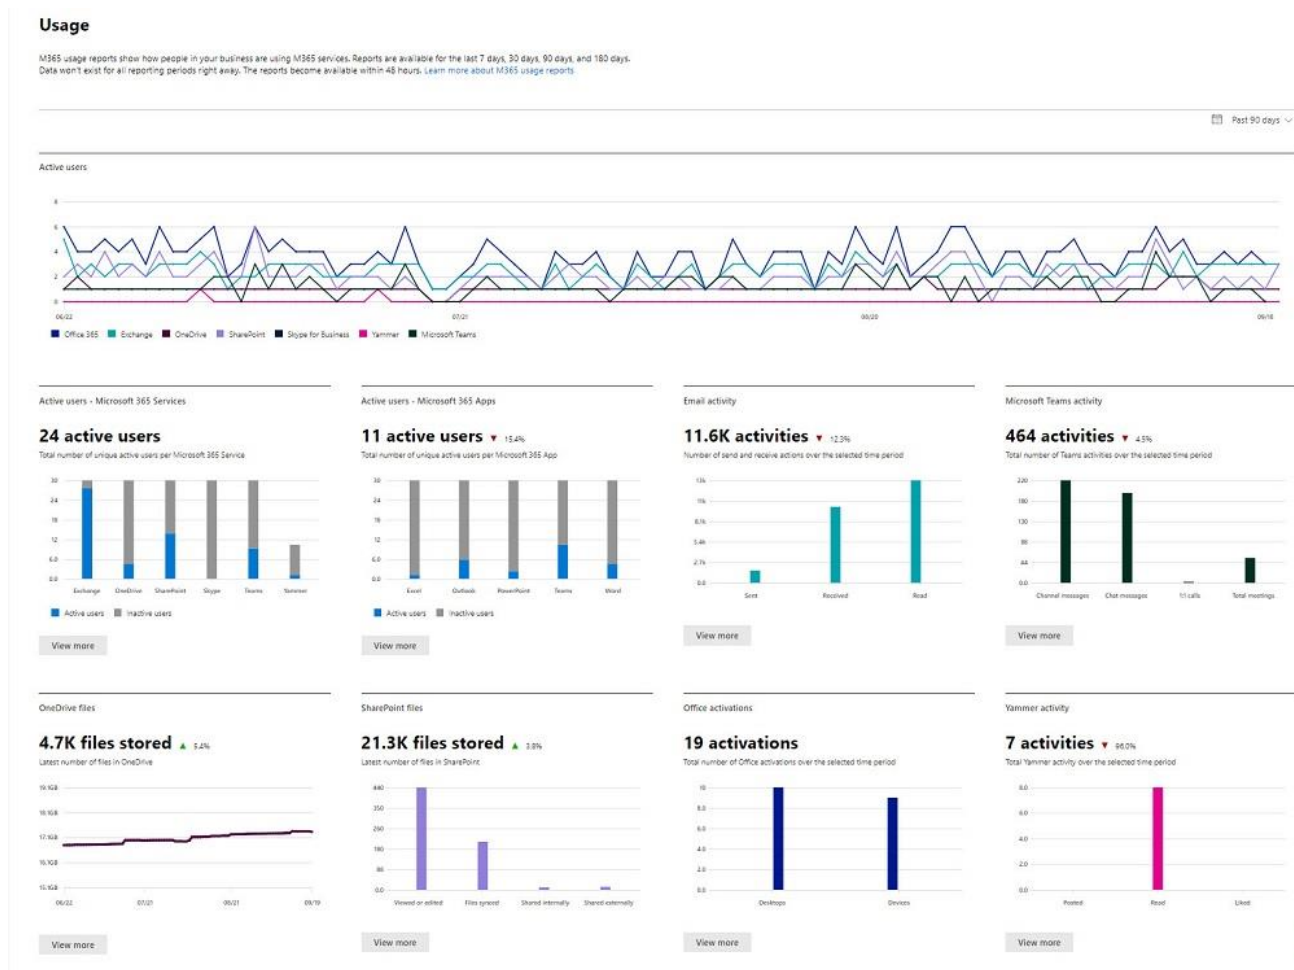


Figure 21-11: Standard workload usage reports for a tenant

The Microsoft 365 admin center features two types of reports:

- **Usage:** How much use people make of different workloads, like how many messages users send and receive using Exchange, meetings attended in Teams, or files viewed or edited with SharePoint. Some

applications, like Planner, do not have usage reports, and some which do, like Teams, report a limited set of actions (mainly in chat and meetings) and don't include any insight into the other ways people interact with the app.

- **Productivity Score** covering remote work elements. Microsoft introduced this section to help tenants understand the impact of employees working from home. The data focuses on collaborative activities (Exchange Online, Teams, Forms, and Yammer) and online document use (SharePoint Online and OneDrive for Business).

The detail and coverage of the reports available in the Microsoft 365 admin center have improved over time and deliver enough information to gain insight into user activity. Microsoft makes up to 180 days of reporting data available to tenants.

The reports for the individual workloads include a grid table containing individual user data. The table holds up to 2,000 entries and can be sorted by the different columns. You can also **Export** to save the report data to a CSV file (you'll need to export the data to sort and analyze user activity if your tenant has more than 2,000 accounts). If you want to preserve the graphics, you can either take a screenshot or use the browser's print to PDF choice to create a PDF file. It is also worth mentioning that you can access additional reports in other admin portals. For example, the Azure AD admin center has reports focusing on directory-centric activities such as password resets, anomalous sign-ins, and rights management, but they are useful in understanding the complete operating environment. The Teams admin center also includes a set of reports about various aspects of Teams usage.

Third-party reporting products make their money by offering a wider range of reports with a more granular level of detail than is available in the standard reports. They also usually give access to information gathered for more than the standard 180-day reporting window or incorporate data gathered from a variety of sources within a customer's IT environment. In comparison to the standard reports, which tend to focus on increasing usage of applications like SharePoint Online or Teams, ISV products often place greater emphasis on understanding what the data means and how to be more effective in areas like license management.

To stay in business, ISV offerings must be better than the reports delivered by Microsoft. Better analysis and insight into business operations are always welcome and this is exactly the advantage that you hope to gain by using a third-party product. To address the need to dive deeper into how workloads and users are performing, ISVs can use the same data from the Microsoft Graph used by the Microsoft 365 admin center to build tailored solutions to satisfy specific customer needs that the standard reports do not address.

## Anonymized Usage Data

Some organizations are uncomfortable with the idea that people holding many administrative roles (including Report Reader) might be able to see sensitive data about user activity. For example, it is easy to discover how many messages the CEO sends and receives or how many Teams meetings he or she attends. Up to September 1, 2021, the default for Microsoft 365 reporting was to show full information about users and groups. From that date, the usage reports in the Microsoft 365 admin center, Teams admin center, and anywhere else which uses the Microsoft Graph usage reports API display anonymized values. Figure 21-12 shows anonymized data in a usage report.

Details						Export
Username	Last activity date (UTC)	Send actions	Receive actions ↓	Read actions		
FE7CC8C15246EDCCA289C9A40...	14 September 2018		1,575	2,692		2,698
784DC33C0AA659D3342FBD75...	08 May 2017		0	233		0
ADEC9FD53F428E35D9028CAF1...	14 September 2018		35	227		574
E0E7E73AB82A0726D7AAA872B...	25 July 2018		0	160		0
E885D8F6B229530E97F74CB5D0...	14 September 2018		9	148		556
8ED6F13F8C692AA248DDB43C3...	05 February 2018		0	111		0
D37C99D58D92387F88D6EDFB...	26 May 2016		0	100		0
B19CECF62B156C7C468E37F5B7...	18 May 2018		0	26		0

Figure 21-12: Anonymized user data in the email activity report

Anonymization (also called deidentification or obfuscation) means that the Microsoft Graph replaces the personal information for user accounts (like display names and user principal names), groups, and sites with system-generated obfuscated values that cannot be associated with the underlying objects, including GUIDs that could be used to identify an object.

If the organization decides that it is acceptable to display full user information, a global administrator can switch the setting in the **Reports** section of **Org settings** in the Microsoft 365 admin center. This action generates an *UpdatedCFRPrivacySettings* audit record in the Office 365 audit log. Because the setting governs the Microsoft Graph usage reports API, the deidentification setting also applies to usage data fetched by scripts and programs which use the API.

Note that some administrative roles such as Usage summary reports reader or Global reader can never see details of user activity, anonymized or not.

## Microsoft 365 Usage Analytics for Power BI

Tenants can install a [Power BI template app](#) (previously known as the Office 365 adoption content pack) to help understand the usage patterns for several Microsoft 365 workloads. Usage analytics is a pre-built dashboard of graphs and charts. You only need a basic Power BI license to access the dashboard. Together, the graphs and charts form a dashboard for a tenant divided into four areas:

- Executive Summary: A general view of workload usage across the tenant (Figure 21-13).
- Activation/Licensing: How many licensed accounts are in the tenant, how many are active, and what devices they use.
- Product Usage: How much use is made of Exchange Online, SharePoint Online, Teams, Yammer, Microsoft 365 Groups, and OneDrive for Business.
- User Activity: Different views of user activity within the tenant.

The current iteration includes good coverage of Exchange Online, SharePoint Online, OneDrive for Business, Teams, and Yammer. It does not cover other applications like Stream or Planner, so the dashboard covers usage of the base workloads rather than giving a comprehensive view of activities across all workloads. Microsoft might include coverage of the missing applications over time.

The dashboard helps tenants understand the adoption rate and usage for various aspects of usage over the prior twelve months. For instance, how many people use SharePoint Online, Teams, and Exchange Online – or all the covered applications. This kind of information can tell you if you have a license management problem where people have licenses for an application that they do not use.

Another use of the dashboard is to look at usage trends, such as whether email traffic increases over time. These insights into what is happening within the tenant can help you to understand changes in user behavior. For instance, if you make a determined effort to convince people to store documents in SharePoint and OneDrive and access the files there rather than attaching them to email for circulation within the company, a steady uptick in file storage should result. If it does not, then the campaign is unsuccessful, and you might need to take a different approach to convince people to use “cloudy attachments.” Another example is in

Teams, where a reduction in email traffic should match an increase in minutes consumed in instant messaging and conversations. Still another is to measure whether the introduction of Teams causes users to move some of their collaboration with fellow employees away from email to conversations within team channels.

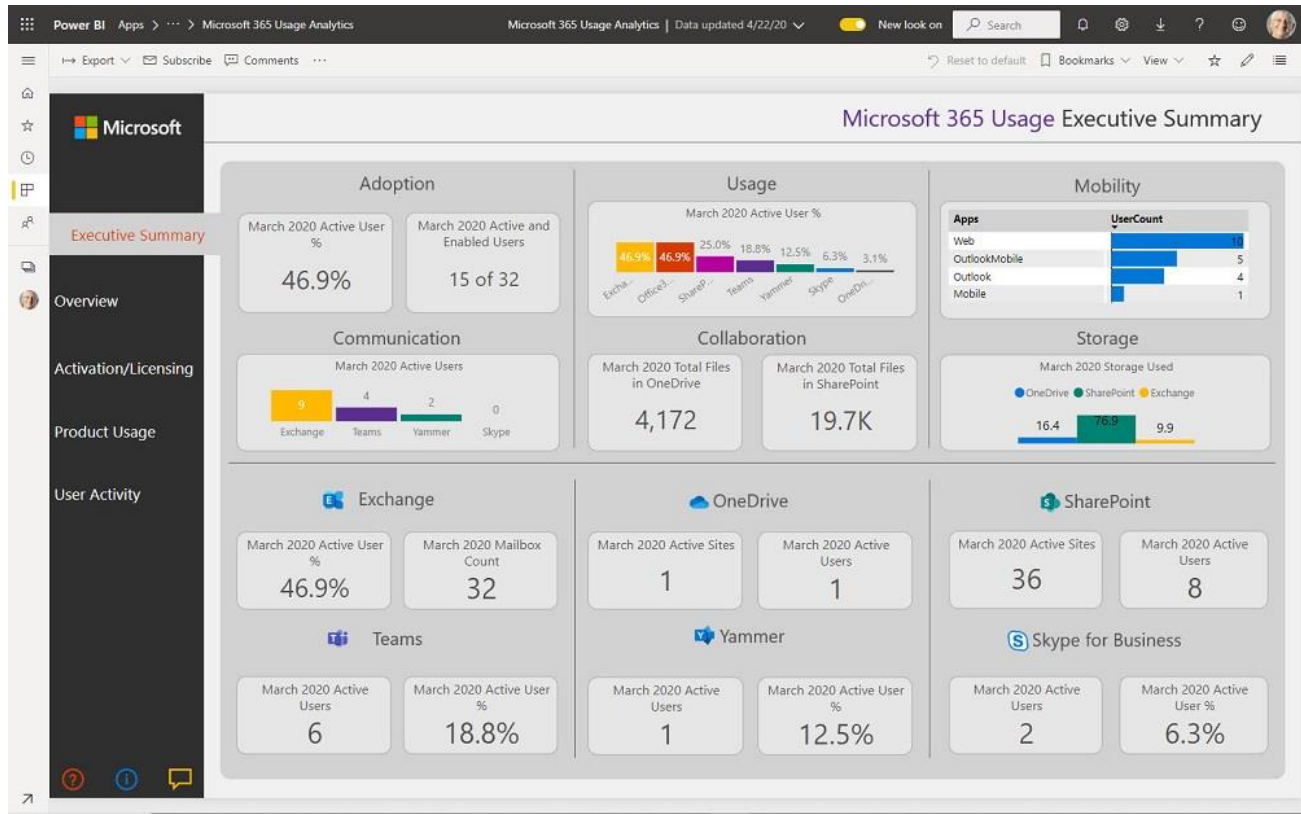


Figure 21-13: Microsoft 365 Usage Analytics dashboard

Microsoft gathers data from across workloads to aggregate and refines the data to refresh the dashboard. The content pack updates weekly by default, but the backend service refreshes data daily, so a tenant can amend the dataset scheduled refresh interval to receive daily updates. However, the actual data lags the current date by between 36 and 48 hours as Microsoft needs this time to gather information and process it for the dashboard. Given the scale of Microsoft 365 and the amount of data that must be processed to generate aggregate statistics for tenants, this delay is understandable. The setting for anonymized data for reports governs whether usernames or anonymous identifiers appear in the dashboard.

## Independent Reporting Products

Few tenant administrators are skilled at generating easy-to-read reports that highlight important data. Programmatic access to the Reporting DataMart is available to allow software developers to use the data belonging to tenants if tenants grant permission for this access. The intention is to allow vendors who specialize in software monitoring and reporting to incorporate usage data into their products to give customers a much better view of what happens inside a tenant than you can get using the basic reports. Apart from the obvious goodness inherent in having a common repository to consult to retrieve operational data about Microsoft 365, the combination of Reporting DataMart and choice of APIs (including PowerShell and the Microsoft Graph) avoids the need for people to create personal methods to extract information.

The number of independent software vendors who create reporting solutions continues to expand. Most of the ISVs that offer reporting solutions focus on pure cloud environments while some include monitoring and reporting for on-premises servers or non-Microsoft cloud workloads. Each product plays to its strengths and employs different approaches to data extraction, analysis, and reporting. Among the criteria you can use to figure out whether you need to use a third-party product instead of the standard usage reports include:

- **Data longevity:** Microsoft stores tenant reporting data for 180 days. This might be insufficient for your needs if you want to chart the usage of applications over a longer period for planning and analysis purposes. Ask for how long the vendor stores tenant reporting data. In addition, ask about the security and privacy controls that are in place to ensure that tenant data cannot be compromised. Ask about how often the product refreshes data from Microsoft 365 workloads to understand how up-to-date the reports are.
- **The number and type of reports:** A limited set of standard usage reports are available in number and in what they report. For example, there are no reports for mobile device usage, perhaps because Microsoft would prefer tenants to use Microsoft Endpoint Manager for mobile device management. A third-party reporting vendor is likely to have hundreds of different reports in their product.
- **Intelligent views:** Looking at a single workload is interesting if you want information about that application, such as the number of active mailboxes in Exchange Online. Things become more interesting when you can compare data extracted from one workload against another. For example, if you have a project to make more use of OneDrive for Business, you will be interested in the number of sites in use and the files and quota used in those sites. It is also interesting to know if the growing use of OneDrive affects email volume and usage.
- **Filtering and Pivoting:** Viewing tables of data through a browser is acceptable for small tenants but rapidly becomes problematic when the number of reported objects grows. For instance, it is difficult to make sense of a table holding details of thousands of mailboxes. The ability to use filters and to view data through pivot tables are huge advantages for large tenants.
- **Access:** The reports available in the Microsoft admin center need some level of administrative access. This is inconvenient when users other than administrators want access too. For instance, a company might want to charge back costs for Exchange Online mailboxes to the operating units of the business. This activity is usually performed by accountants, not tenant administrators, so removing the requirement to have administrative permissions to be able to access reports can be very useful.
- **Delivery:** Not everyone wants to browse a website to access reports. Some products allow reports to be automatically extracted and emailed to users on a scheduled basis.
- **Scale:** How well does the product scale up to deal with the predicted full load of the tenant? Good demos performed against a 5-user tenant might not be so impressive when confronted with the data generated by 20,000 users.
- **Support:** How is support provided? Is support available locally?
- **Cost:** How is the cost of the product calculated?
- **Deployment:** Is any extra software needed to create or view the reports? Is the product fully web-based or do you have to install some software on a workstation or servers to access the reports?

No one product is the best choice for every situation. The best approach is to spend the time to test the software in your environment and make the decision based on what you discover there.

**Gaps and changes:** All third-party reporting products and Microsoft's reports use a single source of truth: the reporting data collected for a tenant. Some third-party products copy that data so that it is retained for longer periods and to perform additional analysis. However, you need to be aware of two issues regarding reporting data. First, the data can change because of a change made by Microsoft. This is what happened when Microsoft changed the way that the *Get-MailboxStatistics* (or *Get-ExoMailboxStatistics*) cmdlet reported system messages stored in user mailboxes (described in the Exchange Online chapter). Second, hiccups and operational glitches sometimes prevent the ingestion of workload data to the reporting data mart, which then creates gaps in reports. Sometimes the gaps are filled in after normal service is restored, sometimes the gaps remain. Keep your eyes open and make sure that if a gap appears, you make Microsoft and/or the third-party reporting vendor aware of the issue.

## Extracting More Detail About Tenant Activity

The standard usage reports include a lot of information about how a tenant consumes resources, but the information is often quite high-level and not as granular as you might like. In addition, the reporting service makes data available for a limited timeframe, usually with a maximum horizon of 180 days, which makes the reporting information provided by the service difficult to use for long-term planning. For example, Figure 21-14 illustrates the pattern of SharePoint Online usage over the last 30 days. The report can show data going back over a maximum of 180 days, but you cannot go back any further to explore questions such as how much growth occurred in the last year or last two years. The bottom part of the report lists the activity of individual users.

Reporting products often use PowerShell to retrieve more information from workloads and Azure AD to build out the data necessary to fill in the knowledge gaps. For example, it is all very well to know that a total of 55 distribution lists exist within a tenant, but it is even better to know the membership of each group and whether the groups are in active use or just taking up dead space in the GAL. To make long-term reports possible and to assure speedy access to information, reporting vendors often download tenant data periodically to repositories that they manage. The data is useful for analyzing trends such as the growth in mailboxes or messaging activity over time. Each product has its set of reports and its unique way to interpret and present information.

Extra reporting capabilities are enabled by Azure AD Premium. For example, if you want detail about a [password-based activity](#) such as user account resets, your account must have an Azure AD Premium license to access that report. See the Azure AD documentation for the [latest information about available reports](#).

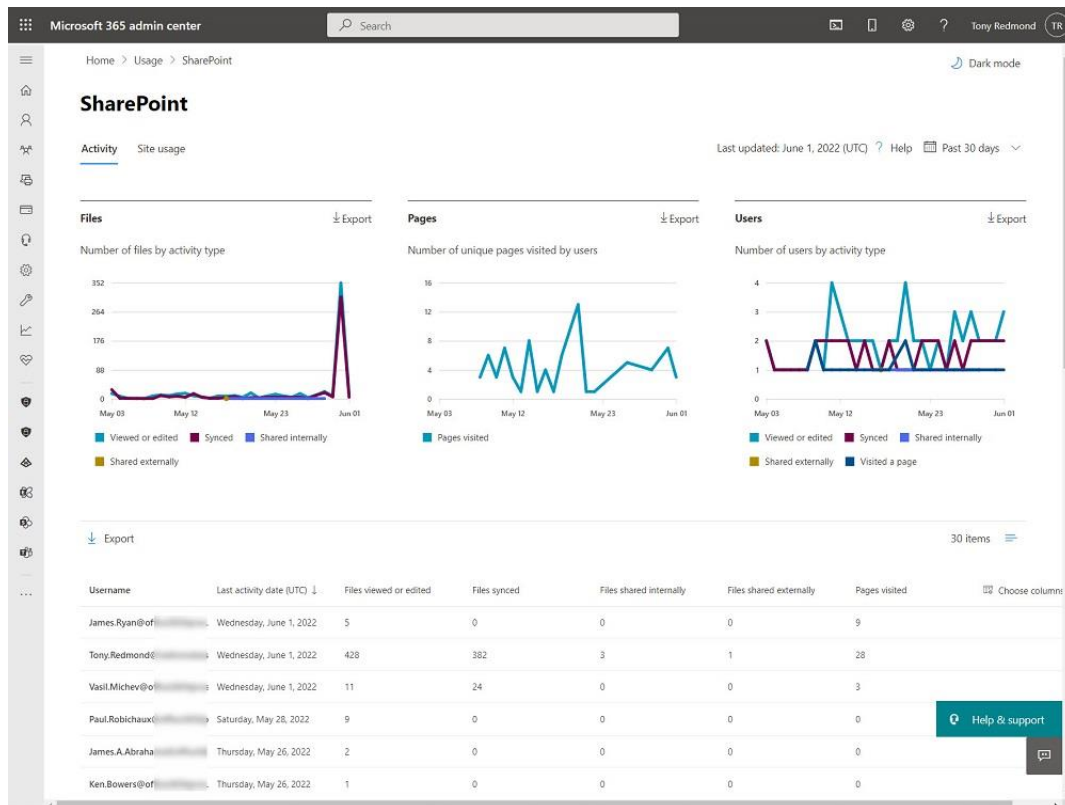


Figure 21-14: Viewing the usage report for SharePoint Online

# Communications Compliance

[Microsoft Purview Communications Compliance](#) is part of the Microsoft 365 Insider Risk solution set. These policies are the third iteration of functionality designed to help organizations monitor communications sent and received by employees. The normal reason why an organization wants to monitor communications is to reduce the potential for risk due to mistakes or deliberate actions by employees in their internal and external communications. The loss caused by inappropriate or illegal communications might be reputational or financial. Either way, it's undesirable, especially for large enterprises, which is why organizations want to proactively detect any potential problems and then prove that they are on top of the situation should the need exist to demonstrate this point to regulators or other corporate bodies.

Although it might seem sneaky and reminiscent of a "big brother" approach to employees, certain industries have regulations that force companies to ensure that they have a supervision policy in place, perhaps only for employees that work in specific areas. For example, [FINRA](#), the U.S. Financial Industry Regulatory Authority, enforces rules governing the activities of brokers and dealers to ensure market transparency and fairness. It was common to find that on-premises organizations created special software called "transport sinks" to capture and examine messages sent between employees. The code ran in the Exchange transport pipeline to examine and capture copies of messages and route the copies to personnel skilled in regulations and compliance, who reviewed the traffic to ensure that no problems exist. Companies often combined transport sinks with transport rules to control how certain groups of users communicated with each other.

Organizations using communication compliance policies must have Office 365 E5, Microsoft 365 E5, or Microsoft 365 E5 Compliance for all accounts covered by the policies.

## Reviewing Microsoft 365 Communications

Communication compliance policies set guidelines for acceptable internal and external communications. In a Microsoft 365 context, communications mean:

- Exchange Online email.
- Teams personal chat and channel communications, including messages sent by users with on-premises mailboxes. The Teams messages reviewed by communication compliance policies are the compliance records captured by the Microsoft 365 substrate and used for eDiscovery (see Chapter 13 for more information).
- Yammer conversations (if the Yammer network runs in Microsoft 365 native mode).
- Messages imported from an external source [through a connector](#) like Bloomberg messaging.

Communications compliance policies depend on messages stored in Exchange Online. Background agents check mailboxes to scan for policy violations in normal messages and the compliance records captured for Teams chat and channel conversations (including records for conversations with external users in other tenants and Skype consumer users).

When the agents detect violations, they copy messages to special mailboxes to make them available for review and resolution. This isn't a real-time process, and you can expect a delay of approximately an hour before policies detect and flag problem messages for review. In most cases, this isn't a problem because compliance checking is a reactive process. In any case, it's usual that policies select only a percentage of detected messages for review, so not every message that could contain a violation will turn up in the portal for review.

## Components of a Communications Compliance Policy

Communication Compliance policies are managed through the Microsoft Purview Compliance portal. A policy includes the following components:

- **Reviewees:** The people (individuals or groups) whose communications come within the scope of the communications compliance policy. You can use email distribution lists or Microsoft 365 groups to define the reviewees for a policy. You can't use dynamic groups or dynamic distribution lists to define reviewees.
- **Locations:** Define the workloads monitored by the policy, including Exchange, Teams, and Yammer.
- **Conditions:** Define the characteristics of communication items policies examine to detect policy violations. The conditions include:
  - Sensitive information types as used elsewhere in Microsoft 365, as in Data Loss Prevention policies.
  - Classifiers, both those provided by Microsoft such as Resumes and Source Code, or those created by the tenant to find items commonly used by the business.
  - Message properties such as message size, attachment type and size, originating domain, and sensitivity label.
- **Sample size:** The percentage of messages which meet the policy conditions captured for review.
- **Reviewers:** The people who will access the items captured for review to decide whether the content of those items complies with the regulations in force. Reviewers process captured items through the Compliance portal and decide whether items are compliant or non-compliant. In the latter case, reviewers can then escalate the situation to the offending user's manager, HR, or some other department for resolution.

An organization can deploy multiple communication compliance policies within their tenant, each of which monitors a set of users. With the structure of a policy in mind, before you can create a new policy, you should answer the following questions:

- What set of users come within the scope of the policy? How is the target set identified?
- Communication between users includes a vast variety of content. How do we focus on the messages for review? For example, are we looking for a specific codeword or terms?
- What is an adequate sample for review? A policy might review 100% of detected messages during the testing or initial deployment phase of deployment to make sure that the conditions specified for the policy are accurate and not too broad. After reaching this point, it's usual to reduce the percentage to 10% or so to lessen the load on reviewers. Experience with supervision policies proves that selecting too high a percentage can generate so many items that reviewers cannot cope with the volume and turnaround reviews in a reasonable time.
- Who will examine the captured traffic? The reviewers need to understand what constitutes a problem when they see it in a message. They also need to understand what is a serious violation that needs immediate action and what is not so serious. Reviewers do not have to be fluent in all the languages used for communications because Microsoft Translator can translate messages into a reviewer's preferred language.
- What happens when a violation occurs? All violations have consequences, but some violations have graver consequences than others. The organization might require an escalation process to allow reviewers to escalate violations to a higher authority for a decision about how to handle a matter, including potentially reporting a case to a regulatory authority.
- User awareness. Before the tenant begins to monitor communications, it is only fair to inform the users whose communication is under review about potential violations and the consequences if a policy detects a problem that proves to be a real issue. This is an area where the company's legal and HR departments have a key role to play.



After determining the answers to these questions, you can go ahead and create a policy, assuming your account holds the right roles.

## Communications Compliance Role Groups

Five role groups support granular management of communications compliance policies and the data captured by policies:

- **Communication Compliance:** Includes all the other communication compliance role groups.
- **Communication Compliance Administrators:** Manages communications compliance policies.
- **Communication Compliance Analysts:** Can view communications compliance message metadata to analyze the effectiveness of policies.
- **Communication Compliance Investigators:** Can view the full body of messages detected by communication compliance policies.
- **Communication Compliance Viewers:** Allows access to communication compliance reports and widgets.

The Communication Compliance role might be the only one used in smaller tenants. The other roles exist to allow granular assignment of permissions to specific people to do certain tasks. For more information, see [this page](#).

## Workflow Considerations

Communications compliance includes simple workflow processing (what Microsoft calls “*flexible remediation workflows*”) to help track and resolve violations, mostly by the dispatch of emails to offenders and their managers and recording the outcome. For the most serious cases, communications compliance is integrated with Microsoft 365 Advanced eDiscovery, allowing for information gathered about violations to be transferred to eDiscovery as prepackaged cases for further investigation and resolution.

Thought must be put into how to integrate what’s available in the application to complement and build on existing HR procedures. For instance, what should happen upon the detection of a violation by an employee? And what escalation steps happen if someone proves to be a serial offender? These are decisions that Microsoft can’t make because every company is different. The implementation of employee monitoring is highly dependent on the industry the organization works in and the applicable regulations.

## Machine Learning and Classifiers

Communication compliance policies can include classifiers to detect types of information. Microsoft 365 uses a variety of classifiers for different purposes and tenants can create their own by going through a training process. During this process, machine learning processes sets of sample documents to build a map of common characteristics that solutions can use to detect the content of the same type. The prepackaged classifiers include:

- Offensive language (now deprecated).
- Profanity.
- Resumes (CVs).
- Source code.
- Targeted harassment.
- Threat.
- Adult images.
- Gory images.
- Racy images.
- Sexual harassment.

Microsoft's text-based classifiers can handle multiple languages (the set grows over time), including English, French, German, Italian, Japanese, Arabic, Dutch, Korean, Spanish, Portuguese, and Chinese (including Chinese traditional). Image-based classifiers are language-independent. Images examined by these classifiers can be in JPEG, GIF, BMP, and PNG formats. Images must be at least 50 x 50 pixels to be assessed. Optical Character Recognition (OCR) scanning for printed or handwritten content embedded in email and Teams messages has [similar size limitations](#).

Supervision policies don't include machine learning and tenants had to define terms of interest to check for. Having classifiers makes it easier for tenants to monitor communications based on well-known and well-understood characteristics. For instance, if someone tells another person that "I hate you," the harassment classifier will recognize the "hate" keyword and the context of its use. Using classifiers and machine learning helps communication compliance policies generate fewer false positives than supervision policies do. We will probably never have zero false positives but eliminating most of these reports eases the load on reviewers.

Even the best machine learning detection experiences problems with the way language flex and evolves. One person's offense is another person's norm, which makes it imperative that reviewers consider messages selected for review in context. For instance, a [scatological reference](#) about someone in an email or Teams chat might be innocuous or offensive, depending on how it is phrased. And calling someone a pile of brown smelly bovine output is likely to pass most machine learning tests. The evolving use of language and the variation of acceptance of different terms can cause many false positives, which is one of the reasons why Microsoft deprecated the offensive language classifier and replaced it with a combination of threat, profanity, and harassment qualifiers, each being more focused and therefore more likely to generate a correct result.

## Anonymized Results

Communication compliance settings include the option to see anonymized usernames (like *Anonyl0-Ibb*) instead of the real display names of users who cause policy matches. You might prefer to select this setting to preserve user privacy during the initial phases of potential violation reviews and investigations.

## Creating a New Communications Compliance Policy

You can create a new policy from scratch, or you can use one of the template policies provided by Microsoft:

- **Regulatory Compliance:** looks for entries in a custom lexicon of words that might indicate violations if present in messages between a specific group of accounts. Internal teams responsible for monitoring regulatory compliance probably deal with violations of this nature.
- **Sensitive Information:** looks for the presence of sensitive information types in messages (Microsoft 365 includes many sensitive information types covering anything from passport numbers to social security numbers). Violations of this policy are like those encountered in Data Loss Prevention processing.
- **Conflict of interest** looks for evidence in communications between users or groups of actions that conflict with the interests of the company.
- **Inappropriate text** uses built-in classifiers to detect text in messages that the organization might consider inappropriate, abusive, or offensive. An example of country-level regulations in this area is [Japan's "power harassment" law](#), which took effect for large companies on 1 April 2020.
- **Inappropriate images** detects adult or "racy" images.

The template policies cover (at least in part) many of the scenarios that create the need to monitor communications, so the quickest way to get going is to create a new policy from a template and amend it afterward to meet your needs. The Inappropriate text template is a good example to start with because it's easy to generate some sample messages for the policy to capture. To create a new policy, click **Create policy** and choose *Detect inappropriate text* from the drop-down list. Several policy settings are pre-populated from

the template, and all you need to do at this point is change the policy name (if you want) and add the users or groups to supervise and the reviewers.

In Figure 21-15, we see the screen used to collect details of the new policy. A distribution list defines the sets of users to monitor, and we've appointed a single reviewer. It's usually best to have more than one reviewer per policy. When you go ahead and create the policy, some background processes run to create the special mailbox used to collect items for review and inform the communications monitoring assistants that a new policy is active. It can take up to an hour for this process to complete setup and a further 24 hours before the assistants begin to monitor messages.

**Monitor communications for conflict of interest**

About this template

Set up a policy to monitor communications between two groups of users across locations like Exchange, Teams, and more. Just choose the two groups whose communications you want to supervise, specify reviewers, and we'll set up the rest.

Settings we need from you

Policy name \*

Conflict of interest (June 2022)

Supervised group A \*

CEO Direct Reports Start typing to find users or groups

Supervised group B \*

Vice Presidents Start typing to find users or groups

Reviewers \*

Tony Redmond Lotte Vettler Start typing to find users

Settings we've filled in for you

You can change these later. Click 'Customize policy' if you want to configure different settings now.

Communications to monitor ⓘ

Monitored locations Exchange, Teams, Yammer

Conditions and percentage

Communication direction Internal

Percentage to review 100

Optical character recognition(OCR) Disabled

More Options

Create policy Customize policy

Figure 21-15: Creating a new communications compliance policy

## Tweaking a Communication Compliance Policy

When you create a policy from a template, you can only change certain policy settings until the setup has finished. If you want access to all policy settings during creation, choose to create a custom policy. Once setup is complete, you can tweak the policy to make it more effective. Common changes include:

- Extend the scope of the policy to include additional users or groups, or to make it apply to all users.
- Exclude selected users from the policy.
- Select the locations to monitor (Exchange, Teams, and Yammer). By default, the policy covers Exchange and Teams. Yammer is only available when the tenant network is in native Microsoft 365 mode.
- Select the direction of communications to monitor. By default, policies check all communications:

- **Inbound:** communications to the recipients specified in the policy from users who do not come under the scope of the policy. Traffic can come from other users within the organization or people outside the organization.
- **Outbound:** communications from the people specified in the policy to anyone else.
- **Internal:** communications between the people included in the policy.
- Amend the policy rules used to detect violations. If you're used to working with data loss prevention rules, you'll find the process very similar.
- Change the percentage of messages to review from the set detected to contain violations.

Figure 21-16 shows the conditions generated for a policy created from the *Inappropriate text* template. As you can see:

- Communications are monitored in all directions.
- Three classifiers are used to identify potential violations. The combination of these classifiers should catch messages containing profane, threatening, harassing, or offensive language subject to the caveats expressed above.
- 100% of all matching messages are selected for review.

## Communication compliance > Edit Identify Bad Words Used by Senior Management

**Choose conditions and review percentage**

**Communication direction \***

- Inbound.** Sent to users you choose to supervise from people not included in this policy.
- Outbound.** Sent from the users you choose to supervise to people not included in the policy.
- Internal.** Sent between the users or groups you identified in this policy.

**Conditions**

By default we will monitor all communications from the users and groups you specified. Add conditions to limit the results to communications matching specific criteria. [Learn more about these conditions](#)

**Content matches any of these classifiers**

Any of these

**Classifiers**

- Targeted Harassment
- Profanity
- Threat
- Add

+ Add condition

**Review percentage**

If you want to reduce the amount of content to review, specify a percentage. We'll randomly select the amount of content from the total that matched any conditions you chose.

100%

Back Next Cancel Need help? Give feedback

Figure 21-16: Amending conditions and sample percentage for a communications compliance policy

You could add more conditions to improve the effectiveness of the checking. Policies focused on language could be improved by adding a condition to check for particular words or terms that the classifiers might not

catch, such as new slang or a word used in known cases of harassment. Other policies will depend on checking for:

- A lexicon of specific words that might indicate suspicious or problematic behavior.
- Microsoft 365 sensitive information types (social security numbers, credit card numbers, and so on) and custom sensitive information types defined by the tenant.
- Message properties. These include:
  - Specified words or phrases appear in the message body (or its subject). Use KQL syntax to define the query, but do not try to get too complex. Alternatively, you can add a condition that checks for messages where the specified words are not present. You will see an error if the specified query is too complex for a supervision policy.
  - Any attachment to the message includes or does not hold the specified words or phrases.
  - The message has an attachment of a specific type or does not have an attachment of a specific type.
  - An attachment is larger than the specified size (in KB, MB, or GB).
  - The overall message size is larger than the specified size.
  - Messages are sent to or received from a specific domain.
  - Messages have a certain sensitivity label.

Once ready, click **Next** and **Save** to save the updates to the policy. Once again, it will take up to a day before the new policy settings are effective.

## Supervision Mailboxes

Communications compliance policies use special supervision mailboxes to store the copies of messages captured by the background agents for review. These hidden mailboxes can't receive email, and don't appear in any Microsoft 365 administrative interface except PowerShell, where you can run the `Get-SupervisoryReviewPolicyV2` cmdlet to reveal the name:

```
[PS] C:\> Get-SupervisoryReviewPolicyV2 "Regulatory Compliance" | Ft reviewmailbox  
  
ReviewMailbox  
-----  
SupervisoryReview{d7c6eb96-20e5-4cbb-9838-2d230f64efb1}@office365itpros.onmicrosoft.com
```

## Reviewing Captured Messages

The task of a policy reviewer is to examine captured items to decide whether any compliance violation exists. Do not underestimate the amount of work involved in processing several hundred review items, a volume that a busy tenant can easily generate daily. Apart from the opening and reading of each item, the reviewer must understand the context of the communication and how the content fits with regulations. Hours of work might be consumed to process items at a high level of accuracy.

To review messages, navigate to the Communication compliance section of the Compliance portal where you can see a list of the policies and statistics for each policy (Figure 21-17). Policies generate an alert any time processing detects four violations within 60 minutes.

To check waiting items, go to the Policies tab, select a policy, and click the *Pending* link. The items pending review can be filtered to focus on specific recipients, senders, domains, item types, subjects, and other properties. Items can then be grouped by family (for instance, show all email together) or by conversation, which assembles the messages in a Teams or email conversation (like a Teams meeting recording transcript) to allow the reviewer to see how an interaction unfolds. This is important because a remark taken out of context can look quite different when considered amid a full conversation.

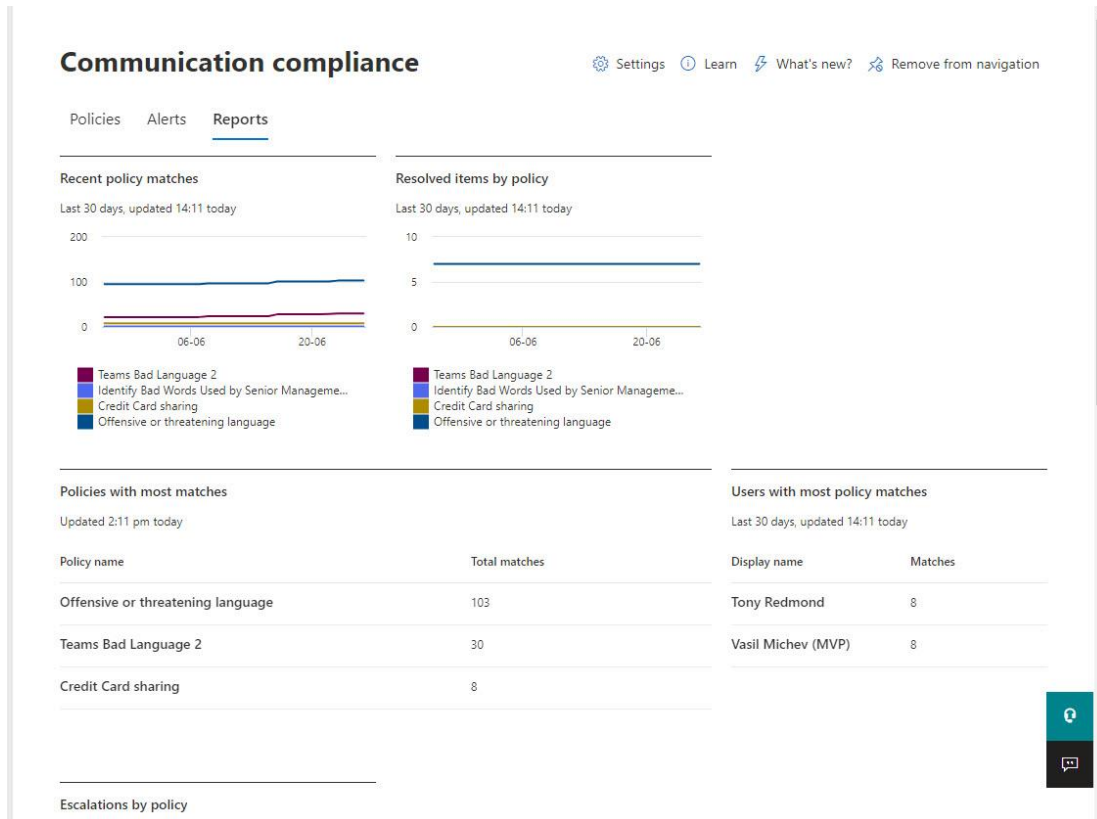


Figure 21-17: Statistics for communications compliance policies

The reviewer can then examine items to decide if a violation is present. The Compliance portal offers three views:

- **Source:** View an item as the user sees it.
- **Plain text:** Strip away all the formatting to show plain text with line numbers. This helps to focus on the words without the distraction of other elements.
- **Conversation:** Show the individual message containing the violation in context within a complete conversation.
- **User history:** Show previous instances when the user had policy violations.

In addition, the reviewer can see the user history to know if the sender of the message has previous violations for the same behavior. First-time offenders often receive more flexibility and understanding than serial offenders do. If the reviewer needs to download an item, they can do so as a message item or PDF.

The item selected for review in Figure 21-18 is a Teams message detected as containing some offensive text. The user history shows that the person sending the message has some other outstanding policy matches awaiting review. This might indicate that some corrective intervention is necessary to help the person modify the text they use in messages. It's likely that reviewers make the decision after checking each of the policy matches to understand the ebb and flow of the conversations and the context of the text detected as policy violations.

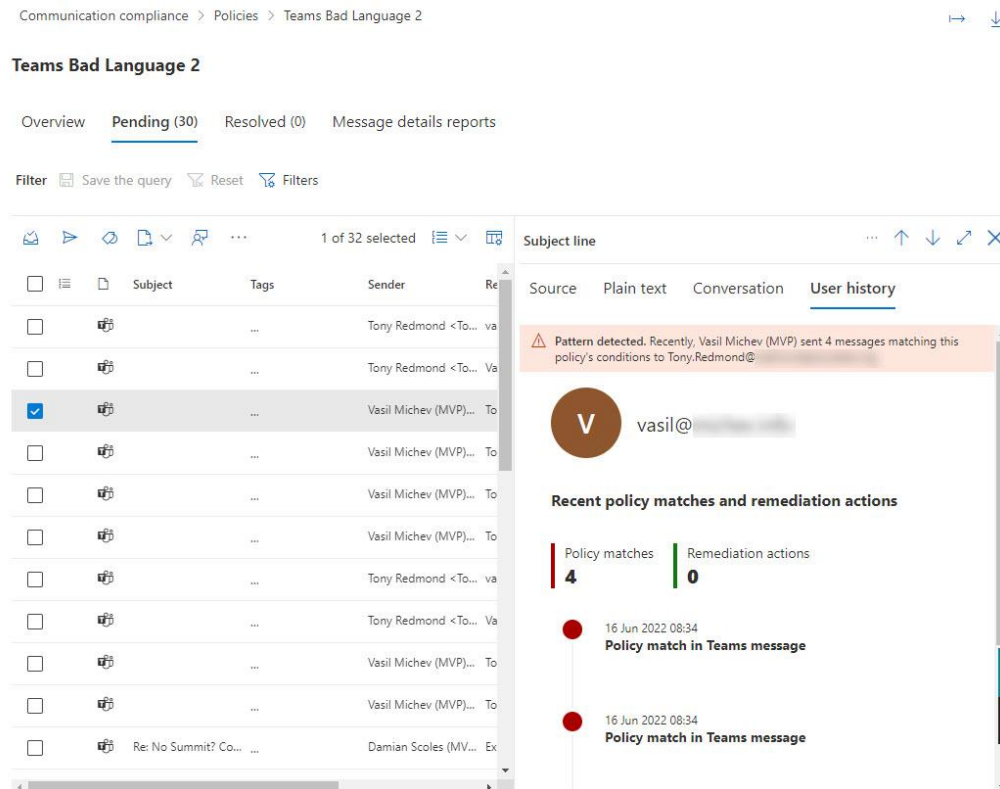


Figure 21-18: Reviewing pending items for a communications compliance policy

The steps that a reviewer can take to progress an item include:

- **Resolve:** Mark the item as resolved, perhaps after some invention. This action removes the item from the pending list. Eventually, the aim is for reviewers to resolve all items on the pending list.
- **Notify:** Send an email to the sender to tell them about the potential violation and (possibly) to ask for their input. Messages created from *Notice Templates* (see below) help ensure that notifications use the right tone and words and include any necessary references to organization policies.
- **Tag as:** Mark the item as compliant, non-compliant, or questionable.
- **Escalate:** Send the message forward to a higher authority for their review and decision. For instance, a message might go to the sender's manager.
- **Escalate for investigation:** Create an advanced eDiscovery case for the item and notify administrators about the new case. This step might be useful when a systematic problem is detected involving multiple people, such as potential insider trading.
- **Report as misclassified:** Help refine the machine learning model by marking an item as not being a problem. This action sends a copy of the misclassified item to Microsoft to help them refine the classifier that marked the item incorrectly.
- **Automate:** Create a [Power Automate flow based on the message](#). You can create a new flow or use a template such as the Notify manager (CC) flow. This flow sends a message to a user's manager when they violate a policy.
- **Show/Hide translation view:** Purview uses Microsoft Translator to translate the plain text of messages from [supported languages](#) into the language of the reviewer.

If the review message comes from Teams and the violation is clear and obvious, the reviewer can remove the message from its source chat or channel conversation.

It's worth emphasizing that these are generic explanations, and a tenant is free to use their interpretation of how to resolve review items based on their business, regulatory environment, and HR policies.

Many of the messages detected by a policy don't need much investigation and reviewers can dismiss them quickly. Reviewers can select multiple messages (or all) and resolve them in a single operation, which speeds up the review process dramatically.

## Notice Templates

Because notifying someone that they might have done something wrong is an exercise fraught with potential problems, compliance communication policies use notice templates to ensure that notifications use the right tone and appropriate wording. Notice templates are available through the Settings section for communications compliance. An organization must add at least one notice template before they can use communications compliance policies.

A template includes email basics such as:

- **The originating mailbox:** It is usually better when notification messages come from a specific mailbox set aside for this purpose instead of an investigator's personal mailbox.
- **Boilerplate CC and BCC recipients:** For instance, notifications might need to be copied to an HR mailbox.
- **Subject:** The subject of notification messages is often dictated by HR or legal guidelines.
- **Message body:** Boilerplate text to inform the recipient why they are being notified and what policy is violated. The text might also include some "next steps" for the recipient to take, such as talking to their manager or contacting HR to arrange for an interview.

When an investigator chooses the *Notify* action, they select a template, and Purview copies the boilerplate settings into a message form. The investigator can then modify the settings as required by taking steps such as including some details about the violation. When everything is ready, they can send the message for delivery as normal to the recipient mailboxes.

## Information Barriers

Exchange has long supported the concept of an [ethical firewall](#), a software barrier to stop defined sets of users from communicating with each other. In the past, organizations wrote bespoke customizations like transport sinks for this purpose. Usually, organizations operating in highly regulated industries deployed ethical firewalls to stop groups of people from communicating. For instance, a bank might stop traders and brokers from communicating. Since the advent of transport rules in Exchange 2007, transport (mail flow) rules are a favorite method to prevent communication between groups. Ethical firewalls continue in Exchange Online, but other methods of communication exist, notably Teams. Information Barriers deliver a cross-workload answer by preventing sets of users from communicating using its supported workloads. Figure 21-19 shows the Information Barrier architecture to control access across Exchange Online, SharePoint Online, Teams, and OneDrive for Business (because of the work needed in the transport system, Exchange Online does not currently support information barriers).



### Information Barrier Architecture

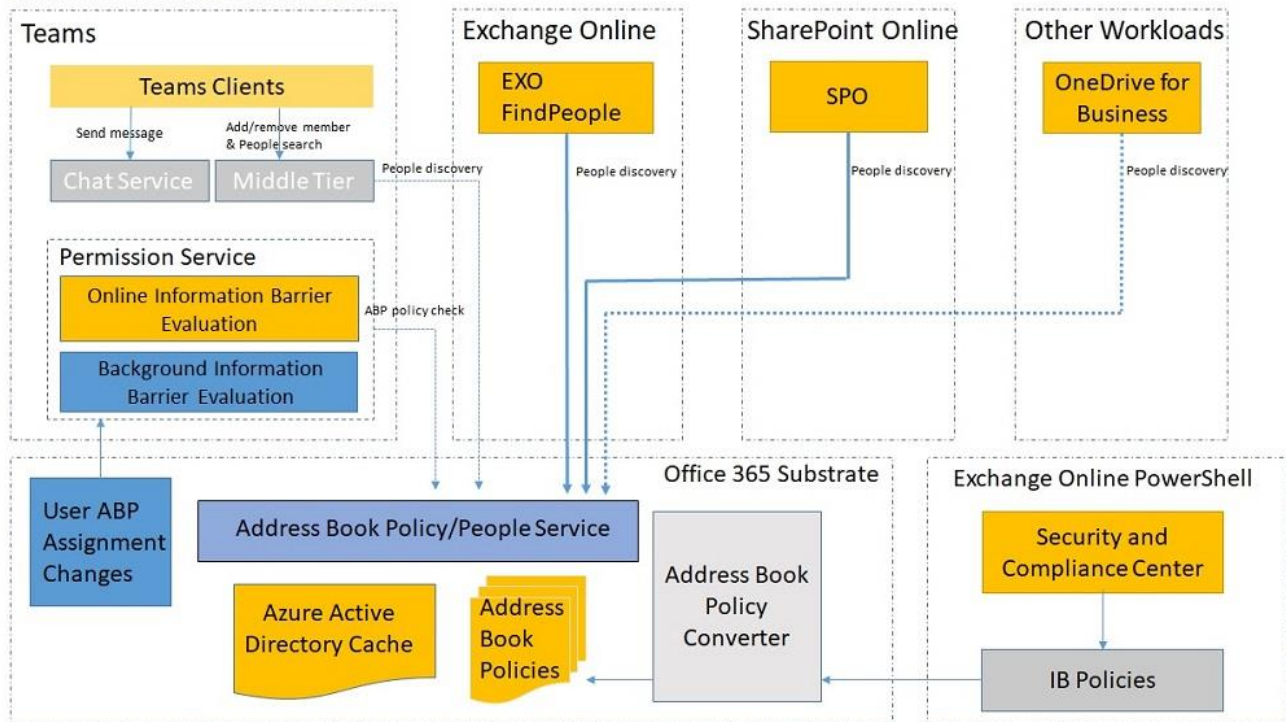


Figure 21-19: Information Barrier architecture (source: Microsoft – Ignite 2018)

Accounts under the control of Information Barrier policies require an Office 365 E5 license or the Microsoft 365 E5 Compliance add-on. An exception exists for education licenses because Microsoft bundles Information Barriers into the A5, A3, and A1 plans. Information Barriers are also part of the Microsoft Insider Risk solution.

The documentation for how to configure Information Barriers is [available online](#). Although global administrators often manage information barrier policies for a tenant, you can restrict permissions by assigning accounts the Information Barrier Compliance Management role as this is all that's needed to manage the policies.

Exchange Online currently doesn't support Information Barriers. However, given that these policies depend on Exchange address book views, it's likely that the eventual implementation of information barriers in Exchange Online will replace Address Book Policies (ABPs), which restrict user access to the directory and have been in use since Exchange 2010. Teams also depends on Exchange to scope directory queries, meaning that you must remove existing ABPs from a tenant before you can define Information Barrier policies. This can be a little tricky if any inactive mailboxes exist with assigned ABPs (assigned before the mailboxes become inactive): you can't remove an ABP if any mailbox (inactive or active) exists with a connection to the ABP, so you need to restore all inactive mailboxes with ABP assignments (restoring an inactive mailbox deletes the assignment) and then remove the ABPs.

## Information Barriers in Teams

As you might expect, workloads implement Information Barriers to match their style of communication. The implementation of Information Barriers in Teams uses several different components, including:

- [Directory scoping](#) to limit the access of users to specific parts of the directory.
- Organization segments (common across all workloads) to define sets of users using queries against Azure AD.
- Information Barrier policies define how segments can or cannot communicate with each other.

- Background processes to apply the settings in Information Barrier policies to personal chats, team memberships, meetings, calls, and (optionally) sharing documents in SharePoint Online and OneDrive for Business.

Each team has a property called *InformationBarrierMode* to control if it comes within the scope of information barrier policy processing. After implementing information barriers for Teams, the property for new teams is set to *Implicit* to mark the team for processing. However, you need to backfill the property for older teams. This is easily done with PowerShell:

```
[PS] C:\> [array]$Teams = Get-UnifiedGroup -Filter {ResourceProvisioningOptions -eq "Team"} -ResultSize Unlimited | ? {$_.InformationBarrierMode -ne "Implicit"}
Write-Host ("{0} teams are being backfilled for Information Barriers" -f $Teams.Count)
ForEach ($Team in $Teams) { Set-UnifiedGroup -Identity $Team.ExternalDirectoryObjectId -InformationBarrierMode "Implicit" }
```

The various modes supported by information barriers are [described here](#).

## Organization Segments

The sets of users used by Information Barrier policies are known as organization segments. Information Barrier policies define which segments can communicate with each other or those blocked from communication. The first step in implementing Information Barriers is to define organization segments for the users whose communication the organization wants to control. You don't have to create segments for the entire organization, but if you don't, the risk exists that some users will be able to communicate where a block should apply.

To define organization segments, you create queries to find Azure AD accounts. Microsoft suggests that you use the same property as the basis for all segments. In other words, if you decide to divide the organization by department, use the department property for all segments. It's also best if the segments don't overlap (multiple segments cover the same set of accounts).

Because Information Barrier policies depend so heavily on Azure AD, it follows that the information in the directory must be accurate. All the properties used to define organization segments must contain appropriate values, so a prerequisite for Information Barriers is to review the state of the directory to ensure that the right data is available. For instance, if you decide to use the department attribute as the basis, you could create a spreadsheet of all accounts including this data to verify its correctness. This PowerShell command does the job:

```
[PS] C:\> Get-MgUser -All -Filter "UserType eq 'Member'" | Select Department, DisplayName, UserPrincipalName, PhysicalDeliveryOfficeName | Sort Department | Export-CSV -NoTypeInfo c:\temp\UsersDepartments.csv
```

## Creating Organization Segments

Information Barrier policies work by either allowing or blocking segments from communicating with each other. To create a policy, we must first define the segments. Each segment has a filter to tell Azure AD what accounts come within its scope. Think of this as similar to the filter used by a dynamic group. In the example scenario, we want to create an ethical firewall for Teams to stop staff in the Trading and Sales departments from communicating with each other. To do this, we need two organization segments: one defines the accounts in the Trading Department, and the other defines the accounts in the Sales Department. Organization segments can include both tenant and guest accounts.

You can manage information barriers through PowerShell or the Microsoft Purview Compliance portal (Figure 21-20). The GUI includes a filter builder that makes it easier to construct complex filters used in organization segments and the conditions for information barrier policies. It's recommended that you start off using the GUI to build segments and policies and use PowerShell whenever necessary for automation.

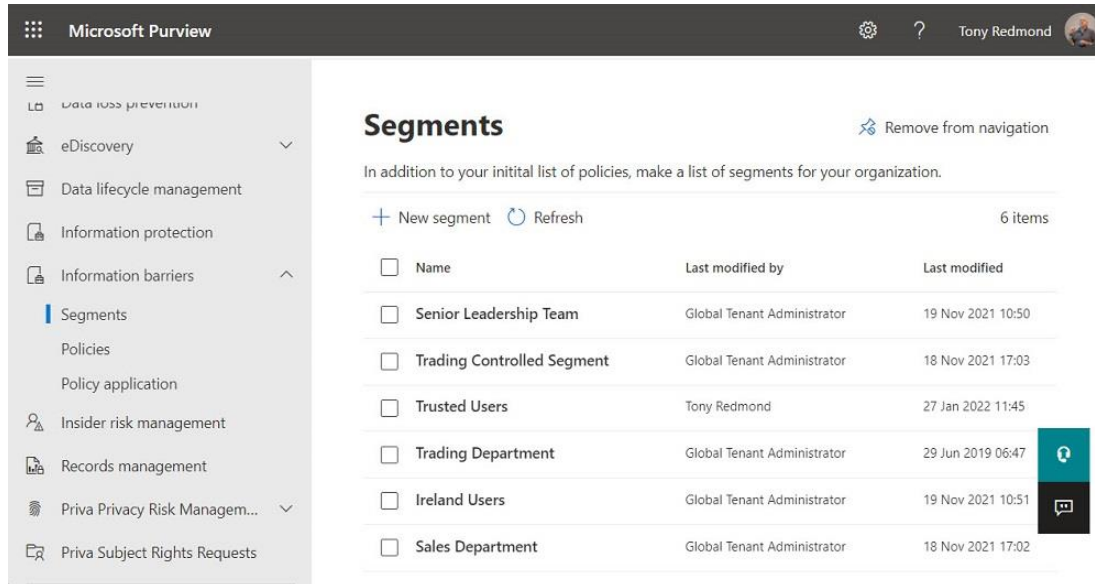


Figure 21-20: Information Barriers in the Microsoft Purview compliance portal.

To use PowerShell, connect to the [compliance endpoint](#). In this example, we create two organization segments.

```
[PS] C:\> New-OrganizationSegment -Name "Trading Department" -UserGroupFilter "Department -eq 'Trading'"
New-OrganizationSegment -Name "Sales Department" -UserGroupFilter "Department -eq 'Sales'"
```

To test the effectiveness of the query used by an organization segment, we can run the query with the *Get-Recipient* cmdlet. For example:

```
[PS] C:\> Get-Recipient -RecipientPreviewFilter {Department -eq 'Sales'}
```

Name	RecipientType
Andy.Ruth	UserMailbox
Eoin.Smith	UserMailbox
JAbrahams	UserMailbox

If you need to replace the filter, use the *Set-OrganizationSegment* cmdlet. To avoid user disruption, it's best to settle on segment definitions before introducing Information Barrier policies into production. However, sometimes organizational changes such as the renaming of a department or other structural updates can force an alteration in a filter. For example:

```
[PS] C:\> Set-OrganizationSegment -Name "HR Department" -UserGroupFilter "Department -eq 'Human Resources and People'"
```

The segments defined above use simple queries. You can construct more complex queries using the [set of filterable attributes supported by information barrier queries](#). The 'co' listed in the set of attributes refers to the country name, like Singapore, Germany, or United Kingdom.

If you can't find a suitable attribute to use, you might be able to use the membership of a distribution list or Microsoft 365 group instead. This example creates a segment based on the membership of a Microsoft 365 group identified by its external directory identifier (the documentation says that you can use the group name, but this doesn't appear to work).

```
[PS] C:\> Get-UnifiedGroup -Identity "Trading Team" | Select ExternalDirectoryObjectId
```

```
ExternalDirectoryObjectId
-----
29b959d6-56ab-4aea-b70d-ffb8397d397f
```

```
New-OrganizationSegment -Name "Trading Controlled Segment" -UserGroupFilter "MemberOf -eq '29b959d6-56ab-4aea-b70d-ffb8397d397f'"
```

After deploying information barrier policies and letting the deployment jobs run to validate segment membership, you can check what policy applies to an individual user with the *Get-InformationBarrierRecipientStatus* cmdlet:

```
[PS] C:\> Get-InformationBarrierRecipientStatus -Identity Andy.Ruth -Identity2 JAbrahams
```

## Creating Information Barrier Policies

After creating the necessary segments, we can create two Information Barrier policies. One policy prevents people in the trading department from communicating with sales, the other blocks reverse traffic. To make it clearer when looking for a policy, it's a good idea to capture the nature of the block in the policy name:

```
[PS] C:\> New-InformationBarrierPolicy -Name "Trading Block Sales" -AssignedSegment "Trading Department" -SegmentsBlocked "Sales Department" -State Inactive
[PS] C:\> New-InformationBarrierPolicy -Name "Sales Block Trading" -AssignedSegment "Sales Department" -SegmentsBlocked "Trading Department" -State Inactive
```

We deliberately create the Information Barrier policies in an inactive state to allow administrators to set everything up (and do some testing) before activating policies across the organization. You can't use a segment in an information barrier policy if it is already in use by another policy (even if that policy is inactive).

Note that you cannot remove an organization segment once an information barrier policy is active. However, you can change its user filter so that the segment no longer applies to any recipients.

## Activating and Debugging Information Barrier Policies

To activate a policy, run the *Set-InformationBarrierPolicy* cmdlet to change the state to Active.

```
[PS] C:\> Set-InformationBarrierPolicy -Identity "Trading Block Sales" -State Active
```

After changing policies to an active state, run the *Start-InformationBarrierPoliciesApplication* cmdlet to begin the process of applying the policy conditions to the affected users.

```
[PS] C:\> Start-InformationBarrierPoliciesApplication
```

It can take some time before the application executes active policies to make them fully effective for users in the relevant organization segments (don't try to start multiple jobs as this generates some horrible errors). The *Get-InformationBarrierPoliciesApplicationStatus* cmdlet returns the status of the information barrier policy application. In this example, the application processed 4,343 recipients.

```
[PS] C:\> Get-InformationBarrierPoliciesApplicationStatus

RunspaceId           : 6f427633-0a48-4f57-9c84-ec5e8bb946f1
Identity             : b4cf044a-e915-4960-9baa-7804415c2a3a
CreatedBy            : Administrator
CancelledBy          :
Type                 : ExoApplyIBPolicyJob
ApplicationCreationTime : 11/16/2021 21:11:54
ApplicationEndTime    : 11/16/2021 22:22:19
ApplicationStartTime  : 11/16/2021 21:11:54
TotalBatches         : 5
ProcessedBatches     : 5
PercentProgress      : 100
TotalRecipients      : 4343
SuccessfulRecipients : 4343
FailedRecipients     : 0
FailureCategory      : None
Status               : Completed
```

By default, the cmdlet returns the status of the last run of the information barrier policy application. To see details of all runs, use:

```
[PS] C:\> Get-InformationBarrierPoliciesApplicationStatus -all:$True
```

The application processes all mail-enabled recipients, including mail contacts. This is to make sure that the address book segmentation functionality works for all recipient types. If the job reports failed recipients, it's probably because those recipients are in multiple organization segments. To resolve the problem:

- Note the identity for the job reported by *Get-InformationBarrierPoliciesApplicationStatus*. In the example shown above, the identity is b4cf044a-e915-4960-9baa-7804415c2a3a.
- Search the audit log to find audit records with the *InformationBarrierPolicyApplication* record type for the date when the application ran and filter the set to find errors belonging to the run.

```
[PS] C:\> [array]$AuditLogs = Search-UnifiedAuditLog -ResultSize 5000 -Formatted -RecordType
InformationBarrierPolicyApplication -StartDate 16-Nov-2021 -EndDate 17-Nov-2021 | Where-Object
{$_ .AuditData.Contains("b4cf044a-e915-4960-9baa-7804415c2a3a")} -and
$AuditData.Contains("IBPolicyConflict") }
```

- Examine the *AuditData* property of each audit record to find the issue reported by the application, noted under *ErrorDetails*:

```
"ErrorDetails": "Status: IBPolicyConflict. Error: IB segment \"352e1fe3-fbee-4980-9208-103ddea4b370\" and IB segment \"93889a38-0cbf-4d0f-a260-005a6a1c4893\" has conflict and cannot be assigned to the recipient \"fdc6b121-44b8-4262-9ca7-3603a16caa3e\".\r\n.",
```

- We can see that the problem is that the conflict is between two segments. To see what the segments are, run the *Get-OrganizationSegment* cmdlet:

```
[PS] C:\> Get-OrganizationSegment | ? {$_ .ExoSegmentId -eq "352e1fe3-fbee-4980-9208-103ddea4b370" -or $_ .ExoSegmentid -eq "93889a38-0cbf-4d0f-a260-005a6a1c4893"} | ft Name
```

```
Name
----
Trading Controlled Segment
Sales Department
```

- To resolve the recipient name, run the *Get-Recipient* cmdlet:

```
[PS] C:\> Get-Recipient -Identity fdc6b121-44b8-4262-9ca7-3603a16caa3e
```

```
Name      RecipientType
----      -
Andy.Ruth UserMailbox
```

You now know why the error happened. The two segments both found the user Andy Ruth. Recipients identified in segments must be unique, so the solution is to update the recipient's properties to move them out of one of the two segments. You can't update a recipient's properties or the filter in an organization segment while the information barriers application is active, so wait until the active job finishes before making any necessary changes.

Microsoft's SLA for the application to detect a change in an account's properties that impact a barrier policy is 24 hours. However, the application usually detects changes like updating someone's department faster (3-4 hours). The thing to realize is that it takes time for workloads to react to changes in the directory and to set expectations accordingly. Updating someone's role at 9 am does not mean that a barrier exists to stop them from communicating with others by 9:30 AM. This is especially true when the directory is in a state of churn due to department reorganizations or other major changes.

If a problem occurs in deploying Information Barrier policies or you make changes to policies, you can run the *Start-InformationBarrierPoliciesApplication* cmdlet to process all active policies.

## Checking User Accounts

After processing an account to activate Information Barrier Policies, the account has an Information Barrier GAL (for the policy applied to the mailbox). This is how Information Barrier policies replace Exchange ABPs.

```
[PS] C:\> Get-Mailbox -Identity Andy.Ruth | Select -ExpandProperty AddressListMembership
\IBPolicyGAL_1bd52480-82f2-4416-814a-022776e46bef
\Default Global Address List
\Mailboxes(VLV)
\All Users
\Offline Global Address List
\All Mailboxes(VLV)
\All Recipients(VLV)
```

The *Get-Recipient* cmdlet also returns organization segment information for the user.

```
[PS] C:\> Get-Recipient -Identity Andy.Ruth | Format-Table InformationBarrierSegments

InformationBarrierSegments
-----
{352e1fe3-fbee-4980-9208-103ddea4b370}
```

We can check the membership of an organization segment by running a query using the GUID for a segment:

```
[PS] C:\> Get-Recipient -RecipientTypeDetails UserMailbox, GuestMailUser | ?
{$_ .InformationBarrierSegments -eq "352e1fe3-fbee-4980-9208-103ddea4b370"} | Format-Table Displayname
```

You'll also see that mailboxes that do not come under the scope of an information barrier policy also receive policy address lists. In effect, although Exchange Online does not support information barrier policies and the mailbox is unrestricted, a policy applies to allow the mailbox owner to communicate with any other user in the organization. Behind the scenes, workloads are aware of policies applicable to mailboxes. Clients use the *FindPeople* API to request this information from the server and receive a restricted or open view in response.

```
[PS] C:\> Get-Mailbox -Identity James.Ryan | Select -ExpandProperty AddressListMembership
\IBPolicyGAL_40ed3838-a6fd-4d1f-aa54-652537d5a708
\IBPolicyGAL_1bd52480-82f2-4416-814a-022776e46bef
\Mailboxes(VLV)
\All Mailboxes(VLV)
\All Recipients(VLV)
\Default Global Address List
\All Users
\Offline Global Address List
```

Outlook desktop users can see the names of the address lists created for Information Barrier policies. As the names are quite obscure, this might cause some questions for the help desk until people realize that they are just internal names.

It might take some time before the organization segments are acceptably accurate and reflect the communication paths allowed by the organization (or by regulation). There's also the small matter of synchronization and client caching to consider as it will take more time for clients to update changes synchronized from Azure AD and respect the blocks imposed by the policy.

## Allowing Communication by Policy

Information Barriers also support the concept of an allow list to allow segments to only communicate with one or more other segments. For example, this command permits communications between accounts in the Group HQ segment with five other segments.

```
[PS] C:\> New-InformationBarrierPolicy -Name "Group HQ Communications" -AssignedSegment "Group HQ"
-SegmentsAllowed "Sales Department", "Marketing", "Research", "Trading Department", "Engineering"
```

If workloads can't find a policy to block or allow communication between two accounts, it assumes that communication is allowed.

## How Teams Uses Information Barriers

Teams imposes blocks to implement Information Barriers at several points.

**Team membership:** If a team owner tries to add someone to a team when a member already exists who is blocked from communicating with the new potential member, Teams won't allow the new member to be added. For existing teams, if a background scan of the membership detects a violation, the Information Barrier Processor removes members to bring the team into compliance. In most cases, removals are processed on a last-in, first-out basis: in other words, the last person who joins a team and causes a violation is removed and existing members are left in place. However, processing depends on when the Information Barrier Processor detects a violation caused by a change to the properties of individual user accounts or updates for organization segments, so removals might happen in a different order.

Any member of a team, including org-wide teams, whose presence causes a policy violation is removed by the Information Barrier Processor. Removal includes owners if they are in violation, meaning that a team can be left ownerless. Audit records for removals are in the audit log where you'll see that the user removing the account from the group is noted as "*ServicePrincipal\_Guid*."

**Information Barriers, Teams, and Guest Users:** One problem that Teams currently has with Information Barriers is that team owners can't add a new guest user account. Existing guest accounts don't cause problems because Teams can check that adding them to a team roster won't violate a policy. However, if you try to add an external person whose account doesn't already exist in the directory, Teams can't validate that the barrier is respected, and the attempt to add the member fails. However, the guest account is created in Azure AD. The workaround is simple: create the guest account through the Azure AD portal or by adding them to the membership of the underlying group using Outlook or OWA. The addition of the new member will be replicated to Teams and any Information Barrier checks will then be imposed. Microsoft is aware that this situation is unsatisfactory and is working on a fix.

**Chats:** Teams won't start the chat if the participants are blocked by policy. At least, Teams will start a chat, but if the chat was created as a 1:1 chat, the only participant will be the person who tries to communicate with the person in violation. In the case of group chats, Teams removes the participants whose presence violates policy and leaves the other participants.

**Joining meetings:** When an account tries to join a meeting, Teams blocks them if other participants blocked by the policy are in the meeting.

**Screen sharing:** Any time someone shares their screen in a meeting, Teams checks for policy violations and won't allow the sharing if a violation is detected.

**VOIP calling:** When someone calls another person or a group, Teams checks the call to make sure that it doesn't violate a policy and terminates the call if a violation is detected.

Users blocked from communicating with each other won't be able to see blocked accounts in organization charts, activity feed, suggested contacts, people cards, and call and chat contacts.

In addition, whenever Information Barrier policies or user accounts are updated, the Information Barrier Policy Evaluation Service evaluates existing chats and team memberships to ensure that no policy violation results from the update. Existing 1:1 chats become read-only if Teams finds that participants are blocked by policy (Figure 17-21). The contents of the chat prior to the detection of the violation remain untouched.

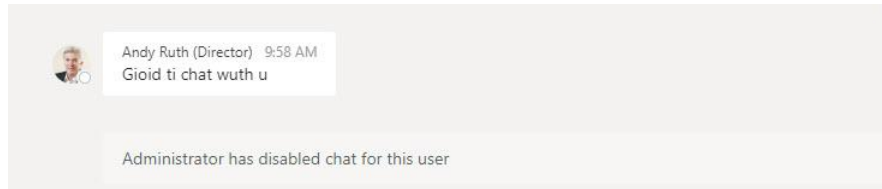


Figure 17-21: An information barrier policy blocks a user from an existing personal chat

## Adding Group Members Using Administrative Interfaces

Although Teams can control the addition of members and how members communicate, it can't control how administrators add members to Groups through administrative interfaces. For instance, you can add a prohibited team member using the Teams admin center or the Microsoft 365 admin center. Likewise, you can do the same by running the *Add-UnifiedGroupLinks*, *Add-TeamUser*, or *Add-MgGroupMember* cmdlets. This is the reason why the Information Barrier Processor exists. It serves as a backstop to eliminate violations introduced through interfaces that do not know (yet) about Information Barrier policies. Here's what happens:

- Someone makes a change to the membership of a group using an administrative interface.
- The change synchronizes to Azure AD.
- The Teams – Azure AD synchronization process detects the violation(s) when it runs to update Teams with changes in Azure AD and blocks the changes. Teams accepts changes that don't violate Information Barrier policies. Clients subsequently download the changes to their local cache.
- The Information Barrier Processor runs later and removes the members in violation from the group.

The Information Barrier Processor serves as a backstop to eliminate violations introduced through interfaces that do not action the settings in Information Barrier policies.

## Information Barriers and SharePoint Online and OneDrive for Business

To use information barriers with SharePoint Online, you associate organization segments with sites, and access to the site is thereafter determined by the segments associated with the site and the information barrier policy assigned to user accounts. Before using information barriers for SharePoint Online and OneDrive for Business, you must first enable the capability by setting tenant properties:

```
[PS] C:\> Set-SPOTenant -InformationBarriersSuspension $False
```

Segment assignment happens automatically for any site connected to Teams. SharePoint admins cannot change the assigned segment for teams-connected sites. For sites not associated with Teams, you can assign segments using the SharePoint admin center or PowerShell. With PowerShell, the first step is to connect to the compliance endpoint and run the *Get-OrganizationSegment* cmdlet to return the set of segments defined in the organization:

```
[PS] C:\> Get-OrganizationSegment | Select Name, EXOSegmentId
```

Name	ExoSegmentId
Trading Department	5517962f-c1e2-4a4b-b905-997d99bdc393
Sales Department	352e1fe3-fbee-4980-9208-103ddea4b370
Group General	c0ccc7f3-c44c-4ef6-a847-e1388bbeclc2

Then assign segments with the *Set-SPOSite* cmdlet:

```
[PS] C:\> Set-SPOSite -id https://office365itpros.sharepoint.com/sites/TradingDepartment
-AddInformationSegment 5517962f-c1e2-4a4b-b905-997d99bdc393
```



A background assistant called the [Information barriers compliance assistant](#) runs every 24 hours to monitor group-connected SharePoint Online sites (excluding those connected to Teams) to detect and remove users who violate IB policies from site membership. Make sure that site owners are part of the segments assigned to the site as otherwise, they won't be able to access the site. In addition:

- SharePoint Online disables the option to use Anyone sharing links (if permitted by the tenant).
- The site and its contents can only be shared with people who match the segments assigned to the site.
- New users can only be added to the site if their segment matches one of those assigned to the site.

If a site has no segments assigned, SharePoint only permits users to share based on the information policy assigned to their accounts. For instance, if a user belongs to the Sales segment and policy blocks interaction with Trading, SharePoint won't allow that user to share content with people in the Trading segment.

The same approach applies to assigning segments to OneDrive for Business sites. In this case, because a OneDrive site is for personal storage, it is even more important to ensure that the assigned segment includes the site owner.

## Audit Events for Information Barriers

Details about the processing of Information Barrier policies are captured in several audit events in the audit log, including:

- *New-ExoInformationBarrierSegment*: when a new organization segment is created.
- *RecipientChange*: when an account is updated because of information barrier settings.
- *ApplyIBPolicy*: when an account is evaluated for information barrier policies.

## The Impossibility of Stopping User Communication

Information Barriers won't stop users from communicating with each other. In fact, in practical terms, no software barrier will stop people from communicating. Information Barriers don't stop users from emailing each other if they know each other's SMTP addresses or decide to use personal email accounts. This is the reason why you might want to use communication compliance policies to check email flowing between different parts of the organization or implement ethical firewalls with transport rules, especially if you want to control communication to specific domains.

It's also true that Information Barriers won't stop users from sending messages to VIP mailboxes, which is why mailbox moderation exists. Because all messages stay within a tenant, Information Barriers are more effective in stopping chat and voice communication with Teams, but again, they won't stop someone using a PSTN number to dial up another Teams user for a voice call.

If blocks stop people from communicating through email and Teams, they will find another route to share information, be it WhatsApp, a simple text message, or a surreptitious note scrawled on a scrap of paper. But that's not the point of Information Barriers or why you would want to deploy these policies. Instead, having policies like this in place helps organizations satisfy regulatory requirements in a demonstrable and provable manner. And if people want to do wrong, HR processes should be defined, available, and communicated to enforce company policy in a way that humans understand.

# Chapter 22: Power Platform

*Christina Wheeler*

## In Search of No-Code/Low-Code Automation

The Microsoft Power Platform is a family of services spanning five main products: Power BI, Power Apps, Power Automate, Power Virtual Agents, and Power Pages (Figure 22-1). These services deliver methods to manipulate, surface, automate and analyze data in a low-code/no-code approach for use with Microsoft 365 and [Dynamics 365](#) (Microsoft's cloud-based customer management system). The Power Platform is based on Dataverse which is the underlying data platform (based on SQL Server) for Dynamics 365. It provides a unified data schema so that applications and services can interoperate.

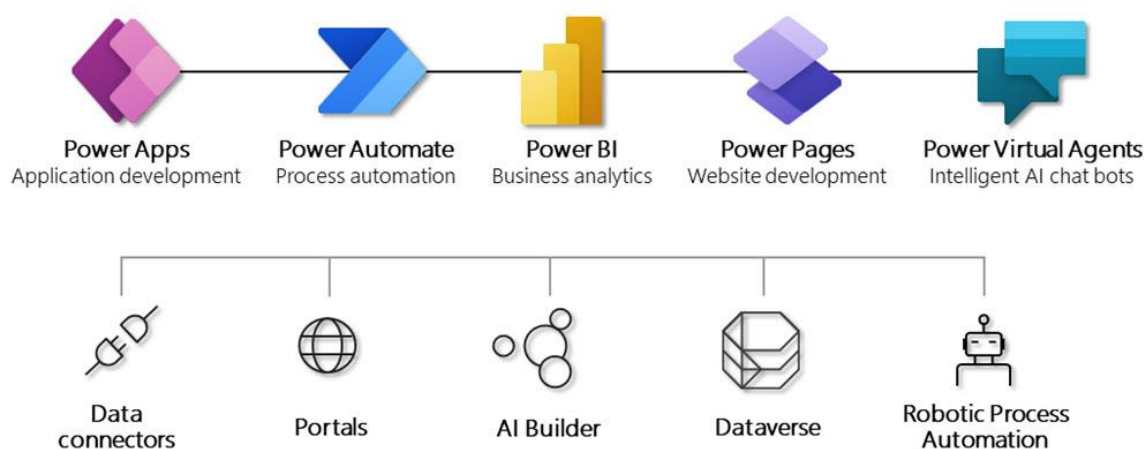


Figure 22-1: Power Platform components

Data Connectors are wrappers around REST APIs (Application Programming Interfaces) to connect other services to interact with the Power Platform components. Currently, there are over 300+ prebuilt connectors available that can connect to Microsoft and non-Microsoft services. These services include SharePoint Online, OneDrive for Business, Twitter, Dropbox, and more. If there's a service you need to communicate with that doesn't have a prebuilt connector, you can create your own custom connector in the Power Platform.

## Microsoft 365 Developer Program

Before we dive into the Power Platform, I want to mention the [Microsoft 365 Developer Program](#). By signing up for the program, you get a free, pre-configured Microsoft 365 tenant (without Windows) that includes 25 E5 licenses. The tenant is renewable every 90 days based on usage. This sandbox tenant is great for learning, building prototypes, or simply testing new features. It includes everything you need for Power Platform development. In addition to the Microsoft 365 Developer Program, you can also sign up for a free [Power Apps Developer Plan](#), which you can add to a free development tenant, to another existing tenant, or can be set up as new. I recommend signing up for both, especially for anyone who is learning and needs an environment to work in separate from production.

# Power Platform Administration

The [Power Platform admin center](#) (Figure 22-2) is the centralized portal for administration tasks for the Power Platform and there are other places in Microsoft 365 where administrators can interact with the Power Platform.

- The [Power Apps Maker portal](#) is where non-administrator accounts access Power Apps.
- The standard Microsoft 365 administrative portal manages aspects that affect Power Platform, such as user accounts, licensing, and auditing.

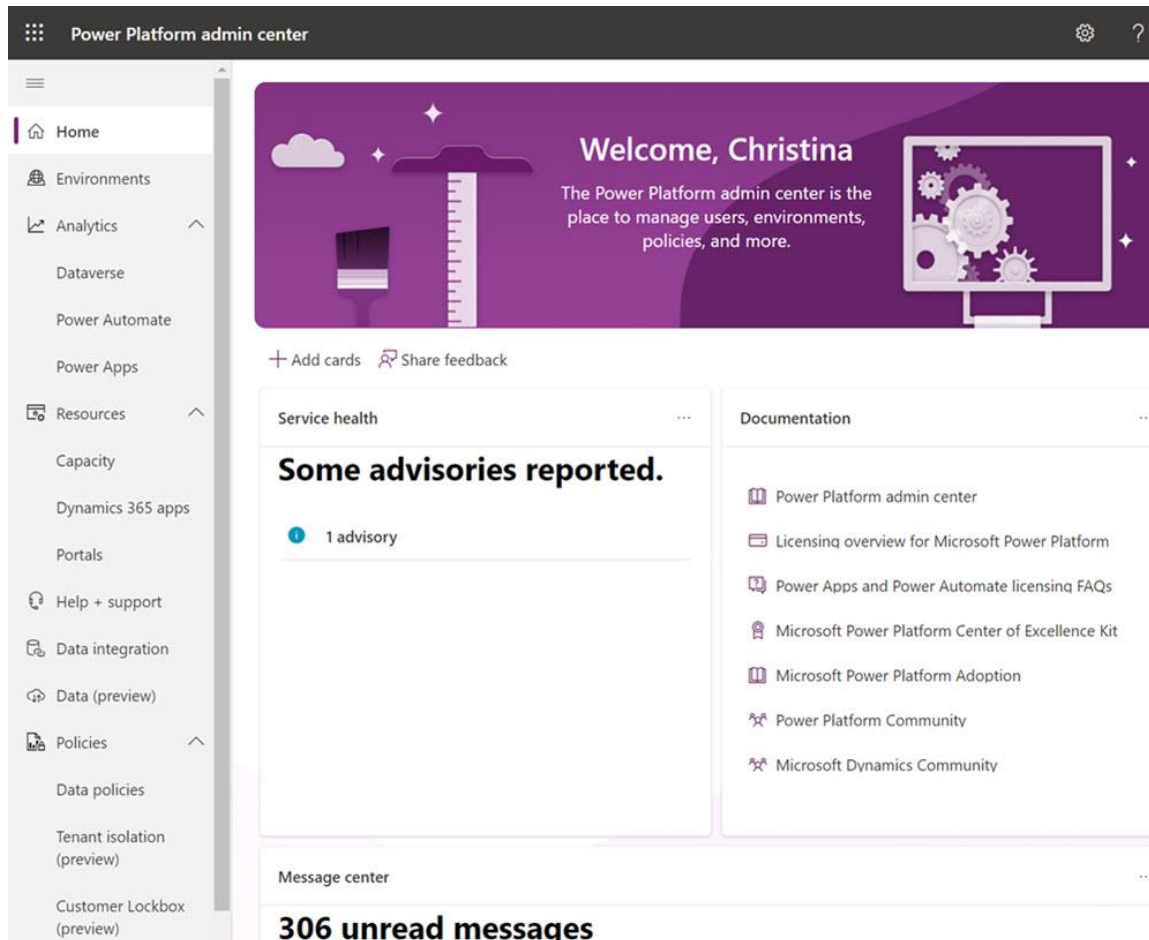


Figure 22-2: The Power Platform admin center

The Power Platform admin center evolves constantly. The capabilities you will find in the current release are:

- **Environments** management. Shows a single list of all your environments, supports sorting on all properties, search across environments, and create/open environments.
- **Analytics** includes detailed views of key metrics for the Dataverse, Power Automate, and Power Apps.
- **Resources** is where you can view and manage the capacity used by Dataverse. It also includes a list of Dynamics 365 apps available to install and configure in the tenant.
- **Help and support**. Provides a list of self-help recommendations and support for Dynamics 365, Power Apps, and Power Automate. Or get real-time support by opening a support ticket.
- The **Data Integration** section is where an administrator can integrate data into Dataverse. It supports integrating data between Dynamics 365 Finance and Operations and Dataverse as well as integrating data into Finance and Operations and Dynamics 365 Sales.

- The gateways (to reach back to on-premises resources and support hybrid integration scenarios) can be managed from the **Data (preview)** section. This includes the ability to set up gateway clusters for load balancing.
- The **Policies** section is an area where new Data Loss Prevention policies can be created and maintained. These policies define what data can be shared with which connector, thus controlling what can be used in flows and apps.

## Environments

If you have a paid plan license for Power Automate or Power Apps, you can use environments to host flows and apps. Simply put, an environment is a space to store and manage the organization's flows, apps, chatbots, connections, custom connectors, and gateways in a geo-located manner. This means the flows and apps that live within an environment are available in the region where the environment is located. Environments are used to separate resources that have different roles, security requirements, or target audiences. This is especially important if you need to apply the EU General Data Protection Regulation (GDPR) in a tenant used worldwide.

When a flow or app is created in a specific environment, it is routed to all data centers in that geographic region which could deliver some performance benefits. You can also create several environments in one physical region to create boundaries inside one geographic area because you can assign different users and policies to environments created in the same region.

Be aware that if the environment is deleted, all resources (flows, apps, connections, custom connectors, and gateways) within the environment are removed. Likewise, because the environment is the isolation boundary for all resources that reside in a specific environment, you can never refer to resources across environments. For example, you cannot create a custom connector in one environment and then create a flow that uses that custom connector belonging to a different environment. However, it is possible to export a flow or app from one environment and import it into a different one.

The type of environments you can create are:

- **Sandbox** - Non-production environment is an isolated environment intended for development and testing separate from production. Requires 1GB of available database capacity.
- **Production** - Intended to use as a permanent environment for production apps, flows, and chatbots. Requires 1GB of available database capacity.
- **Trial (subscription-based)** - Intended for use to develop larger, multiuser solutions and proofs-of-concept. This type of trial has an end date that can be extended.
- **Trial** - This is a standard trial intended for trying new features. Limited to one user and expire after 30 days. The trial plan allows users to extend the trial period two more times after the initial 30 days have ended for a full trial period of up to 90 days. After the full trial has ended Microsoft disables and deletes the environment.

You can create an environment with or without a Dataverse database. To learn more about how to create an environment in the Power Platform admin center, read [this article](#).

### Considerations when creating a new environment

Anytime you create a sandbox or production environment with a Dataverse database, you will have the option to add Dynamics 365 apps (such as Dynamics 365 Sales and Field Services) during the creation process. If you want to add Dynamics 365 apps now or know you will want to add them later, you must select the **Enable Dynamics 365 apps** option during the creation process. Currently, if you don't select this option your database will be provisioned without the support for adding Dynamics 365 apps. This setting cannot be changed later.

Below are some additional points to consider when creating a new environment:

- **When to create an environment with a database** - When creating a new sandbox or production environment with a Dataverse database, you will have the option to add Dynamics 365 apps (such as Dynamics 365 Sales and Field Service) during the creation process by choosing **Enable Dynamics 365 apps**. Currently, this option is only available when creating a new environment and cannot be changed later. Therefore, if you want the ability to install Dynamics 365 apps (even at a later time) then you must enable this on creation. If you do not select Enable Dynamics 365 apps at the time of the database provisioning, you will not be able to make this change later nor will you be able to install Dynamics 365 apps in the environment.
- **When to create an environment without a database** - If you don't need Dynamics 365 apps nor Dataverse and are just creating Power Apps or Power Automate with other data sources, then you can create an environment without a Dataverse database.
- **Dynamics 365 apps and trial environments** - Trial environments currently do not support the ability to enable Dynamics 365 apps. To create a Dynamics 365 trial, you must do so [through this link](#).

## Environment permissions

Power Platform environments without a Dataverse database will automatically have two built-in roles that provide access to permissions within an environment. In the first role, **Environment Admin** can perform all administrative actions in an environment which includes the ability to:

- Add/remove users or groups from the Environment Admin or Environment Maker role.
- Provision Dataverse database for the environment.
- View/Manage all resources created within the environment.
- Set data loss prevention policies.

The second role, **Environment Maker** role can create resources within an environment which includes apps, connections, custom connectors, gateways, and Power Automate flows. Environment Makers can distribute apps they build within an environment to other users within the organization by sharing the app with individual users, security groups, or all users within the organization.

Environments with a Dataverse database will have a **System Administrator** role instead of an **Environment Admin** role for full administrator privileges. For users who need to make apps that connect to a Dataverse database, you must assign them the **System Customizer** role in addition to the **Environment Maker** role.

## Default Environment and Environment Details

Regardless of your license type, there is always a default environment that gets created when the tenant is provisioned. This environment is shared by all users and any licensed user can create flows and apps in it. You cannot delete the default environment, nor can you backup and restore the default environment. The default environment is limited to 32 GB of storage capacity and starts as 1 GB. If you need to store data of more than 32 GB, you should create a production environment instead of using the default environment.

You can view the properties of the default environment or any other environment you created via the Environments page. To open, choose the environment you want to view by selecting it in the Environments list to open its properties page (Figure 22-3). The first panel exposes **Details**. The default environment will not automatically have a Dataverse database. You will know a database hasn't been created if you see the **Add database** section. When a database is not created, the **Security group** section (defining who has access to the environment) offers two options for the environment roles: Environment Admin and Environment Maker. An Environment Maker can create new flows or apps, data connections, and gateways. An Environment Admin can manage data loss prevention policies, add users to environments, and assign admin/maker privileges. An admin can assign users to both roles, allowing them to use the new environment. To create a database for your environment when one does not exist, click on **+ Add database** and proceed to select the desired options, and click the **Add** button.

Environments &gt; Practical 365 Demo

**Details** [See all](#) [Edit](#)

<b>Type</b> Trial (28 days remaining)	<b>Region</b> United States
<b>Refresh cadence</b> Frequent	<b>Purpose</b> Not specified
<b>Environment ID</b> d48a170f-a1e9-ea33-82c9-1a52464c5f0f	

**Add database** [+ Add database](#)

Collect, store, and share your data. Create database for this environment.  
[Learn more about databases.](#)

**Access**

- Environment admin**  
[See all](#)
- Environment maker**  
[See all](#)

**Resources**

- Power Apps
- Flows

Figure 22-3: Environment properties of a Trial environment with Add database section

If you have a database created, you will see the **Version** and **Updates** sections (Figure 22-4). Notice the **Access** section has changed from **Environment admin** and **Environment maker** to **Security Roles, Teams, Users, and S2S Apps**. The third panel (**Resources**) shows Flows, Power Apps, Portals, and Dynamics 365 apps configured in the environment.

**Details** [See all](#) [Edit](#)

<b>Environment URL</b> practical365.crm.dynamics.com	<b>State</b> Ready
<b>Region</b> United States	<b>Refresh cadence</b> Frequent
<b>Type</b> Default	<b>Security group</b> Not assigned
<b>Organization ID</b> ddd8d352-e114-4a4b-a2ee-e27a1613dc9f	<b>Environment ID</b> Default-9e163937-b682-4f9f-95ec-76cbf73d9ea0

**Auditing** [Manage](#)

**Auditing enabled**  
No

**Saving new logs for**  
Forever

**Free up capacity**  
[Delete logs](#)

**Version**

**Database version**  
9.2.22054.000150

**Updates**

**2022 release wave 1**  
On  
[See what's new in the release](#)

**Access**

- Security roles**  
[See all](#)
- Teams**  
[See all](#)
- Users**  
[See all](#)
- S2S Apps**  
[See all](#)

**Resources**

- Dynamics 365
- Portals
- Power Apps
- Flows

Figure 22-4: Environment properties with Version and Updates sections

## Storage

In April 2019, Microsoft introduced Dataverse (*formerly the Common Data Service or CDS*) capacity storage optimized for the Power Platform and Dynamics 365. Dataverse is a secure, cloud-based storage used by the Power Platform. The data itself is stored in a combination of Azure SQL Server and Azure blob storage hosted in the Azure Cloud. Dataverse contains a standard set of tables, columns, and rows (for example, Accounts, Contacts, and various activity types) that are extensible, allowing tenants to add additional data columns to fit business needs. Tables can create relationships with each other, and Business Rules can be created to make fields required, hide fields, and set default values. Each environment can have zero or one Dataverse database.

The Power Platform environment needs at least 1 GB of storage space. Each tenant contains one default environment which includes 1 GB of Dataverse space from the beginning. Additional Environments require extra storage that can be purchased in 1 GB increments. Purchase of extra storage is necessary if the default Environment uses more than the default amount. If the storage capacity is exceeded, the administrator receives a notification about over-capacity usage, and a message is also displayed in the admin center. When a tenant exceeds the storage quota, it is not possible to create a new environment, copy an environment, or restore an environment until additional capacity is purchased.

To review the current storage, sign into the Power Platform admin center and choose **Resources** → **Capacity**. The *Summary* tab provides a tenant-level view of how the organization is using storage capacity. The *Storage capacity* tab gives similar information as found in the Summary tab, but with an environment-level view of where the organization uses capacity. This tab provides statistical details about actual database usage, top database tables and their growth over time, actual file usage, top file tables and their growth over time, actual log usage, and top tables and their growth over time.

## Capacity Limits

Tenants with Power Apps or Power Automate licenses automatically get default capacity. Microsoft increases the capacity for each Power Apps and Power Automate license. Table 22-1 shows the Power Apps and Power Automate capacity limits.

<b>Power Apps capacity Limits</b>	<b>Per license entitlement (Power Apps per-app plan)</b>	<b>Per license entitlement (Power Apps per-user plan)</b>
Dataverse Database Capacity	+ 50 MB	+ 250 MB
Dataverse Log Capacity	+ 0	+ 0
Dataverse File Capacity	+ 400 MB	+ 2 GB
<b>Power Automate capacity Limits</b>	<b>+ Per user</b>	<b>+ Per flow</b>
Dataverse Database Capacity	+ 50 MB	+ 50 MB
Dataverse Log Capacity	+ 0	+ 0
Dataverse File Capacity	+ 200 MB	+ 200 MB

Table 22-1: Power Apps and Power Automate default capacity

## Accessing On-premises Data with the Data Gateway

Often the need to build solutions using Power Platform apps includes a need to integrate data from on-premises data repositories. Traditionally, these repositories include SQL Server databases, documents, and other data stored in SharePoint document libraries and other data stores. To provide connectivity with data stores within a corporate network, the Data Gateway is available (see [this document](#) for software and hardware requirements). The Data Gateway provides one gateway for connecting on-premises data to

multiple cloud services. These cloud services include Azure Analysis Services, Azure Logic Apps, Power BI, Power Apps, and Power Automate. The gateway creates an HTTPS proxy connection that can be accessed securely from these cloud services. It is the same component used by Microsoft's Power BI service. The supported data sources for Data Gateway are:

- SQL Server
- SharePoint
- Oracle
- Informix
- Filesystem (files and file shares)
- DB2

The gateway uses outbound connections, meaning connections are always initiated from the internal network towards the cloud – not the opposite way. Traffic is pushed through Azure Service Bus, which is hidden within Data Gateway's overall architecture. A good practice is to download the Data Gateway executable and deploy it on at least two servers that can access your required data sources. This way Data Gateway provides a highly available setup, as load balancing between the two machines happens automatically. Note that you cannot install Data Gateway on an AD Domain Controller. The machines hosting the Data Gateway can have shared services, but they should be servers, not workstations.

To install Data Gateway, simply download the executable and run through its setup. When you configure the Data Gateway, make sure to store the recovery key in a safe place. You will need it for recovery and when you expand your setup with multiple Data Gateway installations. You also can specify your region, which should be the one where flows are executed to avoid excess latency. If a proxy or a firewall restricts the outbound network connectivity, make certain the ports and destinations are allowed for Data Gateway to operate correctly [following the instructions from Microsoft](#). Once your Data Gateway is successfully installed, you can add data sources in the gateway to make your on-premises data accessible to Power BI, Power Apps, and Power Automate.

## The Microsoft Power Platform CoE Starter Kit

The Power Platform Center of Excellence (CoE) Starter Kit is a collection of tools designed to help drive innovation, improvements, and strategies to adopt and support the Microsoft Power Platform (specifically Power Apps, Power Automate, and Power Virtual Agents). The kit provides tooling and automation (through apps, Power BI analytic dashboards, *and flows*) to help teams build monitoring and automation needed to support a CoE.

The CoE Starter Kit is an open-source initiative that can be downloaded from [GitHub](#). Microsoft provides full [installation and configuration](#) instructions. To use the kit, the administrator accounts must have Premium licenses and Microsoft Power Platform service admin, global tenant admin, or Dynamics 365 service admin rights in the Office 365 tenant.

The Core Components of the Starter Kit (one of the four parts of the Kit) only contains assets relevant to admins. The kit collects many types of information about apps, flows, flow action details, custom connectors, connectors, model-driven apps, shared-with information, chatbots, and logs, before displaying its findings in Power BI dashboards.

## Power Automate

Power Automate is a powerful, built-in workflow engine in the Power Platform that can be used to streamline repetitive tasks and paperless processes. There is a growing demand for retrieving, filtering, accessing, manipulating, and reusing data from Microsoft 365 together with external systems and on-premises data sources without the need for complex custom solutions that can be error-prone and time-consuming to



create. Power Automate aims to provide a user-friendly, powerful engine for automating tasks and integrating data between systems. Power Automate is a fully cloud-based solution of the Power Platform family of services, offering two components: a flow engine and Robotic Process Automation (RPA) capabilities. It allows users to build self-contained workflows that execute when triggered by request (such as when a specific event occurs in a remote system, manually, or as scheduled tasks).

Some might feel that Power Automate is a modern take on the classic Windows Task Scheduler and PowerShell scripting, but with a more polished interface. While that is not an incorrect assumption because both schedule tasks and run scripts for automation, the two are not alike. Task Scheduler is typically used to run backend tasks that users rarely, if ever, need to see, and PowerShell is an administrator tool that requires a solid technical background, while Power Automate is more geared towards delivering value to users – not just administrators or backend services.

Microsoft initially launched Power Automate in 2016 as Microsoft Flow, together with Power Apps. The two share certain capabilities but do not depend on each other. Power Apps is a design and implementation tool for building line of business apps that run in the cloud. Power Automate has matured enormously and now has wide support for integration with many data sources through its use of connectors. It is also one of the few Microsoft services that have remained tightly focused on process automation and integration since its inception.

With Power Automate, the person who builds automation solutions does not need to be a developer. Power Automate does not allow you to call external code unless that code is accessible through an API available to Power Automate through an Open API and REST specified interface. Power Automate is relatively easy for end-users to build small solutions and ramp up to more advanced integration solutions as their skills mature.

Typical solutions power users and administrators might build with Power Automate include:

- Monitor a blog site through its RSS feed for new posts. Once a new post is made (and the RSS feed is updated), send a notification to a mobile phone with links to the new content.
- Free up licenses for disabled accounts.
- Save attachments sent to an unmonitored email address as a file in a SharePoint Online document library.
- Start an approval process when a new request for purchase is made through a SharePoint site.
- Post a weekly status update to a Teams channel detailing changes made to one or more documents.
- Send an email alert to IT administrators when one or more services fail in Microsoft Azure.
- Track changes to an Excel file and create approval requests for certain changes.
- Create to-do items in Microsoft To-Do or Planner for incoming helpdesk requests.

There are three types of flows you can create with the Power Automate service: *Cloud flows*, *Desktop flows*, and *Business Process flows*. A cloud flow is the most used type of flow which enables users to trigger processes automatically based on a trigger, manually, or via a schedule. The three flow types you can create with a cloud flow are:

- **Automated Flows** – An automated flow triggers when a condition is met.
- **Instant Flows** – An Instant Flow is a flow that allows you to start a flow with a click of a button.
- **Scheduled Flows** – A Scheduled Flow is for tasks that need to be automated on a schedule.

In 2020, Microsoft introduced Robotic Process Automation (RPA) capabilities to Power Automate under the name UI Flows (now called Desktop Flows). This feature expands on the automation capabilities of Power Automate, making it possible for bots to perform UI-based tasks centered on a pre-recorded set of UI interactions.

There are two kinds of RPA abilities in Power Automate: attended and unattended. The first means that a bot can record a set of UI actions and then play them back in real-time when necessary. Human interaction must

initiate the process. Unattended RPA requires no human interaction because it is meant for back-end processes.

## Licensing Power Automate

The licensing scheme for the Power Platform is complicated. If you need to manage several types of licenses in a tenant, you might need to consult Microsoft to understand the available options. To help, Microsoft keeps the [Licensing overview for Microsoft Power Platform](#) document updated with detailed information about pricing and licensing different components. The [Power Apps for Microsoft 365 plan](#) is included in most enterprise plans. This license allows users to run a limited number of flows per month with fewer capabilities than found in other plans. For example, it excludes access to premium connectors (all Azure and SQL connectors, HTML connectors, custom connectors, etc.).

Additional premium plans for Power Automate exist:

- **Power Automate per-user plan.** Users with this license can create unlimited cloud flows. Pricing is \$15 per user/month.
- **Power Automate per user plan with attended RPA.** Users with this license can create unlimited cloud flows plus automate legacy applications through attended RPAs and AI. This plan includes 5,000 AI builder credits per month.
- **Power Automate per-flow plan.** Flows implemented with this plan can be run by anyone in the organization. Pricing starts at \$100 per month with a minimum of 5 cloud flows to total \$500 per month.

The premium plans place no limits on trigger frequency or flow runs and instead use a daily capacity limit (5,000 daily API requests). Organizations that need additional capacity for heavy usage scenarios can buy add-on capacity and assign it to specific users or processes.

The UI Flows functionality in Power Automate Attended RPA capabilities costs \$40 per user/month, and Unattended RPA is an add-on costing \$150 per bot/month (on top of the \$40 per/month for the Power Automate plan with Attended RPA).

Unless you block the capability, tenant users can make self-service purchases for Power Platform licenses. The steps needed to prevent this are in Chapter 4. entitlement limitations based on licensing plan.

## Guest Access to Power Automate

You can assign actions in a flow to guest accounts. To use this feature, you must create the guest account and assign it a plan that contains a Power Automate license before assigning it to a flow.

The flow is created as normal and then assigned to the guest account, which receives an email notification to act in the flow. The link in the email routes them to the host tenant. To return to their home tenant, the guest can go to the My Flows page and click the **Go to your org's default environment** link in the notification message or sign out and re-sign to their tenant.

Except for SharePoint flows and Approval flows, connections to any other Azure AD-backed services are created in the context of the user's home tenant, not the tenant that the user is logged into. Currently, there are some limitations for guest access in flows:

- Notification to mobile devices is unsupported for guest users because the Power Automate app for iOS or Android cannot sign users in as guests.
- People picker experiences do not work for guest users because they are not able to query and enumerate members of the tenant but typing the full email address in the selection box works.
- There is no way in the user interface to switch between tenants. Also, the only way to access the flow portal as a guest is through a link in an email.

## Building Automated Solutions with Power Automate

To start to build flows with Power Automate, navigate to <https://powerautomate.microsoft.com> or simply click the Power Automate icon from the app menu. You can create a flow from a pre-defined template or scratch. From the Power Automate front page, click **My flows** in the top navigation bar. This page lists all the existing flows linked to your account with tools to enable or disable, share, and edit each one individually.

### Building Power Automate flows from Scratch

To build a manually triggered, instant cloud flow, follow these steps:

1. From the **My flows** page, click **+ New flow - Instant cloud flow**. You will be prompted with a dialog to define the Flow name and a list of triggers to select from.
2. For this example, set the flow name to **Create calendar event**. Select **Manually trigger a flow** and click **Create**.
3. The Power Automate designer will open with your newly created flow. Here is where you can start building out the actual logic for the flow. Add actions, conditions, and other logical building blocks by clicking **+ New step**. A list of operations will be available for you to choose from (Figure 22-5).

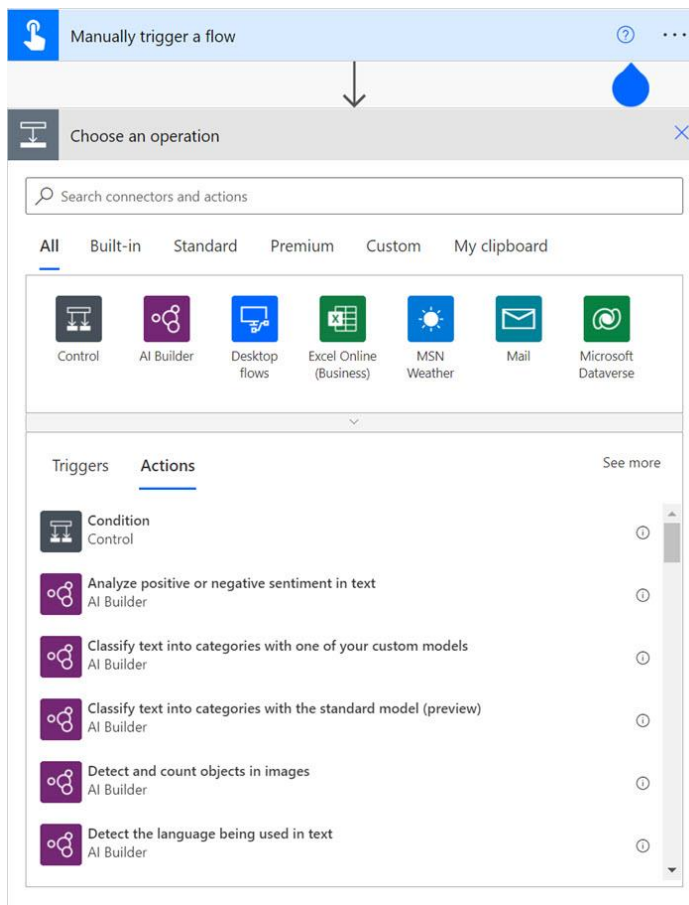


Figure 22-5: Adding a new step for a flow

4. When you scroll down the list of available actions (which numbers are in the hundreds), you can see the vast scope of functionality available for inclusion in a flow. You can connect to external databases, services in Azure, and data stored in on-premises servers. In this case, search for and select **Office 365 Outlook – Create event (V4)** for the first action.
5. Your first action is now transformed into a box with the Outlook icon and some mandatory fields to fill out. Power Automate might prompt you to authenticate as it might not have the authentication token for your account. If prompted, authenticate with an account that has access to the Outlook

calendar you wish to work with. Once authenticated, you can verify the available accounts by clicking the three little dots in the corner of the action.

You can see other options, such as adding a comment for this action (for providing better readability for other editors of this flow), and under *Settings* some specific settings for this action. You can also remove this action should you not need it later. Verify that you are using the desired credentials (in the section **Data - Connections**) to work with the calendar.

6. Next, select the drop-down menu for **Calendar id** and select your calendar. Typically, it is named **Calendar** in English-based interfaces. The name of your calendar might vary if you are using Power Automate with a localized interface. You might also have other calendars visible here, such as those belonging to other users who have granted permission for you to access their calendars. Power Automate also includes an expression language that allows us to specify dynamic values and combine or retrieve data from other actions within the same flow. This way you could dynamically choose a calendar based on data you received earlier in the flow or based on actions the flow is instructed to do.
7. For **Subject**, we can add a fixed string or use the expression language. This would be ideal if you needed to have a date and time or some other dynamic value as part of the event's subject. To keep this example simple type in the following text in the **Subject** field:

```
[Flow] Test event
```

8. For **Start time**, type a date and a time in the future – perhaps the next day. Note the time format, which needs to be in ISO 8601 format, so the 'T' must appear literally to show the beginning of the time element. An example value is:

```
2021-11-01T09:00:00
```

For **End time**, add a similar date and time but a bit later than the start time:

```
2021-11-01T10:00:00
```

In the example, the event starts on the 1<sup>st</sup> of November 2021 at 9 am and ends on the 1<sup>st</sup> of November 2021 at 10 am. Define the Time Zone as well.

You might be wondering what benefits a flow brings if it only supports fixed dates and times for new calendar events. As early said, Power Automate has an expression language that allows us to substitute, generate, and dynamically insert desired values, such as a dynamic date that adds an event 1 day from today. You can access these expressions on the side menu under **Expression**. This provides you with different collections, logical functions, string functions, and similar programming concepts you might need later.

9. Under **Show advanced options** there are many more ways to fine-tune calendar events. For this example, only the mandatory fields (*calendar id, subject, start and end time, and Time zone*) are needed. The event creation action should now look like Figure 22-6

Figure 22-6: Creating the calendar event through Power Automate

10. To change the name of the flow, in the top left corner, click on the original name and replace it with the desired name. Save the flow by clicking **Save** in the top right corner.
11. After saving the flow, you can test it immediately. To test, click on **Test** and then select **Manually** and click **Test**.
12. In the confirmation run flow dialog click **Run flow** to run the flow. After a few seconds, the browser opens directly the flow runs view showing the latest entry in the list.
13. Both the trigger and the action items are visible with a small green checkmark in each corner to indicate the flow worked correctly. Expand the actions to verify the input and output data and troubleshoot for any errors.
14. Now open the calendar where the flow created the event and verify the event is present.
15. At first glance, the calendar event might appear to be an incorrect time. Power Automate uses UTC timestamps while the calendar's owner might configure their calendar to display events in a different time zone. The date and time for the event are correct.
16. Click next on **My flows** within the Power Automate designer browser tab. Find your flow under **Cloud flows**. Click on the more options button (three vertical points) of your flow and click **Turn off** to disable the flow. This will prevent the flow from running every 10 minutes and adding multiple overlapping calendar events.

## Building Power Automate Flows from Templates

Templates are predefined flows you can use in many different scenarios. Via a Power Automate template, you can quickly create new flow instances that only require access to the services defined in the template and additional configuration settings. Microsoft offers several templates to help create a flow in a matter of minutes. You can use the templates as a starting point to implement more refined flows and/or to see how some flow options work. To see the available templates, click on **Templates** in the left navigation bar. You can choose from several hundred examples such as **Create a task in Planner based on Office 365 Outlook calendar event** and **Create a daily summary of Planner Tasks by Bucket**.

## Sharing Flows

If you want to share a flow with a colleague or group in the same tenant, click on the **Share** icon to open the settings page for adding users and groups. Anyone you choose to share your flow with becomes an owner of the flow. They can edit and modify the flow and will also receive access to any connections within the flow. If you chose to authenticate with an Exchange Online mailbox to send emails or create events, those connections become embedded connections when you choose to share the flow with others.

## Building and Running Power Automate Flows in Excel

Excel also supports Power Automate, allowing power users to create new flows within a spreadsheet. This brings the capability to trigger automation from data in Excel and is another way to automate processing. To use Power Automate in Excel, use the **Get add-ins** button in Excel to add the Power Automate Add-in or get it from the [Microsoft Apps Store](#). After installing the add-in, a new Power Automate button in the **Data - Automation** tab or **Insert - Add-Ins - My Add-ins** tab allows the user to select data from a spreadsheet such as a row.

Before Power Automate can interact with data in an Excel file, it must be accessible to Power Automate. You can store the spreadsheet in OneDrive for Business or a SharePoint Online document library and then connect the file to Power Automate.

## Using Data Loss Prevention Policies

It is not a good idea to have your flows retrieve data from a sensitive internal location (such as a SharePoint list that contains salary information) and share the data publicly using something like the Google Drive connector. This is where Data Loss Prevention (DLP) helps. DLP allows you to separate which services can allocate data to the tenant and/or with external sources. [This script](#) creates a report of the flows in use within a tenant and tells you what connectors each flow uses. It is helpful to know what connectors you might need to restrict.

DLP policies allow administrators to set restrictions on usage for Power Automate data handling and data capture. Typically, organizations do not wish to capture or store sensitive information or private data unless it is properly handled and secured. As users build their solutions with plenty of freedom, a reasonable need exists for administrators to set certain boundaries through DLP policies.

To create a new Data Loss Prevention policy, select **Policies > Data policies** in the Power Platform admin center and then click on **+ New policy**. When you create a new policy, you must first give it a **Policy name**, set the **Connectors**, and then set the **Scope**.

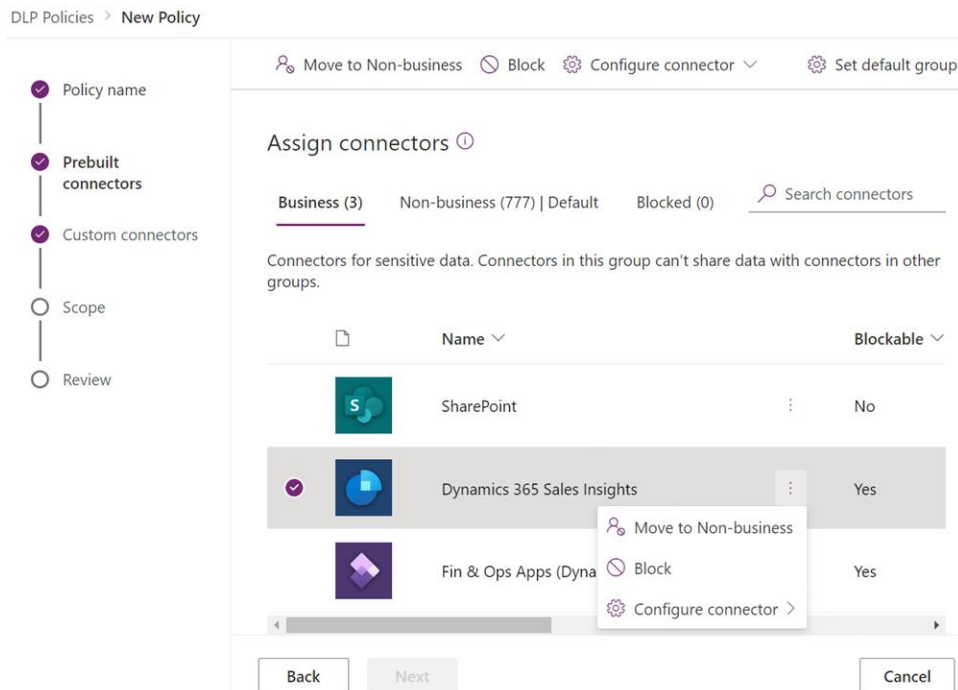


Figure 22-7: Defining data group connectors

**Connectors** are where you define the policy restrictions. In this view, you can define which Power Automate connectors can share business data and which are blocked. You can also set which of these is the default –

can access business data only, or no business data allowed. This model then limits any flows that users create to share data between either group. DLP policies enforce rules which connectors can be used together by classifying the connectors as **Business** or **Non-business**. By default, all services are placed into the "Non-business" classification group. You can freely add and remove connectors from either group simply by clicking on the [...] menu (Figure 22-7) and choosing **Move to Business** or **Move to Non-business**.

If you put a connector in the Business group, data cannot be shared between services in different Data groups and can only be shared with other connectors you added to the same group. One Data group must be designated as the default group. Initially, the "No business data allowed" group is the default group. An administrator can change the default data group using the ellipse button at the top right of each group. New services added to Power Automate will be placed in the designated default group. It is recommended to keep the "No business data allowed" as the default group and manually add services into the Business data after evaluation of the impact of allowing business data to share information with the new service.

**Scope** is where you define the environment's scope. The define scope option defines which **Environments** this policy applies to. An environment is simply an isolation boundary for DLP and data locality. If you have not created any environments yet, you can create a DLP policy that applies to all environments regardless of how many you have in place or plan on creating in the future. You can find extended information about Environments at the start of this chapter.

After you name and save the policy, it is activated automatically for the environment(s) you selected.

Users building flows that violate a DLP policy can create their flows, but Power Automate will put the flows into a suspended state. If a user creates a flow that retrieves data from SharePoint and pushes it to Salesforce (as per the example above), Power Automate will automatically suspend the flow.

When a user flow is suspended, it is not possible to create exceptions within the policy to authorize the use of or more selected flows. If a valid business reason exists to share business data between groups, you can consider editing the DLP policy or asking the user to edit the flow to comply with the DLP policy. Otherwise, the flow will remain irremediably suspended.

## Building Advanced Solutions with Power Automate

Building solutions with Power Automate is straightforward, but things get more complex when flows need to call other flows and when flows must make multiple decisions during its run. It is always best practice to start building a new flow by planning it first:

- How will the flow trigger?
- How many runs do you anticipate for the flow each month?
- Is the flow self-contained or will it call upon other services or third-party APIs?
- What will happen if the flow fails?
- How to log flow activities?

You can always modify flow definitions later. It is also possible to export a flow in a .zip file to keep it as a timestamped copy. This way you can rest assured that you have a copy of your flow before making extensive changes.

Several tools exist to troubleshoot flows. The **Run history** for each flow expands to show each connector input and output data and provides possible exceptions or raw errors the flow encountered. This is the best way to understand why any flow fails or why a single run failed among multiple successful runs. You can also test flows by editing a single flow and clicking on the Test button in the upper right corner. This puts the flow in a test run view and shows a single flow execution in real-time to make debugging easier (Figure 22-8):

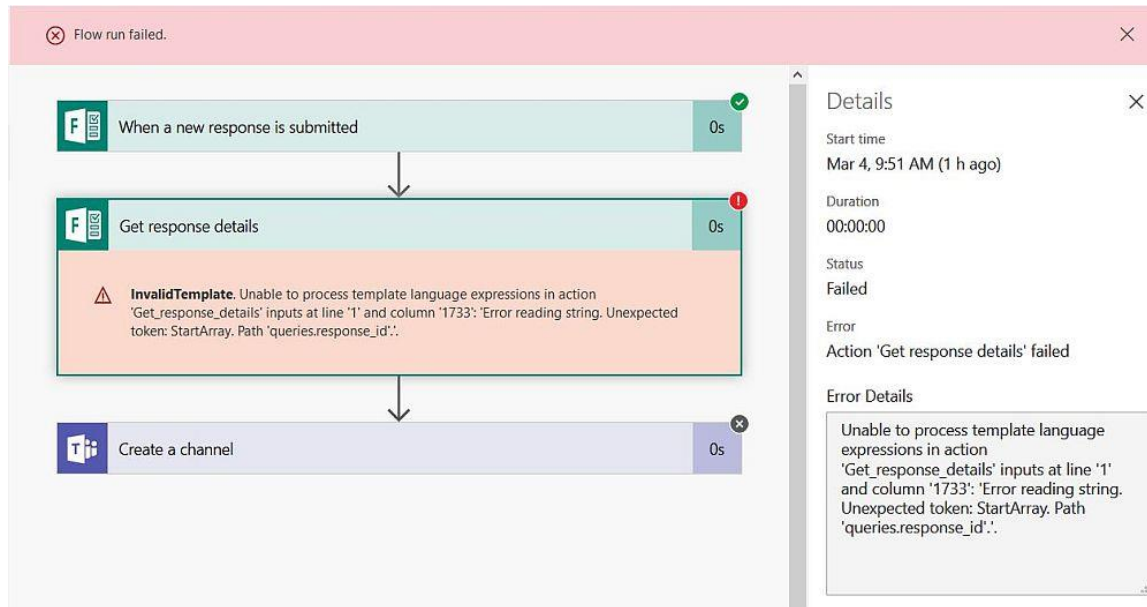


Figure 22-8: Raw output from a flow with execution errors

Flow checker is a tool that can be activated in the upper right corner while editing any flow. It performs an analysis and provides recommendations should the flow have any detected issues.

Certain connectors and triggers also expose settings to allow to further fine-tune how a flow activates and uses these building blocks. For example, when you use the Send Email connector to send email through Exchange Online, you have the option to specify a timeout value and a retry policy.

Some connectors also expose a **Run after** trigger. This can be configured to run if the previous task failed, succeeded, timed out, or skipped. This gives a way to provide further debugging, by sending out data to another API or system if a certain condition was met. If the flow makes a call to a third-party API but it times out, you can alert another system that the API is not answering to avoid further issues in subsequent flow runs.

Over time any account ends up with some flows that are not actively needed or used. Instead of deleting them, disabling allows them to retain their run histories and configuration.

## Extended Examples of Power Automate flows

[Lee Ford's blog](#) explains how to use a Form to allow users to request the creation of a new team and a mixture of Power Automate and Graph API calls to provision the team.

Another example is how to use [Flow to set values for the Asset ID property of documents](#) in SharePoint libraries. The Asset ID property is used for event-based retention processing (see Chapter 18). Like many uses of Power Automate, this example shows how to automate a repetitive task that humans find boring.

See also the article [Combining Microsoft Graph and Flow for Better Administration](#) which explains an example of advanced flows applied to tenant administration.

### API Support for Power Automate

Power Automate supports a Web API to work with flows programmatically. The API is REST-based and supports the ability to list, create, update flows, get all users with whom a flow is shared, share/unshare a flow, and export/import flows.



## Power Automate Desktop and RPA Tools

Power Automate Desktop (PAD) provides RPA (Robotic Process, Automation) capabilities to the Power Platform. RPA attempts to automatically replicate the actions of a user, such as mouse movements, clicks, and keyboard entries, for playback to automate repetitive tasks. RPA tools are often used to build business applications on the fly and for integration with legacy applications.

For a long time, people created "Macros", based on Visual Basic for Applications (VBA) for use in Excel, PowerPoint, and Word to replace repetitive series of keyboard and mouse actions. Macros are potentially dangerous because they have been used as vehicles for malware. PAD flows can be seen as an advanced version of Macros that are more secure because the code is not embedded in Office documents. The resulting flows can be an automation solution for a multitude of processes, not only for Word, Excel, or PowerPoint.

At Ignite 2020, Microsoft announced that PAD is [free to all Windows 10](#) users. Although the Power Automate Desktop app itself is free, the RPA system used to create and save flows to the cloud requires a paid subscription to the Power Platform ([a free trial with limited functionality is available](#)). Power Automate with RPA add-on licensing is only required if you want to publish desktop flows to the cloud.

Power Automate Desktop is already included in Windows 11. If you're running Windows 10, the Power Automate Desktop installer can be [downloaded directly from Microsoft](#).

## Power Apps

Power Apps is a service designed to allow users to build custom business applications without the knowledge of app development or any coding language. Power Apps is essentially a container that allows making apps that can be used across different platforms. Power Apps is designed to work with several different Microsoft data sources such as SharePoint lists and libraries, OneDrive, Excel, Dynamics, and external databases. Some examples of non-Microsoft sources include Dropbox, Box, Twitter, and Facebook.

Historically, the development of apps has involved creating different versions for each operating system they need to run on (one for iOS, one for Android, one for Windows). This essentially triples the development work, triples the support costs, and increases the development resources needed to create business apps. With Power Apps, all the apps run through the app, which takes care of the differences between the operating systems and just allows them to run the apps. There is also a web version of Power Apps, with the same concept but running through any modern web browser instead of a mobile app. Power Apps can also be embedded in other applications, such as SharePoint and Teams.

Typical solutions that power users and administrators might build with Power Apps include:

- Create user interfaces to interact with back-end office solutions to make administrative tasks widely available: hardware and software requests from users, support forms, open new services, and assign licenses.
- Forms to interact with SharePoint Lists and Libraries in a simplified way.
- Requests for provisioning of SharePoint Sites and Site Collections.
- Forms to request guest accounts and other types of administrator tasks.

Power Apps can be easily combined with Power Automate, in such a way that Power Apps works as the user interface and the back-end processes are implemented with Power Automate.

## Licensing Power Apps

The [licensing for Power Apps](#) is like Power Automate where Office 365 users get a non-premium license allowing them to create Apps for Office 365. For Power Apps, there are three additional premium plans:

- **Power Apps per user plan.** It allows users to run an unlimited number of apps, without any feature restrictions. Pricing is \$20 user/month and includes 500 AI Builder service credits per month.
- **Power Apps per app plan.** Allows users to run applications for a specific business scenario based on the full capabilities of PowerApps. Pricing is \$5 user/app/month and includes 250 AI Builder service credits per month. Requires access to the Microsoft 365 admin center with global administrator or billing administrator roles.
- **Pay-as-you-go-plan.** Allows users to run applications without [Power Apps licenses via Azure subscription](#). You only pay for the number of users who used an app in a given month. Pricing is \$10 per active user/app/month. This plan does require an active Azure subscription. This pricing model provides a way for organizations to try out Power Apps without a commitment. It offers the ability to enable additional usage in scenarios where subscription plans are not economical such as in cases where an app is only needed for occasional use for a larger base.

Administrators who need to have access to admin views in Power Automate, Power Apps, and Dataverse do not need a license. Further administration, such as specific Sales and Marketing modules for Dynamics 365 and the use of custom connectors, requires a purchased license.

## Enabling Power Apps for Users

The [Power Platform admin center](#) is the central hub for the administration of Power Apps, Power Automate, and Dataverse. Microsoft enables Power Apps by default for all users in the tenant with E3 or E5 licenses. You need to pay attention to how or when Power Apps is used and the business case(s) to use it. Some organizations disable Power Apps upfront for all users if they do not see an immediate need for the service. Take into consideration that Power Apps is a powerful tool that can improve the acceptance and usability of Office, but users do need to have the necessary training to understand how to use it. See the earlier section about enabling Power Automate to learn how to disable Power Apps for the entire tenant.

Users need a license to access Power Apps. From the Microsoft 365 admin center, navigate to **Users > Active users** and select a user. On the user detail page, select the tab **Licenses and apps**. The list with all product licenses available for the user shows. Expand the **Apps** section, search for **Power Apps for Office 365** and verify it is enabled. If a user account is assigned multiple plans, you only need to enable Power Automate from one plan.

## Building Power Apps

Two types of Power Apps can be created: Canvas apps and Model-driven apps.

- **Model-driven apps** use the Dataverse to configure forms, business rules, and process flows. You create a Model-driven app from the Power Apps site. You start with a data model, building up from the shape of core business data and processes to model forms, views, and other components. Model-driven apps automatically generate a user interface that is responsive across devices. Model-driven apps are used mainly by Dynamics 365.
- **Canvas apps** give the flexibility to arrange the user experience and interface the way you want it. You can start to build Canvas apps from Microsoft tools where the data lives, such as a SharePoint list or a Power BI dashboard, or from the Power Apps site.

Model-driven apps use Dataverse and Canvas apps have the option to use Dataverse as the data source.

**Note:** "Portal" apps can be created from the Power Apps dashboard, but they are not really Power Apps. Portals are internet portals based on Dataverse tables.

It is possible to merge the experiences of canvas and model-driven apps, making hybrid apps where the generated user interface from a Model-driven app can be modified and extended using the Canvas apps possibilities and power.

## Creating a Canvas App

As with Power Automate, there is a web designer for Power Apps. Open [the site](#) or select the Power Apps icon from the app menu. From the site's main page, you can create new apps from scratch or start with a predefined template. Templates help you learn how to build Power Apps and the kind of scenarios to consider using Power Apps. Some applications created using templates are suitable (with some degree of adaption) for production in an enterprise environment.

To see the available templates, click **+ Create** from left navigation menu and scroll down to the **Start from template** section. To follow our example, select **Service Desk**. This application tracks service requests, assignments, and job status, prioritizes jobs, adds notes, and tracks tasks by the assigned technician. The app opens with sample data and is a ready-to-use application deployable to help service desks in any company. Assign the application a name and select to use it in tablet or phone format.

When Power Apps designer studio opens, it shows three panels (Figure 22-9). The left side defaults to the Tree view panel displaying a list of screens in the application. In the center is the canvas with the main app screen. The panel to configure properties and rules for the components on the selected page is on the right.

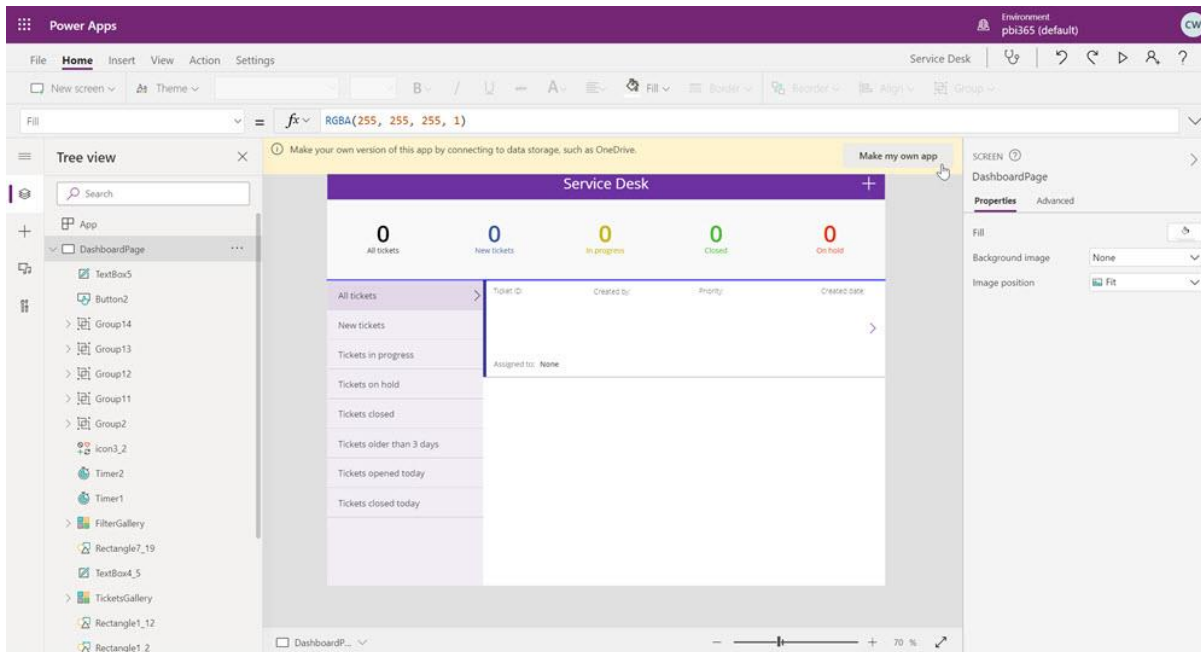


Figure 22-9: Power Apps designer studio

Each screen is composed of controls and there is a plethora of them. Click on the **Insert** tab at the top of the screen and you will see the controls grouped in sections. For example, the **Text** section contains the **Label**, **Text input**, **HTML text**, **Rich text editor**, and **Pen input** controls. There is also a **New screen** button that allows creating new screens for the application with options to start with predefined templates (such as **List**, **Email**, **Tutorial**) or start from **Blank**.

When a control is selected from the left panel of the graphical designer, its properties panel is automatically presented on the panel on the right side. This panel (Figure 22-10) allows the user to configure characteristics of the control (its **Properties**), such as its position in the canvas, its visibility, size, etc. The **Advanced** settings, allow you to configure advanced settings for the app and controls.

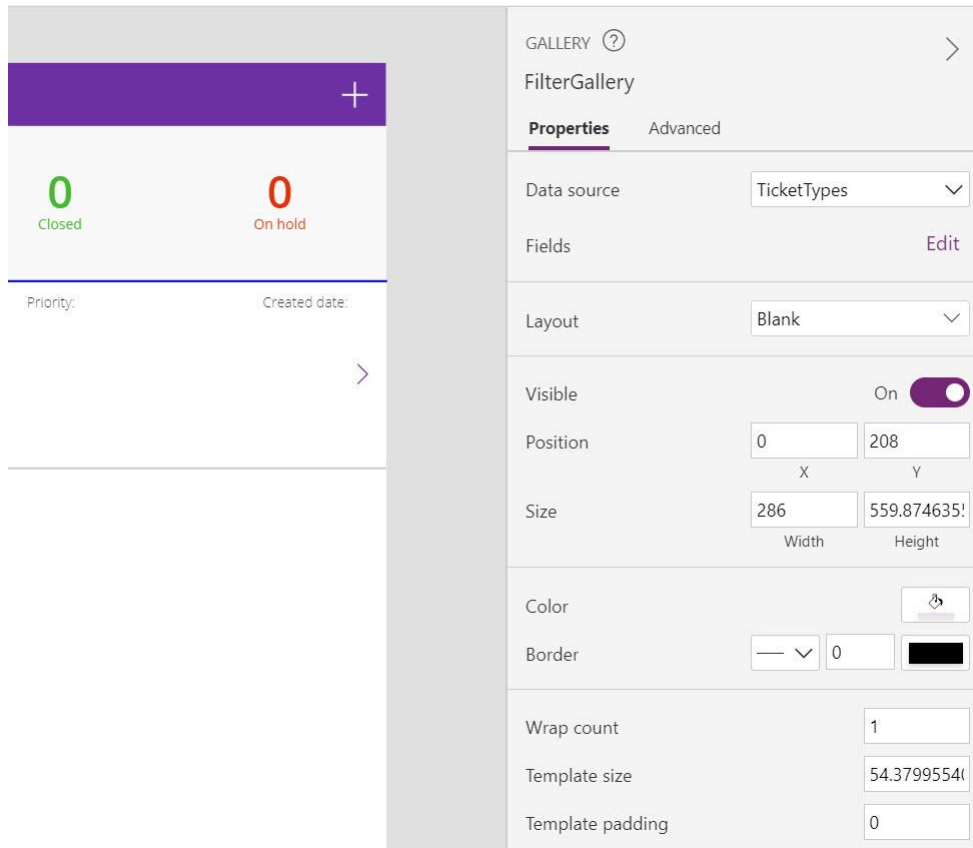


Figure 22-10: The Properties panel with the Advanced link for a selected control

The top-left menu has links to **Insert** controls in the canvas, **View** some essential components that are not reachable from the other windows, and go back to the Advanced properties panel (**Action**).

The **View** menu is important because it gives access to the components that make the application work:

- The **Data sources** are the repositories for the data that the application uses. Several sources can be used, such as Excel, Power Automate, Outlook, OneDrive, and SharePoint. Other services from Azure as the Service Bus and Azure Active Directory; and multiple external sources, such as GitHub, Box, and Dropbox.
- The **Media** menu allows images, videos, and audio to be used in the application.
- **Collections** can be used to store data that users can manage in an app. They are a group of items that can be created manually, or you can put a data source such as a SharePoint list into a collection.
- **Variables** are used to save values through the lifespan of the application, defining its type (text, date, etc.) and initial value.

The top-right menu contains two important buttons: The **App Checker** checks that there are no bugs in the application before it is published. The **Preview** button runs the application for testing before publishing. Use the **Preview** button to run the application (Figure 22-11):

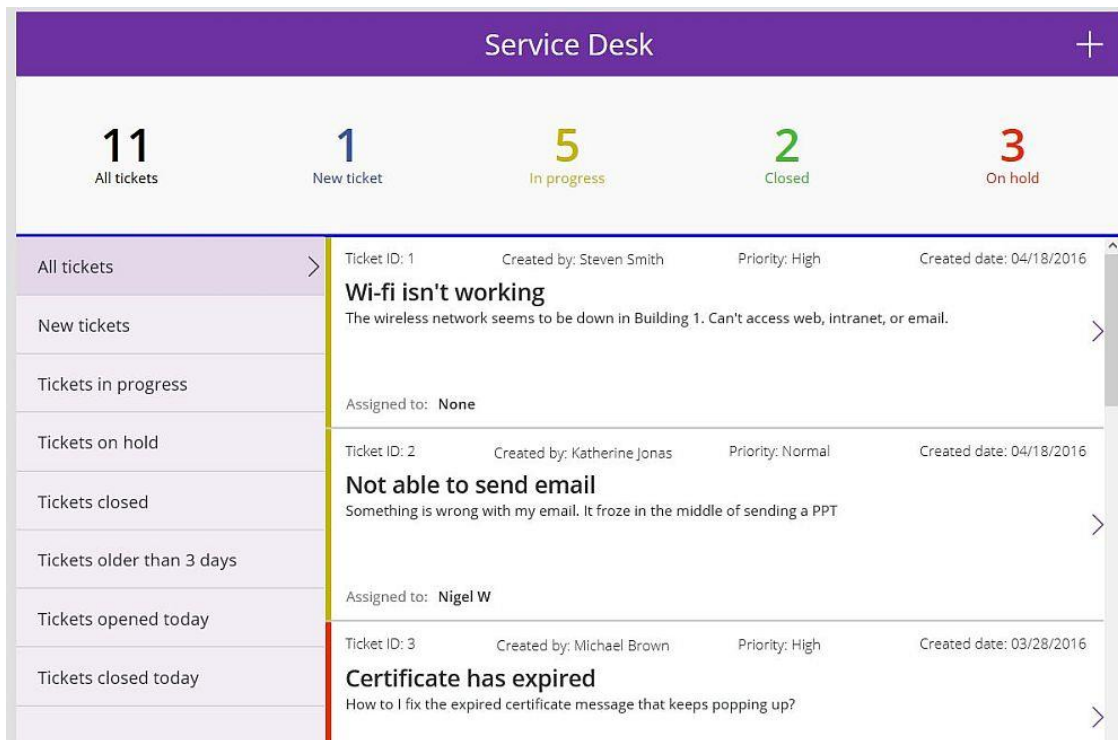


Figure 22-11: The Power Apps application running

The application shows a classification of the tickets saved in the data sources. The application is interactive and fully functional: click on any of the tickets, and the details page opens, allowing the user to modify the state, add comments, etc. New tickets are added using the + button in the upper-right corner. Tickets can be closed, reopened, modified, set on hold, etc., and the totals are shown at the top of the application.

When you are satisfied with the application, go to **File > Save**, define the final name for the application, and use the **Save** button to publish the application and make it available to users.

## Creating a Model-driven App

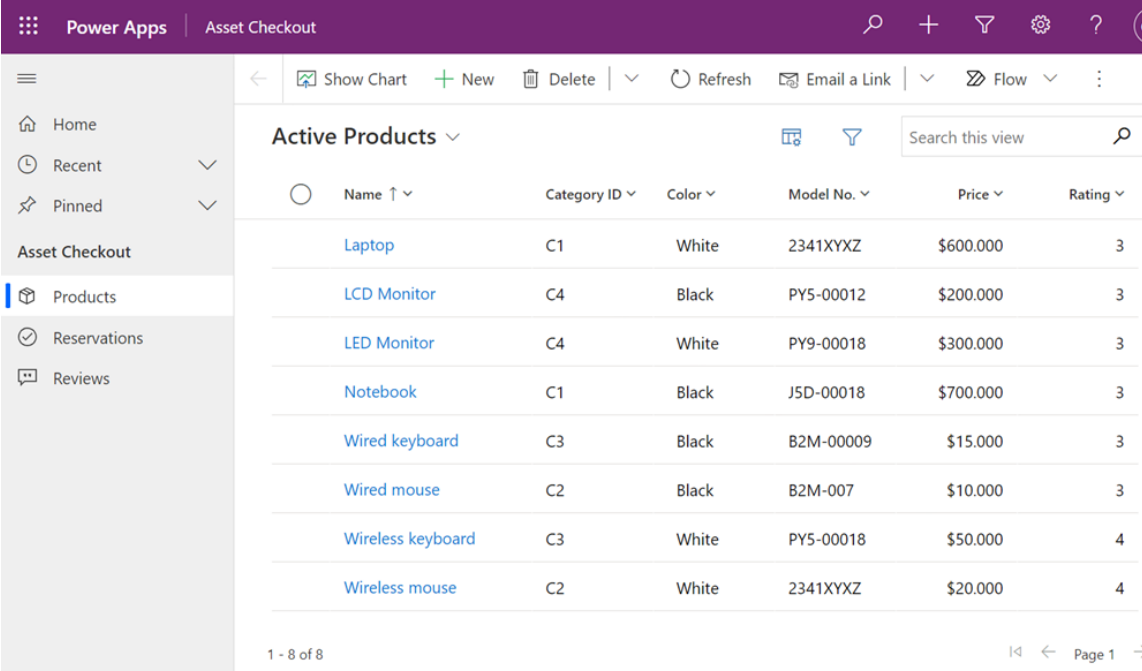
Unlike Canvas apps (where the designer has complete control over app layout), Model-driven app layout is determined by the components added to the app. The creation of Model-driven apps requires three steps:

- Modeling business data - It is important to determine what data the app will need as well as how the data are related to other data. The Model-driven design uses a metadata-driven architecture so designers can customize the application without writing any code.
- Defining business processes - Consistent Business processes are a key aspect of a Model-driven app design. Consistent processes ensure that app users focus on their work and not on remembering to perform a set of manual steps. Processes can be simple or complex and often change over time. To create a process, from the **Model-driven** area select **Settings > Advanced customizations > Open solution explorer**. Then on the left navigation pane in solution explorer select **Processes > New**.
- Composing the app - After modeling data and defining processes you can build the app by selecting and configuring the components needed using the app designer.

**Note:** To work with Model-driven apps you need a database in Dataverse and a separate Environment. To get both, you must have a paid plan account. Remember to assign the license to the users after buying the service from Microsoft.

The apps and user interface for Model-driven apps are based on the Dynamics 365 CE (Customer Engagement) user interface. It's almost always used in combination with Dynamics and seldom-used stand-alone or integrated into other services.

After creating a new Environment, create a new Dataverse database, and be sure you check the option **"Deploy sample apps and data"** at creation time. To start creating a Model-driven app for learning, look to existing templates in the Power Apps gallery. There are three templates for Model-driven apps: Fundraiser, Innovation Challenge, and Asset Checkout. To create a new app based on one of the templates, click on the **Create an app** button and select **Model-driven**. Then assign a name, and description for the app and check **Use existing solution to create the App**. In the following window, you can select an existing solution from a list that includes the three examples. When you are ready to review/modify the copy of the example, use the **Publish** and then the **Play** button to see the app in action, as shown in Figure 22-12.



Name ↑	Category ID	Color	Model No.	Price	Rating
Laptop	C1	White	2341XYZ	\$600.000	3
LCD Monitor	C4	Black	PY5-00012	\$200.000	3
LED Monitor	C4	White	PY9-00018	\$300.000	3
Notebook	C1	Black	J5D-00018	\$700.000	3
Wired keyboard	C3	Black	B2M-00009	\$15.000	3
Wired mouse	C2	Black	B2M-007	\$10.000	3
Wireless keyboard	C3	White	PY5-00018	\$50.000	4
Wireless mouse	C2	White	2341XYZ	\$20.000	4

Figure 22-12: An Asset Checkout model-driven Power App

## Creating a SharePoint Integrated Power Apps app

With Power Apps, you can customize SharePoint Lists and Library forms as was possible with InfoPath before Microsoft deprecated InfoPath. Power Apps has a similar integration with Lists and Libraries and embeds its forms within the SharePoint List or Library.

To start with SharePoint Lists, you can work with an existing modern list or create a new custom list. If you create a new List, ensure it is using the Modern user experience. From the command bar of the list, select **Integrate > PowerApps > Customize forms**. You can then select the fields to show in the list from the available set (Figure 22-13).

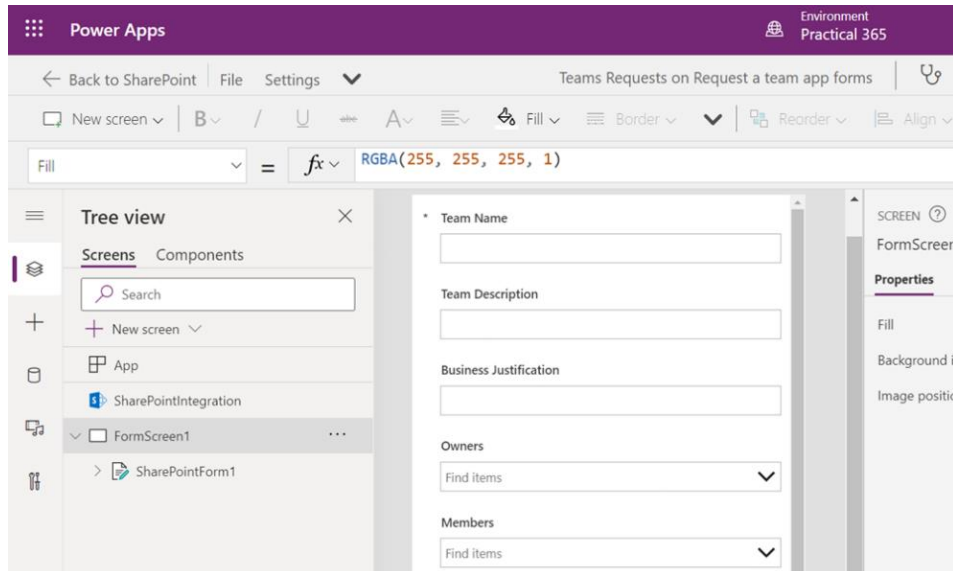


Figure 22-13: A Power Apps Canvas app to replace the default form for a SharePoint List

The complete functionality of Canvas apps is present: you can reorder the fields, change their properties, set conditional formatting to, for example, appear/disappear based on the content of other fields, etc. When you finish, test the app with the **App Checker** button, **Save** and **Publish to SharePoint** the new form.

The Power Apps form replaces the default SharePoint forms automatically. When a user clicks on **+New** to create a new item, the form appears in place of the old form (Figure 22-14):

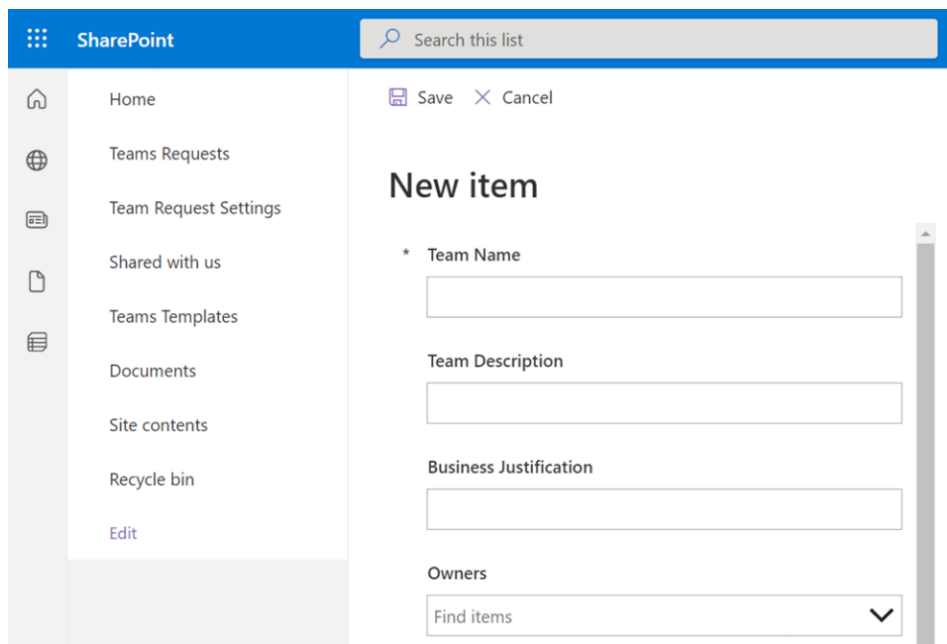


Figure 22-14: A Power Apps Canvas app replacing the default form in a SharePoint List

If changes are needed, the List owner can open the Power App to modify by navigating the same way as when creating a Power App customized form. From the command bar of the list, select **Integrate > PowerApps > Customize forms**.

If you want to change the list back to using the default SharePoint form instead of the Power App customized form, you can re-activate it through the List settings. To re-activate the original form, go to the List in SharePoint, open the **Settings** page by selecting the gear icon near the upper-right corner, and select **List Settings**. On the settings page select the **Form Settings** link under the **General Settings** section. A new page will open. Then select **Use the default SharePoint form** and save the changes. When using the default form,

if you return to the Form Settings page, you will see an extra link **Delete custom form** in the section of Power Apps, under **See versions and usage**. Use this option to remove the Power Apps form.

Take into consideration that if you customize the form for a SharePoint List, the form doesn't appear as an app in Power Apps Studio or Power Apps Mobile. The Power App form can only be opened from the list for which you created it. Anyone with SharePoint permissions (to manage, design, or edit the associated List) can customize the forms. Any guest users that do not have a plan that includes Power Apps will get an error message if they try to access a list form that has been customized using Power Apps.

## Using the Apps

There are several ways to make Power Apps app available to users:

- All Power Apps are mobile-enabled, and Microsoft publishes Power Apps for Android and iPhone. Start the Power Apps phone app and log in with your Microsoft 365 account. A list of apps you have access to appears. Click on any of the apps to open the app and then you can start using it.
- Users can access Power Apps from any modern browser supported by Microsoft 365. When a Power Apps app is published, it gets an identifier that can be used to redirect any browser directly to the app. For example, <https://web.powerapps.com/apps/b0f87da1-1053-4a3f-999f-84a372b1b656>. You can retrieve the identifier and the URL from the app properties (use the details link from the ellipse button in the apps list). The users can bookmark the URL to access the app speedily.
- SharePoint Online has its own Webpart to embed and interact with Power Apps from any modern SharePoint page. Open the SharePoint page where the Power Apps app will be hosted, add the Power Apps Web Part, configure the URL or Identifier of the app, and publish the page.
- Team owners can add Power Apps as a channel tab.

## Moving Power Apps Between Environments

A typical Power Apps administrator activity is moving apps from dev/test to production tenants or between Power Platform Environments. There are two ways to move apps which are via the export/import functionality or by building a Power Apps solution. Canvas apps can be moved using **Export package** under **Apps** and Model-driven apps can be exported using **Export** under **Solutions**. Only the Owner and Co-owner of an app can export a canvas app package. An account must hold the Environment Maker permission to import an app.

To export a Power Apps app, from the main window of the website (<https://web.powerapps.com>) click on the Apps menu on the left side menu to show the list of published apps. Click on the ellipse button ("...") of one app and from the contextual menu select **Export package**.

The **Action** button (Figure 22-15) allows you to select if the packaged solution will be created as a new app or updated as an existing app when imported into a target environment. When you click the **Export** button, Power Apps creates a .zip file and downloads it to the local computer.

To import one package, from the same page with the list of apps use the link **Import package** (menu at the top of the page). A new window opens to upload the .zip file. A second window opens showing the package details; use the **Import** button. If any error occurs, for example, because you try to create an app that already exists, an error message appears indicating the problem found. If the app uses connectors that are not yet configured, the import menu will prompt you to select or create new connections.



## Export package

**Package details**  
Created by Christina Wheeler on 06/27/2022

Name \*  
Space Invaders

Environment  
Christina Wheeler's Environment

Description  
Space Invaders Power App provided by Brian Dang (Power CAT)



**Import setup**

Choose what will happen when each resource is imported to the new environment.

**Create as new**  
This app or flow will be new to the environment when the package is imported.

**Update (default)**  
The app or flow already exists in the environment and will be updated when this package is imported.

**Review Package Content**  
Choose your export options and add comments to provide instruction or add version notes.

NAME	RESOURCE TYPE	IMPORT SETUP	ACTION
Space Invaders	App	Create as new	 

Related resources

NAME	RESOURCE TYPE	IMPORT SETUP	ACTION
No items			

**Export** Cancel

Figure 22-15: Creating an export package for one Power Apps app

## GDPR and Power Apps

The European Union GDPR regulations oblige app developers to manage user personal data in an accountable way. See Chapter 18 for an overview of GDPR and Office 365.

GDPR has two consequences for Power Apps: users can request a data controller (usually a company) to retrieve their data at any moment (“right of data portability”), and they can also require the data controller to remove their data (“right to erasure”) from the system. Power Apps saves personal information in the Environments, Canvas apps, Gateway, Custom Connectors, and Connections.

In general, users can find their information in the Power Apps Portal or using the [App Creator PowerShell cmdlets](#). A user has the right to find and modify their information only. An administrator with Global Administrator or Azure Active Directory Global Administrator rights can do the same for all users. In some cases, a Power Apps paid license plan is required. Microsoft details the actions necessary to [export](#) or [remove](#) personal user data using the Power Apps Portal and/or the PowerShell cmdlets online.

## Power BI

Power BI is a suite of business intelligence (BI) software services, apps, and connectors that work together to provide a way for you to connect, share, and analyse data through reports and dashboards. Power BI is the next-generation business analytics tool comprising several technologies that originated from Excel such as PowerPivot and Power Query. Power BI is packed with collaboration, sharing, and mobile features that are deeply embedded in Microsoft business application solutions. Power BI consists of three main components:

- **Power BI Desktop** – Authoring tool of Power BI used by report designers to access, transform, and model data as well as build the data visualizations on report pages.
- **Power BI Service** – Cloud-based central hub where users can access and interact with dashboards and reports.
- **Power BI Mobile** – Mobile component of Power BI that allows users to interact with published reports using their smartphones and tablets through the Power BI app available on Windows, iOS, and Android devices.

## Licensing Power BI

The licensing options for Power BI include the following:

- **Power BI (free)**. Users can sign up for a free account using a work or school account which provides access to a personal workspace.
- **Power BI Pro**. Users with this standard license can share data, reports, and dashboards with other users who have a Power BI Pro license through shared workspaces. Power BI Pro is included in E5 subscriptions, otherwise the price is \$9.95 per user/month.
- **Power BI Premium**. There are two types of licensing for [Power BI Premium](#) which is capacity-based, and per-user based.
  - **Capacity-based** – This license consists of capacity in the Power BI service exclusively allocated to each organization supported by dedicated hardware fully managed by Microsoft. Capacity can be applied broadly or allocated to assigned workspaces based on the number of users or workload needs with the flexibility to scale up or scale down. Pricing starts at \$4,995 per capacity/month. With this license, report consumers do not require any per-user license however Power BI Pro licensing is needed for report designers who will be designing and publishing content into the Power BI Premium capacity.
  - **Per-User (PPU)** – This license provides the ability for organizations to [license premium features on a per-user](#) basis. Users with this license get Power BI Pro capabilities along with features like paginated reports, AI, and other [premium capacity features](#). Pricing is \$20 user/month.

## Power BI Service

The Microsoft Power BI Services ([app.powerbi.com](http://app.powerbi.com)) is the SaaS part of Power BI. Sometimes referred to as Power BI online, is a cloud-based service where users can view and interact with the reports. Report designers use Power BI Desktop to publish reports to the service and report readers can access the service through shared workspaces. Every user (both free and licensed) receives a personal workspace called *MyWorkspace*. Shared workspaces are available for Power BI Pro or Power BI Premium licensed users.

## Power BI Desktop

Power BI Desktop is a report authoring application used by report designers to build advanced queries, data models, and reports. This application is available to download for free and can be installed on computers running Windows 10 or later. To [get Power BI Desktop](#), you can install it as an app from the Microsoft store or download and install the executable. You do not have to have a Power BI service account to use Power BI Desktop however it is recommended you sign up for the free version if you are not already licensed in Power BI Pro or Premium.

### Using Power BI Desktop

Once you have Power BI Desktop installed you can immediately connect, transform, and build visualizations of your data. Power BI Desktop has three views:

- **Report** – View where you use queries from your data model to build visualizations and arrange them on report pages. In this view, you can add, delete, and manage multiple report pages.
- **Data** – View where you see the data in your report in a data model format and can add measures, create new columns, and manage relationships.
- **Relationships** – View that provides a graphical representation of the relationships established in your data model. In this view, you can add, modify, and manage the relationships.

Power BI Desktop comes with the Power Query Editor. Originating from Excel through PowerPivot, the Power Query Editor is used to connect to one or more data sources. You can shape and transform your data using the Power Query Editor to meet your visualization needs.

While you can create reports directly from the Power BI service, it is recommended you build in Power BI Desktop first and then publish your report to the Power BI service in your individual or shared workspace (Figure 22-16).

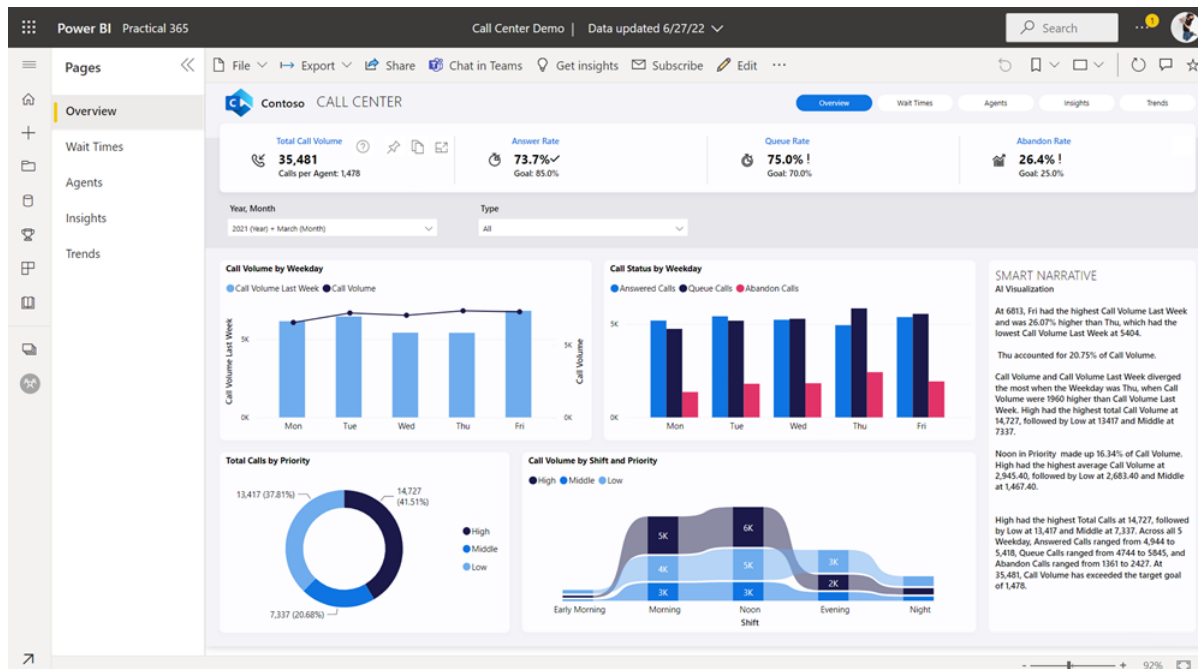


Figure 22-16: Published report displaying in shared workspace called Practical 365

For a getting started guide and tutorials on Power BI Desktop, please visit [Get started with Power BI Desktop](#).

To understand the differences between Power BI Desktop and Power BI Service, please visit [Comparing Power BI Desktop and the Power BI service](#). To learn more about your sharing and collaboration options in Power BI, please visit [Ways to collaborate and share in Power BI](#).

## On-premises Data Gateway

If you need to connect to data within your organization that is not part of the cloud, you can connect using an on-premises data gateway. To learn more about the on-premises data gateway, please visit [this article](#).

## Using Power Apps and Power Automate with Power BI

Power BI is natively a read-only tool, allowing users to see their data but doesn't have built-in capabilities to write back data. Thanks to the integration of Power Apps and Power Automate with Power BI, you can now build Power Apps and Power Automate solutions to provide users the ability to write data back to the source system from within Power BI using the Power Apps and Power Automate visuals.

To learn more, please visit [Power BI data write-back with Power Apps and Power Automate](#).

## Power BI Security

The Power BI Service is built on Azure (Microsoft's cloud computing infrastructure and platform). Power BI uses two primary repositories for storing and managing data that is uploaded from users from Power BI Desktop to the Power BI Services. Data uploaded from users is typically stored in *Azure Blob Storage* and all metadata and system items are stored in *Azure SQL Database*.

For more information on the Power BI architecture, please visit [this article](#) and for more detailed information on Power BI Security, please [read the Power BI Security whitepaper](#).

## Power Pages

At Build 2022, Microsoft announced the evolution of Power Apps portals called Power Pages. Power Pages is the newest member of the Microsoft Power Platform family. It is built on the foundation of Power Apps portals where you can build websites using the new design studio and pre-defined starter templates. Currently, Power Pages is in preview and is available by visiting <https://powerpages.microsoft.com>.

### Design Studio for Power Pages

Power Pages Design Studio is a cloud-based WYSIWYG (what-you-see-is-what-you-get) authoring tool designed for low-code makers to build and style data-centric business sites backed by Dataverse. When you sign up for a trial, an environment will be auto-provisioned for you within your tenant. To build a site simply click on **+ Create a site** to start the creation of your new site.

When first get started with creating a new site, you will be given a list of templates to choose from. Once selected, you will be prompted to provide a site name, web address, and have an option to change the site language. *The site's web address must be unique and will display a message if unavailable and will not let you continue until it's unique.* When ready, click **Done** to proceed to provision the site.

The site will take around 5 minutes to provision and once complete you will be able to see it listed under My sites where you will have the option to preview or edit (Figure 22-17).

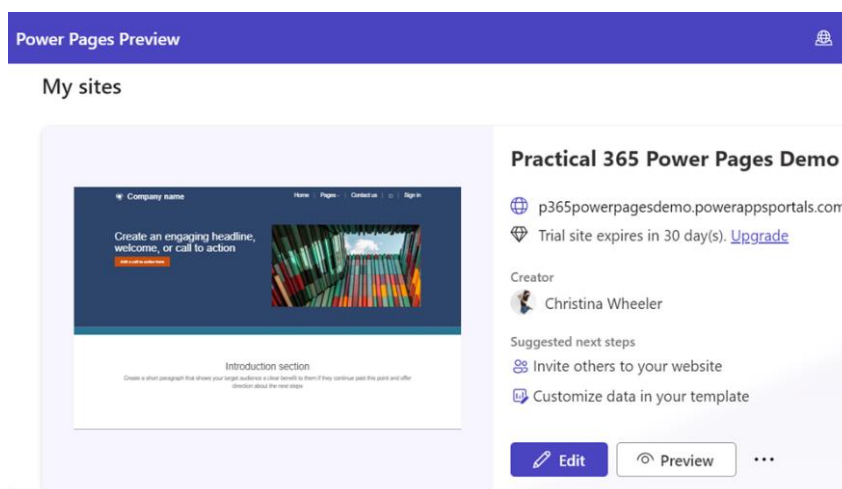


Figure 22-17: Power Pages example website

From here you can click **Edit** to launch design studio. Design studio has four marquee experiences called *workspaces*. The workspaces included are:

- **Pages workspace** - Enables makers to [create and design webpages](#) with no-code and low-code widgets such as lists, forms, text, images, video, and others.

- **Styling workspace** – Enables makers to [apply global site styles](#). There are 13 preset themes available and for each theme, you can customize the color palette, font styles, background color, section margins, and button styles.
- **Data workspace** – Enables makers to easily model, visualize, and manage business data for the site. All data and changes made in the [Data workspace](#) are stored in Dataverse.
- **Set up workspace** – Enables makers and site administrators to [configure site settings](#) such as identity providers for authentication and table permissions.

To dive deeper into Power Pages, please visit [Microsoft Power Pages \(preview\) documentation](#).

To learn more about Microsoft's vision for the future of Power Pages and low-code solutions, please visit [this article](#).

## Power Virtual Agents

"Power Virtual Agents" is the latest service added to the Power Platform. Built on top of Azure's existing Bot Framework Service together with some Artificial Intelligence tools for building bots, Power Virtual Agents enable the development of chatbots with no coding and/or model training requirements. Administrators can take advantage of this service for scenarios such as building a bot to walk a new employee through the onboarding experience.

Power Virtual Agents are created using a guided, no-code graphical interface, which integrates bots with the prebuilt existing Power Platform connectors. Power Automate workflows can be triggered to enable bots to act on behalf of customers. Background processes continuously and automatically monitor the bots to improve their performance using AI and data-driven insights. You can create a bot to be used as a Power Virtual Agents web app (<https://powervirtualagents.microsoft.com>) or as a Power Virtual Agents app in Microsoft Teams (Figure 22-18). If you're new to Power Virtual Agents and want to learn more, I recommend going through the [Quickstart tutorial](#) to create a Power Virtual Agents HR bot in Microsoft Teams.

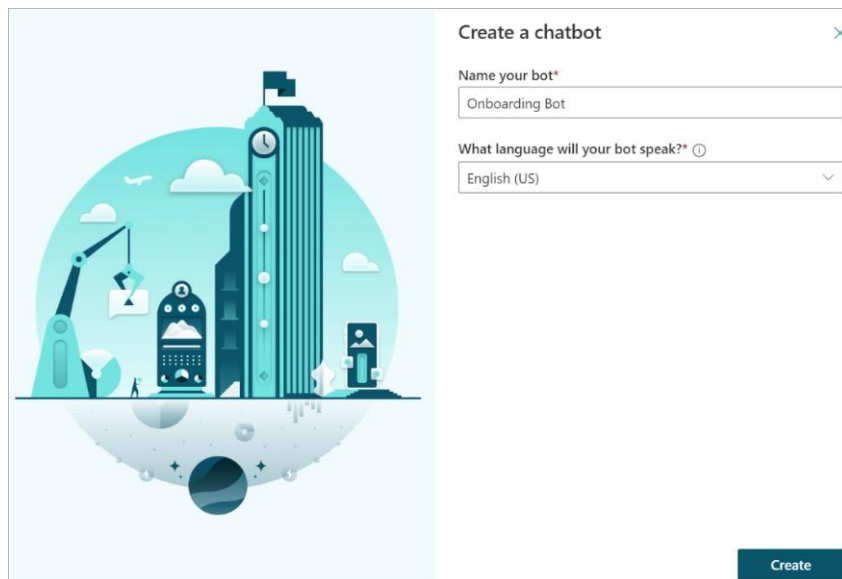


Figure 22-18: Creating a Power Virtual Agents chatbot initial dialog in Teams

## Licensing Power Virtual Agents

To create and manage bots with Power Virtual Agents, you will need to be licensed. The [available licenses for Power Virtual Agents](#) include:

- **Power Virtual Agents per user license.** This license allows you to assign individual licenses to users who need access to create and manage chatbots.
- **Power Virtual Agents per tenant license.** This license allows you to run intelligent chatbots across websites and other channels. You only pay for two-way engagement between users and your chatbots with sessions servicing each end-to-end interaction. This license cannot be assigned to individual users. Price is \$200 per month for 2,000 sessions.

Add-ons available currently are:

- **Power Virtual Agents per tenant (Sessions add-on).** This add-on license allows you to add additional sessions to your Power Virtual Agents plan. Price is \$100 per month for 1,000 sessions and requires a Power Virtual Agent license.

To learn more about Power Virtual agents, please visit [this page](#). Power Virtual Agents' capacity is pooled at the tenant level. To learn more, please visit [Quotas, limits, and configuration values for Power Virtual Agents](#).

## Teams and the Power Platform

Many organizations around the world are using Teams as their central hub for their corporate collaboration. The Power Platform has improved with the integration into Teams. You can leverage the Power Platform and Teams with the ability to do the following:

- Embed a canvas app or a model-driven app as a tab in a team.
- Embed a canvas app or a model-driven app as a personal app.
- Create Power Apps directly in Teams, backed by Dataverse for Teams.
- Create Power Automate flows in Teams.

To embed a Power App as a tab in Teams, you simply click on the + in a team next to the tabs in the channel of the team you want to add it to. Choose Power Apps, then select the Power App you want to embed as a tab in a team (Figure 22-19). For this example, an app called School Transformation Survey is selected to be added as a tab in a team channel.

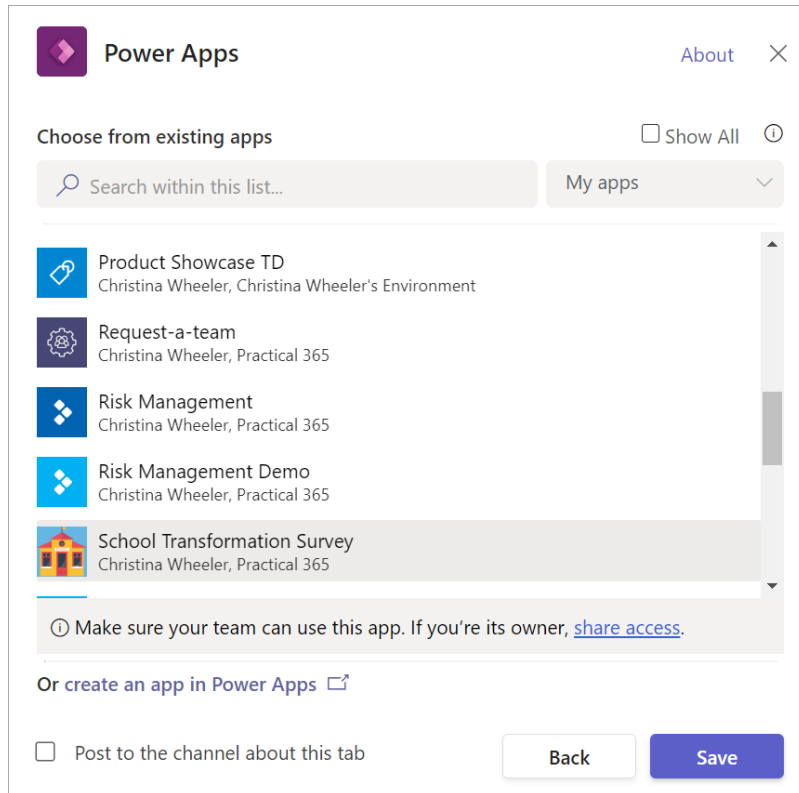


Figure 22-19: Adding a Power App as a tab in a team dialog example

The School Transformation Survey app appears as a tab in a channel (Figure 22-20). This example app was created outside of Teams using Power Apps Studio as a canvas app and is using SharePoint lists as the backend data. You can add any type of canvas app or model-driven app in Teams, or you can create new Power Apps directly inside of Teams using the Power Apps Teams app.

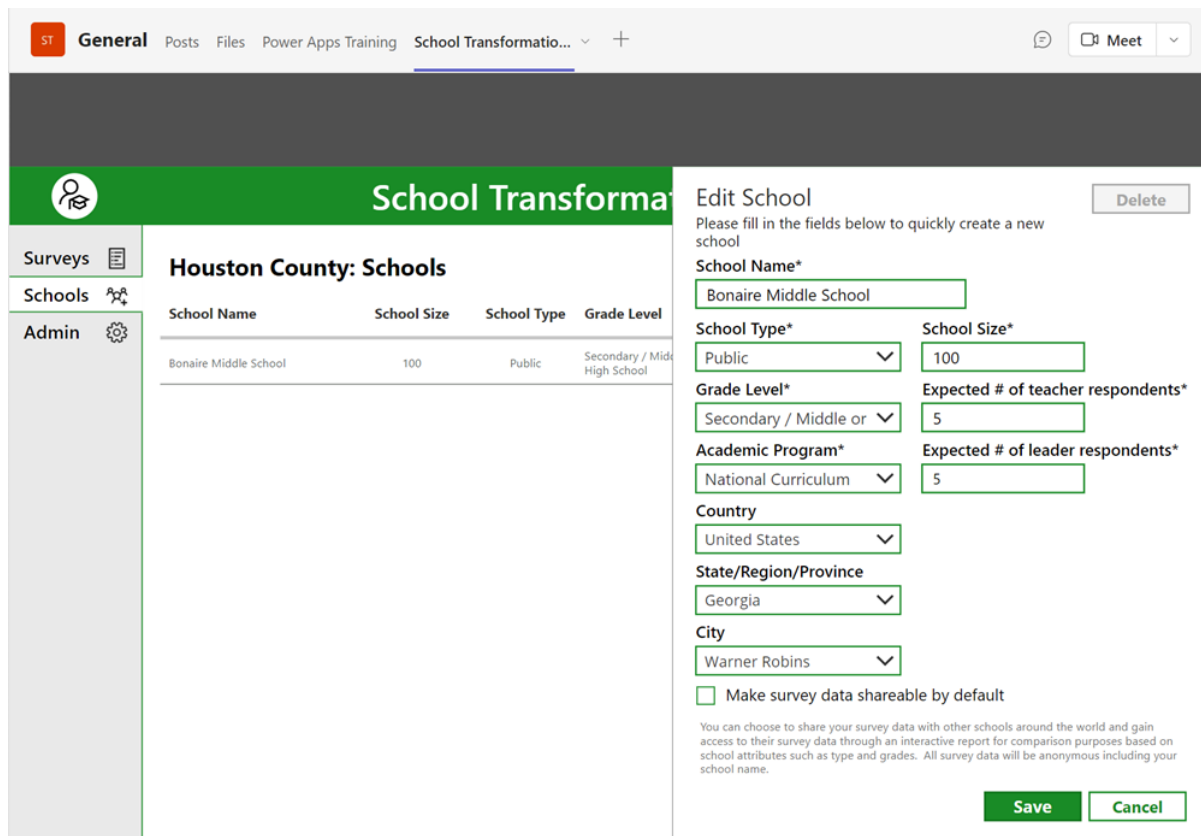


Figure 22-20: Power App embedded as a tab in a team example

For more details on using Power Automate and to see an example of the Request a team Power Platform app template, see the article [Teams + Power Automate: Practical Examples to Automate Tasks](#).

## Dataverse for Teams

When you create a Power App directly in Teams, the backend used is Dataverse. In September 2020, Microsoft introduced Dataverse for Teams which empowers you to create custom apps, bots, and flows in Teams using Power Apps, Power Automate, and Power Virtual Agents. The Dataverse for Teams environment is created automatically when you create a bot or app in Teams or when you install the Power Apps app in Teams and create an app for the first time. Each team can have one Dataverse for Teams environment which is used to store, manage, and share team-specific data, flows, and apps (see the [Dataverse for Teams infographic](#)).

Dataverse for Teams creates a single environment for each team and the capacity is measured with relational, file, and image data. Table 22-4 shows the differences between Dataverse for Teams and Dataverse and Table 22-5 shows Dataverse for Teams service limitations.

<b>Environment lifecycle</b>	<b>Dataverse for Teams</b>	<b>Dataverse</b>
Environments	1 per Team	Unlimited
Maximum size	1 million rows or 2 GB	4 TB or more
Upgrade to Dataverse	Yes	N/A

Table 22-4: Environment lifecycle differences between Dataverse for Teams & Dataverse

<b>Unit</b>	<b>Service limit</b>
Dataverse for Teams environments	5 environments + 1 additional environment for every 20 eligible Microsoft 365 user licenses.
Max storage per tenant (Dataverse for Teams environments)	10 GB + Dataverse for Teams environments × 2 GB (up to a maximum of 19.5 TB).  The 2 GB storage limit can't be extended. If more storage is needed, you can upgrade the environments to Dataverse.

Table 22-5: Environment lifecycle differences between Dataverse for Teams & Dataverse

With the growing amount of data tenants store in Exchange Online, SharePoint Online, OneDrive for Business, and other services, it is natural to consider how best to protect content, especially when the data is confidential or sensitive. Sensitivity Labels and the underlying rights management technology are a good way to protect email and documents, and that is where we go next.



# Chapter 23: Managing Tenants with PowerShell and the Microsoft Graph

*Tony Redmond*

## The Power of Automation

The nature of sharing a massive multi-tenant environment means that Microsoft will never be able to satisfy the unique administrative and automation requirements of every tenant. This chapter discusses how to use PowerShell to manage your tenant. We expect that readers will have some acquaintance with PowerShell, but we'll try to explain how to get work done without plunging too deep into some of the more arcane and complicated aspects of scripting.

Many Microsoft 365 workloads support PowerShell. It's quick and relatively easy to perform an administrative operation in a few lines of script code, which is why we include PowerShell examples to illustrate points in the book. In this chapter, we concentrate on topics that demand more than one or two lines of code, especially for objects that require management on an ongoing basis such as Azure AD accounts, Exchange Online mailboxes, and Microsoft 365 groups. In doing so, we leverage the many examples of how to do different things with PowerShell available in repositories like GitHub, the Microsoft Technical Community, and individual websites. There's no shame in taking code posted by someone and adapting it to solve a problem.

Increasingly, Microsoft 365 applications use Graph APIs to interact with data. We reflect this by covering the basics of running Graph API requests from PowerShell and the Microsoft Graph PowerShell SDK. The combination of easy automation with PowerShell and the power of Graph API requests often delivers the best balance of performance and flexibility, especially in larger tenants.

## Common PowerShell Modules Used with Microsoft 365

Among the modules often used to manage Microsoft 365 tenants are:

- **Azure AD:** Modules are available for both production (AzureAD) and preview (AzureADPreview). The Azure AD Graph API is the foundation of both modules. As noted later, these modules are due for deprecation.
- **Exchange Online:** Manages mailboxes, other mail-enabled objects, and the Exchange Online configuration through the [ExchangeOnlineManagement](#) module.
- **SharePoint Online:** The [Microsoft.Online.Sharepoint.PowerShell](#) module manages sites and other SharePoint administrative objects. If you want access to documents and folders within SharePoint Online sites and OneDrive for Business accounts, use the [SharePoint PnP](#) module.
- **Teams:** The [MicrosoftTeams](#) module manages Teams-specific objects.
- **Microsoft Graph PowerShell SDK:** The cmdlets in this multi-module SDK implement many Microsoft Graph API calls. It is now the preferred option for PowerShell access to Azure AD components.

Full coverage of the Exchange Online, Azure AD, Teams, and Microsoft Graph PowerShell SDK modules is available later. Examples of SharePoint Online management with PowerShell are in the SharePoint chapter.

The modules listed above all support modern authentication, which is what you should use to connect to PowerShell endpoints. This is especially important because of Microsoft's efforts to remove basic authentication from Microsoft 365.

**Other Modules:** Apart from the PowerShell modules created by Microsoft, the PowerShell Gallery holds a variety of third-party modules that might be useful. A good example is the [ImportExcel module](#), which includes cmdlets to interact with Excel workbooks, including the ability to create charts in worksheets.

## Module Installation and Maintenance

To make it easier for ongoing maintenance, install these modules from the PowerShell Gallery rather than by using an MSI file. Do this by running the *Install-Module* cmdlet in a session when signed in as an administrator. For example:

```
[PS] C:\> Install-Module -Name ExchangeOnlineManagement -Force -Scope AllUsers -Repository PSGallery
```

By default, PowerShell considers the PowerShell Gallery to be an untrusted repository (like any other internet source), so PowerShell will prompt you to authorize the installation of modules from the gallery. If you want to make the PowerShell Gallery a trusted resource, connect to PowerShell as an administrator and run the command:

```
[PS] C:\> Set-PSRepository -Name PSGallery -InstallationPolicy Trusted
```

You need to run this command on each workstation where you install PowerShell modules. To see a list of installed modules, use the command:

```
[PS] C:\> Get-InstalledModule
```

Once you've installed a module, you can connect to it using suitable credentials and explore the commands available in the module using the *Get-Command* cmdlet. For example:

```
[PS] C:\> Get-Command -Module Microsoft.Graph
```

Use the *Get-Module* cmdlet to find out the set of modules loaded into a session. Here's how to return the set of loaded modules and their version number:

```
[PS] C:\> Get-Module | Format-Table Name, Version
```

## Updating PowerShell Modules

Microsoft updates PowerShell modules used to manage Office 365 services on an ongoing basis. Although sometimes functionality is dependent on a specific version of a module, you should update the modules used to develop and run PowerShell scripts on an ongoing basis. Most of the modules are now available in the PowerShell Gallery, so you can create a script to check for updates, apply any available updates, and remove old versions from a workstation. The latter step is important to avoid the possibility of errors caused by running obsolete outdated cmdlets.

PowerShellGet, which is the easiest way to work with the Gallery, first appeared in Windows Management Framework (WMF) 5.0, which is essentially PowerShell 5.0. PowerShell 6.0 and later include this functionality too; Windows 10 and Windows Server 2012 R2 or later ship with PowerShell 5.0 or later installed by default, which means that PowerShellGet is already available. For earlier operating systems, you'll either need to upgrade to WMF 5.0 or install PowerShellGet for PowerShell 3.0 or 4.0. Some applications such as Exchange Server are sensitive to changes in the version of WMF installed on the system, so you should not upgrade WMF until you've verified that all your installed software will continue to work.

A script [downloadable from GitHub](#) illustrates how to:

- Define a set of modules to check for updates in the PowerShell Gallery. Tailor the set of modules to meet your needs.
- Use the *Update-Module* cmdlet to check for and apply updates.
- Use the *Get-InstalledModule* cmdlet to check for older versions and the *Uninstall-Module* cmdlet to remove any found.

Ideally, you should check modules for updates monthly. Workload teams update their PowerShell modules at a different pace and it's impossible to predict when a new version will appear.

## Azure Active Directory PowerShell

Many administrative tasks such as user and licensing management, adding and removing domain names, and managing company information in PowerShell can be performed using the Azure AD PowerShell module (referred to by Microsoft as [Azure AD PowerShell for Graph](#)). Microsoft introduced the Azure AD module in 2015 to replace the older Microsoft Online Services ([MSOL module](#)). The MSOL module (created in 2012) is still in use to manage user accounts.

After installing the Azure AD module, connect to your tenant with:

```
[PS] C:\> Connect-AzureAD
```

Microsoft [plans to deprecate the Azure AD and MSOL modules at the end of 2022](#). Their replacement and Microsoft's long-term focus for the development of PowerShell support for Azure AD features is the Microsoft Graph PowerShell SDK. You should use the SDK for any future development and upgrade scripts written using the Azure AD or MSOL modules.

Although Microsoft has scheduled the end of support for the Azure AD and MSOL modules, the cmdlets will continue to work afterward. The exception is license management because, on August 26, 2022, Microsoft 365 adopts a new license management platform. Once the new platform is operational, license information for Azure AD accounts will be accessible only through Graph API requests or the cmdlets in the Microsoft Graph PowerShell SDK.

No automatic migration tool is available to translate code written using cmdlets from the older modules to cmdlets from the SDK. It is a manual process that requires substantial effort to create an inventory of scripts using the older modules, update their code, and test the effectiveness of the new code.

## Microsoft Teams PowerShell

The Teams PowerShell module contains cmdlets to manage elements such as teams, channels, and membership and the ability to create a new session to use the cmdlets for management of Teams policies such as *Get/Set-CsTeamsInteropPolicy*. To manage Teams with PowerShell, there are several sets of PowerShell cmdlets you should be familiar with:

- The [Teams PowerShell module](#) is available in generally available and preview versions. See [the documentation](#) for more information.
- If you need to manage the properties of the Microsoft 365 Groups used by Teams, use the \*-*UnifiedGroup* cmdlets included in the Exchange Online module.
- If you need to manage the properties of the Azure AD group, use the cmdlets in the Microsoft Graph PowerShell SDK.

To connect PowerShell to the Teams endpoint, run the following command.

```
[PS] C:\> Connect-MicrosoftTeams
```

Support for versions of the Teams PowerShell module before 4.0 ceases in mid-June 2022.

## PowerShell for Security and Compliance Features

The easiest way to connect to the compliance endpoint is to run the *Connect-IPPSSession* cmdlet from the Exchange Online management module. For example:

```
[PS] C:\> Connect-IPPSSession
```

The reason why a different PowerShell endpoint exists for compliance is that the functionality managed by the compliance cmdlets supports multiple workloads. Connecting to the compliance endpoint gives access to a range of cmdlets for different tasks such as managing DLP, preservation policies, and eDiscovery cases. For more information, see [Microsoft's documentation for the compliance cmdlets](#).

Many cmdlets with the same name (but slightly different functionality) exist in Exchange Online and the compliance set. For example, if you execute the *Get-RoleGroup* cmdlet against Exchange Online, you see the RBAC role groups defined for Exchange administration. If you connect to the compliance endpoint and run *Get-RoleGroup*, you see a completely different set of role groups. These are the role groups defined for RBAC access to security and compliance features.

## Power Automate and PowerShell

Power Automate support for PowerShell comes in two versions: [Administrator](#) and [Maker](#). The Power Apps PowerShell module also contains the Power Automate cmdlets. Using an administrator session, install the modules as follows:

```
[PS] C:\> Install-Module -Name Microsoft.PowerApps.Administration.PowerShell -Force -Scope AllUsers
Install-Module -Name Microsoft.PowerApps.PowerShell -AllowClobber -Force -Scope AllUsers
```

The current version of the Power Apps module has several dozen cmdlets to monitor Power Automate connector usage, verify what flows exist and run, and delete obsolete flows. Many of these tasks can be performed using the Power Automate portal, but you can automate some of the more tedious tasks with PowerShell. To enumerate the cmdlets, use the command:

```
[PS] C:\> Get-Command *Flow*
```

To test that the installation is successful, run the *Add-PowerAppsAccount* cmdlet to add an account to access Power Automate. This account must be a tenant admin to retrieve information about Power Automate for the tenant. PowerShell prompts for account credentials, which remain valid for up to 8 hours before you need to sign in again.

```
[PS] C:\> Add-PowerAppsAccount
```

If credentials are already available in a variable, you can use them. In this example, the *\$O365Credentials* object holds valid account credentials:

```
[PS] C:\> Add-PowerAppsAccount -Username user@domain.onmicrosoft.com -Password
$O365Credentials.Password
```

Once you're connected with administrative permissions, you can use the *Get-AdminFlow* cmdlet to get a list of all flows currently known in the tenant. Here's an example of how to extract and report the information returned for the flows. Remember to run the *Connect-MgGraph* cmdlet to connect to the Microsoft Graph PowerShell SDK first.

```
[PS] C:\> $Report = [System.Collections.Generic.List[Object]]::new()
$Flows = Get-AdminFlow
ForEach ($Flow in $Flows) {
    $ReportLine = [PSCustomObject] @{
        Created          = Get-Date($Flow.CreatedTime) -format g
        LastModified     = Get-Date($Flow.LastModifiedTime) -format g
        Flow             = $Flow.DisplayName
        CreatedBy        = (Get-MgUser -ObjectId $Flow.CreatedBy.UserId).UserPrincipalName }
}
```

```
$Report.Add($ReportLine) }
```

```
[PS] C:\> $Report | Format-Table Created, CreatedBy, Flow
```

Created	CreatedBy	Flow
10/05/2018 17:04	James.A.Abrahams@office365itpros.com	Create an event in Outlook in the new Planner
09/05/2020 23:30	John.West@office365itpros.com	Sync Microsoft 365 message center to P1...
20/08/2019 15:31	John.West@office365itpros.com	Start approval when a new item is added

The flow identifier is used to retrieve a list of flow runs. In this example, we fetch the identifier for a flow and run the *Get-FlowRun* cmdlet to fetch the run information. You can't fetch flow information for runs executed by another user, so you'll only see runs of the flow executed by your account.

```
[PS] C:\> $FlowId = Get-AdminFlow | ? { $_.DisplayName -Like "*Have I been pwned*" } | Select -
ExpandProperty FlowName
Get-FlowRun $FlowId
```

## Managing Connections to Microsoft 365 Endpoints

With so many different PowerShell management endpoints to connect to, you could easily accumulate a half dozen or more PowerShell functions in your profile to accommodate them all. Fortunately, you can save some time by using the [script written and maintained by MVP Michel de Rooij](#) to connect to the services, including:

- Azure AD.
- Azure Rights Management (Microsoft Information Protection).
- Exchange Online.
- Compliance endpoint.
- SharePoint Online.
- Teams.

The endpoints are liable to change as the service evolves, so you must be prepared to update the script to stay abreast of changes and to tailor its functionality to meet your needs. The script supports connecting to the endpoints with modules that support multi-factor authentication.

## Checking for Connections in Scripts

Before you can run cmdlets, you must connect to the endpoints used by the cmdlets. The number of connections needed depends on the modules that the cmdlets come from. Here's an example of how to check some of the modules used to interact with Microsoft 365 are available in the current session.

```
[PS] C:\> $Modules = Get-Module
If ("ExchangeOnlineManagement" -notin $Modules.Name) {Write-Host "Please connect to Exchange Online
Management before continuing...";break}
If ("MicrosoftTeams" -notin $Modules.Name) {Write-Host "Please connect to Microsoft Teams before
continuing...";break}
If ("Microsoft.Online.Sharepoint.PowerShell" -notin $Modules.Name) {Write-Host "Please connect to
SharePoint Online Management before continuing...";break}
```

**Linux, Mac OS, and Microsoft 365 management:** You can use PowerShell Core V6 (or later) to connect a session to Exchange Online and the compliance endpoint from a Linux or Mac OS X workstation. This is possible because these endpoints support remote sessions. The PowerShell modules for some parts of Microsoft 365, like SharePoint Online or Power Automate, need Windows .NET Core assemblies that do not exist on Linux or macOS, so you cannot connect PowerShell to these services. However, the Microsoft Graph PowerShell SDK modules, which rely only on PowerShell Core, will run just fine on [Mac OS X](#) and [Linux](#).

## Using Azure Cloud Shell with Microsoft 365 PowerShell Modules

Given that Microsoft controls the entire range of Azure compute and storage services, it seems logical that they have introduced products that consume these services to simplify various aspects of IT management. The [Azure Cloud Shell](#) (ACS) is such an example. When you run ACS from an administrative portal or <https://shell.azure.com, a connection occurs to PowerShell Core running on a Linux virtual machine on a Hyper-V host. Commands run on the virtual machine.>

You can connect to ACS from the Microsoft 365 admin center or Teams admin center, both of which include a shell icon in the command bar. ACS signs you in with the same account used to connect to the admin center, and that account must be able to use an Azure subscription to pay for a small amount of Azure storage for temporary files. Naturally, if you plan to do some administrative work, the account must also have the necessary permissions.

Microsoft has updated some PowerShell modules to allow them to run in ACS. You can:

- Connect to Exchange Online by running the *Connect-EXOPSSession* cmdlet. This uses Remote PowerShell to connect, so the newer cmdlets (like *Get-ExoMailbox*) are unavailable.
- Connect to Teams by running the *Connect-MicrosoftTeams -UseDeviceAuthentication* cmdlet. Using device authentication means that you go to <https://microsoft.com/devicelogin> and enter the code shown on the screen to prove your authentication status. Providing the right code allows ACS to connect to a Graph app called MS Teams PowerShell cmdlets and run both the Graph-based cmdlets (like *Get-Team*) and policy management cmdlets (like *Get-CsTeamsMessagingPolicy*).
- Use the *Invoke-RestMethod* or *Invoke-WebRequest* cmdlets to make API calls to the Microsoft Graph. Alternatively, use the Microsoft Graph PowerShell SDK cmdlets, including the *Invoke-MgGraphRequest* cmdlet.

You cannot connect to the SharePoint Online management endpoint via ACS. Other modules might have difficulties unless Microsoft has done the work required to allow them to work with ACS. Given the interconnected nature of Microsoft 365 applications, this can create some insurmountable difficulties with scripts unless developers write the scripts to run with ACS and do not attempt to load modules that cannot run under ACS. For instance, it is possible to write scripts to perform management operations for Teams using the cmdlets in the Teams module along with some Graph API calls (if necessary). However, you need to connect to Exchange Online differently (see above), don't make any calls to SharePoint Online cmdlets, and make sure that you don't use aliases for PowerShell commands like *Sort-Object* which might cause problems for PowerShell Core on Linux.

Although it can be convenient to be able to sign into ACS and run PowerShell cmdlets from administrative portals and browsers, in most situations, it is both more functional and easier to run PowerShell as normal. All modules will load, no special processing occurs, and scripts run as expected.

## Performance

PowerShell commands vary from one-liners to very complex scripts spanning many hundreds of lines. If you only use one-liners and an occasional script, you might not need to worry too much about performance. But once you start with more complicated processing involving multiple steps or thousands of objects, it's wise to consider how to make your PowerShell code run faster. The suggestions made here apply to many types of objects used across Microsoft 365 like mailboxes, groups, teams, distribution lists, or SharePoint sites.

## Object Filtering

PowerShell cmdlets can support server-side and client-side filtering:

- Server-side filtering occurs when you include the *-Filter* parameter in a command. The value passed to the parameter is the filter you want to apply. Server-side filtering is preferable because the remote

server (for example, Exchange Online or Azure AD) applies the filter before it returns a set of objects to the PowerShell session.

- Client-side filtering happens when you use the *Where-Object* (or simply *Where* or the *?* shorthand command) to process an available set of objects. Client-side filtering happens after a server returns a set of objects to extract the objects that you want to work with. Because you must wait for the full set of objects to come from the server before you can apply the filter, the overall time taken to process objects is invariably longer, and often by several factors.

For example, both commands shown below return the same set of objects (Groups with more than 10 guest members). We use PowerShell's *Measure-Command* cmdlet to tell us how long each command takes to execute, and the result is obvious. The server-side filter takes less than 0.3 seconds to return the subset of groups that have more than 10 guest users while fetching all the groups in the tenant and then applying the filter to that set takes about 14 times longer. Your mileage will vary depending on factors such as the current server load, the total number of groups in the tenant, and the complexity of the filter (and the property used). However, server-side filtering is invariably faster and should therefore be used whenever possible when processing objects fetched from servers – mailboxes, groups, sites, teams, and so on.

```
[PS] C:\> Measure-Command {Get-UnifiedGroup -Filter {GroupExternalMemberCount -gt 10}} | Select TotalSeconds
TotalSeconds
-----
0.2947363

[PS] C:\> Measure-Command {Get-UnifiedGroup | ? {$_.GroupExternalMemberCount -gt 10}} | Select TotalSeconds
TotalSeconds
-----
4.2973367
```

## Filtering and Performance with Exchange Online Cmdlets

The newer REST-based Exchange Online cmdlets (discussed later) are an exception to the rule mandating the use of server-side filters whenever possible. According to Microsoft, newer cmdlets like *Get-ExoMailbox* can [return data more slowly using a server-side filter](#) than with a client-side filter. Microsoft is working to address the issue and improve the performance and capabilities of server-side filters for these cmdlets. The root cause for the slowness might be the way that the cmdlets “warm up” during initialization. This is done to ensure speedy access when the cmdlets fetch large numbers of objects. Another downside is that the REST cmdlets do not support the same depth of filtering that their older counterparts do. For example, this query returns a filtered set of mailboxes that do not have mailbox auditing enabled but whose “persisted capabilities” mark them as having an enterprise license (BPOS\_S\_Enterprise) with the advanced auditing feature (M365Auditing):

```
[PS] C:\> Get-Mailbox -Filter "AuditEnabled -ne '$True' -and (PersistedCapabilities -eq 'BPOS_S_Enterprise') -and (PersistedCapabilities -ne 'M365Auditing')" -RecipientTypeDetails UserMailbox | Format-Table DisplayName
```

If you change the code to run *Get-ExoMailbox* instead, you get an invalid filter error. The problem here is likely to be an incompatibility between the way the older cmdlets perform server-side filtering using direct calls to the Exchange directory and the way that the REST cmdlets behave. In the interim, while Microsoft upgrades and improves the REST cmdlets, use these guidelines:

- Use the REST cmdlets whenever possible, especially when it's necessary to retrieve large numbers of mailboxes for processing.
- Try to apply server-side filtering with the REST cmdlets to refine the set of objects returned.
- If the REST cmdlets can't process the filter, use an older RPS cmdlet, or retrieve the data and apply a client-side filter.

In all cases, you should test filters when upgrading from the older cmdlets to the newer cmdlets to make sure that the filters work as expected.

## Using Server-Side Filters with Mail-Enabled Objects

Not every property of an object supports server-side filtering. Cmdlets in the Exchange Online module to process objects like mailboxes, distribution lists, mail contacts, and groups share many [common filterable properties](#). Here are some examples:

- Filter on the custom attributes (1 to 15 for single-value attributes, 1 to 5 for the extension attributes, which accept multiple values).

```
[PS] C:\> Get-ExoMailbox -Filter {CustomAttribute3 -ne $Null}
Get-UnifiedGroup -Filter {ExtensionCustomAttribute5 -eq $Null}
```

- Filter on the member counts for a Microsoft 365 group (total member count and guest member count). Distribution lists and security groups don't maintain counts as part of group properties. If you want this information, you must access the group and count its members.

```
[PS] C:\> Get-UnifiedGroup -Filter {GroupMemberCount -gt 20}
```

- Filter on the display name, name, or alias of mail-enabled recipients.

```
[PS] C:\> Get-DistributionGroup -Filter {DisplayName -like "*Manager*"}
Get-UnifiedGroup -Filter {Name -like "*Office*"}
Get-Recipient -Filter {Alias -like "Jeff"}
```

- Filter on whether groups are hidden from Exchange address lists. Groups like this are usually team-enabled.

```
[PS] C:\> Get-UnifiedGroup -Filter {HiddenFromAddressListsEnabled -eq $True}
```

- Filter on Email addresses. As noted above, *Get-ExoMailbox* supports a limited set of filterable properties. This is one example as *Get-ExoMailbox* doesn't support filtering by email addresses (all addresses or the primary address), so you should use *Get-Mailbox* instead. Note how the cmdlets use an extra filter to return just user mailboxes.

```
[PS] C:\> Get-Mailbox -Filter {EmailAddresses -like "*Office*"} -RecipientTypeDetails UserMailbox
Get-Recipient -Filter {PrimarySmtpAddress -like "Vasil*"} -RecipientTypeDetails UserMailbox
```

- Filter on the created date for a group, or the date when its properties were last updated.

```
[PS] C:\> Get-UnifiedGroup -Filter {WhenCreated -lt "01/01/2016 00:01"}
Get-DistributionGroup -Filter {WhenChanged -gt "01/01/2018 00:01"}
```

- Find the set of groups not enabled for Microsoft Teams.

```
[PS] C:\> Get-UnifiedGroup -Filter {ResourceProvisioningOptions -ne "Team"}
```

- If you want to use a calculated date from the *Get-Date* cmdlet, you need a slightly different format. In this example, we fetch all groups created in the last year.

```
[PS] C:\> [string]$CheckDate = (Get-Date).AddDays(-365)
Get-UnifiedGroup -Filter "WhenCreated -gt '$CheckDate'"
```

When filtering mailboxes and other mail-enabled objects, you can focus on a set of specific recipient types using the *RecipientTypeDetails* parameter. In this example, Exchange Online returns only user mailboxes and not shared mailboxes, room mailboxes, group mailboxes, and so on:

```
[PS] C:\> Get-Mailbox -Filter {EmailAddresses -like "*Office*"} -RecipientTypeDetails UserMailbox
```

While this example returns only distribution groups and ignores mail-enabled security groups and room lists.



```
[PS] C:\> Get-DistributionGroup -RecipientTypeDetails MailUniversalDistributionGroup
```

The *Get-Team* cmdlet returns a set of team objects. It supports some basic filtering on properties like user and privacy type (public or private). However, *Get-Team* does not support server-side filtering. We discuss this topic further in the section about using the Teams PowerShell module.

**Typed Variables Matter:** When you write PowerShell and use the Microsoft 365 modules to access data, consider using typed variables used to receive data fetched by cmdlets. Usually, PowerShell assigns the variable type automatically based on the data in use, but it's often better to be precise. For instance, declaring a variable to be an array is best when cmdlets return single items. One advantage of using an array is that you can be sure the count property works. For example, instead of using:

```
$Variable = Get-ExoMailbox
```

Use:

```
[Array]$Variable = Get-ExoMailbox
```

## Returning All Data

By default, many cmdlets retrieve a limited number of objects from the service. To make sure that you retrieve all the data you need, specify the number of objects to return as a value passed to the *ResultSize* parameter. Alternatively, you can use *Unlimited* to instruct the cmdlet to retrieve as many objects as exist. For example, this command retrieves all user mailboxes:

```
[PS] C:\> Get-ExoMailbox -RecipientTypeDetails UserMailbox -ResultSize Unlimited
```

The more objects returned, the longer it takes to process a command against those objects. Some of the examples explained here include the *ResultSize* parameter and some do not. This is not a general recommendation, and the correct usage depends on how many objects of a certain type exist in a tenant. In some cases, so many objects exist that it is best to retrieve them for processing in batches. The custom attributes available for mail-enabled recipients are useful here because you can assign one of the attributes to hold a batch number and then fetch objects for that batch using a filter. For example:

```
[PS] C:\> Get-UnifiedGroup -ResultSize 2000 -Filter {CustomAttribute15 -eq "Batch7"}
```

As we'll discuss later, the Microsoft Graph PowerShell SDK cmdlets use a different method to retrieve specific amounts of data. Because Microsoft generates these cmdlets from Graph APIs using a process called *AutoRest*, they use the *Top* parameter when it's important to fetch a precise number of records. Some cmdlets also support the *All* parameter to fetch all available data.

## Process Data Remotely

In a small tenant, it is usually unimportant if you run a cmdlet like *Get-UnifiedGroup* without any filters to return all available groups. However, as you scale up to deal with many thousands of groups (or teams), the time necessary to run the command increases. For Exchange Online objects like groups, we can speed things up by having PowerShell run commands on the remote Exchange server to which our session is connected. Things speed up because the remote server has local access to the data and is likely to have more resources available to process the data. By piping the output to the *Select-Object* cmdlet, the amount of data returned is reduced to just the requested properties. To ask the remote server to process a command, we call the *Invoke-Command* cmdlet.

For example, to retrieve the full set of groups and store them in a variable, we can run:

```
[PS] C:\> [array]$Groups = (Get-UnifiedGroup -ResultSize Unlimited | Select-Object DisplayName, Alias, GroupMemberCount, WhenCreated)
```

Another method to fetch data uses the *Invoke-Command* cmdlet. The advantage of this approach is that the server processes all data. This speeds things up by restricting the amount of data that must travel between server and client. In this example, the `$Session` variable points to the remote session connected to Exchange Online. The session identifier can be found with a command like:

```
[PS] C:\> $Session = Get-PSSession -InstanceId (Get-OrganizationConfig).RunspaceId.Guid
```

Now that we know the session, we can send the command to the server:

```
[PS] C:\> [array]$Groups = (Invoke-Command -Session $Session -ScriptBlock {Get-UnifiedGroup -ResultSize Unlimited | Select-Object DisplayName, Alias, GroupMemberCount, WhenCreated})
```

The `$Groups` variable ends up with the same data as in the first example. Running *Get-UnifiedGroup* directly is simpler but slower than running it via *Invoke-Command*. The difference is small (probably 10%) when dealing with 100 groups, but the performance gap widens as more objects are processed, so this technique is something to consider when you use PowerShell to process large numbers of groups.

## Where Method and Where-Object

It's common to form a collection of objects in an array with cmdlets like *Get-UnifiedGroup* or *Get-ExoMailbox* and then filter those objects to find a subset. You can do this by piping the objects through the *Where-Object* cmdlet or by using the *Where* method for the array. For example, we create an array of groups with:

```
[PS] C:\> $Groups = Get-UnifiedGroup -ResultSize Unlimited
```

To find the set of groups that mention "Office 365" in the display name, we can do:

```
[PS] C:\> $Groups | Where-Object {$_.DisplayName -Like "*Office 365*"}
```

Or:

```
[PS] C:\> $Groups.Where({$_ .DisplayName -Like "*Office 365*"})
```

Usually, the *Where* method is 10%-15% faster than using *Where-Object* (or the `?` shortcut). Note that the *Where-Object* cmdlet returns `$Null` if it finds no data with the filter while the *Where* method returns an empty array.

**Don't Limit Your PowerShell Output:** You might notice that sometimes PowerShell limits the output of values and uses an ellipsis (...) to show that some added information is available. The `$FormatEnumerationLimit` preference variable, which has a default value of 4, controls how many items PowerShell displays when a property has multiple items. If you want to be sure to see all available items, set `$FormatEnumerationLimit` to -1 (minus 1). Be aware that changing this value might affect the appearance of some of your reports.

## Select Properties

When you run a cmdlet like *Get-UnifiedGroup* or *Get-Team* to fetch information about a set of groups or teams, the cmdlet returns a set of properties for all matching objects. You might need all properties, but it is more likely that you only need a few. To [speed up processing](#) (because you don't need to loop through the set of objects to extract the needed properties), use the *Select-Object* (or just *Select*) cmdlet to pick the properties you need to process. For example, this command returns four properties for all the Groups in the tenant.

```
[PS] C:\> $Groups = (Get-UnifiedGroup | Select DisplayName, Alias, GroupMemberCount, WhenCreated)
```

Even better, define the set of properties that you want in a variable and use that with the command:

```
[PS] C:\> $Properties = 'DisplayName', 'ExternalDirectoryObjectId', 'Alias', 'PrimarySmtpAddress'
$Groups = (Get-UnifiedGroup | Select $Properties)
```

Note that *Get-UnifiedGroup* returns blank values for some properties unless you explicitly request values to be returned. To ensure that property values are up-to-date and complete, use the *IncludeAllProperties* parameter. This slows processing a little as Exchange Online must calculate some properties before they can be returned.

```
[PS] C:\> Get-UnifiedGroup -IncludeAllProperties
```

If you use the *Invoke-Command* cmdlet to send the command to the server for processing, *Select-Object* runs on the server. Otherwise, PowerShell filters the returned items using the *Select-Object* on the client.

The *Get-UnifiedGroupLinks* is another example of how selecting properties can help performance. This cmdlet returns the members or owners of a group, including a bunch of properties for each object. The set properties returned for local users are roughly equivalent to those returned by *Get-Mailbox* while those returned for guest accounts are the equivalent of *Get-MailUser*. Most people don't realize that calling *Get-UnifiedGroupLinks* results in the retrieval of quite so many properties as you only ever see the properties if you go looking, but it does mean that this is a "heavy" cmdlet in terms of processing overhead because of the work necessary to fetch the properties available through the cmdlet. If you only need a few properties for a group member, use *Select-Object* to fetch those properties after expanding the membership. This technique invariably performs better than using a *ForEach-Object* loop to process each item to extract properties.

**Running PowerShell against Large Numbers of Exchange Online Objects:** If you need to run scripts against large numbers of Exchange Online objects, consider using the [RobustCloudCommand](#) module. Developed by a Microsoft Exchange escalation engineer, the module supports running long-lasting scripts with automatic renewal of Azure AD access tokens.

## Avoid Throttling

Microsoft 365 is a multi-tenant environment and applies a throttle to stop users or processes from consuming more than their fair share of resources. When a throttling condition exists for PowerShell, the service signals a "micro delay" warning. Typically, this only happens when you try to process thousands of objects with resource-intensive cmdlets like *Get-MailboxFolderStatistics* (or *Get-ExoMailboxFolderStatistics*).

If you run into this condition, you can try to avoid throttling by introducing a short delay between each call to the resource-intensive cmdlet by calling the *Start-Sleep* cmdlet. In this example, we loop through a set of groups to find out how many conversation items exist in the mailbox belonging to each group. We report the number and the date the last conversation occurred and then sleep for 250 milliseconds.

```
[PS] C:\> ForEach ($G in $Groups) {
    $Conversations = Get-MailboxFolderStatistics -FolderScope Inbox -IncludeOldestAndNewestItems
                  -Identity $G.Alias | Select ItemsInFolder, NewestItemReceivedDate
    Write-Host $G.DisplayName "has" $Conversations.ItemsInFolder "conversations in the Inbox. The
last is dated" $Conversations.NewestItemReceivedDate
    Start-Sleep -Milliseconds 250 }
```

Like most tuning exercises, it might take a couple of attempts to calculate the right delay needed to avoid throttling.

## Speed Group Retrieval with Get-Recipient

Sometimes you only need to retrieve a list of groups for subsequent processing. In these scenarios, the *Get-Recipient* cmdlet (or the *Get-ExoRecipient* cmdlet from the Exchange Online Management module) is much faster at retrieving objects than *Get-UnifiedGroup* is. You can prove this theory by using the *Measure-Command* cmdlet to report how fast each cmdlet returns information.

```
[PS] C:\> Measure-Command {Get-Recipient -RecipientTypeDetails GroupMailbox}
Measure-Command {Get-UnifiedGroup}
```

*Get-Recipient* doesn't return any group properties, so you need to call *Get-UnifiedGroup* for each group to retrieve these properties and make them available for processing. Note that when we call *Get-Recipient*, we

select the *ExternalDirectoryObjectId* property. This is because a group's *ExternalDirectoryObjectId* (a GUID pointing to the group object in Azure AD) is unique whereas an alias or display name might not be, so when referencing an individual group in a loop, it's best to use *ExternalDirectoryObjectId* because you'll always be sure to access the group you want.

```
[PS] C:\> $Groups = (Get-Recipient -RecipientTypeDetails GroupMailbox | Select DisplayName, ExternalDirectoryObjectId)
ForEach ($G in $Groups) {
    $Site = (Get-UnifiedGroup -Identity $G.ExternalDirectoryObjectId).SharePointSiteURL
    Write-Host "The URL for the SharePoint site for the Group is" $Site }
```

You can use the identifier with the *Set-UnifiedGroup* cmdlet too:

```
[PS] C:\> ForEach ($G in $Groups) {
    Write-Host "Setting value for" $G.DisplayName
    Set-UnifiedGroup -Identity $G.ExternalDirectoryObjectId -CustomAttribute12 "Some value"}
```

*Get-Recipient* supports filtering, but you can only use properties common to all recipient types, like *DisplayName*, the custom attributes, or the alias. Note that because *Get-Recipient* returns all recipient types unless instructed otherwise, you must include a filter for group mailboxes.

```
[PS] C:\> Get-Recipient -Filter {DisplayName -Like "*Sanjay*" -and RecipientTypeDetails -eq "GroupMailbox"}
```

If you need to filter the set of groups using group-specific properties (like the number of guest members or the *HiddenFromExchangeClientsEnabled* flag), use *Get-UnifiedGroup* as described earlier. *Get-Recipient* does not support the group-specific properties.

An alternative approach uses the *Get-MgGroup* cmdlet from the Microsoft Graph PowerShell SDK. *Get-MgGroup* is even faster than *Get-Recipient*. However, it requires you to load an extra module and you must create a registered Azure AD app to run the SDK cmdlets in non-interactive jobs. In addition, *Get-MgGroup* retrieves basic group objects, and you'll still need to run *Get-UnifiedGroup* to retrieve properties that make the group a Microsoft 365 group. More details about running *Get-MgGroup* are in the section covering the Microsoft Graph PowerShell SDK.

## Use PowerShell Lists

Often scripts need to store some data. It's easy to use arrays for storage with statements like:

```
[PS] C:\> $MyMailboxes = @()
$MyMailboxes = "Vasil Michev", "Paul Robichaux"
$MyMailboxes += Get-ExoMailbox -Identity Tony.Redmond | Select -ExpandProperty DisplayName
$MyMailboxes
Vasil Michev
Paul Robichaux
Tony Redmond
```

This works, but it's better to use PowerShell list objects. The reason is that when you use the += operator to add a new item, PowerShell discards and rebuilds the array. This overhead is acceptable for small arrays, but not when dealing with large numbers of objects. PowerShell lists deliver better performance, and you can easily add multiple properties for each item. For example:

```
[PS] C:\> $MyMailboxes = [System.Collections.Generic.List[Object]]::new()
$Mailbox = Get-ExoMailbox -Identity Vasil.Michev -Properties WhenCreated
$MailboxData = [PSCustomObject][Ordered]@{
    UPN = $Mailbox.UserPrincipalName
    Name = $Mailbox.DisplayName
    Age = (New-TimeSpan($Mailbox.WhenCreated)).Days }
$MyMailboxes.Add($MailboxData)
```

## Dates

Properties holding dates for Microsoft 365 objects are often interesting pieces of data to report or query. For instance, you might want to find out the oldest distribution list in the tenant, which we can do for mail-enabled recipients using the *WhenCreated* property. For example:

```
[PS] C:\> Get-DistributionGroup | Sort WhenCreated | Select -First 1 | Format-Table DisplayName, WhenCreated
```

DisplayName	WhenCreated
Exchange MVPs (Only)	27/01/2014 20:37:00

The same technique works for mailboxes, mail contacts, mail users, Microsoft 365 groups, and so on. As we'll discuss when reviewing how to use the *Get-ExoMailbox* cmdlet, Exchange doesn't return date properties like *WhenMailboxCreated* unless you ask for them. In this example, we use that property to report the last user mailbox created.

```
[PS] C:\> Get-ExoMailbox -Properties WhenMailboxCreated -RecipientTypeDetails UserMailbox | Sort WhenMailboxCreated | Select -Last 1 | Format-Table DisplayName, WhenMailboxCreated
```

The examples above depend on the *WhenCreated* or *WhenMailboxCreated* properties being sortable as a date/time value. Sometimes, cmdlets return dates as strings, or a script stores a date as a string. In these cases, you need to tell PowerShell to regard the value as a date for sorting. For instance, let's assume that you want to create a list of guest accounts where the guest has not accepted their invitation. These accounts might be candidates for removal as it's obvious that the guest doesn't want to collaborate with our tenant. Our code grabs the set of guests that haven't accepted their invitation from Azure AD, calculates how many days it is since the account's creation, and stores the report data in a PowerShell list. The dates extracted from Azure AD are strings, so when we examine the data and want to sort the data to show the most recent guest accounts first, we must tell PowerShell to treat the property as a date.

```
[PS] C:\> $Report = [System.Collections.Generic.List[Object]]::new()
[array]$Guests = Get-MgUser -Filter ("UserType eq 'Guest'") -All | ? {$_.ExternalUserState -eq "PendingAcceptance"}
ForEach ($Guest in $Guests) {
    $AccountAge = ($Guest.CreatedDateTime | New-TimeSpan).Days
    If ($AccountAge -gt 365) {
        $ReportLine = [PSCustomObject]@{
            Guest      = $Guest.Mail
            Name       = $Guest.DisplayName
            Date       = $Guest.CreatedDateTime
            Age        = $AccountAge
        }
        $Report.Add($ReportLine) }
} # End ForEach
$Report | Sort {$_.Date -as [DateTime]} -Descending | Format-Table Guest, Date, Age
```

## Error Handling

In general, cmdlets return an error status that a script can interpret to handle different conditions. For example, if you try and access a group that doesn't exist and suppress the output of the error message to the screen, you can examine the error detail logged by PowerShell (in this case the reason is *ManagementObjectNotFoundException*, meaning that the system can't find the object) and then decide what to do next:

```
[PS] C:\> Get-UnifiedGroup -Identity NotThere -ErrorAction SilentlyContinue
$error[0].CategoryInfo.Reason
ManagementObjectNotFoundException
```

The cmdlets in the Teams PowerShell module are based on the Microsoft Graph API and sometimes behave differently compared to how cmdlets from other modules function. First, they don't support the ability to suppress error messages on the screen. Second, the reasons logged for the error conditions are sometimes indistinguishable across different conditions. For instance:

```
[PS] C:\> Get-Team -GroupId Rubbish -ErrorAction SilentlyContinue
get-team : BadRequest in /groups/ endpoint
At line:1 char:1
+ get-team -GroupId test -ErrorAction SilentlyContinue
+ ~~~~~
+ CategoryInfo          : NotSpecified: (:) [Get-Team], HttpRequestException
+ FullyQualifiedErrorId : System.Net.Http.HttpRequestException,Microsoft.TeamsCmdlets.PowerShell.Custom.GetTeam
[PS] C:\> $Error[0].CategoryInfo.Reason
HttpRequestException
```

[PowerShell's Try and Catch construct](#) can be used to handle error conditions more elegantly.

```
[PS] C:\> Try {Get-Team -GroupId 92e84aba-f3f7-4078-9f53-241575ce4d01}
Catch { Write-Host "No team found!"}
```

## Dual Writes for Exchange Online and Azure AD

When cmdlets in the Exchange Online management module update the properties of an object that also exists in Azure AD (like a Microsoft 365 group), they perform dual writes against the Exchange Online directory and Azure AD. This removes the need for synchronization between the two directories. In the case of Groups, it also means that Azure AD can go ahead and replicate updates to workloads like Teams or SharePoint Online.

The dual-write nature of these updates means that a write operation might not work because Azure AD is unavailable (usually momentarily). When this happens, the cmdlets receive a *UnableToWriteToAadException* error. This fact should be considered when building error handling into your scripts, especially when cmdlets update groups.

## Deciding Which PowerShell Cmdlets to Use

With so many PowerShell modules available, it can be confusing to understand the best cmdlet to choose to perform a task. Some choices are easy because modules like Exchange Online management, SharePoint Online, and Teams have clear purposes. Some choices are less clear. For instance, these cmdlets all return some information about a mailbox:

- *Get-MgUser* (from the Microsoft Graph PowerShell SDK).
- *Get-User* (from Exchange Online).
- *Get-Mailbox* or *Get-ExoMailbox* (from Exchange Online).
- *Get-CsOnlineUser* (from Microsoft Teams).

Depending on what task you're trying to perform, you might also need to use one or two of the cmdlets listed above to return some desired information. For example, the *Get-MgUser* cmdlet reveals the licenses held by an account but doesn't return the set of mail-related attributes available through the *Get-Mailbox* or *Get-ExoMailbox* cmdlets. There's currently no single cmdlet that will return every available property for a user account. Over time, you will develop an understanding of which cmdlets are best suited for different tasks. As a rule, use the cmdlets that are most related to the nature of the change or information you are working with. When working with PowerShell for Microsoft 365, remember:

- Use the cmdlets from the Microsoft Graph PowerShell SDK whenever possible for Azure AD account management, license assignment, and tenant configuration tasks of the type performed in the Microsoft 365 admin center. In this book, we use PowerShell SDK cmdlets whenever possible.

Organizations should upgrade scripts that include the Azure AD and MSOL cmdlets (like *Get-AzureADUser* and *Get-MsOUser*) and replace these cmdlets with Microsoft Graph API calls or cmdlets from the Microsoft Graph PowerShell SDK.

- If you work with on-premises accounts synchronized or federated to Azure AD, you'll need to manage most settings for those accounts using cmdlets from the on-premises Active Directory module.
- Use the workload-specific cmdlets (such as *Get-ExoMailbox* or *Get-User* for Exchange Online, or *Get-Team* and *Get-TeamChannel* for Teams) to perform the tasks that you would normally perform in those admin portals.

## Managing Exchange Online with PowerShell

PowerShell is critical to the smooth operation of Exchange Online. Administrators can transfer the skills that they accumulate working with the Exchange Management Shell used for on-premises server management to Exchange Online. However, do not expect that every piece of script that you've developed to help manage on-premises Exchange will move across because a reduced set of cmdlets is available in the cloud. The reason is logical – Microsoft takes care of a lot of the mundane work of operating servers and databases, so you do not need access to the cmdlets used for these purposes. The cmdlets that are available reflect the actions that you can take against different objects. Exchange Online uses the same RBAC mechanism as found in on-premises deployments to control the set of Exchange cmdlets and parameters available to a user. You can check the definitions of management roles to see what cmdlets are available for each role.

[The Exchange Online management module](#) supports Ubuntu 18.04 or later to allow Exchange Online management from Linux devices. The same version supports macOS devices.

### Exchange Online Management Module

Remote PowerShell (RPS) is the default mechanism for administrators to run cmdlets against a server. Although the mechanism works well on-premises, it is less successful in the cloud. To address the problem, Microsoft developed the [Exchange Online Management module](#), which is now the preferred method to run PowerShell against Exchange Online.

When they developed the module, Microsoft optimized several cmdlets for bulk retrieval operations. These cmdlets don't exist in on-premises Exchange PowerShell (see Table 23-1). Microsoft chose to create these cmdlets because their RPS equivalents are the most heavily used and error-prone within Exchange Online. You'll notice that the cmdlets all retrieve objects or settings, and none amend or create settings. This is because *Get-* cmdlets often deal with multiple objects and therefore run for much longer than the short interaction with the server needed by *Set-*, *New-*, or *Remove-* cmdlets.

<b>Remote PowerShell cmdlet</b>	<b>REST-based cmdlet</b>
Get-Mailbox	<i>Get-ExoMailbox</i>
Get-MailboxStatistics	<i>Get-ExoMailboxStatistics</i>
Get-MailboxFolderStatistics	<i>Get-ExoMailboxFolderStatistics</i>
Get-CasMailbox	<i>Get-ExoCasMailbox</i>
Get-Recipient	<i>Get-ExoRecipient</i>
Get-RecipientPermission	<i>Get-ExoRecipientPermission</i>
Get-MailboxPermission	<i>Get-ExoMailboxPermission</i>
Get-MailboxFolderPermission	<i>Get-ExoMailboxFolderPermission</i>
Get-MobileDeviceStatistics	<i>Get-ExoMobileDeviceStatistics</i>

Table 23-1: Remote PowerShell Cmdlets and REST equivalents

Microsoft says that they plan to add more REST-based cmdlets in future releases. For example, the cmdlets Microsoft created to control user access to Cortana briefing emails (*\*-UserBriefingConfig*) are REST-based. In addition, Microsoft is upgrading 800+ older RPS-based cmdlets to:

- Remove dependencies such as the Windows Remote Management (WinRM) service.
- Support modern authentication and remove the requirement to allow accounts to use basic authentication to connect to Exchange Online with PowerShell.
- Improve stability and error handling.

At the time of writing, Microsoft has upgraded approximately all its target cmdlets (in a preview module). To make it easy to run existing scripts, when you connect to Exchange Online, the module loads the older RPS cmdlets as well as the upgraded and new REST-based cmdlets. For version 2.0.6 and above, connecting to the module loads only the set of upgraded cmdlets. To load all cmdlets including those that Microsoft has not upgraded, you must connect using the *UseRPSsession* switch. This connection requires the user account to be enabled for basic authentication for PowerShell:

```
[PS] C:\> Connect-ExchangeOnline -UserPrincipalName Ken.Bowers@office365itpros.com -UseRPSsession
```

The upgraded cmdlets are more robust at handling network or server errors and include support for pagination and automatic retry. Because of multithreading, the REST cmdlets are faster than their RPS equivalents and can run at speeds of up to 8-10 times better. The exact speed increase is highly dependent on the number of objects processed. Small batches of mailboxes or other objects might see an improvement of 2-3 times while commands that process higher numbers (more than a thousand objects) should experience up to ten times better performance.

As you start working with the cmdlets optimized for bulk retrieval, you might notice that the first time you run a cmdlet against a set of objects, it's not quite as fast as you might expect. The impression is especially strong with small sets of objects. The initial slowdown is because the cmdlets support pagination, so when a cmdlet begins to process a set of objects, it sorts the objects to allow the cmdlet to resume processing from a retry point should the need occur. During this period, Exchange evaluates the cmdlet against RBAC to ensure that the user can perform the requested action. I call this the "warm-up" phase, and once it's finished, access to the objects is fast and robust. You don't notice this happening when dealing with large sets of objects because the time used to paginate is a much smaller percentage of the overall run.

**Specify the Right Domain:** If you use the *Organization* parameter with the *Connect-ExchangeOnline* cmdlet, make sure that you specify the service domain for the tenant (the onmicrosoft.com sub-domain) as otherwise, you might [experience some performance problems](#).

## Upgrading Scripts

Converting scripts to use the REST cmdlets is reasonably straightforward. Two big things to remember are that the REST cmdlets optimize how they communicate with Azure AD and minimize the properties retrieved for Exchange objects. Although you can use display names, alias, and even distinguished names to identify mailboxes, it's suggested that you use the mailbox's external directory object identifier (*ExternalDirectoryObjectId*) whenever possible. This is the Azure AD object identifier for the owning account, and it's guaranteed to be unique.

The REST-based cmdlets use the concept of property sets to group properties together into collections that calls can request. You'll always get the minimum property set (15 properties of 255) returned, but in many situations, you'll need to ask Exchange Online to return more information. For example, if you want to work with mailbox quotas, you'll specify that you want the Quota property set in a call like this:

```
[PS] C:\> $Mbx = Get-ExoMailbox -RecipientTypeDetails UserMailbox -ResultSize Unlimited -PropertySets Quota
```



If you need specific properties from a set, specify their name like this:

```
[PS] C:\> Get-ExoMailbox -AuditEnabled, WhenCreated, WhenChanged
```

If you want every available property, fetch the All property set.

```
[PS] C:\> Get-ExoMailbox -PropertySet All
```

## Script Conversion Checklist

Here is a checklist to consider when reviewing code to use the REST-based cmdlets:

1. Make sure that the script connects to the Exchange Online REST endpoint by running the *Connect-ExchangeOnline* cmdlet. This will load both the REST cmdlets and the older RPS cmdlets.
2. Change the cmdlet names from old RPS names to REST cmdlet names. For example, *Get-Mailbox* becomes *Get-ExoMailbox*.
3. Update the Identity parameter used in the scripts to find mailboxes using the object identifier or user principal name.
4. For best performance, only request the properties the script needs to process by specifying the correct property set.
5. Check that the REST cmdlets support all the parameters used in your script. Check that filters used to reduce the number of returned objects work with the REST cmdlets (see earlier discussion).
6. The REST cmdlets return data ordered alphabetically while the older cmdlets order objects by creation date. This might affect how scripts work.
7. Use piping whenever possible to take advantage of the multi-threading built into the REST cmdlets. Unfortunately, in most instances, if you want to manipulate or generate data about objects in the pipeline a *ForEach* loop is necessary.
8. Check your error handling. The complex error handling sometimes created in scripts to deal with RPS issues might not be necessary or need to be adapted for the REST cmdlets.
9. Check your performance after the upgrade. Hopefully, you'll be pleased. You should see at least a 2x-4x performance increase over equivalent RPS cmdlets.
10. Test, test, and test again to ensure the functionality of the script is as expected.

## Certificate-Based Authentication

The Exchange Online Management module supports certificate-based authentication using an app registered in Azure AD to allow unattended execution of the type needed to run scripts as background processes. This is a technique often used for long-running scripts to create reports, perform maintenance, and so on. The basic principles are:

- The app registered in Azure AD is assigned the *Exchange.ManageAsApp* permission and an Azure AD administrative role such as Exchange Service Administrator. This gives the app permission to access Exchange Online objects.
- A self-signed X.509 certificate is uploaded to the Azure AD app. This generates a thumbprint (unique value) usable for authentication. For testing, the easiest way to generate the certificate is by running the *New-SelfSignedCertificate* PowerShell cmdlet.
- When the background process connects to Exchange Online, it passes the object identifier for the Azure AD app, the thumbprint for the certificate, and the tenant's service domain. Together, these elements are enough for Exchange Online to allow authenticated access corresponding to the role assigned to the app. For instance, if the app is assigned the Exchange Service Administrator role, it can interact with Exchange like an administrator signed into Exchange interactively.

See the [Microsoft documentation](#) for more information about how to use the Exchange Online management module cmdlets in unattended scripts.

## Limited Sessions

An account can have up to three active sessions to a remote PowerShell endpoint for Exchange Online. The `Get-PSSession` cmdlet lists the current sessions in use at any time. For example:

```
[PS] C:\> Get-PSSession
```

Id	Name	ComputerName	ComputerType	State	ConfigurationName	Availability
1	ExchangeOnli	outlook.offi...	RemoteMachine	Opened	Microsoft.Exchange	Available

Once three sessions are present, if you need to create a new session, you must either remove a session with the `Remove-PSSession` cmdlet or wait 15 minutes for one of the sessions to time out. For example:

```
[PS] C:\> Remove-PSSession -Id 1
```

After you remove a session, the cmdlets loaded into the session are unavailable. Automatic management of connections to Exchange is a big advantage gained when you connect using the Exchange Online Management module rather than a Remote PowerShell connection.

## Exchange Online Organization Configuration

One reason why administrators use PowerShell to manage Exchange Online is to update settings in the organization configuration. These settings include those to control distribution list naming, public folders, enable or disable connectors for Groups and Teams, use of MailTips, and controls for Exchange web services. Although some of the settings are updatable using a management GUI, most are not because they are controlled by an application. In these instances, the settings are present because the configuration is a convenient place to store and read settings. The documentation for the [Set-OrganizationConfig](#) cmdlet mentions some settings that are only available for on-premises organizations; on the other hand, some settings found in the configuration (like `DirectReportsGroupAutoCreationEnabled`) are artifacts of defunct software that no longer exists.

The `Set-OrganizationConfig` cmdlet sets values while the `Get-OrganizationConfig` cmdlet retrieves values. For example, to enable the Bookings app for the tenant, you run the command:

```
[PS] C:\> Set-OrganizationConfig -BookingsEnabled $True
```

**Throttling PowerShell:** Microsoft throttles the use of PowerShell to ensure that commands do not absorb excessive resources. Throttles are a form of “micro delays” that cause processing to pause for a short period. Although throttles exist for a good reason, they can be problematic when the need arises to process large object sets, such as if you wanted to run the `Get-ExoMailboxStatistics` cmdlet for every mailbox. Two techniques help to avoid throttling. First, use the [Invoke-Command method](#) when you need to process large numbers of objects with PowerShell. Instead of running commands locally, `Invoke-Command` passes the commands to Exchange Online servers for the servers to process. Second, use filters to limit the number of objects requested from the server and the number of properties for each object. Server-side filtering (the cmdlet provides a filter to allow the server to return only the required objects) is much more efficient than client-side filtering (the items are all returned and then filtered with the `Where-Object` cmdlet) and should be used whenever possible. See [this Microsoft article](#) for details of Exchange cmdlets and how to use filters when requesting objects.

Note: The new REST-based cmdlets have very different performance, which means that these cmdlets do not suffer as much from throttling as the older Remote PowerShell cmdlets do. The examples in this book use the REST cmdlets whenever possible.

## Organization Configuration Hydration

Microsoft applies a standard configuration to Exchange Online for new tenants and most tenants use a standard configuration. Microsoft refers to the settings in the standard configuration as *dehydrated*, meaning that although the organization is free to customize standard settings such as role definitions or OWA mailbox policies before this can happen, the settings must be prepared for customization by copying them from the standard configuration. This process is called *rehydration* and tenants with non-standard configurations are *hydrated*.

Rehydration is a one-time operation. It can happen explicitly when an administrator runs the *Enable-OrganizationCustomization* cmdlet, or implicitly when EAC or another administrative interface runs the cmdlet when an administrator updates a setting for the first time. When the cmdlet runs successfully to prepare the organization for customization, it updates the *isHydrated* property in the organization configuration. You can check the property to know if an organization uses a customized configuration:

```
[PS] C:\> Get-OrganizationConfig | Select isDehydrated
```

```
IsDehydrated
```

```
-----  
False
```

## Reporting Mailbox Quota Usage

One of the first PowerShell scripts created after the launch of Exchange 2007 was to report the quotas assigned to mailboxes and the amount of quota consumed by each mailbox. Over time, many variations on this report have appeared. Most of the scripts available on the internet are relatively simple and depend on the *Get-ExoMailbox* and *Get-ExoMailboxStatistics* cmdlets. This book wouldn't be complete if we didn't include an example.

Two things are notable in the code below. First, if you want to do computations with the mailbox quota and usage information returned by the cmdlets, you should convert the data into numbers, which is what the script does. Second, the script checks for mailboxes that are above a threshold of quota used (in this case 85%) to highlight mailboxes that might need to have their quota increased (or remove some items to free quota). The output is a CSV file sorted by the mailbox display name.

```
[PS] C:\> # Set threshold % of quota to use as warning level
$Threshold = 85
Write-Host "Finding mailboxes..."
$Mbx = Get-ExoMailbox -RecipientTypeDetails UserMailbox -PropertySet Quota -Properties DisplayName -
ResultSize Unlimited
$Report = [System.Collections.Generic.List[Object]]::new() # Create output file for report
ForEach ($M in $Mbx) {
    # Find current usage
    Write-Host "Processing" $M.DisplayName
    $ErrorText = $Null
    $MbxStats = Get-ExoMailboxStatistics -Identity $M.UserPrincipalName | Select ItemCount,
TotalItemSize
    # Return byte count of quota used
    [INT64]$QuotaUsed = [convert]::ToInt64((((($MbxStats.TotalItemSize.ToString()).split("(")[-
1]).split(")")[0]).split(" ")[0]-replace '[,]', ''))
    # Byte count for mailbox quota
    [INT64]$MbxQuota = [convert]::ToInt64((((($M.ProhibitSendReceiveQuota.ToString()).split("(")[-
1]).split(")")[0]).split(" ")[0]-replace '[,]', ''))
    $MbxQuotaGB = [math]::Round(($MbxQuota/1GB),2)
    $QuotaPercentUsed = [math]::Round(($QuotaUsed/$MbxQuota),4).ToString("P")
    $QuotaUsedGB = [math]::Round(($QuotaUsed/1GB),2)
    If ($QuotaPercentUsed -gt $Threshold) {
        Write-Host $M.DisplayName "current mailbox use is above threshold at" $QuotaPercentUsed -
ForegroundColor Red
        $ErrorText = "Mailbox quota over threshold" }
    # Generate report line for the mailbox
```

```

$ReportLine = [PSCustomObject]@{
    Mailbox           = $M.DisplayName
    MbxQuota          = $MbxQuotaGB
    Items             = $MbxStats.ItemCount
    MbxSizeGB         = $QuotaUsedGB
    QuotaPercentUsed = $QuotaPercentUsed
    ErrorText         = $ErrorText}
$Report.Add($ReportLine) }
# Export to CSV
$Report | Sort Mailbox | Export-csv -NoTypeInformation MailboxQuotaReport.csv

```

## Figuring Out the Last Login Time for a Mailbox

The last login time is a property reported for a mailbox by the *Get-MailboxStatistics* and *Get-ExoMailboxStatistics* cmdlets. Many assume that this is a timestamp noting the last time that the mailbox owner logged onto their mailbox. Indeed, this is exactly what the property was when *Get-MailboxStatistics* first appeared in Exchange 2007, which is why so many PowerShell scripts use the *LastLogonTime* property to calculate the last login time for a mailbox. For instance, the code below is an example of what is often run for on-premises servers to find user mailboxes and output the *LastLogonTime* and *LastLoggedOnUserAccount* properties.

```

[PS] C:\> Get-Mailbox -RecipientTypeDetails UserMailbox -ResultSize Unlimited | Get-MailboxStatistics | Format-Table DisplayName, LastLogonTime, LastLoggedonUserAccount -AutoSize

```

DisplayName	LastLogonTime	LastLoggedOnUserAccount
Deirdre Smith	22/06/2020 08:48:35	
Tony Redmond	21/06/2020 15:16:03	
TempAdmin	21/06/2020 15:16:03	
Jeff Guillet	19/06/2020 07:14:52	

When *Get-MailboxStatistics* first appeared, it was reasonable to expect that only users connected to mailboxes, so the timestamp was accurate. Over time, the number of mailbox assistants running inside Exchange on-premises and Online servers expanded to perform different background processing and mailbox maintenance. Because multiple processes, like the Managed Folder Assistant, connect to a mailbox daily, using *Get-MailboxStatistics* or *Get-ExoMailboxStatistics* to determine the last login time became less and less accurate.

### Using Mailbox Diagnostics

Exchange Online captures diagnostic information for mailboxes on an ongoing basis. Other values are computed by background assistants using signals logged in the Microsoft Graph. The data is accessible using the *Export-MailboxDiagnosticLogs* cmdlet. Not every action causes Exchange to update a counter. For example, deleting an item in a mailbox folder doesn't update the *LastEmailTimeCurrentValue* count, but creating and sending a message does. Background assistants and other tasks continue to update *LastLogonTime*, which means that the value doesn't represent the last time the user logged onto the mailbox. For example:

```

[PS] C:\> $Log = Export-MailboxDiagnosticLogs -Identity TRedmond -ExtendedProperties
$xml = [xml]($Log.MailboxLog)
$xml.Properties.MailboxTable.Property | ? {$_.Name -like "Last*time*"}

```

Name	Value
LastRecordIdentifiedTime	18/06/2022 05:04:26
LastLogonTime	20/06/2022 10:55:37
LastSharingPolicyAppliedTime	12/04/2020 05:30:55
LastScheduledTimerChangeToken	0xCF2E68845952DA0800000000
LastUserActionTime	18/06/2022 17:38:53
LastUserActionUpdateTime	19/06/2022 05:22:04
LastContactsTimeCurrentValue	16/06/2022 13:50:34
LastEmailTimeCurrentValue	19/06/2022 19:51:36
LastFileTimeCurrentValue	19/06/2022 17:06:00
LastCalendarTimeCurrentValue	16/06/2022 10:44:47

```
LastTasksTimeCurrentValue      18/06/2022 10:32:46
LastProfileTimeCurrentValue    16/02/2019 19:20:25
LastUserActionWorkloadAggregateTime 19/06/2022 19:51:36
```

The existence of these properties means that we can extract and use the information to gain an insight into how people use Exchange Online. A script [described in this article](#) describes how to scan user mailboxes to extract mailbox statistics from the `Get-ExoMailboxStatistics` and `Export-MailboxDiagnosticLogs` cmdlets and Azure AD sign-in information using the `Get-MgAuditLogSignIn` cmdlet to create a report highlighting potentially underused mailboxes. Including Azure AD sign-in information is important. The influence of Teams on communications means that fewer internal messages arrive in user mailboxes, and people might work predominantly in Teams and not check their mailboxes for a few days. Therefore, someone might be very active across Microsoft 365 while their mailbox shows signs of inactivity.

The `LastLogonTime` property outputs the same value as its counterpart in mailbox diagnostics but displays the time in local time rather than UTC. For example:

```
[PS] C:\> Get-ExoMailboxStatistics -Identity TRedmond -Properties LastLogonTime, LastLogOffTime |
Format-List DisplayName, LastLogonTime, LastLogOffTime

DisplayName      : Tony Redmond
LastLogonTime    : 11/04/2022 13:03:10
LastLogoffTime   : 11/04/2022 13:03:59
```

## Using Audit Records to Track Logins

Events captured in the Office 365 audit log (see the auditing chapter) are useful when tracking user mailbox logins. For instance, the code shown below searches the audit log for the last 90 days (you can search back 365 days with Office 365 E5) to retrieve two types of records:

- `UserLoggedIn` comes from Azure AD and tells us when the user account last signed into a service. If someone only uses a single app (like Exchange) with a desktop client, the client only needs to execute a log-in after their token expires. Logins from browser clients (and Teams) are more accurate because they refresh their access tokens every hour. You can also fetch Azure AD sign-in data for accounts [using the Graph API](#) or with the `Get-MgAuditLogSignIn` cmdlet (only for the last 30 days).
- `MailboxLogin` comes from Exchange Online mailbox auditing and tells us the last time the user signed into the mailbox. Two problems make these records less dependable than user logins. First, mailbox auditing must be configured for mailboxes (Exchange Online sets mailbox auditing on by default) and the `MailboxLogin` event must be included in the Owner configuration. Second, the process to ingest Exchange Online audit data into the audit log is much slower than Azure AD events, which means that you might not see the events for some time after they occur.

Combining user login and mailbox login events in the search gives us more confidence that we can find when users last logged in, especially if mailbox auditing captures mailbox login events. While the code is slow to process large numbers of mailboxes and not guaranteed to generate 100% accurate results, the results give us better data to work with.

```
[PS] C:\> [array]$Mbx = (Get-ExoMailbox -RecipientTypeDetails UserMailbox -ResultSize Unlimited -
Properties DisplayName | Select PrimarySmtpAddress, DisplayName, UserPrincipalName)
$StartCheckDate = (Get-Date).AddDays(-90)
ForEach ($M in $Mbx) {
    [array]$AuditRecs = (Search-UnifiedAuditLog -StartDate $StartCheckDate -EndDate (Get-Date) -
UserIds $M.UserPrincipalName -Operations UserLoggedIn, MailboxLogin -ResultSize 1)
    If ($AuditRecs) {
        Write-Host "Last Login date for" $M.DisplayName "is" $AuditRecs[0].CreationDate }
    Else {
        Write-Host "No logins found for" $M.DisplayName "since" $StartCheckDate }
}
```

Records sent by Exchange Online to the audit log originate through user or system activity. They do not record access to non-user mailboxes, such as shared mailboxes or group mailboxes. To retrieve the last activity in a shared mailbox, you could use the *Get-ExoMailboxFolderStatistics* cmdlet to return the last modified date for an item found in the Inbox. Here's an example of how to scan shared mailboxes:

```
[PS] C:\> Get-ExoMailbox -RecipientTypeDetails SharedMailbox -ResultSize Unlimited | Get-ExoMailboxFolderStatistics -IncludeOldestAndNewestItems -FolderScope Inbox | Format-Table Identity, LastModifiedTime -AutoSize
```

Identity	LastModifiedTime
Customer Services\Inbox	25 May 2019 07:45:47
DLPIncidents\Inbox	14 Jan 2019 02:37:28
Redirect for Removed Mailboxes\Inbox	26 May 2019 00:44:09
Redmond Shared Events\Inbox	3 Jul 2019 07:20:19
Company Information\Inbox	22 May 2019 23:37:38
Book Comments\Inbox	10 Aug 2019 16:05:54

The same technique works for group mailboxes if you pipe the output from the *Get-UnifiedGroup* cmdlet to *Get-ExoMailboxFolderStatistics*.

## Managing Microsoft 365 Groups with PowerShell

A broad range of policy-driven features are available to control aspects of Microsoft 365 Groups like naming, creation, and expiry, but you might not like how those features work or want to pay for the extra licenses that they need. You might decide that you prefer to do things differently and that the best way to reach your goal is to write some PowerShell.

Because Microsoft 365 Groups are mail-enabled and enhanced versions of Azure AD groups, standard group management with cmdlets from the Microsoft Graph PowerShell SDK (covered later) can be used to manage many aspects of Groups, including basic properties like display names and group membership. However, if you want to manage email-related properties, such as proxy addresses, you must do this using the Groups cmdlets contained in the Exchange Online management module. Some of the basic cmdlets to work with Groups are:

- *New-UnifiedGroup*: Create a new group.
- *Get-UnifiedGroup*: Return the properties of one or more groups. As noted earlier, sometimes it is faster to use *Get-Recipient* to fetch information about groups.
- *Set-UnifiedGroup*: Manipulate the properties of a group.
- *Remove-UnifiedGroup*: Remove a group and remove all the content contained in the group mailbox, calendar, document library, and notebook. You can run the *Undo-SoftDeletedUnifiedGroup* cmdlet to restore a soft-deleted group within 30 days of its deletion.
- *Add-UnifiedGroupLinks*: Add a user object to the membership of a group. You can add users as a group owner, a member, or a subscriber.
- *Get-UnifiedGroupLinks*: Return information about the membership of a group.
- *Remove-UnifiedGroupLinks*: Remove a member, owner, or subscriber from a group.

Updates applied with the Groups cmdlets use a dual-write process to make sure that changes flow through to the underlying Azure AD group objects. If the target object cannot be updated in both EXODS and Azure AD, the write fails. After successful updates, a synchronization process pushes the changes to applications such as Teams and SharePoint Online.

## Creating New Microsoft 365 Groups

The *New-UnifiedGroup* cmdlet creates a new group. This operation can sometimes take a little time to complete because it creates a new group object in Azure AD followed by synchronization of the new object to EXODS to force the creation of the group mailbox. Only users with Exchange Online mailboxes can run the *New-UnifiedGroup* cmdlet.

In this example, we create a new private group called "Corporate Banking Team." If we want the group to be public, we'd set the *AccessType* parameter to "Public." Remember, you can change a group's access type afterward if needed. Note that the *DisplayName* is a required parameter. In this example, we also set:

- The owner (if you do not specify an owner, the user who runs the cmdlet becomes the group owner).
- The classification of the group to mark it as private. If you use sensitivity labels, specify the identifier (GUID) of the label to apply in the *SensitivityLabel* parameter and do not include the *Classification* parameter.
- *AutoSubscribeNewMembers* so that members are also added to the group's subscriber list to receive updates via email.
- Set the *RequireSenderAuthenticationEnabled* flag to *False* to allow people outside the tenant to communicate with the group.
- Create the membership list. You can add as many members as you like by specifying their alias or primary email address (or their distinguished name, if you like).
- Set the alias of the new group.

```
[PS] C:\> New-UnifiedGroup -DisplayName 'Corporate Banking Team' -AccessType Private -Alias BankingTeam -RequireSenderAuthenticationEnabled $False -Classification "External Access" -Owner TRedmond -Members "Kim Akers", "Ben Owens" -AutoSubscribeNewMembers
```

If you want to add multiple owners, you can do this after creating the group by running the *Add-UnifiedGroupLinks* cmdlet to add the extra owners. Before you can add someone as an owner, you must add them as a group member.

### Groups and the Alias Property

*New-UnifiedGroup* fails if you try to create a new group with an alias already assigned to another group. Although Exchange Online allows mail-enabled objects to have duplicate aliases (or *MailNickname*), it checks the alias value when *New-UnifiedGroup* runs to ensure that the alias is unique and fails if a match exists with another group. However, *New-UnifiedGroup* allows you to create a group with an alias already used for another mail-enabled object, such as a mail contact. The *Set-UnifiedGroup* cmdlet flags an error if you try to update a group with a duplicate alias. Having objects with duplicate aliases is a bad idea. You should always ensure that the alias is unique, which is one of the reasons why Exchange Online changed to create the *Name* and *DistinguishedName* properties of new mailboxes based on the Azure AD object identifier of the owning account.

Because each object has a unique SMTP address, the presence of objects with duplicate aliases does not interfere with mail flow. But because directory synchronization can use the alias as a unique identifier for objects, having objects with duplicate aliases can cause problems, including when you synchronize groups in hybrid configurations. Overall, it is best practice to check whether another mail-enabled object already has the alias that you want to use by running the *Get-Recipient* cmdlet. For example, this command checks whether another mail-enabled recipient has the alias "BankingTeam."

```
[PS] C:\> Get-Recipient -Filter {Alias -eq "BankingTeam"}
```

Because multiple mail-enabled objects can exist with the same alias, it's a bad idea to use the alias as an identity to find groups in scripts. Instead, use a property that is guaranteed to be unique, like the primary SMTP address or the external directory object identifier. For example:

```
[PS] C:\> Get-UnifiedGroup -Identity "BankingGroup@office365itpros.com"
```

## After Creating a Group

After the creation of a group, we can examine its properties with the *Get-UnifiedGroup* cmdlet. Some of the properties that you will see are shown below. Note that if you want to see all available properties for a group, specify the *-IncludeAllProperties* switch to force Groups to return properties such as the *InboxURL* (a URL pointing to the Inbox folder in the group mailbox).

```
[PS] C:> Get-UnifiedGroup -Identity BankingTeam | Format-List

AccessType                : Private
AutoSubscribeNewMembers  : True
ConnectorsEnabled        : True
Alias                     : BankingTeam
ManagedBy                : {Kim Akers, Ben Owens}
EmailAddresses            : {SMTP:BankingTeam@Office365ITPros.onmicrosoft.com}
PrimarySmtAddress        : BankingTeam@Office365ITPros.onmicrosoft.com
Name                     : Corporate Banking Team_89b6fbd74a
DisplayName               : Corporate Banking Team
RequireSenderAuthenticationEnabled: True
MaxSendSize               : 35 MB (36,700,160 bytes)
MaxReceiveSize           : 36 MB (37,748,736 bytes)
RecipientType             : MailUniversalDistributionGroup
Language                  : en-US
SharePointSiteUrl        : https://office365itpros.sharepoint.com/sites/BankingTeam
SharePointDocumentsUrl   : https://office365itpros.sharepoint.com/sites/BankingTeam/Shared
Documents
SharePointNotebookUrl    :
(etc.)
```

Many of the properties supported by Groups are familiar because they are also found in Exchange distribution lists. Groups determines the values of some of these properties automatically and you do not have any control over their values. Several properties are worthy of comment:

- The account that runs the *New-UnifiedGroup* cmdlet becomes a group owner (the *ManagedBy* property) and a group member (the *Members* property). If you do not want this to be the case, you should remove the account from these properties afterward. However, you must ensure that at least one group owner always exists, and that account must also be a group member.
- The recipient type is *MailUniversalDistributionGroup*, which is the same as an Exchange distribution list. However, the *RecipientTypeDetails* property for a group is *GroupMailbox*.
- The URLs for the SharePoint Online resources are not available at once for a new group. These values become available after SharePoint Online provisions the resources.
- Depending on what email address policies are in force, Exchange Online creates one or more email addresses for a new group. The first, or primary address (shown by the SMTP prefix in capitals), uses the alias of the group together with the default SMTP domain for the tenant (as in *BankingTeam@Office365ITPros.com*). The second uses the tenant's service "onmicrosoft.com" domain. If you do not have a vanity domain or do not use this domain as the default, the "onmicrosoft.com" domain becomes the primary address. If you need to change things later, it is easy to add other email addresses or completely rewrite the set of email addresses as needed by running the *Set-UnifiedGroup* cmdlet.
- The language set for the group comes from the tenant. In this case, it is "en-US," meaning U.S. English. If you want users to receive group notification messages in other languages, you can change this value with the *Set-UnifiedGroup* cmdlet.



- The values of the *MaxSendSize* and *MaxReceiveSize* properties control the largest message size that the group mailbox can send and receive. You cannot change these values.
- The *AutoSubscribeNewMembers* property controls whether new members added to the group will receive email notifications for new conversations. In other words, they join the subscriber list. The default is *\$False*. You can set this to *\$True* with the *Set-UnifiedGroup* cmdlet.
- The *RequireSenderAuthenticationEnabled* property works in the same way as a regular distribution list. When set to *\$True*, the group will only accept messages from accounts known to the tenant. Set the property to *\$False* when creating the group or with the *Set-UnifiedGroup* cmdlet afterward if you want to allow external users to be able to send messages to the group.

## Finding Groups

Several ways are available to find a list of groups with PowerShell. You can use:

**Get-UnifiedGroup:** Because this cmdlet returns all available information about groups, this is the preferred cmdlet when dealing with Groups.

**Get-Mailbox:** Group mailboxes are mailboxes, so it follows that you can use the *Get-Mailbox* cmdlet to find groups.

```
[PS] C:\> Get-Mailbox -GroupMailbox
```

The *Get-ExoMailbox* cmdlet does not currently support the *-GroupMailbox* switch.

**Get-Recipient:** Because a group mailbox is a mail-enabled object, it is a recipient. You can find group mailboxes with:

```
[PS] C:\> Get-Recipient -RecipientTypeDetails GroupMailbox
```

*Get-Recipient* is faster than either the *Get-UnifiedGroup* or *Get-Mailbox* cmdlets because it retrieves less information for each object. In many situations, the best approach for performance is to use *Get-Recipient* to fetch a set of groups for processing followed by calling other cmdlets to do whatever's necessary thereafter. As noted in the section about performance, you might need to make calls to *Get-UnifiedGroup* to return specific properties for a group. If preferred, you can use the *Get-ExoRecipient* cmdlet in the same manner.

Always use server-side filters when available to ensure the fastest performance when retrieving groups.

## Creating a List of Groups

This command creates a CSV file listing the alias, display name, group owner, access type (private or public), the counts for group members and external members, and the creation date, and then calls Excel to open the file. You can then sort and order the information as you need.

```
[PS] C:\> Get-UnifiedGroup -ResultSize Unlimited | Sort DisplayName | Select Alias, DisplayName, ManagedBy, AccessType, GroupMemberCount, GroupExternalMemberCount, WhenCreated | Export-CSV -Path C:\temp\TenantGroups.CSV -NoTypeInfo -Encoding Ascii
[PS] C:\> Invoke-Item C:\temp\TenantGroups.CSV
```

In this example, we limit the set of groups to report on to those created within the last 90 days.

```
[PS] C:\> Get-UnifiedGroup |? {$_.WhenCreated -gt (Get-Date).AddDays(-90)} | Format-Table DisplayName, WhenCreated
```

A variant of the above which uses server-side filtering for better performance is:

```
[PS] C:\> Get-UnifiedGroup -Filter ([scriptblock]::create("WhenCreated -gt '$((Get-Date).AddDays(-90))'"))
```

Many variants exist for reporting groups. For example, you might want to look for groups or teams that are inactive. As explained later, we can do this by looking for signs of activity in the conversations stored in group inboxes, team compliance records, or SharePoint documents.

## Setting Group Properties

The *Set-UnifiedGroup* cmdlet manipulates group properties. In this example, we use it to update some of the properties of our newly created group. The example changes the default language for group notifications, changes the email address list, tells Groups that we want new members to automatically join the subscriber list, sets up a MailTip, and updates the display name for the group. Just like any other mail-enabled object, the MailTip will be displayed by Outlook and OWA when a user addresses a message to the group.

```
PS: C:\> Set-UnifiedGroup -Identity BankingTeam -Language 'fr-FR' -PrimaryEmailAddress
'BankingTeam@Office365ITPros.com' -AutoSubscribeNewMembers $True -MailTip 'Corporate Banking is
great fun' -DisplayName 'Fun with corporate banking'
```

The language selected for a group does not affect user contributions. Instead, the choice of group language only affects the information contained at the bottom of notification messages.

## Calendar Access

In some cases, you might not want to allow all the group members to be able to edit items posted to the group calendar. The most common situation is in education when a class shares a group but only the teacher should be able to add or update calendar items. You can set the *CalendarMemberReadOnly* property for a group to restrict add and update access to group owners. For example:

```
[PS] C:\> Set-UnifiedGroup -Identity Classroom1 -CalendarMemberReadOnly $True
```

Currently, the feature only restricts access to group calendar items through OWA.

## Setting Email Addresses

Like any other mail-enabled recipient, a group has at least a primary SMTP address used by Exchange to route email to the group. In addition, a group can have several secondary proxy addresses that Exchange can also use in routing. The difference between a primary and secondary address is that the primary address is stamped on outgoing messages from the group and is therefore the reply address for the group. However, Exchange can route messages using any of the addresses assigned to a group. To set the full set of addresses for a group, run the *Get-UnifiedGroup* cmdlet:

```
[PS] C:\> Get-UnifiedGroup -Identity BankingTeam | Select -ExpandProperty EmailAddresses
smtp:BankingTeam@Office365itpros.com
SMTP:CorporateBanking@Office365ExchangeBook.com
SPO:SPO_b4f8244d-d4ee-46b8-96d5-21d2b6c44c81@SPO_b662313f-14fc-43a2-9a7a-
d2e27f4f3478smtp:BankingTeam@office365itpros.onmicrosoft
smtp:BankingTeam@office365itpros.onmicrosoft.com
```

This group has a primary SMTP address (it has a capitalized SMTP: prefix), two proxy SMTP addresses, one of which is for the service domain (the last in the list), and a proxy address used by SharePoint Online (SPO). To change the primary SMTP address, run the *Set-UnifiedGroup* cmdlet and pass the new address in the *PrimarySMTPAddress* parameter. All proxy addresses must be in one of the domains belonging to the tenant:

```
[PS] C:\> Set-UnifiedGroup -Identity BankingTeam -PrimarySMTPAddress
BankingGroup@Office365itpros.com
```

Setting a new primary address demotes the previous primary address but keeps it as a secondary proxy address. Therefore, in this example, we start with three SMTP addresses and end up with four. If you need to rewrite the set of addresses for a group, use the *EmailAddresses* parameter. For example, here we replace the existing addresses with a new set. Exchange Online assumes that an address is SMTP if you don't specify its

type, and if you don't specify which proxy address is the primary, Exchange Online assumes that it is the first in the list.

```
[PS] C:\> Set-UnifiedGroup -Identity BankingTeam -EmailAddresses  
"SMTP:BankingTeam@Office365itpros.com", "CorporateBanking@office365itpros.com",  
"Banking.Team@office365itpros.onmicrosoft.com"
```

Note that the new set includes an address in the tenant's service domain (onmicrosoft.com). This address is a Microsoft Online Email Routing Address (MOERA) and Exchange Online requires all Microsoft 365 groups to have a MOERA address. Attempts to remove the MOERA address for a group fail if you do not replace the MOERA address at the same time with a new proxy address from the service domain.

To add a new address without affecting the set of current addresses, specify the new address in a hash table. You can add or remove several addresses at one time by including the addresses in a comma-separated list.

```
[PS] C:\> Set-UnifiedGroup -Identity BankingTeam -EmailAddresses  
{Add="BankingGroup@Office365itpros.com"}
```

Similar syntax is used to remove an address. In this case, the command will not work because we are trying to remove the primary address for the group. To remove the primary address, you need to first assign a new primary address to the group before proceeding to remove its old primary address.

```
[PS] C:\> Set-UnifiedGroup -Identity BankingTeam -EmailAddresses  
{Remove="BankingTeam@Office365itpros.com"}
```

## Setting up Private Groups

Some groups exist to share sensitive information between their members. It might be the case that you do not want these groups to be generally known or accessible to non-members, so you can take steps to prevent users from knowing about their existence. The first step is to make sure that the group is private. It does not make sense to have public access to a sensitive group and making the group private prevents any chance that non-members will discover documents belonging to the group through a search performed with SharePoint or Delve. You set the access type for a group by editing its properties or by running the *Set-UnifiedGroup* cmdlet:

```
[PS] C:\> Set-UnifiedGroup -Identity "Senior Managers" -AccessType Private
```

Next, we can hide the group's existence from address lists to prevent it from appearing to users who browse the GAL or another address list to look for groups that they might like to join. This is an important step to take for groups that might have sensitive or controversial words in their name, such as "HR Task Group on Layoff Planning" or "Merger with the Contoso Corporation". To hide a group, run the *Set-UnifiedGroup* cmdlet to set the *HiddenFromAddressListsEnabled* property as follows:

```
[PS] C:\> Set-UnifiedGroup -Identity "Senior Managers" -HiddenFromAddressListsEnabled:$True
```

Hiding a group from the address lists is done by Exchange Online and is only effective for clients that connect to an Exchange mailbox. For instance, if you hide a group, it will still appear in Stream or Planner but not in OWA or Outlook. Members of a hidden group will not be able to see it if they browse an address list. However, they can access the group through a client. For example, if they expand the Groups list in the Outlook desktop client, they can see the hidden group because this list shows all groups to which the user belongs. The usual approach is to have members of hidden groups mark them as favorites so that the groups are easily accessible.

## Groups Hidden from Exchange

Teams uses the *HiddenFromExchangeClientsEnabled* property to hide the Microsoft 365 groups that it creates from Exchange clients like Outlook. If the *HiddenFromExchangeClientsEnabled* property for a group is *\$True*,

Exchange clients do not try to access the group, even if the user is in its membership. If `$False`, the group is available to Exchange clients.

```
[PS] C:\> Set-UnifiedGroup -Identity GroupForTeams -HiddenFromExchangeClientsEnabled
```

Because *HiddenFromExchangeClientsEnabled* is a switch, you do not need to pass `$True` if you want to enable it, but when the time comes to disable the switch, you do have to pass `$False`.

```
[PS] C:\> Set-UnifiedGroup -Identity GroupForTeams -HiddenFromExchangeClientsEnabled:$False
```

## Hidden Membership

The *New-UnifiedGroup* cmdlet supports the *HiddenGroupMembershipEnabled* switch, which conceals the membership of a group from everyone except group members and tenant administrators, who can always see and amend the membership using the Microsoft 365 admin center, EAC, or PowerShell. The *HiddenGroupMembershipEnabled* switch is only available when you create groups through PowerShell, and you cannot change the hidden status of membership afterward by running *Set-UnifiedGroup*. The intention is to block casual viewing of group membership for sensitive groups like those who are working on special projects or (as needed by law in some countries) the students in a certain class. An example of creating a group with hidden membership appears below. Remember to set the group access type to private.

```
[PS] C:\> New-UnifiedGroup -Alias QTX -HiddenGroupMembershipEnabled -DisplayName "Secret Planning Team" -AutoSubscribeNewMembers -AccessType Private
```

To list the groups with hidden membership, use the command:

```
[PS] C:\> Get-UnifiedGroup | ? {$_.HiddenGroupMembershipEnabled -eq $True} | Format-Table DisplayName, Notes, AccessType
```

Teams cannot use groups with hidden membership. The presence of the flag blocks their inclusion in the set of groups presented to an owner to be potentially enabled for Teams.

## Controlling Who Can Send Email to a Group

Like Exchange distribution lists, you can control the individual users and groups who can send messages to a group. To do this, use the *AcceptMessagesOnlyFromSendersOrMembers* parameter to create the set of individual senders and groups from whom the group will accept emails. You might want to do this to protect sensitive groups from unwanted communications. For instance, you might want to allow external users to email the group but restrict this access to some specific email addresses.

The list of individual senders can include mailboxes belonging to the tenant, mail contacts, mail users, and guests. Valid groups include distribution lists, dynamic distribution lists, and Groups. By default, the parameter value is *\$Null*, meaning that anyone can send messages to the group. Exchange Online rejects messages from external senders if the *RequireSenderAuthenticationEnabled* flag for the group is set to *\$True*.

The following example creates a set of allowed senders for a group using a mixture of email addresses, mailbox names, and the group itself. Adding the group to its allowed sender list means that Exchange delivers messages sent by members to the group. You might want to prevent members from sending messages to a group, but this is not usually the desired outcome.

```
[PS] C:\> Set-UnifiedGroup -Identity "Banking Communications" -AcceptMessagesOnlyFromSendersOrMembers Flay@outlook.com, "Ben Owens", "Banking Communications"
```

When Exchange Online creates the allowed sender list for the group, it resolves the addresses specified against Azure AD and writes the values into three group properties:

- **AcceptMessagesOnlyFromDLMembers:** Stores any reference to the groups on the allowed sender list.

- **AcceptMessagesOnlyFrom**: Individual senders on the allowed sender list.
- **AcceptMessagesOnlyFromSendersOrMembers**: The complete list of senders and groups allowed to send to the group.

We can see these properties with the *Get-UnifiedGroup* cmdlet:

```
[PS] C:\> Get-UnifiedGroup -Identity "Banking Communications" | Format-List accept*
AcceptMessagesOnlyFrom           : {flay_outlook.com#EXT#, Ben Owens}
AcceptMessagesOnlyFromDLMembers  : {BankingCommunications_a6275ce562}
AcceptMessagesOnlyFromSendersOrMembers : {flay_outlook.com#EXT#, Ben Owens,
                                         BankingCommunications_a6275ce562}
```

As you can see, the addresses used for guests are obvious because of their unique format.

## Using Exchange Online Cmdlets with Groups

The general rule is that the set of cmdlets created for Groups should be used to interact with groups. However, many other cmdlets in the Exchange Online and Exchange Online Management modules can be used to manage group mailboxes. The presence of a *GroupMailbox* parameter means that a cmdlet supports Groups. The current set that supports this parameter includes:

- *Add-MailboxFolderPermission*.
- *Add-MailboxPermission*.
- *Get-Mailbox*.
- *Get-MailboxFolderPermission*.
- *Get-MailboxPermission*.
- *Get-UserPhoto*.
- *New-MailboxAuditLogSearch*.
- *New-SyncRequest*.
- *Remove-MailboxFolderPermission*.
- *Remove-MailboxPermission*.
- *Remove-UserPhoto*.
- *Search-MailboxAuditLog*.
- *Set-Mailbox*.
- *Set-UserPhoto*.

Although these cmdlets give access to properties that aren't supported by the Groups cmdlets, you will rarely need to use them. The preference is to always manage Groups through the purpose-designed cmdlets instead of looking to use the general-purpose cmdlets.

## Mailbox Identities

Even in cases where a cmdlet does not support the *GroupMailbox* parameter, the *Get-UnifiedGroup* cmdlet uses the same kind of input formats for the Identity parameter (for example, alias or distinguished name) as do the other cmdlets in the Exchange Online module. This means that you can pipe the identity to these cmdlets to retrieve information about group mailboxes. For example, you can use a command like this to discover the number of messages in conversations stored in group mailboxes:

```
[PS] C:\> Get-UnifiedGroup -ResultSize Unlimited | Get-MailboxStatistics | Format-Table DisplayName,
ItemCount, LastLogonTime -AutoSize
```

DisplayName	ItemCount	LastLogonTime
Ignite 2016	235	18 Nov 2019 01:43:34
Managers	252	18 Nov 2019 01:13:52
Board Members (Secret)	262	17 Nov 2019 18:26:59
HR Working Group	265	18 Nov 2019 01:01:35

Mountaineering committee	227 17 Nov 2019 16:43:23
Interesting places to go	244 12 Apr 2018 03:43:17
Sanjay Patel's Favorite Places	222 18 Nov 2019 03:03:36

As pointed out earlier, using *Get-Recipient* to fetch a list of group mailboxes will make this code run faster (and there's an even faster way explained below). For example:

```
[PS] C:\> Get-Recipient -RecipientTypeDetails GroupMailbox -ResultSize Unlimited | Get-MailboxStatistics | Format-Table DisplayName, ItemCount, LastLogonTime -AutoSize
```

The last logon time shown for each group seems interesting. However, it is of little value because no one logs onto group mailboxes and the time reported is based on the last access of a background assistant to the mailbox. Using the *Get-MailboxFolderStatistics* cmdlet, we can also see that items accumulate in other mailbox folders such as:

- **Inbox:** conversation items posted by Outlook clients.
- **Calendar:** items in the group shared calendar.
- **Sent Items:** copies of items posted to the group by some clients. All the other default folders that you would expect to see in a user mailbox are also present.
- **TeamsMessagesData** (in the non-IPM part of the mailbox): compliance records captured for team channel conversations, if the group is team-enabled.
- **Junk Email:** potential spam and malware delivered to the group mailbox.
- **Recoverable items:** items held for eDiscovery and other compliance purposes.

This command reports all the folders in a mailbox (or group mailbox) that hold some items:

```
PS C:\> Get-MailboxFolderStatistics -Identity BankingTeam | ? {$_.ItemsInFolder -gt 0} | Format-Table Name, ItemsInFolder -AutoSize
```

## Using Exchange Online cmdlets with Groups

The [REST-based Exchange Online cmdlets](#) in the Exchange Online Management module support the processing of group mailboxes. These cmdlets don't support the *GroupMailbox* parameter but will process a set of mailboxes retrieved by the *Get-ExoRecipient* cmdlet when specifying a value of *GroupMailbox* in the *RecipientTypeDetails* parameter.

In this example, we use *Get-ExoRecipient* to fetch the set of groups and then pipe each mailbox through *Get-ExoMailboxFolderStatistics* to return the number of items and the date of the latest item in the Inbox folder. Calculated properties are used to create more readable output.

```
[PS] C:\> Get-ExoRecipient -RecipientTypeDetails GroupMailbox -ResultSize Unlimited | Select-Object -Property @{Name = 'UserPrincipalName'; Expression = {$_.PrimarySmtpAddress}} | Get-ExoMailboxFolderStatistics -FolderScope Inbox -IncludeOldestAndNewestItems | Format-Table @{"Name"="Group";"Expression"={$_.Identity.Split("\")[0]}}, ItemsInFolder, @{"Name"="Size";"Expression"={$_.FolderSize.ToString()}}, @{"Name"="Newest Item";"Expression"={$_.NewestItemReceivedDate }}
```

Looking at the date of the latest list in the Inbox tells us the last time an Outlook client posted an item to a conversation in the group. The same technique can be used with the *TeamsMessagesData* folder to discover the last time someone posted a message to a channel conversation in a team-enabled group.

## Managing Group Membership

A newly created group is a lonely place because it has no group members apart from the group owner. We can fix the problem by running the *Add-UnifiedGroupLinks* cmdlet to add some members to the group. In this example, we add Kim Akers to the group.

```
[PS] C:\> Add-UnifiedGroupLinks -Identity BankingTeam -Links Kim.Akers@Office365ITPros.com -LinkType Members
```

In this instance, the user principal name specifies the new member, but you can also use a mailbox alias or display name. The important thing is that the provided value is resolvable to a unique account. Although you can add unlicensed accounts to a Group, users will not be able to access group resources unless they have the right license. For instance, you cannot access the files in a document library unless you have a license to access SharePoint Online.

To add a manager (owner) to the group, we change the *LinkType* to "Owners". A user must be a member of the group before they can become an owner. To add a subscriber to the group, change the *LinkType* to "Subscribers." Members of the subscriber list receive email updates for new messages sent to conversations. Logically, a mailbox must be a member of a group before they can become a subscriber.

If you want to add a large set of members to a group, it is always both more efficient and less liable to throttling to include multiple users in a single call to *Add-UnifiedGroupLinks*. The easiest way to do this is to create an array of the users that you want to add to the group and then use the array as an input to the cmdlet. In this example, we use a variable to hold the mailboxes in the membership of a distribution list. The aliases for each of the accounts in the array can become an input for the *Add-UnifiedGroupLinks* cmdlet to add the users to the group:

```
[PS] C:\> $Members = (Get-DistributionGroupMember -Identity MyDL |
Where-Object {$_.RecipientTypeDetails -eq 'UserMailbox'})
[PS] C:\> Add-UnifiedGroupLinks -Identity MyNew0365Group -LinkType Members -Links
$Members.PrimaryStmpAddress
```

Another example scans for all user mailboxes in the tenant not marked with a value in one of the customized attributes. Again, the resulting array becomes the input to the cmdlet:

```
[PS] C:\> $Users = (Get-Recipient -RecipientPreviewFilter {RecipientTypeDetails -eq "UserMailbox"
-and CustomAttribute3 -ne "N"})
Add-UnifiedGroupLinks -Identity 0365Group -LinkTypeMembers -Links $Users.PrimarySmtpAddress
```

## Creating an All Users Group

It is common for organizations to have an "All Users" distribution list to allow for the distribution of company communications to everyone in a tenant. We can mimic the same approach with Groups by creating a group holding all the mailboxes in the tenant. The steps are simple:

- Create a variable holding the details of all user mailboxes in the tenant.
- Create the group with the *New-UnifiedGroup* cmdlet.
- Populate the membership of the group with *Add-UnifiedGroupLinks*. First, we add users as members of the group and then add them as subscribers.

The array holding the mailbox data (*\$Mailboxes*) contains many different properties for each mailbox. To make sure that we pass unique values, we tell Exchange Online to use the distinguished name property for each mailbox when adding it to the group membership. You could also use the primary SMTP address. Here is the code:

```
[PS] C:\> [array]$Mailboxes = Get-ExoMailbox -RecipientTypeDetails UserMailbox -ResultSize Unlimited
New-UnifiedGroup -Alias AllMailboxes -DisplayName 'Company Communications'
Add-UnifiedGroupLinks -Identity AllMailboxes -LinkType Members -Links $Mailboxes.DistinguishedName
Add-UnifiedGroupLinks -Identity AllMailboxes -LinkType Subscribers
-Links $Mailboxes.DistinguishedName
Get-UnifiedGroupLinks -Identity AllMailboxes -LinkType Members | Format-Table DisplayName
```

In large tenants, you might want to use batches of input objects to update group membership. PowerShell does not flag an error if a user already exists in a member list. Later, we describe a script to create a team for all employees. The script uses similar steps, adjusted to accommodate cmdlets in the Teams PowerShell module.

## Checking Group Membership

The *Get-UnifiedGroupLinks* cmdlet reports the membership of a group. The value passed to the *LinkType* parameter is either *Members*, *Owners*, or *Subscribers* to tell the cmdlet what membership information to display. For example:

```
[PS] C:\> Get-UnifiedGroupLinks -Identity BankingTeam -LinkType Members

Name           RecipientType
----           -
TRedmond       UserMailbox
Kim Akers      UserMailbox
```

As described earlier, *Get-UnifiedGroupLinks* returns an extensive set of properties for an individual member. In other words, when you use the cmdlet to interrogate group membership, Exchange Online returns a lot of useful information about each member including their display name, proxy email addresses, and custom attributes.

## Reporting the Membership of Groups and Teams

Many people ask for a report detailing the membership of groups. You can [download a script from GitHub](#) showing how to create an HTML report with the membership of a group sorted into members and owners. The script also generates a CSV file.

Another common request is to report the set of groups (or teams) someone belongs to. This example uses the *Get-MgUserMemberOf* cmdlet to find the set of groups the requested person belongs to. By default, this cmdlet returns details of membership of any group, so we filter the list to find just Microsoft 365 Groups. The cmdlet returns a set of group identifiers, so we run the *Get-MgGroup* cmdlet to retrieve details for each group. The output is a list holding details for each group usable as an input to another function or to create a CSV.

```
[PS] C:\> $Mailbox = Read-Host "Enter User to check"
$Mbx = (Get-ExoMailbox -Identity $Mailbox)
If (!$Mbx) { Write-Host ("Can't find a mailbox for {0}" -f $Mailbox) }

Write-Host "Checking Microsoft 365 Groups membership for" $Mbx.DisplayName
$FoundGroups = [System.Collections.Generic.List[Object]]::new()
[array]$Groups = (Get-MgUserMemberOf -UserId $Mbx.ExternalDirectoryObjectId | ?
{$_ .AdditionalProperties["groupTypes"] -eq "Unified" } | Select -ExpandProperty Id)
ForEach ($G in $Groups) {
    $GroupDetails = Get-MgGroup -GroupId $G
    $FoundGroup = [PSCustomObject]@{
        GroupName = $GroupDetails.DisplayName
        Mail      = $GroupDetails.Mail
        Id       = $G }
    $FoundGroups.Add($FoundGroup)
} #End For

Write-Host ("{0} is a member of {1} Microsoft 365 Groups" -f $Mbx.DisplayName, $FoundGroups.Count)
$FoundGroups | Select GroupName, Mail
```

To discover the set of groups owned by an account, use the *Get-MgUserOwnedObject* cmdlet to interrogate Azure AD, such as in the example shown below. As in the previous example, the *Get-MgUserOwnedObject* cmdlet returns more than Microsoft 365 groups, so we apply the same filter.

```
[PS] C:\> $Mailbox = Read-Host "Enter User to check"
[array]$Mbx = (Get-ExoMailbox -Identity $Mailbox | Select DisplayName, ExternalDirectoryObjectId)
If (!$Mbx) { Write-Host ("Can't find a mailbox for {0}" -f $Mailbox) }
$FoundGroups = [System.Collections.Generic.List[Object]]::new()
[array]$Groups = (Get-MgUserOwnedObject -UserId $Mbx.ExternalDirectoryObjectId | ?
{$_ .AdditionalProperties["groupTypes"] -eq "Unified" } | Select -ExpandProperty Id)
ForEach ($G in $Groups) {
    $GroupDetails = Get-MgGroup -GroupId $G
    $FoundGroup = [PSCustomObject]@{
```



```

    GroupName = $GroupDetails.DisplayName
    Mail       = $GroupDetails.Mail
    Id        = $G }
    $FoundGroups.Add($FoundGroup)
} #End For

```

```

Write-Host ("{0} is an owner of {1} Microsoft 365 Groups" -f $Mbx.DisplayName, $FoundGroups.Count)
$FoundGroups | Select GroupName, Mail

```

Another approach to finding out what groups a member belongs to is to use the distinguished name of the mailbox in a server-side filter. This approach works for both the *Get-Recipient* cmdlet and the *Get-UnifiedGroup* cmdlet, and it's useful when a script already has the Exchange Online management loaded. In this example, *Get-Recipient* finds the set of group mailboxes and the filter checks the membership of each group to see if the member exists. Apart from not having to call a separate cmdlet to fetch and check group membership, this is a good example of how server-side filtering speeds up processing.

```

[PS] C:\> $Mailbox = Read-Host "Enter User to check"
$SDN = (Get-Mailbox -Identity $Mailbox).DistinguishedName
$Groups = (Get-Recipient -ResultSize Unlimited -RecipientTypeDetails GroupMailbox -Filter "Members -
eq '$SDN'" | Select DisplayName, Alias)
Write-Host $Mailbox "is a member of" $Groups.count "Groups"
$Groups | Select DisplayName, Alias

```

The technique of checking group membership on a per-user basis can be exploited to create a report detailing the membership of every Microsoft 365 group (and team) in a tenant. The common approach to this problem is to create an array of groups and loop down through each group to extract and report its membership. This approach works, but as discussed above, using *Get-UnifiedGroup* and *Get-UnifiedGroupLinks* can be slow when many groups are to be processed. Checking group membership on an account-by-account basis delivers faster results, as evident in [this script from GitHub](#), which you can download to see the technique in action. The script generates three files: an HTML report and two CSV files holding summary per-user data (how many groups a user belongs to and their names) and full membership information for each group.

## Checking if Someone is in a Group

No off-the-shelf cmdlet exists to allow you to check whether a user is already a member of a group. For small groups, it is possible to check by iterating through the membership to look for a match, but once group membership grows past a hundred or so entries, that kind of processing becomes very slow. One way to check is to use code like that shown below to use a filter with the *Get-Mailbox* cmdlet to check the distinguished name of a user against those in the membership of a group. If a value is returned to the *\$Status* variable, the mailbox is in the group's membership. If not, it isn't. This is rough and ready code that does not include error checking.

```

[PS] C:\> $Check = (Read-Host "Enter User to check")
$User = (Get-Mailbox -Identity $Check | Select DisplayName, DistinguishedName)
$CheckGroup = (Read-Host "Enter name of group to check")
$TargetGroup = (Get-UnifiedGroup -Identity $CheckGroup)
$GroupDN = $TargetGroup.DistinguishedName
$Status = (Get-Mailbox -Identity $User.DistinguishedName -Filter "MemberOfGroup -eq '$GroupDN'"
-ErrorAction SilentlyContinue)
If ($Null -eq $Status) { Write-Host ("{0} is not a member of the {1} group" -f $User.DisplayName,
$TargetGroup.DisplayName) }
Else { Write-Host ("{0} is a member of the {1} group" -f $User.DisplayName,
$TargetGroup.DisplayName) }

```

## Dealing with Users and Guests

We can make things a little more efficient by building a hash table from the group membership and then checking a user against the hash table. This is faster than using a simple array and it allows us to deal with

groups that include guest members. In this example, we create a hash table from the membership of the All Employees group. The table consists of keys (user display names) and values (their email addresses), and looks something like this:

Name	Value
Sanjay Patel	Sanjay.Patel@office365itpros.com
Tony Redmond	Tony.Redmond@office365itpros.com
Jack Healy	Jack.Healy@office365itpros.com
Imran Khan	Imran.Khan@office365itpros.com

Guest users have the string "#EXT#" in their name, so we need to do some processing to normalize the names in the table for tenant users and guests. If you want to exclude guests, comment out the lines that process these members. After we build the table, the next thing is to ask for an SMTP address to check for in the group and then look for that address in the hash table. If the code finds a match in the table, we report success.

```
[PS] C:\> $GroupMembers = @{}
# Populate the array with current group membership
$OrgName = "@" + (Get-OrganizationConfig).Identity
$TargetGroup = "All Employees"
Get-UnifiedGroupLinks -LinkType Member -Identity $TargetGroup | % {
    $User = $_.Name.toString()
    # Handle guests
    if ($User -like "*#EXT*") {
        $GuestUPN = $User+$OrgName
        $GuestUser = Get-MgUser -UserId $GuestUPN -ErrorAction SilentlyContinue
        $GroupMembers.Add($GuestUser.DisplayName, $GuestUser.Mail) }
    else {
        # local mailbox
        $LocalUser = (Get-ExoMailbox -Identity $User)
        $GroupMembers.Add($LocalUser.DisplayName, $LocalUser.PrimarySmtpAddress) }
}
$Check = (Read-Host "Enter SMTP address to check")
If ($GroupMembers.Values -Match $Check) {
    Write-Host "User" $Check "is in the group" $TargetGroup }
```

Another approach to the problem is to create an array holding the group identifier and display name of the groups to which someone belongs and then check the object identifier of the target group against the array.

```
[PS] C:\> $Groups = @{}
$Check = (Read-Host "user to check")
$Dn = (Get-Mailbox -Identity $Check -ErrorAction SilentlyContinue).DistinguishedName
If (!(($DN)) { Write-Host "Can't find" $Check; break }
$TargetGroup = (Get-UnifiedGroup -Identity "Office 365 for IT Pros")
# Populate the array with current membership of the target group
Get-Recipient -Filter "Members -eq '$Dn'" -RecipientTypeDetails GroupMailbox | % {
    $GroupName = $_.DisplayName
    $GroupId = $_.ExternalDirectoryObjectId
    $Groups.Add($GroupId, $GroupName) }

If ($Groups.ContainsKey($TargetGroup.ExternalDirectoryObjectId) -eq $True) {
    Write-Host "User" $Check "is in the group" $TargetGroup.DisplayName }
```

## Removing Members from Groups

To remove a member from a group, run the *Remove-UnifiedGroupLinks* cmdlet and follow some simple rules:

- Only a group owner, a tenant administrator, or an Exchange administrator can run the cmdlet. Holders of other roles (like a Teams administrator) can use different cmdlets to update group membership.
- You must specify the link type of the list from which you want to remove the user – Owners, Members, or Subscribers.

- If a user is a group owner, you must remove their owner status first and then remove them as a group member. Remember that if you're removing the only remaining group owner, you should promote another member to be the group owner so as not to leave the group without an owner. Cmdlets stop the removal of the last group owner.
- If a group member has subscribed to the group for email updates, you do not have to remove this link separately as the removal of their membership also removes their subscription.

As an example, here is how to remove Kim Akers as a member of the Banking Team group:

```
[PS] C:\> Remove-UnifiedGroupLinks -Identity BankingTeam -LinkType Members -Links 'Kim Akers'
```

Exchange Online caches link information to improve performance so changing the links for a group might take some time before the change is fully effective across the tenant.

## Removing Guest Users from Groups

As discussed earlier, you might want to remove a certain account from all groups in a tenant. Removing the guest account from Azure AD is a straightforward brute-force solution, but it means that the guest loses any access they have to SharePoint Online and OneDrive for Business gained through sharing links. Removing a guest account is probably a good idea if you discover that some team owners invite a guest from a competitor company to join their teams. However, if you want to remove guests from groups (teams) while leaving the guest account in place, a more surgical approach is needed. Let's discuss what you could do.

This code depends on an exact check against the display name of either an Azure AD member or guest account. If an account is found, the code returns the set of groups the account is a member of.

```
[PS] C:\> $$CheckUser = (Read-Host "user to check")

[array]$User = Get-MgUser -Filter "displayName eq '$($CheckUser)'"
If ($User) {
    Write-Host "Fetching details of groups..."
    [array]$Groups = (Get-MgUserMemberOf -UserId $User.Id | ? {$_.AdditionalProperties["groupTypes"]
    -eq "Unified" } | Select -ExpandProperty Id)
    Write-Host $User.DisplayName "found in" $Groups.Count "Groups"}
Else { Write-Host ("Can't find {0} in Azure AD" -f $CheckUser) }
```

The next step is to remove the membership link from the groups where the user is present. This code asks the administrator if they want to go ahead and remove the user from the groups. Upon confirmation, the code first checks whether a user is a group owner and removes that link before it removes the member link.

```
[PS] C:\> $NextStep = Read-Host $User.DisplayName "found in" $Groups.Count "Groups. Do you want to
remove the user from the groups?"
If ($NextStep.SubString(0,1).ToUpper()-eq "Y")
{
    Write-Host "Removing" $User.DisplayName "from Groups..."
    ForEach ($G in $Groups) { # Process each group the user has been found in and remove
    $GroupDetails = Get-MgGroup -GroupId $G
    # Have to remove them as an owner first
    [array]$Owners = Get-UnifiedGroupLinks -Identity $G -LinkType Owners
    If ($User.Mail -in $Owners.PrimarySmtpAddress) {
        Write-Host "Removing" $User.DisplayName "as owner of group" $GroupDetails.DisplayName
        Remove-UnifiedGroupLinks -LinkType Owners -Links $User.Mail -Identity $G -Confirm:$False -
ErrorAction SilentlyContinue
    } # End Remove as Owner
    # Now remove their membership link
    [array]$Members = Get-UnifiedGroupLinks -Identity $G -LinkType Members
    If ($User.Mail -in $Members.PrimarySmtpAddress) {
        Write-Host "Removing" $User.DisplayName "as member of group" $GroupDetails.DisplayName
        Remove-UnifiedGroupLinks -LinkType Members -Links $User.Mail -Identity $G -Confirm:$False -
ErrorAction SilentlyContinue
    } # End remove as member
    } # End For
} # End If
```

Although this code does an effective job of cleaning up member and owner links for a selected user, it will hit a problem if it tries to remove the last owner for a group because *Remove-UnifiedGroupLinks* will signal an error and fail to remove the owner. Some extra error checking will solve the issue.

## Checking Groups with Guests

Another common reason to examine groups is to understand how many groups have guests and how many guests are in the membership of each group. Remember, it doesn't matter whether guests were added through Groups, Teams, or Planner – the guests all end up in the same membership. We can carry out our task by checking the membership count data held for each group. In this case, we fetch all the groups in the tenant and sort them by total member count. We then output the total member count, the count of guests, and a computed value for the number of users in the tenant.

```
[PS] C:\> Get-UnifiedGroup -Filter {GroupExternalMemberCount -gt 0} | Sort-Object GroupMemberCount -Descending | Format-Table -AutoSize DisplayName, GroupMemberCount, GroupExternalMemberCount, @{Name="Tenant Users"; Expression = {$_.GroupMemberCount - $_.GroupExternalMemberCount}}
```

DisplayName	GroupMemberCount	GroupExternalMemberCount	Tenant Users
Exchange's Grumpy Old Men	62	57	5
Microsoft 365 Discussion	61	55	6
Microsoft 365 Engage 2020 Speakers	42	41	1
Company Communications	28	0	28

The member counts recorded in group properties should only ever be treated as guidance. The actual number might be off by 1 or 2 in some groups before the background process which updates these numbers catches up with recent changes. If you want to be absolute about member counts, you must check the individual members in each group.

## Finding Inactive Guest Accounts

Another aspect of guest management is the removal of accounts after they are no longer required. [Azure AD Access Reviews](#) attempt to solve the problem by requiring group owners to review and approve continued access for guests in the groups that they manage. Access reviews can also look for inactive guest accounts using Azure AD sign-in data. The downside is that Access Reviews require Azure AD Premium P2 licenses.

If you don't have Azure AD Premium P2 licenses, it's still possible to look for potential signs of aging or inactive accounts. One approach is to write a script to look for guest accounts in a tenant directory over a certain age (in this example, a year). The check is based on the account's *CreatedDateTime* value. After finding the set of guest accounts, you can retrieve the set of Microsoft 365 groups that each guest belongs to. The code below creates a report of guest accounts older than one year together with their group membership.

```
[PS] C:\> Select-MgProfile Beta
[array]$GuestUsers = Get-MgUser -Filter "usertype eq 'Guest'" -All
$Today = (Get-Date); $StaleGuests = 0; $Report = [System.Collections.Generic.List[Object]]::new()
# Check each account and find those over 365 days old
ForEach ($Guest in $GuestUsers) {
    $AccountAge = ($Guest.CreatedDateTime | New-TimeSpan).Days
    If ($AccountAge -gt 365) {
        $StaleGuests++
        Write-Host "Processing" $Guest.DisplayName
        $i = 0; $GroupNames = $Null
        # Find what Groups the guest belongs to... if any
        $DN = (Get-Recipient -Identity $Guest.Id).DistinguishedName
        # The distinguished name for some accounts might contain an apostrophe, so we need to handle this
        If ($Dn -like "*'*)" {
            $DNNew = "" + "$($dn.Replace("'", ''))" + ""
            $Cmd = "Get-Recipient -Filter 'Members -eq '$DNNew'' -RecipientTypeDetails GroupMailbox |
Select DisplayName, ExternalDirectoryObjectId"
            $GuestGroups = Invoke-Expression $Cmd }
        Else {
```

```

    $GuestGroups = (Get-Recipient -Filter "Members -eq '$Dn'" -RecipientTypeDetails GroupMailbox
| Select DisplayName, ExternalDirectoryObjectId) }
    If ($GuestGroups -ne $Null) { $GroupNames = $GuestGroups.DisplayName -join ", " }
    $ReportLine = [PSCustomObject]@{
        UPN      = $Guest.UserPrincipalName
        Name     = $Guest.DisplayName
        Age      = $AccountAge
        Created  = $Guest.CreatedDateTime
        Groups   = $GroupNames
        DN       = $DN    }
    $Report.Add($ReportLine) }
}
$Report | Sort Name | Export-CSV -NoTypeInformation c:\Temp\OldGuestAccounts.CSV

```

The code uses distinguished names to perform a server-side check to find groups a guest account belongs to. We might run into a situation where a distinguished name contains special characters like the apostrophe, which is why the code checks for this condition. If the script finds an apostrophe in a distinguished name, it “escapes” the character before attempting to fetch the set of groups the guest account belongs to.

Azure AD sign-in logs are a good source to find the last sign-in date for an account. You can search the Azure AD sign-in logs using the *Get-MgAuditLogSignIn* cmdlet to fetch the information for guest users based on their account identifier. However, it’s easier to request the *Get-MgUser* cmdlet to return sign-in activity data along with the other properties it returns. Here’s an example. Starting off by using the *Get-MgUser* cmdlet to search to find accounts matching whatever the user enters:

```

[PS] C:\> Select-MgProfile Beta
$Guest = Read-Host "Enter name of guest to search for"
$Search = 'DisplayName:' + $Guest
[array]$Guests = Get-MgUser -Search $Search -ConsistencyLevel Eventual -Property SignInActivity
ForEach ($G in $Guests) {
    If ($G.UserType -eq "Guest") {
        If (!(($string)::IsNullOrEmpty($G.SignInActivity.LastSignInDateTime))) {
            $Days = New-TimeSpan($G.SignInActivity.LastSignInDateTime)
            Write-Host ("Guest member {0} last signed in on {1} or {2} days ago" -f $G.DisplayName,
$G.SignInActivity.LastSignInDateTime, $Days.Days ) }
            Else { Write-Host ("No recent Azure AD sign-in data available for {0} ({1})" -f
$G.DisplayName, $G.Mail) }
        }
    }
}

```

## Finding Inactive Guest Accounts Based on Activity

Reviewing guest accounts based on account age or their last sign-ins is certainly a valid approach. However, as Microsoft discovered in the Azure AD group expiration policy, age-based expiration is not necessarily the best approach: activity-based expiration is better. To establish if a guest account is in active use, we can check user activity in:

- The audit log to check if an auditable action happened for the account within the last 90 days.
- The message trace data gathered by Exchange Online to check if anyone sent email to the guest account over the last 10 days.

Writing a script to gather activity information is more complex than the previous example. Our example ([available from GitHub](#)) searches the audit log for three actions (you can add more if you like) that are important to guest accounts:

- *UserLoggedIn*: The guest signs into the tenant. For example, when a guest account signs in to SharePoint Online to edit a shared document.
- *SecureLinkUsed*: A guest accepts the invitation in a secure link to access a document.
- *TeamsSessionStarted*: A guest starts a Teams session. Due to token expiration, this event is logged every hour for the duration of the session.

In addition to checking the audit log, the script also runs a message trace to see if anyone sent messages to the guest's email address. An online check of message trace data can only go back 10 days, but it's a good way to discover if a guest account receives copies of Outlook group conversations via email. Finally, instead of using the *RefreshTokensValidFromDateTime* attribute of the guest account as the basis for calculating its age, this script uses the *creationDateTime* for the account. Because guest accounts authenticate against their tenant's directory, the value of *RefreshTokensValidFromDateTime* is usually the same as the *creationDateTime* for guest accounts, but *creationDateTime* returns an accurate creation timestamp for all Azure AD accounts.

As before, the script captures the names of any groups the guest belongs to. After processing all guest accounts, the script outputs the results to a CSV file. The script flags guests with no evidence of activity for review and potential deletion. One way to automate the deletion is to delete the records for guests that you want to keep from the CSV file and then use the updated file as input to the *Remove-MgUser* cmdlet.

## Removing Guest Accounts

If you find problematic guest accounts, you can remove them selectively or in bulk. This is not something to do without thinking (or testing) as removing a guest account ends any shared access it has to SharePoint Online and OneDrive for Business documents or folders in the tenant as well as its membership of Groups, which then renders access void to Groups, Teams, Planner, Yammer, and any other application which depends on Groups to manage its membership. With that caution in mind, to remove an individual guest account, run the *Remove-MgUser* cmdlet and pass either the account's object identifier or user principal name. For example:

```
[PS] C:\> Remove-MgUser -UserId JohnSmith_yandex.com#EXT#@office365itpros.onmicrosoft.com
Remove-MgUser -UserId 78cbdb6d-6a59-4b45-9925-c78f28dfefal
```

Azure AD doesn't prompt for confirmation before it removes the account. This is one reason why searching for an account using *Get-MgUser -Search* can be so dangerous as it is all too easy to remove the wrong account from the set returned by the search. Following deletion, the account stays for 30 days in a soft-deleted state and can be restored from this state using the *Restore-MgUser* cmdlet or the Azure AD admin center.

On the other hand, if you want to permanently remove a deleted guest account before its 30-day retention period expires, run the *Remove-MgDirectoryDeletedItem* cmdlet. Be careful as this cmdlet doesn't prompt either and once an object is permanently removed from the directory it cannot be recovered, meaning that the guest account must be recreated if it is required afterward. A recreated account has a different object identifier than the deleted account, so it will not regain access to resources that the deleted account had. Here's how to completely remove a guest account from a tenant directory.

```
[PS] C:\> $Id = (Get-MgUser -UserId JohnSmith_yandex.com#EXT#@office365itpros.onmicrosoft.com).Id
Remove-MgUser -UserId $Id
Remove-MgDirectoryDeletedItem -DirectoryObjectId $Id
```

## Bulk Removal

As an example of bulk account removal, this snippet looks for guest accounts belonging to a certain domain (in this case, Badpeople.com) and removes them from Azure AD.

```
[PS] C:\> [array]$Users = (Get-MgUser -Filter "UserType eq 'Guest'" -All | Select DisplayName,
UserPrincipalName, Id)
Foreach ($U in $Users) {
  If ($U.UserPrincipalName -Like "*BadPeople.com*") {
    Write-Host "Removing" $U.DisplayName
    Remove-MgUser -UserId $U.Id }}}
```

Alternatively, if you do not want to remove the guest account from the tenant, you can change the loop to run the *Remove-UnifiedGroupLinks* cmdlet to remove the guest from the membership of groups they belong to. If you only want to remove the guest from or Teams, run the *Remove-TeamUser* cmdlet.

## Synchronizing Groups with Security Groups

You might have a set of security groups that you use to control access to information and want to use the same membership of those security groups with Groups or Teams. No synchronization exists within Microsoft 365 today to keep the membership of a security group aligned with a group or vice versa, but it's reasonably easy to do with PowerShell.

Let's take an example where you have a security group and want to have an equivalent group (or team). In this case, the "eDiscovery Admins" security group is the master, and the "eDiscovery Administrators" group is the replica. Synchronization is done from master to replica. Assuming the group already exists, we fetch the membership of the security group and add any user account found there to the group.

```
[PS] C:\> $M365Group = (Get-UnifiedGroup -Identity "eDiscovery Administrators")
$SecurityGroupId = (Get-MgGroup -Filter "displayName eq 'eDiscovery Admins'").Id
# Grab list of security group members
[array]$SecurityGroupMembers = (Get-MgGroupMember -GroupId $SecurityGroupId)
# Populate the Group with the members of the security group
ForEach ($Member in $SecurityGroupMembers) {
    $MemberDetails = Get-MgUser -UserId $Member.Id
    If ($MemberDetails.UserType -eq "Member") {
        Add-UnifiedGroupLinks -Identity $M365Group.ExternalDirectoryObjectId -LinkType Member -Links
        $MemberDetails.UserPrincipalName }}}
```

This process doesn't handle nested groups, so if the security group includes a nested group, you'll have to extract the membership from that group and process those entries.

After we write the security group membership into the group, we need to do some added processing to fully synchronize the two groups. The security group is the master for membership, and some other members previously added to the group previously might not belong to the security group. To complete synchronization, we need to compare the two lists and remove anyone found in the group who doesn't exist in the security group. Again, this is a straightforward operation:

```
[PS] C:\> $GroupMembers = (Get-UnifiedGroupLinks -Identity $M365Group.ExternalDirectoryObjectId
-LinkType Member)
# Check if any group members do not exist in the security group
ForEach ($i in $GroupMembers) {
    If ($SecurityGroupMembers -Match $i.UserPrincipalName)
        {Write-Host $i.DisplayName "is in security group" }
    Else
        { Write-Host "Removing" $i.DisplayName "from group because they are not in the security group"
-ForegroundColor Red
Remove-UnifiedGroupLinks -Identity $M365Group.ExternalDirectoryObjectId -Links $i.Alias -
LinkType Member -Confirm:$False}
}
Write-Host "Current Membership of" $M365Group.DisplayName
Get-UnifiedGroupLinks -Identity $M365Group.ExternalDirectoryObjectId -LinkType Member | Select
DisplayName
```

To be fully effective, you would need to set up a periodic batch job to synchronize the two memberships. Some more robust error checking and handling of nested groups would improve the script too.

**The GMM Tool:** The Group Membership Management (GMM) tool is [available on GitHub](#) to allow tenant administrators to synchronize the membership of Microsoft 365 groups from source groups at regular intervals.

## Flagging Unowned Groups

All Microsoft 365 groups (and teams) should have at least one owner. If groups end up lacking owners, it means that no one except administrators can add new members or perform other group management functions and that administrators must process group expiry notifications. Here's some PowerShell to flag any Microsoft 365 groups with missing owners:

```
[PS] C:\> $Groups = (Get-UnifiedGroup -ResultSize Unlimited | Select DisplayName, ManagedBy)
$Groups | ? {$_.ManagedBy.Count -eq 0}
```

Another way of looking at group ownership is to check what groups a certain user owns. You can do this by running the *Get-Recipient* cmdlet to return a list of group mailboxes and using a filter on the *ManagedBy* property to extract the groups owned by the person that you want to check. The only complicating factor is that you must use the user's distinguished name in the filter. For example, this PowerShell finds all the groups where Kim Akers is an owner.

```
[PS] C:\> $Dn = (Get-Mailbox -Identity "Kim Akers").DistinguishedName
Get-Recipient -RecipientTypeDetails GroupMailbox -Filter "ManagedBy -eq '$Dn'" | Format-Table
DisplayName
```

Another way is to query Azure AD for information about the objects it believes someone owns:

```
[PS] C:\> $UserId = (Get-ExoMailbox -Identity "Kim Akers").ExternalDirectoryObjectId
[array]$Groups = (Get-MgUserOwnedObject -UserId $UserId | ? {$_.AdditionalProperties["groupTypes"] -
eq "Unified"}) | Select -ExpandProperty Id )

ForEach ($Group in $Groups) {
    Get-MgGroup -GroupId $Group | Select DisplayName, Description }
```

## Checking for Unused Groups

The group expiration policy gives tenants an automated method to age out groups after a set period. However, you might want to check how active groups are before they become candidates to expire to understand the reasons why users create groups and what happens to those groups in the months afterward. To achieve this goal, we can use PowerShell to find any groups where no activity has occurred recently. Depending on their use, groups display different signs of activity. A group used by people who center their activity around Outlook is likely to use conversations heavily while another team might use Teams for their communications. Table 23-2 describes how users interact with groups and the consequent dependency on the two primary workloads.

<b>Type of group</b>	<b>Files stored in SharePoint Online</b>	<b>Items stored in Exchange Online</b>
Groups: documents	Yes	N/A
Groups: conversations and shared calendar	N/A	Yes
Yammer Groups: documents	Yes	N/A
Yammer Groups: discussions and shared calendar	N/A	Yes (calendar)
Group created by Microsoft Planner	Yes	Yes (discussions)
Group created by Microsoft Teams	Yes	Yes (calendar and compliance records)

Table 23-2: Exchange and SharePoint usage in various group types

You can check the potential underuse of groups used for document generation and management by looking for evidence of SharePoint file activity for the site as recorded in the audit log. After connecting a PowerShell session to SharePoint Online and Exchange Online, the code shown below checks to see whether any audit records generated by SharePoint file operations over the last 90 days exist for the site belonging to each



group. If no audit records for SharePoint file operations for a site are available, it's reasonable to assume that not much document-centric activity has taken place for the site, and it is therefore potentially obsolete.

```
[PS] C:\> $WarningDate = (Get-Date).AddDays(-90)
$Today = (Get-Date)
$Groups = Get-UnifiedGroup -ResultSize Unlimited
# This is faster: [array]$Groups = Get-MgGroup -Filter "groupTypes/any(c:c eq 'unified')" -All
# but then you need to fetch the details of each group
$ObsoleteGroups = 0
ForEach ($G in $Groups) {
    If ($G.SharePointSiteUrl) {
        $SPOSite = (Get-SPOSite -Identity $G.SharePointSiteUrl)
        Write-Host "Checking" $SPOSite.Title "..."
        $AuditCheck = $G.SharePointDocumentsUrl + "/*"
        [array]$AuditRecs = (Search-UnifiedAuditLog -RecordType SharePointFileOperation `
            -StartDate $WarningDate -EndDate $Today -ObjectId $AuditCheck -ResultSize 1)
        If (!$AuditRecs) {
            Write-Host ("No audit records found for {0}. The site is potentially obsolete!" -f
                $SPOSite.Title) ; $ObsoleteGroups++ }
        Else
            { Write-Host ("Recent audit record found for {0}. Last access dated {1}" -f $SPOSite.Title,
                $AuditRecs.CreationDate[0]) }
        }
    Else
        {
            Write-Host "SharePoint has never been used for the group" $G.DisplayName
            $ObsoleteGroups++
        }
    }
}
Write-Host $ObsoleteGroups "obsolete group document libraries found out of" $Groups.Count "checked"
```

Microsoft 365 Groups that use emails for communication (Outlook groups) require a different approach because we need to track the activity level for conversations. In this case, we use the *Get-ExoMailboxFolderStatistics* cmdlet to check the date and time for the last item (conversation) added to the Inbox. If that conversation occurred within the last year, all is well, and we can go on. If not, we flag the group as potentially obsolete. Checking the date of the last item is a simple sign of activity. We can improve it by checking how many items are in the Inbox. If just a few items exist, it is a sign that the group processed the first notification message sent following its creation and a couple of other conversations, and then gently fell into decay.

```
[PS] C:\> $Groups = Get-Recipient -RecipientTypeDetails GroupMailbox -ResultSize Unlimited
$ObsoleteGroups = 0; $WarningDate = (Get-Date).AddDays(-365)
ForEach ($G in $Groups) {
    Write-Host "Checking Inbox traffic for" $G.DisplayName
    $Data = (Get-ExoMailboxFolderStatistics -Identity $G.Alias -IncludeOldestAndNewestItems -
        FolderScope Inbox)
    If ($Data.NewestItemReceivedDate -le $WarningDate) {
        Write-Host ("Last Inbox item found in {0} was {1}. Not much going on here!" -f $G.DisplayName,
            $Data.NewestItemReceivedDate); $ObsoleteGroups++ }
    Else {
        Write-Host ("{0} has {1} items in the Inbox. Folder size is {2}" -f $G.DisplayName,
            $Data.ItemsInFolder, $Data.FolderSize) }
    }
}
Write-Host $ObsoleteGroups "Obsolete Groups found out of" $Groups.Count
```

The same approach can be used to check whether groups enabled for Teams or Yammer are active. For Teams, Microsoft 365 saves compliance records for regular channel conversations in the *TeamsMessagesData* folder, while the Yammer folder stores compliance records for conversations in Yammer-enabled groups. This means that we can check the usage of a group by looking at the newest item in the relevant folder. Here's an example for Teams:

```
[PS] C:\> $G = Get-UnifiedGroup -Identity 0365ITPros
```

```
$TeamsData = (Get-ExoMailboxFolderStatistics -Identity $G.Alias -FolderScope NonIpmRoot
-IncludeOldestAndNewestItems | ? {$_.FolderType -eq "TeamsMessagesData"})
Write-Host ("Last Teams conversation created in {0} was {1}." -f $G.DisplayName, $TeamsData.
NewestItemReceivedDate)
```

Note that if you want to process all team-enabled groups, the recommended approach is to use the *Get-Team* or *Get-MgGroup* cmdlets to create a collection of teams and then process each group in the set.

These examples use different methods to find potentially underused groups. It would not take a lot of added effort to add some code to combine the results and make a better decision as to which groups are obsolete and those that only use either SharePoint Online or Exchange Online. An example script that combines checks of Teams, SharePoint Online, and Exchange Online to find and report potentially obsolete Groups and Teams [is described in this post](#). A [Graph-based version of the script](#) is also available and is up to four times faster than the cmdlet equivalent.

You can use the Groups activity report available in the Microsoft 365 admin center to identify inactive groups based on the last activity date, but the nice thing about approaching the problem through PowerShell is that you have a chance to tailor the assessment and the output to meet your needs, such as including details of the group access type or using one of the custom properties that can be assigned to groups to hold information about its status.

A problem shared by both the Groups activity report and any attempt to detect inactive groups via PowerShell is that the focus is on activity created in Exchange and SharePoint.

## Blocking Guest Access to Sensitive Groups

Chapter 11 discusses how to use a policy to disable guest access for a selected group. If you use classifications in your tenant, you probably have a classification like “Company Confidential” or “Secret” to mark groups (or teams) that hold sensitive information. In some cases, you might want to make sure that these groups do not have guest users. Another example is when a school wishes to block guest access to any group used for student communications.

It’s easy to solve the problem with PowerShell if a suitable classification is assigned to sensitive groups. With the classifications in place, we can scan for those groups and then block them all for guest access. In this example, the first step is to collect a set of groups classified as “Confidential.” The code then loops through each group to discover whether group-specific policy settings are in place. If so, the code updates the settings to block guest access. Groups that don’t have a policy setting are controlled by the tenant policy, so the first step is to create policy settings for the group. We can then update the setting to block guest access.

```
[PS] C:\> $GroupTemplate = (Get-MgDirectorySettingTemplate | ? {$_.DisplayName -eq
"Group.Unified.Guest"})
[array]$Groups = (Get-UnifiedGroup -ResultSize Unlimited | Where {$_.Classification -eq
"Confidential"})

ForEach ($Group in $Groups) {
    $GroupSettings = Get-MgGroupSetting -GroupId $Group
    if($GroupSettings) { # Policy settings already exist for the group - so update them
        Update-MgGroupSetting -GroupId $Group.ExternalDirectoryObjectId -TemplateId $GroupTemplate.Id
-Values (@{'name'='AllowToAddGuests';'value'='false'} | ConvertTo-Json) -DirectorySettingId
$GroupSettings.Id
        Write-Host "External Guest accounts prohibited for" $Group.DisplayName
    }
    Else
    { # Settings do not exist for the group - so create a new settings object and update
        New-MgGroupSetting -GroupId $Group.ExternalDirectoryObjectId -TemplateId $GroupTemplate.Id -
Values (@{'name'='AllowToAddGuests';'value'='false'} | ConvertTo-Json
        Write-Host "External Guest accounts blocked for"$Group.DisplayName
    }
}
```

The update is to the Azure AD group. Synchronization to the workload directories must occur before clients learn about the new status for a group and the block becomes effective.

To check that the block works, we can create a list of the groups blocked from having guest members. To do this, run the *Get-UnifiedGroup* cmdlet to check the *AllowAddGuests* property, which is *\$False* for blocked groups. For example, this command reports the display names and classification for all blocked groups. Remember that the block is effective for all clients that populate group membership, including Teams.

```
[PS] C:\> Get-UnifiedGroup -ResultSize Unlimited | ? {$_ .AllowAddGuests -eq $False } | Format-Table
DisplayName, Classification
```

## Using Sensitivity Labels to Block Guest Access

If you use Sensitivity Labels to assign a classification to Groups, you can block guest access to a group by choosing a label that restricts external guest access. After an administrator or owner assigns a label blocking guests from a group, Exchange Online sets the *AllowAddGuests* property of the group to *\$False* and updates the *AllowToAddGuests* setting in the Azure AD policy for the group to block guest access.

You can assign a sensitivity label to a group with the *Set-UnifiedGroup* cmdlet using the label identifier. To find the identifier for labels, run the *Get-Label* cmdlet (from the Security and Compliance module):

```
[PS] C:\> Get-Label | Format-Table DisplayName, Guid
```

Once you know the identifier for the label to apply, run *Set-UnifiedGroup*. For example:

```
[PS] C:\> Set-UnifiedGroup -Identity "Banking Group" -SensitivityLabelId "1b070e6f-4b3c-4534-95c4-08335a5ca610"
```

After a label is assigned with PowerShell, background synchronization updates SharePoint Online and Teams so that these workloads apply the label settings.

## Checking Guests in Sensitive Groups or Teams

When you assign a sensitivity label to a group or team that blocks guest access, any existing guests in the membership are not removed. Unless someone goes and checks, this can lead to a situation where unwanted guests linger in group memberships when they shouldn't. It's relatively straightforward to check the membership of all groups tagged with a specific label to see if they have guest members. Here's some code to do the job:

```
[PS] C:\> Write-Host "Finding confidential Groups..."
[array]$Groups = Get-UnifiedGroup | ? {$_ .SensitivityLabel -eq "c29e68f9-bc4f-413b-a741-6db8e38ad1c6" -and $_.GroupExternalMemberCount -gt 0}
If (!$Groups.Count) { Write-Host "No Groups found with that label"}
Else {
    $Report = [System.Collections.Generic.List[Object]]::new(); $NumberGuests = 0
    Write-Host "Now examining the membership of" $Groups.Count "groups to find guests..."
    ForEach ($Group in $Groups) {
        Write-Host "Processing" $Group.DisplayName
        [array]$Users = Get-UnifiedGroupLinks -Identity $Group.Alias -LinkType Members
        ForEach ($U in $Users) {
            If ($U.Alias -Match "#EXT#" -and $U.ExternalEmailAddress -NotLike "*teams.ms*") {
                $Domain = $U.ExternalEmailAddress.Split("@")[1]
                $ReportLine = [PSCustomObject]@{
                    Email           = $U.Name
                    User            = $U.DisplayName
                    Domain          = $Domain
                    Group           = $Group.DisplayName
                    Site            = $Group.SharePointSiteURL }
                $Report.Add($ReportLine)
                $NumberGuests++ }
        }
    }
}
$Domains = $Report.Domain | Sort -Unique; $DCount = $Domains.Count; $Domains = $Domains -join ", "
```

```
Write-Host ("{0} guests found in {1} confidential groups from these {2} domains: {3}" -f
$NumberGuests, $Groups.Count, $DCount, $Domains)
$Report | Sort Email | Out-GridView
$Report | Export-CSV -NoTypeInformation C:\temp\UnwantedGuests.csv
```

## Archiving Inactive Groups

Projects tend to have a natural lifetime. For example, let's assume that you spin up a group to support the planning and coordination of a financial project. After the project finishes, the group holds conversations about project issues, the calendar of meetings, and all the documents related to the project. Depending on the compliance regime that applies to the company operations, you might have to keep this information for an extended period. Unlike mailboxes, which administrators can put into an inactive state to preserve their content in a state that is inaccessible to users, we need a different approach to archive group content.

Conceptually, the steps to archive a group are easy. The aim is to put the group into a state where we hide the content from general view while the content stays indexed and discoverable for compliance purposes. After finding the set of groups to archive, here's what we might do to make the groups read-only to all but a new group owner:

- Add a new group owner. Ideally, this should be a special compliance administration account instead of a tenant administrator. Remember, before you can add an owner, they must first be a member of the group. Note that Teams cmdlets are used to update membership rosters when a group is team-enabled. This is to ensure that the new roster is active as quickly as possible.
- Check if the group is team-enabled and if so, check for private and shared channels so that you can add the compliance administrator as the owner for each channel.
- Remove all owners from the group's membership list.
- Remove all users from the group's membership list.
- Set *RequireSenderAuthenticationEnabled* property to *\$True* to stop Exchange Online from accepting any external email sent to the group. Guest members are not subject to this check as they can always email a group until their membership is removed.
- Assign a new primary SMTP address to the archived group and remove the old address to stop the delivery of new messages from users inside the tenant.
- Hide the group from Exchange clients (so it doesn't appear in Teams) and from address lists.
- Assign an appropriate sensitivity label to make the archived group private (if your tenant doesn't use sensitivity labels yet, you can change the access type of the group to be private instead).
- Update one of the custom properties available for group mailboxes to record some information about the group's archived status as well as information about who archived the group.

Although Microsoft doesn't offer an out-of-the-box method to archive inactive Groups, (Teams offers the option to archive a team by setting it to read-only mode), we can do the job with PowerShell. You can download a [script from GitHub](#) demonstrating how to archive groups using the steps described above.

A short time after the script runs, the group will disappear from clients. The exact time depends on the client. OWA is the fastest because it reads information from the online directory. It is slowest for Teams because of the need to synchronize the membership changes with Teams. Because we remove people from the group membership, they lose access to any resource available to the group like Planner or SharePoint.

## Archiving Hundreds of Groups

If you're in a situation where you need to archive hundreds of groups (or teams), it might be inconvenient to update and use a group attribute as the indicator that it should be archived. This situation often occurs in education, where large numbers of teams need to be archived at the end of the academic year. In the

enterprise, it might be the case that you scan for inactive groups and teams annually (perhaps using the [Groups and Teams Activity Report](#) PowerShell script) and archive those that aren't in use.

An alternative approach is to:

- Extract details of groups and write the data to a CSV file.
- Review the CSV file in Excel and remove groups that you want to keep (not archive). The advantage here is that someone who isn't skilled in PowerShell but knows how the groups and teams are used can do the review.
- Use the updated CSV file as input to the script.

For example, this code finds groups and creates a CSV file:

```
[PS] C:\> $Groups = Get-UnifiedGroup -ResultSize Unlimited | Select DisplayName, Notes, DistinguishedName, Alias, PrimarySmtpAddress, SensitivityLabel, ExternalDirectoryObjectId
$Groups | Sort DisplayName | Export-CSV -NoTypeInformation c:\temp\GroupsForReview.CSV
```

After editing the file to remove the groups you don't want to archive, we can then replace the call to *Get-UnifiedGroup* in [the script](#) with:

```
$ArchiveGroups = Import-CSV c:\temp\GroupsForReview.CSV
```

The rest of the code in the script is unaltered.

## Group Reactivation

To reactivate the group, all we need to do is to add a new owner, who can then add new members to allow them access to group resources. We should also reverse some of the settings made to group properties to make the group available to clients. The following command shows how to reverse some of the settings:

```
[PS] C:\> Set-UnifiedGroup -Identity ArchivedGroupAlias -HiddenFromAddressListsEnabled $False
-HiddenFromExchangeClientsEnabled:$False -CustomAttribute1 $Null
```

You should also change the assigned email address so that it doesn't include "archived" and update the display name. You don't need to change the setting to require email authentication unless you want to allow external users who are not guests to be able to email the group. Finally, if you use sensitivity labels, you should update the label assigned to the group to be whatever is appropriate for its new role.

## Tracking Archived Groups

As mentioned above, an optional but useful step is to record information about the archival of a group. Two custom properties store the status and information about when the archival occurred together with the username of the account that ran the script. The username comes from the credentials used to sign onto Exchange Online with the PowerShell session. The variable used will be different depending on how you connect. Marking archived groups in this manner makes it easier to find the groups if needed. Here's how we can find the groups archived with the script:

```
[PS] C:\> $Output = @{{Expression = {$_.DisplayName}; Label = "Group name"; width=30}, @{{Expression =
{$_.CustomAttribute2}; Label = "Archived on and by"}}
Get-UnifiedGroup -Filter {CustomAttribute1 -like "*Archived*"} | Format-Table $Output
```

Group name	Archived on and by
Sales Department team	Archived 07/15/2017 15:14:01 by Tony.Redmond@office365itpros.com
Nikon Camera Club	Archived 09/09/2017 23:15:12 by Administrator@office365itpros.com
Tips and Tricks for 365	Archived 01/10/2018 16:18:47 by Administrator@office365itpros.com

## Working with the Group Expiration Policy

Two sets of cmdlets are available to work with the group expiration policy:

- The *\*-MgGroupLifecyclePolicy* cmdlets manipulate the expiration policy settings.
- The *Add-MgGroupToLifecyclePolicy* and *Remove-MgGroupFromLifecyclePolicy* cmdlets are used when you want the expiration policy to process selected groups rather than apply to every group in the tenant. These cmdlets add or remove groups to the policy to bring them under the scope of the policy.

Because all Groups in a tenant can come within the scope of the group expiration policy, you should be aware that the settings you apply through the policy affect all applications that use Groups. In other words, the expiration of a group affects Teams, Planner, Yammer, Power BI, and other applications that use Groups to control their membership.

To retrieve the settings for the current expiration policy, you run this command:

```
[PS] C:> Get-MgGroupLifecyclePolicy | Format-List

AlternateNotificationEmails : Tony.Redmond@office365itpros.com
GroupLifetimeInDays       : 750
Id                        : 23036082-ac71-49ed-a8fd-24db5ad38379
ManagedGroupTypes       : Selected
AdditionalProperties      : {}
```

To update a policy setting, run the *Update-MgGroupLifecyclePolicy* cmdlet. In this case, we set the expiration interval to 750 days and update the backup group owner (the person who receives notifications about expiring groups when no group owner is available).

```
[PS] C:\> Update-MgGroupLifecyclePolicy -GroupLifecyclePolicyId (Get-MgGroupLifecyclePolicy).Id -
GroupLifeTimeInDays 750 -AlternateNotificationEmails Ben.Owens@Office365itpros.com
```

As an example of applying the policy to a selected set of groups instead of every group in the tenant, here is how to make the policy only process groups assigned a specific sensitivity label (identified by its GUID). We need to create a collection of groups marked as such and then add each group to the policy. The following code does the trick if you set the policy to process selected groups rather than all. Note that you cannot assign different expiration periods to individual groups. All groups covered by the policy use the same expiration period.

```
[PS] C:\> $Groups = (Get-UnifiedGroup | ? {$_.SensitivityLabel -eq "d6cfd185-f31c-4508-ae40-
229ff18a9919"})
$Id = (Get-MgGroupLifecyclePolicy).Id
$PolicyId = (Get-MgGroupLifecyclePolicy).Id
ForEach ($G in $Groups) {
    Add-MgGroupToLifecyclePolicy -GroupLifecyclePolicyId $Id -GroupId $G.ExternalDirectoryObjectId }
```

A more developed example scans the set of groups used by Teams and adds any which are not included in the expiration policy.

```
$PolicyId = (Get-MgGroupLifecyclePolicy).Id
$TeamsCount = 0
Write-Host "Fetching list of Teams in the tenant..."
[array]$Teams = Get-MgGroup -Filter "resourceProvisioningOptions/Any(x:x eq 'Team')" -All
ForEach ($Team in $Teams) {
    $Uri = " https://graph.microsoft.com/beta/groups/" + $Team.Id + "/groupLifecyclePolicies"
    $CheckPolicy = Invoke-MgGraphRequest -Uri $Uri -Method Get
    If ($CheckPolicy.Value.Id -eq $PolicyId) {
        Write-Host "Team" $Team.DisplayName "is already covered by the expiration policy" }
    Else {
        Write-Host "Adding team" $Team.DisplayName "to group expiration policy"
        Add-MgGroupToLifecyclePolicy -GroupLifecyclePolicyId $PolicyId -GroupId $Team.Id -ErrorAction
        SilentlyContinue
        $TeamsCount++ }}
Write-Host "All done." $TeamsCount "teams added to policy"
```

To remove a group from the expiration policy, run the *Remove-MgGroupFromLifecyclePolicy* cmdlet and pass the policy identifier and group identifier:

```
[PS] C:\> Remove-MgGroupFromLifecyclePolicy -GroupLifecyclePolicyId $PolicyId -GroupId $GroupId
```

## Reporting Groups Covered by The Expiration Policy

The set of groups covered by the expiration policy is visible in the Azure portal. Although no cmdlet is available to return the set of groups, you can find this information in two ways.

First, you can check each group in the tenant using a Graph request to confirm if a group is covered by the expiration policy:

```
[PS] C:\> [array]$Groups = (Get-Recipient -RecipientTypeDetails GroupMailbox -ResultSize Unlimited | Sort DisplayName)
ForEach ($G in $Groups) {
    $Uri = " https://graph.microsoft.com/beta/groups/" + $G.ExternalDirectoryObjectId +
"/groupLifecyclePolicies"
    $Status = Invoke-MgGraphRequest -Uri $Uri -Method Get
    If ($Status.Value) {
        $Days = (New-TimeSpan -Start $G.WhenCreated -End (Get-Date)).Days
        Write-Host "Group:" $G.DisplayName "Date created:" $G.WhenCreated "days old:" $Days }
}
```

An easier method is to check the group's *ExpirationTime* property, which Groups updates when it assesses a group against the expiration policy. We can therefore find the set of Groups subject to the expiration policy by checking the property. In this example, we find the set of groups with an expiration time set and report them in descending order:

```
[PS] C:\> [array]$ExpirationPolicyGroups = Get-UnifiedGroup | ? {$_.ExpirationTime -ne $Null} | Sort
{$_.ExpirationTime -as [DateTime]}
$ExpirationPolicyGroups | Format-Table DisplayName, ExpirationTime
```

DisplayName	ExpirationTime
GDPR Planning Mark II	28 May 2022 22:23:02
Office 365 for IT Pros	28 Jun 2022 09:42:57
Launch Party	11 Jul 2022 14:36:54
All-Employees	20 Aug 2022 20:23:38

Equipped with this knowledge, we can develop the code to report more interesting information. This version:

- Fetches the set of groups with an expiration time set.
- For each group, fetches the last renewed date from Azure AD.
- Calculates the days remaining before expiration.
- Calculates the age of the group in days (not used in the report, but it could be).
- Outputs a line into the report (a PowerShell list).

After processing all the groups, the report displays sorted by the number of days remaining before the group expires. You could pipe the output to the *Out-GridView* cmdlet to view on screen or to the *Export-CSV* cmdlet to export the data as a CSV file.

```
[PS] C:\> $Report = [System.Collections.Generic.List[Object]]::new(); $Today = (Get-Date)
Write-Host "Finding Microsoft 365 Groups to check.."
[array]$ExpirationPolicyGroups = (Get-UnifiedGroup -ResultSize Unlimited | ? {$_.ExpirationTime -ne
$Null} | Select DisplayName, ExternalDirectoryObjectId, WhenCreated, ExpirationTime )
If (!$ExpirationPolicyGroups) { Write-Host "No groups found subject to the expiration policy -
exiting" ; break }
Write-Host $ExpirationPolicyGroups.Count "groups found. Now checking expiration status."
ForEach ($G in $ExpirationPolicyGroups) {
    $Days = (New-TimeSpan -Start $G.WhenCreated -End $Today).Days # Age of group
    $LastRenewed = (Get-MgGroup -GroupId $G.ExternalDirectoryObjectId).RenewedDateTime
    $DaysLeft = (New-TimeSpan -Start $Today -End $G.ExpirationTime).Days
```

```

    $ReportLine = [PSCustomObject]@{
        Group      = $G.DisplayName
        Created     = Get-Date($G.WhenCreated) -format g
        AgeinDays  = $Days
        LastRenewed = Get-Date($LastRenewed) -format g
        NextRenewal = Get-Date($G.ExpirationTime) -format g
        DaysLeft   = $DaysLeft}
    $Report.Add($ReportLine)
} # End Foreach
CLS;Write-Host "Total Microsoft 365 Groups covered by expiration policy:"
$ExpirationPolicyGroups.Count
Write-Host ""
$Report | Sort DaysLeft | Select Group, @{n="Last Renewed"; e= {$_.LastRenewed}}, @{n="Next Renewal Due"; e={$_.NextRenewal}}, @{n="Days before Expiration"; e={$_.DaysLeft}}

Total Microsoft 365 Groups covered by expiration policy: 55

Group                                Last Renewed                Next Renewal Due            Days before
-----                                -
GDPR Planning Mark II                28 May 2022 22:23          16 Jun 2021 23:23           18
Office 365 for IT Pros                28 Jun 2022 09:42          17 Jul 2021 10:42           48
Launch Party                          11 Jul 2022 14:36          30 Jul 2021 15:36           62
All-Employees                         20 Aug 2022 20:23           8 Sep 2021 21:23           102

```

Two potential issues exist in this code. First, the *Get-UnifiedGroup* cmdlet is slow. Although it will fetch thousands of groups, you'll wait for this to happen. Second, a separate call to the *Get-MgGroup* cmdlet is needed to retrieve the last renewed date for a group because this property is not among those available through *Get-UnifiedGroup*. The solution to speed up the code is to replace these cmdlets with Graph API requests. An explanation of [how to upgrade the code is in this article](#).

## Managing Teams with PowerShell

The Teams PowerShell module is available in the [PowerShell Gallery](#). Microsoft updates the module every few months and you can install or update the module from the PowerShell Gallery. In terms of the administrative control over Teams, the module is capable of handling many of the operations available in the Teams admin center. [Cmdlets are available](#) to create teams (including a team for an existing group), list all teams in the tenant, update settings for teams, and create new channels.

It is best to use the most recent version of the module when working with Teams objects. This is especially true for Teams due to the ongoing work by Microsoft to refresh and improve cmdlets inherited from Skype for Business Online. To update to the latest release, run PowerShell as an administrator and use the following command to download and install the latest module.

```
[PS] C:\> Install-Module -Name MicrosoftTeams -Repository PSGallery -Force -Scope AllUsers
```

After installing the module, you run the *Connect-MicrosoftTeams* cmdlet to connect to the Teams service. To validate that you have the latest module, run the *Get-InstalledModule* cmdlet to check the version number.

```
[PS] C:\> Get-InstalledModule -Name MicrosoftTeams -AllVersions | Format-Table Name, Version
```

```

Name                Version
-----
MicrosoftTeams     4.4.1

```

### Preview Teams PowerShell Module

Microsoft publishes both the production (general availability) and public preview (beta) versions of the Teams module in the PowerShell Gallery. To install the preview release, check the PowerShell Gallery to find the latest



available version (or check the Teams PowerShell [release notes](#)) and then run the *Install-Module* cmdlet. To support the *AllowPreRelease* switch, you might have to update the *PowerShellGet* module as shown below:

```
[PS] C:\> Install-Module -Name PowerShellGet -Repository PSGallery -Force
Install-Module -Name MicrosoftTeams -RequiredVersion "4.5.9-preview" -AllowPreRelease -Scope
AllUsers
```

You can't run the generally available and preview modules together. If you need to switch, run the *Remove-Module* cmdlet to remove the cmdlets and functions for the currently loaded module from memory and then run the *Import-Module* cmdlet to load the desired version of the module (specified in the *Version* parameter).

## Permissions

Your account must have permissions to be able to run some of the Teams cmdlets.

- To create a new team: Tenant administrator, Teams service administrator, or be allowed to create new Groups.
- To change team settings, including updating team membership: Tenant administrator, Teams service administrator, or owner of the target team.
- To list team settings, including team membership: the account must be a member of the team.

## Finding Teams

Many scripts start by fetching the set of teams in a tenant and then go on to process each team. To return the set of teams, run:

```
[PS] C:\> $Teams = Get-Team
```

The *Get-UnifiedGroup* supports a server-side filter for the *ResourceProvisioningOptions* property, which should be set to "Team" when a group is team-enabled. Thus, we can also find the set of team-enabled groups with:

```
[PS] C:\> $Teams = Get-UnifiedGroup -Filter {ResourceProvisioningOptions -eq "Team"}
```

Although you can use *Get-UnifiedGroup* to find the set of team-enabled groups, it's often better to use the *Get-Team* cmdlet when working with teams. Because it does not return the large set of properties processed by *Get-UnifiedGroup*, *Get-Team* is usually faster to fetch objects. However, this approach has some limitations. For example, let's assume that you want to return the list of SharePoint Online sites used by Teams. There's no easy way to get this using the properties returned by *Get-Team*, but information about SharePoint is in the set of properties retrieved by *Get-UnifiedGroup*, so you can run a command like this:

```
[PS] C:\> $Teams = Get-UnifiedGroup -Filter {ResourceProvisioningOptions -eq "Team"}
$Teams | Format-Table DisplayName, SharePointSiteUrl
```

The most important properties returned by *Get-Team* are:

- Group identifier.
- Display Name.
- Description.
- MailNickName (otherwise called Alias).
- Visibility (called access type for Groups).
- Archived (either True or False to show if the team is in an archived state).

The group identifier is the most important value. It is a GUID that uniquely identifies the Azure AD group object for the team and is used as the input to most of the other cmdlets in the Teams module. In addition, you can use the group identifier to fetch information about the underlying group from Azure AD or the mail-enabled properties using the Groups cmdlets.

The list of teams is returned in order of creation date. To view the list in another order, use the *Sort-Object* cmdlet to sort by the *DisplayName*, *Description*, or *GroupId* properties.

```
[PS] C:\> Get-Team | Sort DisplayName
```

To list the teams that the currently logged-in user is a member of, specify the *User* parameter and pass their full User Principal Name:

```
[PS] C:\> Get-Team -User Kim.Akers@Office365itpros.com
```

You cannot use the *Get-Team* cmdlet to fetch details of the teams that another user belongs to. If you want to find the teams that a user belongs to, use *Get-Team* to fetch the full list of teams in the tenant and then *Get-TeamUser* to discover whether the user is a member of each team. The need to check the members of every team ensures that this will not be a fast process! Here's an example of the type of command needed:

```
[PS] C:\> Get-Team | ? {Get-TeamUser -GroupId $_.GroupId | ? {$_ .User -eq "Kim.Akers@office365itpros.com"}}}
```

This is a good example where a Microsoft Graph API delivers much better performance than PowerShell because the Teams endpoint supports the *joinedTeams* API specifically to facilitate fast access to the set of teams a user belongs to. Thus, we can use the *Get-MgUserJoinedTeam* cmdlet to do the job:

```
[PS] C:\> Get-MgUserJoinedTeam -UserId Kim.Akers@office365itpros.com
```

## Combining Teams Cmdlets with Other Modules

The limited properties of team objects returned by *Get-Team* mean that you often need to use other cmdlets, including those in other modules, to solve problems. For example, to get information from Azure AD about the team membership, use the group identifier as the input object identifier for *Get-MgGroupMember*:

```
[PS] C:\> [array]$Members = Get-MgGroupMember -GroupId 129bfa64-3d11-4d6f-b9a9-6c20ec03bd8d
ForEach ($Member in $Members) { Get-MgUser -UserId $Member.Id }
```

Another example is using the group identifier with the *Groups* cmdlets from the Exchange Online management module. For example, here's how to get the membership for the same team using *Get-UnifiedGroupLinks*:

```
[PS] C:\> Get-UnifiedGroupLinks -LinkType Member -Identity 129bfa64-3d11-4d6f-b9a9-6c20ec03bd8d
```

In a more complex example, we use *Get-Team* to update an array with details of the teams in a tenant and then use the *Get-UnifiedGroup* cmdlet to fetch some interesting properties about each group, and the *Get-TeamUser* cmdlet to fetch details of the team owners.

```
[PS] C:\> $Report = [System.Collections.Generic.List[Object]]::new()
[array]$Teams = Get-Team | Sort DisplayName
Write-Host "Reporting" $Teams.Count "teams"
ForEach ($T in $Teams) {
# Fetch information about the team - mostly by getting it from the Unified Group Object
Write-Host "Processing" $T.DisplayName
$G = (Get-UnifiedGroup -Identity $T.GroupId | Select GroupMemberCount, GroupExternalMemberCount,
PrimarySmtpAddress, AccessType, Classification, WhenCreated)
$Owners = (Get-TeamUser -Role Owner -GroupId $T.GroupId | Select Name)
$OwnerNames = $Null
$First = $True
ForEach ($O in $Owners) {
If ($First -eq $True) {
$OwnerNames = $O.Name
$First = $False}
Else {
$OwnerNames = $OwnerNames + ", " + $O.Name }}
$ReportLine = [PSCustomObject]@{
Team = $T.DisplayName
```

```

    Email           = $G.PrimarySmtpAddress
    Members         = $G.GroupMemberCount
    Guests          = $G.GroupExternalMemberCount
    Owners          = $OwnerNames
    Access          = $G.AccessType
    Created         = $G.WhenCreated
    Classification  = $G.Classification }
    $Report.Add($ReportLine) }
$Report | Export-Csv -NoTypeInformation c:\temp\Teams.csv

```

If your tenant uses sensitivity labels instead of classifications, remember to replace *Classification* with *SensitivityLabel*.

## Filtering Get-Team

The *Get-Team* cmdlet does not include a *Filter* parameter to ask the server to return a subset of teams. Instead, the *Get-Team* cmdlet supports several parameters to select common subsets. For instance, you can return a set of teams that are archived by running the command:

```
[PS] C:\> Get-Team -Archived $True
```

To return just the set of private or public teams, use *Private* or *Public* in the *Visibility* parameter. For instance:

```
[PS] C:\> Get-Team -Visibility Public
```

You can also combine these parameters to arrive at a more select subset of teams. For instance, this command returns the set of private teams that are archived that have Tony Redmond in their membership:

```
[PS] C:\> Get-Team -Visibility Private -Archived $True -User Tony.Redmond@office365itpros.com
```

Although server-side filtering isn't available with *Get-Team*, you can apply a client-side filter to the display name, classification, or description properties to find specific teams. For example:

```
[PS] C:\> Get-Team | ? {$_.DisplayName -Like "*Office*"} | Format-Table DisplayName, Description
```

We can combine a filtered lookup for a specific team to return its identifier and use that value to fetch information from Azure AD or the group:

```
[PS] C:\> $Group = (Get-Team | ? {$_.DisplayName -eq "Microsoft 365 Questions"} | Select GroupId)
Get-UnifiedGroup -Identity $Group.GroupId
```

The lack of good filtering for *Get-Team* doesn't matter as much when your tenant only has a few hundred teams. It becomes more of an issue as the number of teams increases.

## Creating New Teams

You can create a new team in two ways: either create a brand-new team from scratch or team-enable an existing group. If you add a new team from scratch, you also need to add some users.

The *New-Team* cmdlet applies in both cases. In this example, we create a new team from scratch and set its access type to *Private*, and apply one of the classifications defined for the tenant. To specify the owner of the new team, pass their User Principal Name in the *Owner* parameter. In this example, we also set the value of the *RetainCreatedGroup* parameter to *True* to tell PowerShell to keep the newly created Microsoft 365 group even if it's unable to create the new team:

```
[PS] C:\> $TeamId = (New-Team -DisplayName "Planning Co-Ordinators" -MailNickName Coordinators
-Description "Team for the folks who make sure we turn up in the right place at the right time"
-Visibility Private -Classification Confidential -Owner James.Ryan@office365itpros.com
-RetainCreatedGroup:$True)
```

**No sensitivity label:** The *New-Team* and *Set-Team* cmdlets do not support a method to assign a sensitivity label to a team. The example shown above assigns an old-style text-only classification. Until Microsoft updates the PowerShell module, you can create teams with sensitivity labels using the Graph API, or you can edit the team using the desktop or browser client to assign a label there.

The *\$TeamId* variable holds the GUID for the newly created team. We can use the GUID to identify the newly created team when we go to populate the team with members and owners. When you specify an owner for a new team, *New-Team* adds them as both a member and as an owner. The same occurs when you run the *Add-TeamUser* cmdlet. This code adds Brian Weakliam as a team member to the team and Don Vickers as both an owner and a member.

```
[PS] C:\> Add-TeamUser -GroupId $TeamId.GroupId -User Brian.Weakliam@Office365itpros.com -Role Member
Add-TeamUser -GroupId $TeamId.GroupId -User Donald.Vickers@Office365itpros.com -Role Owner
```

It's wise to add a delay of a few seconds before proceeding to populate membership for a new team. Waiting for a short period allows the provisioning process between Teams and Azure AD and Teams to complete and Teams to be ready to build the team roster. After populating the team membership, you can proceed to amend group settings and create some channels – or consider posting a welcome message to the General channel in the new team.

## Posting a Welcome Message to a New Team

A new team begins with no content. Sometimes it's nice to give members of the team some guidance about the purpose of the team and how to use it. PowerShell is used for administrative operations and the Teams module doesn't include a cmdlet to create and post messages to a Teams channel, but this can be done by calling the Graph API from PowerShell.

The basic requirement is to register a new app with Azure to act as the entry point for read-write access to Teams via the Graph. The details of how to register an app are described in the Microsoft Graph section. After registering an app, you can use its service principal to authenticate with the Graph to get an access token. The access token contains the permissions which allow the app to execute Graph API requests via PowerShell. In this example, we have created a new team and added the logged-in user to the team (otherwise they won't be able to post to a channel). The identifier for the new team is stored in the *\$GroupId* variable. We find the identifier for the General channel in the target team, use an array to compose some HTML text that we want to post, convert the array to JSON, and then post the message to the channel using the Graph.

```
[PS] C:\> $DisplayName = "Project Condor"
$GroupId = (Get-Team |? {$_.DisplayName -eq $DisplayName}).GroupId
$Description = (Get-Team |? {$_.DisplayName -eq $DisplayName}).Description
$MessageLine = "Team description:" + $Description
$ChannelId = (Get-TeamChannel -GroupId $GroupId |?{$_.DisplayName -eq "General"}).id
$apiUrl = "https://graph.microsoft.com/beta/teams/$GroupId/channels/$ChannelId/chatThreads"
# End up with something like https://graph.microsoft.com/beta/teams/41fd1ff0-099b-46b9-8b55-
e49c8e378383/channels/19:ba56cf956b0a47d991cb953906f605ee@thread.skype/chatThreads
$body = @{}
"rootMessage" = @{}
    "body" = @{}
        "contentType" = 1;
        "content" = "

# Welcome to $DisplayName



<b><i>Welcome to your new Team.</i></b><p>We have some basic rules of Teams etiquette that we would like you to follow to make everyone's life easier: </p><ul><li>Always start a new topic with a <strong>message subject</strong>. </li><li>When you have something to say about a topic, <strong>reply to that topic</strong> and don't start another thread.</li><li>Use <strong>@mentions </strong>to bring something important to the attention of individuals, teams, or a channel.</li></ul><p>$MessageLine</p>" }
}
$bodyJSON = $body | ConvertTo-Json
$msg = Invoke-RestMethod -Headers @{Authorization = "Bearer $accessToken"} -Uri $apiUrl -Method Post -Body $bodyJSON -ContentType 'application/json'


```

Figure 23-1 shows the welcome message generated by the code, which demonstrates the principle of what needs to be done to post to a team using PowerShell and needs some more work before it can be incorporated into the workflow to create a new team.

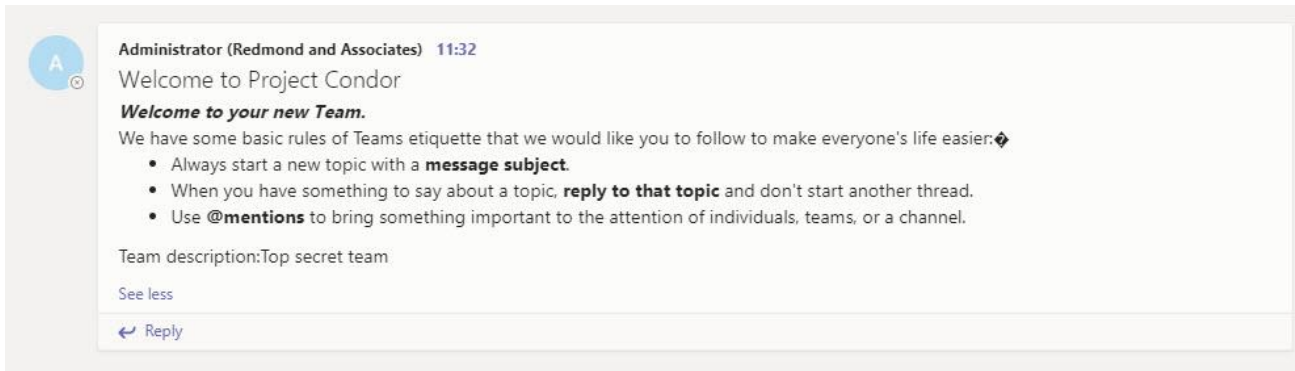


Figure 23-1: A welcome message for a new team generated by a Graph call through PowerShell

## Creating a Team from an Existing Group

To create a new team based on an existing group, run the *New-Team* cmdlet with the identifier (GUID) for the source group passed in the *Group* parameter. In this example, we fetch the group identifier and check if a group description (Notes) is available. If it is, a simple *New-Team* call is sufficient to ensure that the new team inherits the group description. If a description is not available, we create a description and include it in the call to *New-Team*:

```
[PS] C:\> $SourceGroup = Get-UnifiedGroup -Identity "BoardMembers"
If ($SourceGroup.Notes -eq $Null) {
    New-Team -Group $SourceGroup.ExternalDirectoryObjectId -Description "Team created from group" }
Else {
    New-Team -Group $SourceGroup.ExternalDirectoryObjectId }
```

You can only team-enable a group if your account is an owner of the group, a global administrator, or a Teams Service administrator.

## Controlling the Group Alias

The *Alias* property for the group created for a team can be specified in the *MailNickname* parameter for the *New-Team* cmdlet. This is an important property because it controls the generation of the SMTP primary address for the group (and team). If you don't specify a value to the *MailNickname* parameter, Teams generates unique values for the *Name*, *Alias*, and *PrimarySMTPAddress* properties (when the Teams client creates a new team, it creates the alias based on the team's display name). For example, here's what Teams generates for a new team as viewed with the *Get-UnifiedGroup* cmdlet when no value is passed for the *MailNickname*:

```
Name           : msteams_3e8d41_245f9292-29ca-4773-9240-de5fcee03853
Alias           : msteams_3e8d41
PrimarySmtAddress : msteams_3e8d41@office365itpros.com
```

The name given to the team is composed of "msteams," an underscore, and the team's Azure AD object identifier. The alias is derived from the name and the primary SMTP address uses the alias and the default tenant domain. Users only ever see the display name of a team, so having such a value in its name doesn't make any difference.

If you use Teams messaging for communication, you probably aren't concerned about the email address assigned to the group, so generating automatic values for these properties might not be an issue, but if you want to control the email address assigned to a new team, include the name you want to use in the *MailNickname* parameter. For instance, if I pass "Ignite2019" in the *Alias*, the following values are used:

```
Name : Ignite2019_bb966263-a263-4731-8454-afbd7be22b81
Alias : Ignite2019
PrimarySmtpAddress : Ignite2019@office365itpros.com
```

Make sure that the *MailNickName* you specify for *New-Team* is unique as the cmdlet will fail if you try to use a value that already exists for another object in the tenant.

## Creating an All-Employees Team

For tenants with 10,000 or fewer accounts, Teams supports org-wide teams with automatic maintenance of membership based on Azure Active Directory. The advantage of these teams is that Teams updates their membership in the background. We can do the same with PowerShell, with the advantage that we can impose whatever filter or other qualification we want to build the team membership. For example, the team might be built from all employees in one or more countries.

The group identifier is needed to add members to a team. Usually, we would fetch a group identifier using either the *Get-UnifiedGroup* or *Get-MgGroup* cmdlets, but in this case, we put the result of the call to the *New-Team* cmdlet in a variable so that the identifier is available for later use. Another interesting aspect is how we use *Get-Recipient* to fetch a set of current user mailboxes because it's the fastest method to get this data. You could also use *Get-Mailbox*, *Get-Recipient*, or *Get-MgUser* to create a set of input members. In all cases, remember to fetch the right property to get the user's primary email address as this is needed to identify the user when you add them to the team. No attempt occurs to check whether a user has a valid license to use Teams before we add them to the membership.

The last command in this code hides the new team from Exchange clients and sets a sensitivity label to control settings like guest access. Assigning a sensitivity label to a new team is only needed when you use labels for container management. You can't assign a sensitivity label with the *New-Team* or *Set-Team* cmdlets.

```
[PS] C:\> $Tenant = (Get-OrganizationConfig).Name
$Tenant = $Tenant.Split(".")[0]
$TeamName = "All-Employees ("+$Tenant+)"
$NewTeam = (New-Team -DisplayName $TeamName -MailNickName AllEmployeesTeam -Description "All-
employee team" -AllowCreateUpdateChannels $False -AllowDeleteChannels $False -
AllowCreateUpdateRemoveTabs $False -AllowCreateUpdateRemoveConnectors $False -AllowAddRemoveApps
$False)
$Mailboxes = (Get-Recipient -RecipientTypeDetails UserMailbox -ResultSize Unlimited)
$Owner = (Get-TeamUser -GroupId $NewTeam.GroupId -Role Owner)
# Add all the mailboxes except the owner, who's already there
ForEach ($M in $Mailboxes) {
    If ($M.WindowsLiveId -ne $Owner.User) {
        Add-TeamUser -GroupId $NewTeam.GroupId -User $M.WindowsLiveID -Role Member }
}
Set-UnifiedGroup -Identity AllEmployeesTeam -HiddenFromExchangeClientsEnabled -SensitivityLabelId
e42fd42e-7240-4df0-9d8f-d14658bcf7ce
```

The last two commands in the script hide the new team from Exchange clients and restrict the ability of team members to add or remove channels, connectors, and tabs from the team as it's unlikely that you would want to allow members to change the structure of an all-employees team. The resulting team is a skeleton that needs to be built out with channels, tabs, and apps. In addition, you should restrict posts to the general channel and add some owners to moderate what is often a busy forum. Finally, if you add new employees to the company after running the script, you must add them to the membership manually.

If your company has more than 10,000 employees, Yammer might be a better solution for an all-employees forum.

## Viewing and Updating Team Settings

Team settings control how users interact with a team.

- What members can do to affect content within a team.

- What guests can do inside a team.
- What use can be made of “fun” objects like GIFs and stickers inside conversations.

To see what members can do, run *Get-Team* as follows:

```
[PS] C:\> Get-Team -GroupId 72ee570e-3dd8-41d2-bc84-7c9eb8024dd4 | Format-List
AllowCreateUpdateChannels, AllowDeleteChannels, AllowAddRemoveApps, AllowCreateUpdateRemoveTabs,
AllowCreateUpdateRemoveConnectors, AllowOwnerDeleteMessages, AllowUserEditMessages,
AllowUserDeleteMessages, AllowTeamMentions, AllowChannelMentions

AllowCreateUpdateChannels      : True
AllowDeleteChannels            : True
AllowAddRemoveApps             : True
AllowCreateUpdateRemoveTabs    : True
AllowCreateUpdateRemoveConnectors : True
AllowOwnerDeleteMessages       : True
AllowUserEditMessages          : True
AllowUserDeleteMessages        : True
AllowTeamMentions              : True
AllowChannelMentions           : True
```

These values are all true, so we know that members of this team can:

- Create and update channels (*AllowCreateUpdateChannels*).
- Delete and restore channels (*AllowDeleteChannels*).
- Add or remove apps (*AllowAddRemoveApps*).
- Create, update, and remove tabs within channels (*AllowCreateUpdateRemoveTabs*).
- Create, update, and remove connectors within channels (*AllowCreateUpdateRemoveConnectors*).
- Team owners can delete all messages (*AllowOwnerDeleteMessages*).
- Members can delete their messages (*AllowUserDeleteMessages*).
- Members can @mention the team (*AllowTeamMentions*).
- Members can @mention a channel (*AllowChannelMentions*).

Guest members have the same rights as other members. However, you can control the ability of guests to add and remove channels with two settings. In this case, the settings are False, so guests cannot interfere with channels in this team.

```
[PS] C:\> Get-Team -GroupId 72ee570e-3dd8-41d2-bc84-7c9eb8024dd4 | Format-List
AllowGuestCreateUpdateChannels, AllowGuestDeleteChannels

AllowGuestCreateUpdateChannels : False
AllowGuestDeleteChannels       : False
```

To see the “fun stuff” settings for a team, use this command:

```
[PS] C:\> Get-Team -GroupId 72ee570e-3dd8-41d2-bc84-7c9eb8024dd4 | Format-List AllowGiphy,
GiphyContentRating, AllowStickersAndMemes, AllowCustomMemes

AllowGiphy          : True
GiphyContentRating  : moderate
AllowStickersAndMemes : True
AllowCustomMemes    : True
```

These settings tell us that:

- Giphy is enabled for the team (*AllowGiphy*).
- Moderate is selected for Giphy content (*GiphyContentRating*).
- Stickers and Memes are enabled for the team (*AllowStickersAndMemes*).
- Members can upload custom memes and use the memes in team conversations (*AllowCustomMemes*).

## Updating Team Settings

Apart from restricting channel creation for guests, the settings for the team that we've been looking at are quite liberal. You can update the settings for a team using the *Set-Team* cmdlet. Let's assume that we want to exert tighter control over channels, tabs, apps, and connectors and don't want to see custom memes used.

```
[PS] C:\> Set-Team -GroupId 72ee570e-3dd8-41d2-bc84-7c9eb8024dd4 -AllowCustomMemes $False
-AllowCreateUpdateChannels $False -AllowDeleteChannels $False -AllowAddRemoveApps $False
-AllowCreateUpdateRemoveTabs $False -AllowCreateUpdateRemoveConnectors $False
```

Remember that org-wide messaging and other policies applied to user accounts also affect how users interact with Teams and override the settings of an individual team.

## Updating Team Photos

Team owners can run the *Set-TeamPicture* cmdlet to update the image used for the team in listings. You can use PNG, JPEG, or GIF images, but make sure to size the image appropriately (100 x 80 pixels works well). Only owners can update the picture for a team. For example:

```
[PS] C:\> Set-TeamPicture -GroupId 34d68904-9d7c-4ef7-b715-eed283e80243 -ImageFile
c:\temp\TeamFile.jpg
```

If administrators need to update team pictures without becoming a team owner, they can run the *Set-UserPhoto* cmdlet to update the picture for a Microsoft 365 group. This also updates the team picture, so it's the best way to process team pictures centrally. For example:

```
[PS] C:\> Set-UserPhoto -Identity "Privacy Advocates" -PictureData
([System.IO.File]::ReadAllBytes("c:\temp\privacy.jpg")) -GroupMailbox
```

## Working with Team Membership

If you don't pass a value in the user parameter, the *Get-Team* cmdlet returns a list of the teams to which the logged-in user belongs. The cmdlet only returns three properties, the most important being the group identifier, or *GroupId*, which we use with most of the other cmdlets. This is the key to finding the group in Azure AD.

```
[PS] C:\> Get-Team | Format-Table GroupId, DisplayName, Description
```

GroupId	DisplayName	Description
8836e3ca-6531-402a-ac5f-b71c83d6affa	Mail Filter Patent	Discussing the Mail Filter patent
84fe8ec1-d85f-4564-a786-1f7bedd9862c	Engineering Testers	People who do testing for engineering
72ee570e-3dd8-41d2-bc84-7c9eb8024dd4	Office Chat	What's happening in the Office

If you pass the user principal name for an account to *Get-Team*, the cmdlet returns the set of teams that the user is a member of:

```
[PS] C:\> Get-Team -User Jane.Sixsmith@office365itpros.com | Format-Table GroupId, DisplayName,
Description
```

GroupId	DisplayName	Description
e7c971e0-653d-48ed-9611-4935bfebb6ca	Ongoing Budget Discussions	Budget planning team
9367a82e-9c65-4258-86fb-c2eb304c79ef	Privacy Advocates	Privacy in all its shape
15fe15f7-c2b9-4a2c-9ad8-bbd6a11d2d7f	Infrastructure and Technology Plans	Information relating to the

To list the membership of a team, use the *Get-TeamUser* cmdlet. This example outputs the display name (Name), role, and Azure AD account (User) for each member. Guest members are clearly marked.

```
[PS] C:\> Get-TeamUser -GroupId 72ee570e-3dd8-41d2-bc84-7c9eb8024dd4 | Format-Table Name, Role, User
```



Name	Role	User
----	----	----
Tony Redmond	owner	Tony.Redmond@office365itpros.com
Jeff Guillet	member	Jeff.Guillet@office365itpros.com
Administrator	member	Administrator@office365itpros.com
Vasil Michev (MVP)	guest	Vasil@Michev.info#EXT#@Office365itpros.onmicrosof...

You can output members with a specific role by including the `-Role` parameter and passing either *Member* or *Owner*:

```
[PS] C:\> Get-TeamUser -GroupId 72ee570e-3dd8-41d2-bc84-7c9eb8024dd4 -Role Owner
```

`Get-TeamUser` can return the membership of any team if you know the group identifier, which you can find out from Azure AD. For example:

```
[PS] C:\> Get-TeamUser -GroupId (Get-MgGroup -Filter "displayName eq 'Office 365 for IT Pros'").Id
```

Alternatively, use the `Get-UnifiedGroup` cmdlet to retrieve the group identifier:

```
[PS] C:\> Get-TeamUser -GroupId (Get-UnifiedGroup -Identity Hydra).ExternalDirectoryObjectId
```

`Get-TeamUser` returns the group membership for any kind of Azure AD group as the cmdlet does not check that the target group is a Microsoft 365 group. This proves that `Get-TeamUser` is simply another way to call the `Get-MgGroupMember` cmdlet. Likewise, the other `-TeamUser` cmdlets are proxies for the underlying `-MgGroupMember` cmdlets.

## Adding and Removing Team Members

The `Add-TeamUser` cmdlet adds a new member or owner to a team. Although Teams and Groups share a common membership list, you do not have to add an owner as a member of the team first as you do with the `Add-UnifiedGroupLinks` cmdlet. The `-Role` parameter tells `Add-TeamUser` what type of member the user is.

```
[PS] C:\> Add-TeamUser -GroupId 72ee570e-3dd8-41d2-bc84-7c9eb8024dd4 -User
"Don.Vickers@Office365itpros.com" -Role Member
Add-TeamUser -GroupId 72ee570e-3dd8-41d2-bc84-7c9eb8024dd4 -User Tony.Redmond@office365itpros.com -
Role Owner
```

Adding individual members to a team soon becomes a pain when there are more than a couple of new members to process. A simple loop through source data is a good way to add multiple users, as in this example which adds all the members of an email distribution list to a team. The code needs to check whether users in the distribution list are already members of the team, so we create a hash table holding the team members and use that to check a user's status. Building the table is a one-time operation that consumes some resources but checking an item against a hash table is much faster than an iterative loop through team membership. The key to the hash table is the display name of the mailbox or guest; the value is their email address and includes both local and guests. If you do not need to include guests, comment out the lines that create these entries. The table looks like this:

Name	Value
----	----
Robert Wille	rw99@outlook.de
Paul Cunningham	Paul.Cunningham@office365itpros.com
Stanislav Buldakov	sbuldakov@office365itpros.com
Mahmoud Magdy	memagdy@outlook.com

We can use the email address of mailboxes retrieved from the distribution list to check against the hash table to see if a member of a source distribution list is already part of the team. One issue always faced when using distribution lists as a source is how to handle the different recipient types that can be in the list. User mailboxes are straightforward because we can check whether they are in the team and add them if not. But then you must be able to deal with mail contacts, mail users, nested distribution lists, guest accounts, public

folders, and groups. This code only processes mailboxes and ignores the other recipient types. See [this post](#) to get an idea of how complex it can be to create a comprehensive script to report the membership of nested distribution lists.

```
[PS] C:\> $TargetGroupId = "8836e3ca-6531-402a-ac5f-b71c83d6affa"
$SourceDL = "GDPR Working Group"
#Create hash table to look up existing team members
$TeamMembers = @{}
Get-TeamUser -GroupId $TargetGroupId | % {
    $User = $_.User.toString()
    # Handle guests
    if ($User -like "*#EXT#") {
        $GuestUser = Get-MgUser -UserId $User
        $TeamMembers.Add($GuestUser.DisplayName, $GuestUser.Mail) }
    else {
        # local mailbox
        $Email = $_.User
        $User = $_.Name
        $TeamMembers.Add($User, $Email) }
}
$DLUsers = Get-DistributionGroupMember -Identity $SourceDL
ForEach ($U in $DLUsers) {
    Switch ($U.RecipientType){
        "UserMailbox" { # Mailbox - check and add if not found
            If ( $TeamMembers.ContainsValue($U.PrimarySmtpAddress)) -eq $False)
            { Add-TeamUser -GroupId $TargetGroupId -User $U.PrimarySmtpAddress -Role Member
              Write-Host "Mailbox" $U.Name "added to team" }
        }
        "MailUser" { # Could be a guest user
            Write-Host "Can't add Mail User" $U.Name "to the team. Please add them as a guest
user"; Break }
        "MailContact" {
            Write-Host "Teams doesn't support Mail contacts" $U.Name "not added"; Break }
        "MailUniversalDistributionGroup" {
            Write-Host "Need to do something with this nested DL" $U.Name; Break }
    }
}
```

The *Remove-TeamUser* cmdlet removes a member from a team. You do not have to specify the user's role. If you remove a team owner, the account is also removed as a member. The exception is when the owner being removed is the only owner for the team. In this scenario, Teams requires you to assign someone else (a current user or a new member) to be an owner to ensure that the team is not left ownerless.

```
[PS] C:\> Remove-TeamUser -GroupId 72ee570e-3dd8-41d2-bc84-7c9eb8024dd4 -User
"Don.Vickers@Office365itpros.com"
```

A tenant administrator or an account holding the Microsoft 365 User Management admin role can add a member to any group (and team) by fetching the group identifier with the *Get-UnifiedGroup* cmdlet and using it with *Add-TeamUser* to update membership for the selected team. This works even if a group is not team-enabled. For example:

```
[PS] C:\> Add-TeamUser -GroupId (Get-UnifiedGroup -Identity "Stock Market
Club").ExternalDirectoryObjectId -User Donald.Vickers@office365itpros.com -Role Member
```

Because a team always has an underlying group, you also can use the cmdlets that work against groups to view or modify team membership. However, any changes made through cmdlets like *Add-UnifiedGroupLinks* take extra time to be effective within Teams as the change must synchronize from Azure AD to Teams, so overall it is best to use the Teams cmdlets to manipulate team membership with PowerShell.

## Working with Channels

Channels can be standard (regular), private, or shared. To return the channels in a team, use the *Get-TeamChannel* cmdlet:

```
[PS] C:\> Get-TeamChannel -GroupId 72ee570e-3dd8-41d2-bc84-7c9eb8024dd4
```

Id	DisplayName	Description
19:534ebe8882364e68acd46935b126fdca@thread.skype	General	Exchange Grumpy Old ...
19:f1dc5034b5d54d01b97b14f2d37dfc31@thread.skype	The EHL0 blog	Posts from the EHL0 ...
19:d286c7126d6a406f84fb74981e9276c2@thread.skype	MVP Tweets	Tweets from MVPs
19:96a1c26a207642b8bbf7822cfd606f75@thread.skype	Bug Tracking	

This output doesn't show the channel type. To see this information, run *Get-TeamChannel* and specify the channel type (Standard, Private, or Shared) in the Membership parameter:

```
[PS] C:\> Get-TeamChannel -GroupId $GroupId -MembershipType Private
```

You cannot update the settings of the General channel. For other channels, the only settings updateable with the *Set-TeamChannel* cmdlet are the channel name and description. In this example, we update the name and description of a channel:

```
[PS] C:\> Set-TeamChannel -GroupId 72ee570e-3dd8-41d2-bc84-7c9eb8024dd4 -CurrentDisplayName "Service Updates" -NewDisplayName "Updates for Service Incidents" -Description "Incident information extracted from the Service Management API"
```

The *New-TeamChannel* cmdlet adds a channel to a team. For instance:

```
[PS] C:\> New-TeamChannel -GroupId 72ee570e-3dd8-41d2-bc84-7c9eb8024dd4 -DisplayName "Bug Tracking" -Description "A place to track bugs as we discover them in our favorite product"
```

If you don't specify the membership type for a new channel, Teams creates a standard channel. To create a private or shared channel, set its *MembershipType* to the appropriate type and make sure to include a channel owner, who must be a member of the team.

```
[PS] C:\> New-TeamChannel -GroupId $GroupId -DisplayName "Project Hydra" -Description "Discussions about the Hydra Project" -MembershipType Private -Owner Tony.Redmond@office365itpros.com
```

Once set, you can't change a channel type.

The *Remove-TeamChannel* cmdlet deletes a channel from a team. When you remove a private or shared channel, you also remove the SharePoint site belonging to the channel.

```
[PS] C:\> Remove-TeamChannel -GroupId $GroupId -DisplayName "Project Hydra"
```

## Private Channel Membership

You can populate the membership of a private channel with the *Add-TeamChannelUser* cmdlet. If you want to add a new owner for the channel, make sure that you add them as a member first. Here's an example:

```
[PS] C:\> Add-TeamChannelUser -GroupId $GroupId -DisplayName "Legal Discussions" -User Oisin.Johnston@Office365itpros.com
Add-TeamChannelUser -GroupId $GroupId -DisplayName "Legal Discussions" -User Oisin.Johnston@Office365itpros.com -Role Owner
```

To list the members in a private channel, run the *Get-TeamChannelUser* cmdlet.

```
[PS] C:\> Get-TeamChannelUser -GroupId $GroupId -DisplayName "Legal Discussions"
```

To remove a member from a channel, run the *Remove-TeamChannelUser* cmdlet. For example:

```
[PS] C:\> Remove-TeamChannelUser -GroupId $GroupId -DisplayName "Legal Discussions" -User
James.Ryan@office365itpros.com
```

Teams won't allow you to remove the last owner from a channel. If you need to remove someone who is the last owner, you'll have to add a member as the owner and then remove the old owner.

## Shared Channel Membership

Adding a member to a shared channel uses the same syntax as a private channel, but only if the member you add has a tenant account. You cannot invite an external member from another tenant using the *Add-TeamChannelUser* cmdlet, nor can you use the cmdlet to invite a team from the home tenant or an external tenant.

The *Get-TeamChannelUser* cmdlet doesn't list teams that are members of a shared channel. The cmdlet only returns individual members from the home and external tenants.

## Checking for Active Teams

Administrators often ask for a way to know what groups are enabled for Teams or the other resources available to groups. The *Get-UnifiedGroup* cmdlet returns the basic properties of groups, including the number of members, the URL for SharePoint resources, and whether connectors are enabled, including if a group is team-enabled. If you're already working with the Exchange Online PowerShell module and don't want to load another module just to fetch a list of teams, it's reasonable to use *Get-UnifiedGroup* for this purpose. However, in most situations, the *Get-Team* cmdlet is the right way to find the team-enabled groups in a tenant.

Knowing the set of teams in a tenant is one thing. Figuring out if the team is active is another. One approach to test for team activity is to use the *Get-ExoMailboxFolderStatistics* cmdlet to check the group mailbox to see if compliance records are present. As explained earlier, the substrate generates compliance records for messages in all standard channel conversations within a team. If the folder in a team's group mailbox holds some items, it is a good sign that the team has been active. The caveat is that the group mailbox does not hold compliance records for conversation activity in private and/or shared channels, so this activity does not show up when you check the number of compliance records in the group mailbox.

The script ([downloadable from GitHub](#)) makes a collection of all teams in the tenant and runs *Get-ExoMailboxFolderStatistics* for each team to check its activity level. The script generates a report with details of each team, including the last time something happened in a channel and the number of conversations per day. The script also assigns an active status based on the number of conversations per day. Again, it is easy to adjust the code to process a team differently. The code extract below shows how to extract the conversation statistics for a team:

```
[PS] C:\> $TeamsData = (Get-ExoMailboxFolderStatistics -Identity $G.Alias -FolderScope NonIpmRoot -
IncludeOldestAndNewestItems | ? {$_.FolderType -eq "TeamsMessagesData"})
If ($TeamsData.ItemsInFolder) {
    Write-Host "Processing" $G.DisplayName
    $TimeSinceCreation = (Get-Date) - $TeamsData.CreationTime
    $Count++
    $ChatCount = $TeamsData.ItemsInFolder
    $NewestChat = $TeamsData.NewestItemReceivedDate
    $ChatsPerDay = $ChatCount/$TimeSinceCreation.Days
    $ChatsPerDay = [math]::round($ChatsPerDay,2)
} #End if
If ($TeamsData.ItemsInFolder -eq 0) {
    Write-Host "No Teams compliance records found for" $T.DisplayName -foregroundcolor Red
    $ChatsPerDay = 0
    $NewestChat = "N/A"
    $ChatCount = 0
}
If ($ChatsPerDay -gt 0 -and $ChatsPerDay -le 2) { $ActiveStatus = "Moderate" }
```

```
Elseif ($ChatsPerDay -gt 2 -and $ChatsPerDay -le 5) { $ActiveStatus = "Reasonable"}
Elseif ($ChatPerDay -gt 5) { $ActiveStatus = "Heavy" }
```

The output of the script is a CSV file, which can be opened with Excel to format the data to meet your requirements. Several properties included in the report are extracted from groups. If you want to include extra group properties, you can add them to the *Select* statement piped from the call to *Get-UnifiedGroup* and include the properties in the output line generated for each group. It's also easy to get a quick oversight of the activity level within the tenant by grouping the assigned activity status like this:

```
[PS] C:\> $Report | Group ActiveStatus | Sort count -Descending | Format-Table Name, count
```

Name	Count
Moderate	38
Inactive	27
Reasonable	1

The *Get-ExoMailboxFolderStatistics* cmdlet is expensive in terms of the system resources it consumes, so running it to check hundreds of groups is impractical if you intend to create the report very regularly. The solution explored here is imperfect and flawed, but it shows what is possible by combining Teams cmdlets with cmdlets from other modules.

## Archiving a Team

The *Set-TeamArchivedState* cmdlet archives a team. As explained in the Managing Teams chapter, the archive process puts the messaging part channels into a read-only mode. Team members can access conversations, but they cannot add new topics. Team owners and administrators can add or remove members and can unarchive the team when they want. Optionally, the SharePoint Online site belonging to the team can also be put into a read-only state. In this code snippet, we put the group identifier for the team to archive into a variable and use it as input to *Set-TeamArchivedState*. The *SetSPOSiteReadOnlyForMembers* parameter controls if the SharePoint Online site is read-only (\$True) or not (\$False). The default is \$False.

```
[PS] C:\> $TeamId = "bd2c9ef7-909f-4466-8375-2467eb826c63"
Set-TeamArchivedState -GroupId $TeamId -SetSPOSiteReadOnlyForMembers:$True -Archived:$True
```

## Cmdlets to Manage Teams Policies

Some of the cmdlets to manage Teams policies originated in Skype for Business Online. Microsoft has modernized these cmdlets to remove the need to use basic authentication and incorporated the cmdlets in the Teams PowerShell module. In most cases, you won't need to use the cmdlets to create a new policy or tweak policy settings because it's easier to do this through the Teams admin center. However, if you want to apply a new or updated policy to many accounts, it's easiest and much faster to do this with PowerShell. Table 23-3 lists the seven most important cmdlet sets for policy management.

<b>Cmdlet Set</b>	<b>Purpose</b>
*-CsTeamsAppPermissionPolicy	Controls which apps Teams users can access. See <b>Permissions policies</b> under Teams apps in the Teams admin center.
*-CsTeamsAppSetupPolicy	Controls the set of apps pinned to the Teams app bar. See <b>Setup policies</b> under Teams apps in the Teams admin center.
*- CsTeamsCallingPolicy	Controls the calling policies available within a tenant. See <b>Calling policies</b> under Voice in the Teams admin center.
*- CsTeamsChannelsPolicy	Controls the org-wide settings for channels. See <b>Teams policies</b> under Teams in the Teams admin center.
*- CsTeamsFeedbackPolicy	Controls the display of feedback surveys and the Give feedback option for Teams clients. This policy isn't available in the Teams admin center.

*- <i>CsTeamsMeetingPolicy</i>	Controls the type of meetings users can create and the features they can use in meetings. <b>See Meeting policies</b> under Meetings in the Teams admin center.
*- <i>CsTeamsMessagingPolicy</i>	Controls the messaging features available to users. See <b>Messaging policies</b> in the Teams admin center.

Table 23-3: Important Teams policy cmdlets

Each set has a *Get-* cmdlet (fetch policy settings), *Set-* cmdlet (update policy settings), *New-* (create new policy of the type) and *Grant-* cmdlet (assign a policy to a user). For example, to assign a new Teams app permissions policy to all mailboxes in a tenant, you could run this code:

```
[PS] C:\> $Mbx = Get-Mailbox -RecipientTypeDetails UserMailbox -ResultSize Unlimited | Select
DisplayName, UserPrincipalName
ForEach ($M in $Mbx) {
    Write-Host "Processing" $M.DisplayName
    Grant-CsTeamsAppPermissionPolicy -PolicyName "New Teams App Permissions Policy" -Identity
    $M.UserPrincipalName }
```

The identity used as the input for the *Grant-* cmdlets is the account's user principal name. In the example above, Exchange Online mailboxes are used as the source of the identity. Another way of finding a set of accounts to use as input for granting a policy is to use the *Get-CsOnlineUser* cmdlet to fetch the set of users enabled for Teams. For example:

```
[PS] C:\> Get-CsOnlineUser | Grant-CsTeamsMessagingPolicy -PolicyName "Advanced Users"
```

## Update Teams Client Settings

The *Set-CsTeamsClientConfiguration* cmdlet updates the general client settings policy to control features such as email integration and whether the organization tab appears in channels. For example, this command updates the global client settings policy to disable Google Drive as a valid cloud storage location:

```
[PS] C:\> Set-CsTeamsClientConfiguration -AllowGoogleDrive $False -Identity Global
```

## Update Guest Settings for Teams

The settings that control the set of Teams functionality available to guests are managed with the *Set-CsTeamsGuestMessagingConfiguration* cmdlet. For example, to allow guests to chat with other users:

```
[PS] C:\> Set-CsTeamsGuestMessagingConfiguration -AllowUserChat $True
```

You can also disable chat during meetings for all or some users through the *MeetingChatEnabledType* setting in the *CsTeamsMeetingPolicy*. By default, the setting is *Enabled* to allow chat during meetings. To disable chat, change the setting to *Disabled* in a meeting policy and assign that policy to the target set of users.

## Using Cmdlets from Other Modules to Manage Teams

As obvious from the examples here and elsewhere in the book, other cmdlets in the Exchange Online and Azure AD modules are often useful when dealing with Teams. For instance, every group has twenty custom attributes for tenant-specific use. Fifteen (*CustomAttribute1* through *CustomAttribute15*) support single-value properties and five (*ExtensionCustomAttribute1* through *ExtensionCustomAttribute5*) support multi-value properties. You can populate these attributes with whatever values you need. For instance, you might want to add a department name to each team to track how many teams each department uses. In this example, we use the first single-value attribute to hold a department name.

```
[PS] C:\> Set-UnifiedGroup -Identity Team1 -CustomAttribute1 Marketing
```

Updating a multi-value attribute uses a different syntax:

```
[PS] C:\> Set-UnifiedGroup -Identity Team1 -ExtensionCustomAttribute1 @{Add="IT"}
```

When the attributes are set, you can easily retrieve just the teams for a specific department or generate reports sorted by department. For example, to process the teams that belong to the IT department, we can do something like this to fetch a list of groups stamped with a departmental value and then extract the membership for each team.

```
[PS] C:\> $Teams = (Get-Recipient -RecipientTypeDetails GroupMailbox -Filter  
{ExtensionCustomAttribute1 -eq "IT"})  
ForEach ($Team in $Teams) {  
    Write-Output $Team.DisplayName  
    Get-TeamUser -GroupId $Team.ExternalDirectoryObjectID  
    Write-Output "" } }
```

## Sending Messages with PowerShell

When you strip everything away, mail servers like Exchange Online are all about sending messages. Normally, you use a client to compose and send messages, but it is possible to do the job with PowerShell. Being able to create and send messages in code is often useful when the need exists to dispatch notification messages.

Three methods are available:

- **The *Send-MailMessage* cmdlet:** This is the most common method in use today. *Send-MailMessage* uses the SMTP AUTH protocol to connect to Exchange Online with basic authentication.
- **The *.Net SmtplibClient Class*:** You might come across older scripts that use a combination of the [.NET SmtplibClient](#) and [MailMessage](#) classes to send emails from PowerShell. This method exploits the .NET mechanism available to any programming language to create and send SMTP messages. Microsoft has deprecated this method and you should replace the code in any script which uses it to send an email.
- **The *SendMail* call from the Outlook Graph API:** Like other Graph APIs, *SendMail* uses modern authentication. As explained in the section about the Graph, it is relatively straightforward to convert a script using *Send-MailMessage* to use the Graph.

In the future, Microsoft might upgrade the *Send-MailMessage* cmdlet to force the SMTP AUTH protocol to use modern authentication. Until then, if you want to move away from basic authentication, you must use the Graph.

### Using Send-MailMessage

The code needed to create and send a message using the *Send-MailMessage* cmdlet is very simple. This example uses what's known as the authenticated client submission method (because it connects to a mailbox) and the following steps occur:

- Make sure that TLS 1.2 is used by the workstation where the script is run to negotiate a connection to Exchange Online. If this isn't done, Exchange will refuse the connection.
- Check whether suitable credentials to allow the script to send the message through an Exchange Online mailbox are available. If not, call the *Get-Credential* cmdlet to collect the credentials.
- Set up some basic parameters such as the SMTP server and port number to use. For Exchange Online, the server is smtp.office365.com and the port is 587.
- Create a basic HTML style for the body of the message.
- Add some text for the body. The example shows using text that includes HTML formatting commands. This approach works well for simple text, but for more complicated messages it's easier to create the HTML content in a text file and then read it into a variable using the *Get-Content* cmdlet. Apart from anything else, using a text file means that the script code doesn't need to be edited every time you want to change the message text.

- Populate the message parameters in an array and call *Send-MailMessage* to dispatch the message via the nominated mailbox. Exchange won't allow you to connect to a mailbox and messages with *Send-MailMessage* if the mailbox is blocked for SMTP AUTH submissions.

```
[PS] C:\> [Net.ServicePointManager]::SecurityProtocol = [Net.SecurityProtocolType]::Tls12
If (-not $SmtpCred) {
    $SmtpCred = (Get-Credential)}
$MsgFrom = $SmtpCred.UserName
$SmtpServer = "smtp.office365.com" ; $SmtpPort = '587' ; $MsgSubject = "Action Required"

$htmlhead="<html>
<style>
    BODY{font-family: Arial; font-size: 10pt;}
    H1{font-size: 22px;}
    H2{font-size: 18px; padding-top: 10px;}
    H3{font-size: 16px; padding-top: 8px;}
</style>"

$htmlBody = "<body>
<h1>Group Non-Activity Notification</h1>
<p><strong>Generated:</strong> $Date</p>
Please review the activity in the $GroupName group as it doesn't seem to have been used too
much recently. Perhaps we can remove it?"
$htmlMsg = $htmlHead + $htmlBody
# Construct the message parameters and send it off...
$msgParam = @{
    To = "Vasil.Michev@office36itpros.com"
    From = $MsgFrom
    Subject = $MsgSubject
    Body = $htmlMsg
    SmtpServer = $SmtpServer
    Port = $SmtpPort
    Credential = $SmtpCred }
Send-MailMessage @msgParam -UseSSL -BodyAsHTML
```

This code could be enhanced to include multiple addressees (separate the addresses with commas), including BCC and CC recipients, or by adding one or more attachments. If you want to send messages to more than a few recipients, remember that Exchange Online limits the number of messages any individual mailbox can send to 30 per minute. With this limit in mind, you might need to include a delay of a second or two after sending each message.

Very importantly, Exchange Online sends the messages created using this method using the credentials of a mailbox-enabled account. The mailbox must exist, and it must have enough remaining quota to store the copies of the messages in the Sent Items folder.

Exchange Online also supports direct send and connector-based methods to send messages from SMTP-enabled appliances. Often appliances use these methods to send notification messages when they finish processing a job. See [this article](#) for more information on these methods.

## Sending Multiple Messages

You can use the techniques explained above to send a message to multiple recipients. For example, you could read a list of addresses from a CSV file and loop through the set to create and send a message to each address. When you generate multiple messages like this, you need to be aware of two important points:

1. If you send more than 200 or so messages from a mailbox in a short period, Exchange Online will suspect that either the mailbox is sending spam or an attacker has compromised the account owning the mailbox. In either case, Exchange Online will block the mailbox from sending messages and flag a high-priority alert to the administrator. You can release mailboxes from a restricted status by [going to this page](#).



2. When sending a continuous stream of messages, a single SMTP connection is used. Exchange Online allows the connection to be open for a maximum of ten minutes before it drops. If you build a pause between messages into your script (with the *Start-Sleep* cmdlet), make sure that the number of messages processed multiplied by the number of seconds paused between each message does not exceed 600. If it does, reduce the number of messages or reduce the pause between each message.

# Understanding How to Use the Microsoft Graph to Manage Microsoft 365

The basic workloads originally brought the tools used on-premises to cloud automation. After its introduction in Exchange 2007, the focus for Exchange was PowerShell, which quickly became part of the Exchange administrator's toolkit as a powerful way to automate common operations. SharePoint also supports PowerShell, but the coverage of the SharePoint PowerShell module is less comprehensive and powerful than its Exchange counterpart. For this reason, developers usually prefer the SharePoint PnP ([Patterns and Practice](#)) initiative when they want to extend SharePoint functionality. For more information about SharePoint PnP, see the discussion in Chapter 8.

## The Microsoft Graph APIs

The services composing Microsoft 365 create a different environment from the on-premises world. For programmatic access to both user and system data, Microsoft's focus for development is a common method for programmatic access to all manner of entities taken from workloads drawn from across Microsoft 365. This is the [Microsoft Graph](#), a set of REST-based API that interact with endpoints for different workloads drawn from across Microsoft 365. An endpoint is something like mail, contacts, calendar, or groups. The Graph connects the different endpoints and the entities hosted by the endpoints. For example, interacting with a team means using data from:

- Azure AD (user and group).
- Exchange Online (mailbox and team calendar).
- SharePoint Online (site).
- Teams (rosters, channels, tabs, and other team-specific information).

Once a program successfully authenticates, it can consume any of the Graph APIs using standard HTTP requests without needing to reauthenticate. A program can fetch some data from Azure AD (like the membership of Microsoft 365 group for a team), discover who the team owners are, and retrieve details of events in the team calendar and the SharePoint Online sites used by shared channels. Despite the variation in the entities the program interacts with, the format of the requests are similar, and all the requests use the same access token granted by the authentication process.

Microsoft has deliberately made the Graph approachable by including support for a wide choice of development technologies. Programmers can use their preferred language to issue HTTP requests against the Graph endpoints. For example, GET requests fetch data while POST requests write data. Most of the time, requests follow the same approach whether they deal with user, group, or device data. However, because different product groups are responsible for different APIs, it's inevitable that some differences in implementation exist, even if the approach appears to be unified.

The Microsoft Graph is a critical part of Microsoft's platform. Microsoft developers use Graph APIs to create new applications like Teams and the mobile apps for many workloads within Microsoft 365. The Microsoft Graph also includes data from other services, such as the Microsoft partner portal and Microsoft's consumer

cloud services, meaning that programmers can mix and match data drawn from multiple sources within the same application. By using the Graph APIs, programmers can do the following:

- Access data from multiple applications.
- Traverse information users access to track their usage of the service.
- Understand data trends to be able to find data that is important to the tenant or groups of users.
- Build new tools and applications using Microsoft 365 data.

As programmers start to use the Microsoft Graph to build applications on top of Microsoft 365, Microsoft hopes that they will view Microsoft 365 as a complete fabric rather than limiting their vision in the same way as they would when working with a single application. This is the approach taken to create applications like Teams.

Although you can grab a PowerShell script from the Internet and run it straight away (subject to the execution policy on a workstation), the additional complications which exist, notably the need to create a registered app and set up whatever authentication the application uses, means that the same is not true for Graph-based applications. Once someone knows the process, taking code from elsewhere and adapting it to run in a tenant is reasonably straightforward.

## V1.0 and Beta

Workloads often support two endpoints. The V1.0 endpoint is the production-ready fully-supported version. It is the endpoint that Microsoft wants people to use. The beta endpoint is under active development and API requests executed against this endpoint or the results they return can differ over time. Microsoft can change the beta endpoint without warning and without any documentation, so some risk exists in using this endpoint over V1.0. However, sometimes developers need to use the beta endpoint because it supports specific functionality required by their code. Accessing license information for Azure AD accounts is a good example of functionality that was only available through the beta endpoint for a long time.

Although it is preferable to use the V1.0 endpoint whenever possible, it's certainly permissible to use the beta endpoint, even in production, once you understand the risk. Indeed, many of the Microsoft administrative portals have used the beta endpoint to make new features available to customers. However, after Microsoft upgrades the V1.0 endpoint to support whatever feature you need, it's wise to redo your code to use it rather than the beta endpoint.

## Graph Permissions

Not all is rosy and perfect with the Graph. In exchange for its power, some conditions exist to gain entry. One example is the requirement that Graph applications must authenticate against Azure AD; to make your code do this you have to have a good understanding of how Modern Authentication (OAuth 2.0) works and how to request the necessary permissions to run the API requests necessary to interact with data. Likewise, your Graph applications must gain consent to use (the right) permissions in the target Azure AD tenant to work. Graph permissions break down into two sets:

- **Delegated permissions** operate as if a signed-in user runs the query. Some Graph APIs, like Planner, only support delegated permissions. An app using delegated permissions can never gain more permission than the signed-in user. For example, even if the app is granted the *User.ReadWrite.All* permission (the ability to read and write the user profile of every account in the organization), the app can only access the accounts that the signed-in user can. If the signed-in user holds the global administrator or user administrator role, the app can use the administrative permissions assigned to their account to access all accounts. If not, the app can only update the profile of the signed-in user's account. See [this article](#) for more information about how delegated permissions combine with those held by an account to create the set of effective permissions for use with data.

- **Application permissions** run without a signed-in user. These permissions are suitable for unattended automation such as apps running as background services or as daemons. Because apps receive the full benefit of application permissions, only administrators can grant consent when an app receives these permissions. Unlike the limitation placed on delegated permissions, an app can use an application permission like *User.ReadWrite.All* to access every user account in the tenant.

Apart from guessing what Graph permissions are necessary for an action, three methods are available to find this information:

1. Use the Graph Explorer (explained below). After running a query, the Modify permissions tab shows what permissions allow the query to run.
2. Read the Graph documentation. The documentation for each API lists the delegated and application permissions it supports.
3. Use the *Find-MgGraphPermission* and *Find-MgGraphCommand* cmdlets from the Microsoft Graph PowerShell SDK. These cmdlets list the permissions necessary to perform actions against different objects (like users or groups) with SDK cmdlets. See the later explanation for more detail.

Sometimes these methods will list several permissions which support an action. The Graph operates on a least permission model, meaning that you should select the permission with the least capabilities to perform an action. For example, to list users, use the *Users.Read.All* permission instead of the *Users.ReadWrite.All* permission. As we'll discuss later, the danger exists that apps used to run Graph API requests can become over-permissioned over time, so it's wise to keep an eye on the set of permissions an app has.

## Registered Azure AD Apps

To execute Graph API requests, a PowerShell script must prove it has the permissions to run the requests. In many situations, administrative processing requires permission to access data across a range of accounts using application permissions as described above. In these scenarios, it is necessary to register an app in Azure AD. A registered app acts as a container to hold permissions and secrets (see below). It also has a service principal that can authenticate with Azure AD to secure an access token confirming the permissions held by the app. According to a Microsoft Graph architect, you can think of a service principal as a convenient object to hold consents granted to an app. The service principal takes the same name as the registered app. You can use the service principal to track sign-ins by the app through the Azure AD admin center or via a query to the Graph *SignIn* API.

To register a new app, go to the App registrations section of the Azure AD admin center. You'll then need to add:

- The name of the app. The name doesn't matter as only administrators will likely see this detail. However, it is a good idea to give apps names that indicate their function.
- The supported account types. In most cases, this choice is *accounts in the organization only*. However, you can create apps that can run in any Azure AD tenant.

You can then add app secrets and certificates to the app for use in authentication and define the set of permissions from the Microsoft Graph and other APIs that you want the app to use. Figure 23-2 shows the properties of an app called *GetTeamsList*. The application (client) identifier and the directory (tenant) identifier are critical elements needed for authentication. We can also see that the application has an app secret and that its use is limited to the organization.

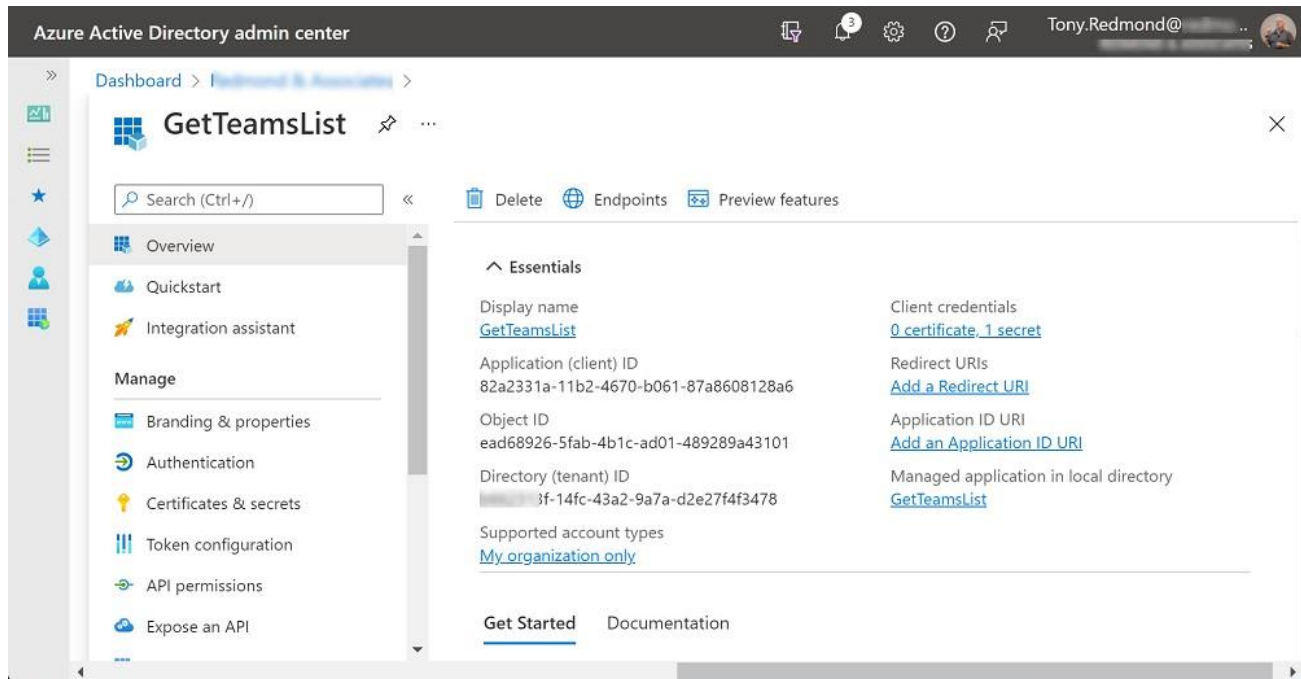


Figure 23-2: Properties of a registered application in the Azure AD admin center

The API permissions option in the menu allows administrators and app owners to define permissions in the set that an app can use. You can browse through the set of permissions available for the Graph and other APIs and mark those needed by the app (Figure 23-3). As noted above, before it can use an application permission, an administrator must grant consent to the app (see the Identities chapter). If limited to their data, a user can sometimes grant approval for an app to use a permission. The important point is that you should limit the permissions granted to an app to those necessary for the app to function.

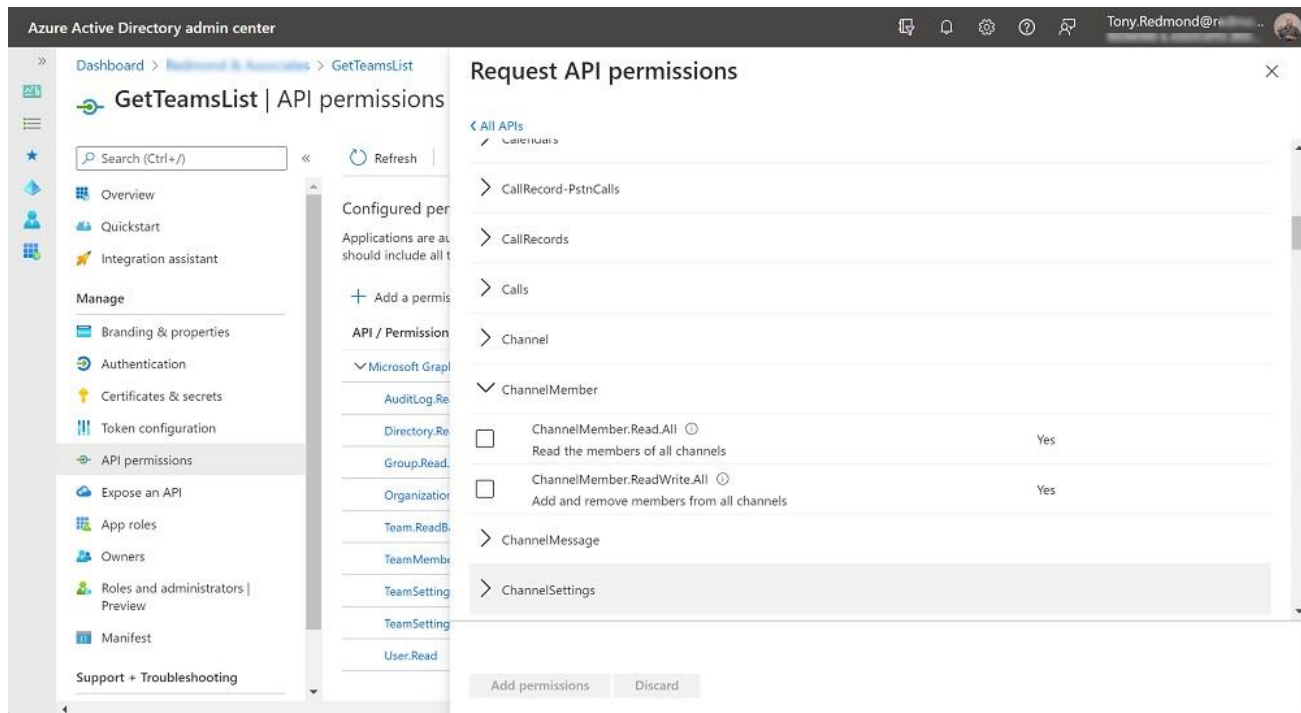


Figure 23-3: Adding permissions to a registered application

If you're unsure about the permissions needed by an app, you can use the Graph Explorer to run a test query and see what permissions the Explorer reports, and then grant consent for those permissions to the app. Sometimes an app needs unexpected permissions to access data. For instance, if a group has a hidden

membership, an app needs the *Member.Read.Hidden* permission to read its membership. Hackers often try to have elevated permissions assigned to apps that they use to retrieve data, so it's a good idea to [keep an eye on consents](#) granted to apps within the organization. It's also wise to limit the permissions granted to individual apps to make sure that they do not become targets for compromise.

## Client Secrets and Certificates

For unattended operations, apps can authenticate with client secrets or certificates. A client secret is a string generated by Azure AD for a registered app. The developer wishing to use the client secret must note it down at the time of creation as thereafter Azure AD obscures the string. A client secret can last up to 24 months. Using client secrets is convenient during the development and testing of a script intended for unattended operation, such as a script executed on a scheduled basis by an Azure Automation runbook. However, when the time comes to deploy the script in production, client secrets can represent a potential security vulnerability because if an attacker secures the application identifier for a registered app, the tenant identifier, and the client secret, they can use the three pieces of information to authenticate using the service principal for the app and access whatever data the app has permission to access. Although attackers can also compromise certificates, it requires more effort and certificates are safer. For this reason, you should use certificates to authenticate unattended jobs.

Equipped with:

- The application (client) identifier.
- The tenant (organization) identifier.
- A certificate or app secret to prove its identity.
- Consent to use permissions necessary to access data.

We can proceed to ask Azure AD for an access token to use in a PowerShell script.

## Getting an Access Token

Azure AD issues [OAuth 2.0 and OpenID Connect bearer tokens](#) to applications when they authenticate. Three types of tokens are available:

**ID tokens** prove a user's identity. For example, after a user enters their credentials into a client to sign in to Azure AD successfully, Azure AD issues an id or authentication token. The client can then use the id token to access resources as the signed-in account.

**Access tokens** obtain access to additional resources over and beyond what is available through an authentication token. When you create a registered app in Azure AD and use the service principal of the app to authenticate against Azure AD, the list of permissions granted to the app allows access to the additional resources such as data accessed through Graph API requests.

**Refresh tokens** are obtainable using either an Id or access token. Tokens have fixed lifetimes and expire (become invalid) at the end of the time. The usual lifetime for an access token issued to an app to access resources via the Graph API is one hour. Apps that support Azure AD continuous access evaluation (CAE) extend token validity to 28 hours because the apps can revoke access immediately when necessary. An app that makes just a few Graph API requests through PowerShell and then exits might only need the original access token. However, if the app executes Graph API requests for longer than the lifetime of the access token, it must use the refresh token (from the original access token) to acquire a new access token before the original token expires. If the app doesn't refresh its access token, any Graph API requests run after the access token expires will fail.

Sample PowerShell code to obtain an access token using an Azure AD registered app is below. You can see definitions for the application identifier, tenant identifier, and client secret in variables used to construct a URI

and body for input to the *Invoke-WebRequest* cmdlet. The URI points to the endpoint responsible for issuing the access tokens.

```
[PS] C:\> # Define the values applicable for the application used to connect to the Graph
# These variables vary from tenant to tenant and app to app
$AppId = "82a2331a-11b2-4670-b061-87a8608128a6"
$TenantId = "b662313f-14fc-43a2-9a7a-d2e27f4f3478"
$AppSecret = 'Mco7Q~D-IQTsaQYSRPLo01ctw4b1Hsw0D_zQ9'

# Construct URI and body needed for authentication
$uri = "https://login.microsoftonline.com/$tenantId/oauth2/v2.0/token"
$body = @{
    client_id      = $AppId
    scope          = "https://graph.microsoft.com/.default"
    client_secret  = $AppSecret
    grant_type     = "client_credentials"
}
# Get OAuth 2.0 Token
$tokenRequest = Invoke-WebRequest -Method Post -Uri $uri -ContentType "application/x-www-form-urlencoded" -Body $body -UseBasicParsing
# Unpack Access Token
$Token = ($tokenRequest.Content | ConvertFrom-Json).access_token
```

The *\$Token* variable stores the access token issued by Azure AD. If you examine the token, it won't make much sense because the token is a structured JSON web token object encoded in [base64url](#) and signed. However, if you copy the variable and paste it into [the jwt.ms site](#), the site decodes the token, and you can see that it contains a list of claims:

```
{ "typ": "JWT",
  "nonce": "gq3zmJhybfXGDGqt6R02PX9s0cimmRpSRrT090sQ4w4",
  "alg": "RS256",
  "x5t": "Mr5-AUibfBii7Nd1jBebaxboXW0",
  "kid": "Mr5-AUibfBii7Nd1jBebaxboXW0"
}.
{ "aud": "https://graph.microsoft.com",
  "iss": "https://sts.windows.net/a662313f-14fc-43a2-9a7a-d2e27f4f3478/",
  "iat": 1644833772,
  "nbf": 1644833772,
  "exp": 1644837672,
  "aio": "E2ZgYJif1+eocvtzqRIrgDGA2V3AQ==",
  "app_displayname": "GetTeamsList",
  "appid": "82a2331a-11b2-4670-b061-87a8608128a6",
  "appidacr": "1",
  "idp": "https://sts.windows.net/a662313f-14fc-43a2-9a7a-d2e27f4f3478/",
  "idtyp": "app",
  "oid": "4449ce36-3d83-46fb-9045-2d1721e8f032",
  "rh": "0.AVwAPzFitvwUokOaetLif080eAMAAAAAAAAAwAAAAAAAAABcAAA.",
  "roles":
[ "TeamSettings.ReadWrite.All", "TeamMember.Read.All", "Group.Read.All", "Directory.Read.All",
  "Team.ReadBasic.All", "TeamSettings.Read.All", "Organization.Read.All", "AuditLog.Read.All"],
  "sub": "4449ce36-3d83-46fb-9045-2d1721e8f032",
  "tenant_region_scope": "EU",
  "tid": "a662313f-14fc-43a2-9a7a-d2e27f4f3478",
  "uti": "BU1RVc7mHkmBq2FMcZdTAA",
  "ver": "1.0",
  "wids": [ "0997a1d0-0d1d-4acb-b408-d5ca73121e90" ],
  "xms_tcdt": 1302543310
}
.[Signature]
```

The deciphered token divides into three parts: header, payload, and signature. The aim of a token is not to hide information, so the signature is not protected by encryption. Instead, it's signed using a private key by the issuer of the token. Details of the algorithm and private key used to sign an access token are in its header. An application can [validate the signature of an access token](#) if necessary, but this is not usually done when

running a PowerShell script. The payload is the location for the claims made by the token and is the most interesting place to check.

The list of claims in the payload of an access token includes simple claims and scopes (groups of claims). A claim is an assertion about something related to the token. In this case, the claims tell us about the tenant, the app name, the app identifier, the security token service (STS) which issued the token, and the issuance time for the token and its expiration time (both in Unix epoch time, where 1644837672 means 11:21:12 GMT on February 14, 2022). Scopes are a logical grouping of claims, and they can serve as a mechanism to limit access to resources. The roles claim contains a scope of Graph API permissions starting with *TeamSettings.ReadWriteAll* and ending with *AuditLog.Read.All*. Therefore, we know this app has consent from the organization to use the permissions stated in the scope when it executes Graph API requests.

After obtaining an access token, the script can then use the token to make Graph APIs requests. For instance, here's how to retrieve the set of user accounts from Azure AD (note the escaping required by PowerShell before the *\$filter* operator). The access token is in the *\$Headers* variable included in the *Invoke-RestMethod* command.

```
$Headers = @{Authorization = "Bearer $token"}  
  
Write-Host "Fetching user information from Azure AD..."  
$Uri = "https://graph.microsoft.com/v1.0/users?&`$filter=userType eq 'Member'"  
[Array]$Users = (Invoke-RestMethod -Uri $Uri -Headers $Headers -Method Get -ContentType  
"application/json")  
$Users = $Users.Value
```

## Summarizing Graph Access for PowerShell Scripts

In summary, to use Graph API requests in a PowerShell script, you need to:

- Register an app in Azure AD. As explained earlier, the app is the point of interaction with Graph APIs.
- Know the tenant identifier. As mentioned above, several ways exist to retrieve the tenant identifier. The important thing is that you know the identifier when the time comes to authenticate the app to Azure AD.
- Grant consent for the permissions needed by the Graph API requests the app runs to access data. [Microsoft documentation](#) lists the available permissions and those needed to perform different operations. Ensure that an administrator grants the app consent only for the permissions required to access the data it processes and isn't over-permissioned.
- Use the app identifier, app secret, and tenant identifier to [request an access token](#) from the token endpoint (other methods can be used with PowerShell to authenticate an app against the Graph – [here's an example](#)). The response to a successful request is an access token that attests to the right of the app to use its assigned permissions to access data. The authorization header for Graph API calls includes the access token. Access tokens expire after an hour, so the app must renew the token periodically during long-running programs.
- Use the PowerShell [Invoke-WebRequest](#) or [Invoke-RestMethod](#) cmdlets to send HTTPS queries to run against the Graph. For example, you use a GET command to request information. *Invoke-RestMethod* understands the format of the output generated by REST APIs and parses the JSON output to create PowerShell objects. We recommend using *Invoke-RestMethod* with Graph API calls whenever possible.
- To use [advanced queries against Azure AD](#) (for objects like groups, members, and owners), the header passed in the query must include *ConsistencyLevel = eventual*. Consistency is a concept used by distributed databases to describe the accuracy of data given that the query must fetch from multiple places (as in a Graph database). [Eventual means that the data is highly available but might suffer from accuracy](#) due to write latency or other influences.
- The script processes the data returned by the Graph. It's usually best to convert the data from JSON format to make it more accessible in PowerShell.

- If the app receives consent for the necessary permission, it can add data by sending a POST request to the Graph API. You can update data with a PATCH command and remove information with a DELETE command.

Using Graph APIs in PowerShell scripts is not difficult to master or hard to code. Practice makes perfect.

# Approaching the Use of Graph API Requests with PowerShell

Traditionally, the role of PowerShell was to automate administrative operations. You can also use Graph APIs to automate administrative operations and access content stored in different repositories, like Teams, Outlook, and Planner, making it possible to read and write conversations, files, and shared calendars, none of which PowerShell can access. To close the gap between PowerShell and the Microsoft Graph, developers can combine PowerShell and the Graph by making Graph API calls using the *Invoke-RestMethod* and *Invoke-WebRequest* cmdlets. As described later, the [Microsoft Graph PowerShell SDK](#) makes it easier to include Graph API calls in scripts using PowerShell cmdlets which act as a wrapper around Graph API calls.

Although it might seem a good idea to always use Graph API calls to access Microsoft 365 data through PowerShell, we recommend a more measured four-step approach:

- **Sketch out the solution:** Understand what source data is available and how to access it. Define the expected output and the processing needed to achieve the result. Make an initial selection of PowerShell modules and Graph APIs which might be useful, understanding that some data is only accessible to the Graph (and might need a beta API). In addition, some Graph APIs have limitations, such as the lack of administrator access to Planner data. Do an internet search to see if anyone has already written code to do what you want or something similar. Never reinvent the wheel if someone else has one to use.
- **Code in PowerShell first:** It's often wise to write the initial code in PowerShell before introducing any Graph APIs. The code you write might work well enough to be the solution you need without doing any further work. This is often the case when a small amount of data is involved (say, fewer than a couple of thousand mailboxes or groups), in which case PowerShell can do the job in a reasonable timeframe without any need to introduce Graph API calls. Staying with PowerShell cmdlets avoids the overhead of creating a registered app in Azure AD, maintaining client secrets, including API calls in scripts, and so on, all of which are important factors to consider. The caveat here is that scripts written for background processing that use certificate-based authentication may need to use registered apps.
- **Speed Things Up:** Usually, the biggest advantage gained through using Graph APIs is speed, especially when fetching large numbers of objects. The next step is to find places in the code where cmdlets like *Get-UnifiedGroup* take a long time to execute and consider replacing those cmdlets with Graph APIs.
- **Adjust for Production:** Every tenant has their view on how to run PowerShell scripts in production. After developing a script that runs interactively, you might need to change it to run as a background process using Azure Automation. During this process, you'll probably need to change the authentication method to use certificates instead of client secrets or user/password credentials. Because people adjust scripts for production usage, the code explained in books and articles is often intended to demonstrate principles rather than be a fully worked-out solution.

The most important point in the checklist is the internet search for code. For example, many working examples of using Graph APIs in scripts are available in the [Office 365 for IT Pros GitHub repository](#). If you don't find a suitable script to remove the need to create anything new, you'll probably find the basis or starting point for what you want to do. Given the rich search facilities available, it is astonishing that people



post questions in forums when it is perfectly obvious that they haven't done the basic research to uncover details that can help solve their problem.

## Comparing PowerShell and the Graph

Table 23-4 compares PowerShell cmdlets and Graph API calls. This is not an extensive comparison. Instead, it serves to underline the point that PowerShell is easier to work with data, but Graph API calls are faster. In general terms, a good approach to take is to prototype and prove that a solution works with PowerShell, and then attempt to replace PowerShell cmdlets with Graph API calls to increase performance where necessary.

	<b>PowerShell</b>	<b>Graph API calls</b>
<i>Access</i>	Must load appropriate module before using cmdlets. Interaction through command-line, ISE, or scripts.	Must use a registered Azure AD app to connect to Graph API endpoints after obtaining an access token. Interaction through programs (including scripts).
<i>Data returned</i>	PowerShell objects.	CSV, JSON, or OData.
<i>Filters</i>	Depends on the cmdlet. Supports both server-side and client-side filtering.	Depends on the API call. Graph APIs return data (filtered if necessary) and the client can apply extra filters.
<i>Processing</i>	PowerShell supports extensive methods to process objects of different .NET types.	Graph API calls fetch and update data. Any language capable of calling the Graph API can process the data returned by the API.
<i>Use for</i>	Ad-hoc queries and maintenance. Scripts that run against a few hundred objects.	Structured regular maintenance involving anything up to tens of thousands of objects.

Table 23-4: Comparing PowerShell cmdlets and the Graph API

## Graph Pagination

To maintain performance and conserve resources, the Graph returns a limited set of objects when you make calls to fetch information about Groups, Teams, and other types of data. This is referred to as a page of data and several pages might need to be retrieved to fetch the full set of objects you want to process. [The Graph documentation](#) says "When a result set spans multiple pages, Microsoft Graph returns an `@odata.nextLink` property in the response that contains a URL to the next page of results." In effect, the nextlink is a URL to tell the Graph the next set of data to return if an application wishes to request that data.

Therefore, when a script uses the `Invoke-WebRequest` cmdlet to fetch Graph data, the code should check the presence of the nextlink to see if more data is available and continue doing so until the nextlink is null. Here's an example that fetches all the teams in a tenant. The code:

- Fetches the first page of data using `Invoke-WebRequest`. Notice the use of the escape (backtick) character when forming the Uri because a \$ sign precedes the Graph filter qualifier.
- Details of the retrieved teams are inserted into a [PowerShell hashtable](#). A hashtable can only contain two properties (key and value), so it's suitable for storing the group identifier (GUID) and the display name for each team. If you want to store more information, create and populate an array or use a [generic list](#) instead.
- Examine the nextlink and if it is not null, make a further call to fetch the next page of data. As the code fetches each page, it updates the data in the hashtable.
- This process continues until the nextlink is null. The hashtable now holds details of all teams.

```
[PS] C:\> $headers = @{Authorization = "Bearer $token"}
$uri = "https://graph.microsoft.com/V1.0/groups?`$filter=resourceProvisioningOptions/Any(x:x eq 'Team')"
```

```
$Teams = Invoke-WebRequest -Method GET -Uri $Uri -ContentType "application/json" -Headers $headers |
ConvertFrom-Json
$TeamsHash = @{}
$Teams.Value.ForEach( {
    $TeamsHash.Add($_.Id, $_.DisplayName) } )
$NextLink = $Teams.'@odata.NextLink'
While ($NextLink -ne $Null) {
    $Teams = Invoke-WebRequest -Method GET -Uri $NextLink -ContentType $ctype -Headers $headers |
ConvertFrom-Json
    $Teams.Value.ForEach( {
        $TeamsHash.Add($_.Id, $_.DisplayName) } )
    $NextLink = $Teams.'@odata.NextLink' }
```

To process the teams, we can loop down through the information in the hashtable and extract the group identifier and the display name for each team. The GUID can be used to fetch more data using cmdlets like *Get-Team*, *Get-UnifiedGroup*, and so on.

```
[PS] C:\> ForEach ($Team in $TeamsHash.Keys) {
    $TeamId = $($Team); $TeamDisplayName = $TeamsHash[$Team] #Populate variables for the team
    Write-Host "Team" $TeamDisplayName "now being processed" }
```

The same technique works with other Graph endpoints, including those used to retrieve documents from a SharePoint library, or sign-in information for Azure AD accounts. [This article](#) includes a useful PowerShell function to retrieve information from the Graph using pagination to fetch all available data.

## Graph Speed

Although it is more complicated to create a PowerShell script that uses Graph API calls, the Graph API calls are much faster at fetching data. This isn't so important when only a tenant only has a few hundred mailboxes, accounts, or groups; it becomes more interesting once several thousand objects are involved.

A good example of where the combination of PowerShell and Graph APIs work well is to retrieve and process large collections of groups. PowerShell can do this using the *Get-UnifiedGroup* cmdlet, but progress is slow in an organization where large numbers of groups exist. The same is true for the *Get-Team* cmdlet in the Teams module (very slow in the past but improved in recent versions of the Teams module). In these scenarios, fetching information using the [Graph API for Groups](#) with PowerShell offers the prospect of much better performance. Depending on the processing done by a script, it can be between four and ten times faster to use Graph calls instead of PowerShell cmdlets.

Although faster, it can be less convenient to retrieve group information with the Graph because:

- The script must do some up-front processing to connect to the Graph and secure an access token.
- The Groups Graph API works across all group types, including distribution lists and security groups. Several calls are often necessary to retrieve the information delivered by *Get-UnifiedGroup* such as the URLs for the SharePoint Online team site belonging to the group, the sensitivity label assigned to the group, or the count of members and guest members.
- Some attributes stored in the Exchange Online Directory such as the set of fifteen custom attributes, five extension attributes, mailbox quota settings, and MailTips aren't available through the Graph.
- Graph APIs use pagination to control the amount of data returned for a call. Scripts must retrieve pages of available data until no more is available.

Leaving the complexities aside, the sheer speed of access is the trump card for the Graph. Although some additional work might be necessary to fetch all the data a script needs, access via the Graph APIs is always faster than pure PowerShell. As an example, using the *Get-UnifiedGroup* cmdlet to retrieve details of 200 groups might take 20 seconds while the initial Graph query to fetch the set of available Microsoft 365 groups using a filter takes less than half a second. The query is:

```
[PS] C:\> $Uri = https://graph.microsoft.com/v1.0/groups?`$filter=groupTypes/any(a:a eq 'unified')
```

If you just need the set of Teams-enabled groups, the filter is:

```
[PS] C:\> $Uri =
'https://graph.microsoft.com/v1.0/groups?' $filter=resourceProvisioningOptions/Any(x:x eq 'Team')
```

An alternative is to use the beta version of the Teams List API:

```
[PS] C:\> $Uri = "https://graph.microsoft.com/beta/teams"
```

In all cases, once you know the query to run, we execute it as follows:

```
[PS] C:\> [array]$Data = Invoke-WebRequest -Method GET -Uri $Uri -ContentType "application/json" -
Headers $headers | ConvertFrom-Json
```

The array returned by the query contains another array called *Value*. The data for the groups or teams is in this array. For example, to fetch the first element in the array, we use *\$Data.Value[0]*.

The Teams and Groups activity report ([downloadable from GitHub](#)) contains examples of retrieving group information such as SharePoint information, member counts, and attributes specific to Microsoft 365 Groups. Another example showing how to report Teams settings not available through the Teams PowerShell cmdlets [is available on GitHub](#). The script checks the channels for all teams in a tenant and reports the email addresses of those that are mail-enabled.

## Advanced Queries Against Azure AD Objects

Many Graph API requests are straightforward. Some are more complicated and require the Graph to perform more processing to satisfy the request. Given the fundamental role played by Azure AD, it makes sense for Microsoft to optimize queries against Azure AD directory objects perform as well as possible. To do this, Azure AD indexes properties often used in queries and holds the index in a separate store. To use these properties, you run [advanced Azure AD queries](#) (some Microsoft Graph PowerShell SDK cmdlets like [Get-MgGroup](#) also support advanced queries). Advanced queries have the *ConsistencyLevel* header set to *Eventual* and use the *\$count=true* query string.

An example of a simple query is to find Azure AD group objects:

```
[PS] C:\> $uri = "https://graph.microsoft.com/v1.0/groups"
[array]$Data = Invoke-WebRequest -Method GET -Uri $Uri -ContentType "application/json" -Headers
$headers | ConvertFrom-Json
[array]$Groups = $Data.Value
```

The data returned in the *\$Data* array contains Microsoft 365 groups, security groups, and distribution lists. To make it easier to process, we extract the array containing the group information and store it in *\$Groups*. To refine the set returned to a specific type, we need to apply a filter. This example of an advanced query fetches the set of distribution lists by applying a filter to check that the *MailEnabled* property is True and that the group type is not "Unified" (a Microsoft 365 group):

```
[PS] C:\> $Uri = "https://graph.microsoft.com/v1.0/groups?' $filter=MailEnabled eq true and not
groupTypes/any(c:c eq 'unified')&' $count=true"
[array]$Data = Invoke-WebRequest -Method GET -Uri $Uri -ContentType "application/json" -Headers
$headers | ConvertFrom-Json
[array]$DLs = $Data.Value
```

The query returns all distribution lists except room lists (unlike regular distribution lists, these objects usually don't exist in Azure AD). The set includes mail-enabled security groups (the *SecurityEnabled* property is True). To find just mail-enabled security groups, the query is:

```
[PS] C:\> $Uri = "https://graph.microsoft.com/v1.0/groups?' $filter=Mailenabled eq true and
SecurityEnabled eq true and not groupTypes/any(c:c eq 'unified')&' $count=true"
```

Not all properties of Azure AD objects support all operators. For instance, to find the set of Microsoft 365 groups with an expiration date set, you'd assume that you could filter based on the property holding the expiration date not being null. However, that's not the case (see [the limitations in queries for various Azure AD objects](#)), and you end up checking against an arbitrary date in the past using the greater than or equal to operator, which is supported:

```
[PS] C:\> $uri = "https://graph.microsoft.com/beta/groups?$`filter=ExpirationDateTime ge 2014-01-01T00:00:00Z AND groupTypes/any(a:a eq 'unified')&`$count=true"
```

An alternative approach is to fetch the complete set of objects and use a client-side filter to extract the objects you want to process. Although effective, this method could be slow.

## Graph Filters and OData Syntax

PowerShell developers who start using Graph calls might miss some of the flexibility of PowerShell when it comes to applying filters to Graph API requests.. Two examples using the Office 365 Service Communication API demonstrate the point. First, let's assume that you want to find the set of issues in a tenant for the last seven days. To do this, you include [a filter in OData syntax](#) to tell the Graph to return only incidents with a start date within the last week. OData filters use sortable dates. We can get a sortable date seven days back from the current date with:

```
[PS] C:\> $DaysRange = (Get-Date).AddDays(-7)
$DaysRangeZ = Get-Date($DaysRange) -format s
```

*\$DaysRangeZ* now contains a value like 2021-07-29T14:41:09, which is a sortable date. However, it won't work if used in a query like this:

```
[PS] C:\> $Uri =
"https://graph.microsoft.com/beta/admin/serviceAnnouncement/issues?`$filter=startDateTime ge
$DaysRangeZ"
```

The reason is that the Graph requires a Z at the end of a sortable date. You can see this in the example of looking for groups with an expiration date described above. After adding a Z, this URI works:

```
[PS] C:\> $Uri =
"https://graph.microsoft.com/beta/admin/serviceAnnouncement/issues?`$filter=startDateTime ge
$DaysRangeZ" + "Z"
```

Now let's assume that we want to add another property to the filter so that the query retrieves open incidents only. The *isResolved* property is True or False to indicate the status of an incident. PowerShell is happy checking against False with either a lower- or upper-cased initial letter. The Graph demands lower case, meaning that we need to use true or false in Graph API requests, like this:

```
[PS] C:\> $Uri =
"https://graph.microsoft.com/beta/admin/serviceAnnouncement/issues?`$filter=startDateTime ge
$DaysRangeZ" + "Z and IsResolved eq false"
```

The above illustrates that the Graph is more particular with the syntax of its requests than PowerShell is. It can be infuriating at times to have code that appears perfectly good be rendered unacceptable due to a minor imperfection. The Graph Explorer is very helpful at resolving syntax issues, so be sure to use it to test Graph API requests before you include them in scripts.

## Lambda Qualifiers

In some of the examples used to date, you might have noticed that the construct of the filter includes /any. This means that the filter operates against a multi-value property, like the set of email addresses assigned to a

mail-enabled object or the licenses assigned to a user account. For example, this query uses a [lambda qualifier](#) to find Azure AD accounts assigned an Office 365 E3 license:

```
[PS] C:\> Uri = https://graph.microsoft.com/beta/users?`$filter=assignedLicenses/any(s:s/skuId eq 6fd2c87f-b296-42f0-b197-1e91e994b900)
```

The **/any** operator in the filter *"iteratively applies a Boolean expression to each member of a collection and returns true if the expression is true for any member of the collection, otherwise, it returns false"*. In this case, the filter examines the *assignedLicenses* property of each account to find instances where the Skuld (one of the values stored in the property) matches the GUID for Office 365 E3. Using [Microsoft's documentation](#) to interpret the components of the lambda /any qualifier, we can see that:

- **assignedLicenses** is the parameter, or the property to which the filter is applied. The property can contain a collection of values or a collection of entities. In this case, the *assignedLicenses* property for an account contains a collection of one or more license entities. Each license is composed of the Skuld and any disabled plans unavailable to the license holders.
- **s:s** "is a range variable that holds the current element of the collection during iteration." The interesting thing is that you can give any name you like to the range variable. In this case, it could be *license:license* or even *rubbish:debris*. It's just a name for a variable.
- **Skuld** is the subparam, or value within the property being checked. When there's only one value in a collection (as when you check for team-enabled groups), you don't need to specify a subparam. In the case of *assignedLicenses*, it is needed because we want to match against the Skuld within the collection of entities in the *assignedLicenses* property.
- **6fd2c87f-b296-42f0-b197-1e91e994b900** is the value to match against items.

Lambda qualifiers exist because it's more difficult to filter against complex properties than it is to filter against simple strings, like this straightforward filter looking for any accounts with a display name starting with Tony:

```
[PS] C:\> $Uri = https://graph.microsoft.com/v1.0/users?`$filter=startswith(displayName, 'Tony')
```

Knowing when a lambda qualifier is needed in a query is a matter of figuring out if a property is multi-valued, checking Microsoft's documentation, or finding an example of a working query created by someone else. Or a combination of all the above.

## Sending Email with the Graph

Earlier we covered how to use the *Send-MailMessage* cmdlet to send messages through PowerShell. *Send-MailMessage* depends on the SMTP AUTH protocol and basic authentication. The Graph *SendMail* API call can replace *Send-MailMessage* in scripts. *SendMail* can be used in two ways, [depending on the permissions](#) held by the app which calls it:

- If the app has the *Mail.Send* application permission, the app can send email as any user or shared mailbox (but not a group mailbox) in the organization. An administrator must give consent to an app to use application permissions.
- If the app has the *Mail.Send* delegate permission, the app can send mail as the signed-in user.

To convert a script to use the *SendMail* API, we need:

- A registered app in Azure AD.
- Assignment of the *Mail.Send* permission to the app.
- Script code to call the app, get an access token, and build and send the message.

Code to get an access token is like that used by any script wishing to use the Graph APIs. From a conversion perspective, the work is to remove the code used to create and send a message with *Send-MailMessage* and replace it with the format used to construct a message for *SendMail*. The following points should be noted:

- The SMTP address for the sender must be resolvable against the tenant directory. The *To* and *CC* addresses must be valid SMTP addresses.
- Attachments must be Base64-encoded before being added to the message. You can attach a file of up to 3 MB and send it in one operation. Larger attachments [need a different approach](#). Remember to assign the correct MIME type to the attachment.
- The message structure is described in JSON. Here's an example of a message with one *To* and two *CC* recipients and an attachment. The HTML content for the message is defined in the `$htmlMsg` variable.

```
# Create message body and properties and send
$MessageParams = @{
    "URI" = "https://graph.microsoft.com/v1.0/users/$MsgFrom/sendMail"
    "Headers" = $Headers
    "Method" = "POST"
    "ContentType" = 'application/json'
    "Body" = @(
        @{
            "message" = @{
                "subject" = "A test message"
                "body" = @{
                    "contentType" = 'HTML'
                    "content" = $htmlMsg }
            }
        }
    )
    "attachments" = @(
        @{
            "@odata.type" = "#microsoft.graph.fileAttachment"
            "name" = $AttachmentFile
            "contentType" = "application/vnd.openxmlformats-officedocument.wordprocessingml.document"
            "contentBytes" = $ContentBase64 })
    "toRecipients" = @(
        @{
            "emailAddress" = @{ "address" = "Chris.Bishop@office365itpros.com" }
        } )
    "ccRecipients" = @(
        @{
            "emailAddress" = @{ "address" = "Vasil.Michev@office365itpros.com" }
        },
        @{
            "emailAddress" = @{ "address" = "James.Ryan@office365itpros.com" }
        } )
    }
} | ConvertTo-JSON -Depth 6
} # Send the message
Invoke-RestMethod @MessageParams
```

A full script to illustrate these concepts is [available from GitHub](#).

## Scoping Graph Applications to Specific Exchange Online Mailboxes

If applications built with the Graph APIs have the necessary permissions, they can read and write information into many types of data, including Exchange Online mailboxes. Because mailboxes can store confidential information, Exchange Online supports [application access policies](#) to limit access to sets of mailboxes defined through membership of a mail-enabled security group. When an application access policy is in place, apps covered by the policy can access only the mailboxes that are members of the security group. See [this page](#) for more information.

Application access policies are protocol agnostic. [They also govern the use of the Exchange Web Services \(EWS\) API](#), including when programs use EWS for impersonation (to behave as if the program was a mailbox's owner) to access mailboxes. If a policy blocks access to a set of mailboxes, a program cannot use EWS to impersonate those mailboxes and access the data within the mailboxes.

Using an application access policy to control access to apps that use the *SendMail* API is a good use case. As you'll recall, if an administrator grants consent to an app to use the application *Mail.Send* permission, the app

can send mail from any user or shared mailbox in the organization. This is a dangerous situation as it creates the possibility that someone could use an app to send unauthorized messages. An application access policy can restrict access to the app so that messages can be sent only from certain mailboxes. You can only create an application access policy in PowerShell. For example:

```
[PS] C:\> New-ApplicationAccessPolicy -AppId 970e01d1-ce75-46ba-a054-4b61c787f682 -  
PolicyScopeGroupId SendMailApp.Control@office365itpros.com -AccessRight RestrictAccess -Description  
"Restrict access to app allowed to send email using the Graph SendMail API"
```

This policy:

- Applies to the registered app with identifier 970e01d1-ce75-46ba-a054-4b61c787f682.
- Restricts access to the mailboxes defined in the mail-enabled security group SendMailApp.Control@office365itpros.com.

In other words, the app can only use the *Mail.Send* API to create and send messages using the mailboxes of a member of the security group. This [blog post](#) explains how to check service principals in a tenant to discover those which use Exchange Online mailbox permissions. It's wise to perform checks like this periodically to ensure that apps don't abuse permissions.

## Tenant Privacy Controls

The [Microsoft Graph Insights resource type](#) is used by applications to access *insights* such as the set of documents viewed, modified, or shared by a user. Insights appear in different Microsoft 365 apps including Viva Insights, Workplace Analytics, and the user profile card. The Microsoft Graph calculates insights using signals gathered from apps about a user's activity and spanning information in their mailbox, OneDrive for Business account, and files they have access to in SharePoint Online or other users' OneDrive for Business accounts. Three types of insights are available:

- **Item insights:** Recommendations about documents available to a user which they might be interested in. This information shows up in the profile card in the "Files" section.
- **Meeting insights:** When users open a meeting in their calendars, Outlook displays [insights relevant to the meeting](#), such as documents and emails covering the content of the meeting.
- **People insights:** Information about [people deemed to be relevant to an individual](#) and who have a "public relationship" with that individual. For example, they have the same manager (as noted in Azure AD) or share a common membership of a public group or distribution list with fewer than 30 members. People insights show up on the profile card in the "works with" section, with the people who appear ranked in order of the public and private communication (direct email or meetings rather than group-based communications) between the listed individuals and the owner of the profile card.

Tenants control the display of insights through organization settings. Meeting insights are either enabled or disabled for the entire organization, while tenants can disable item and people insights for a specific Microsoft 365 group or security group. Administrators can modify settings for item and meeting insights in the Search & intelligence section of Settings in the Microsoft 365 admin center. Users can disable item insights for their account through the Privacy tab of their [MyAccount page](#) (this action does not add them to the group used by the tenant to disable item insights for selected users). Controls for people insights are only available through the Microsoft Graph.

The people insight control is available through <https://graph.microsoft.com/beta/organization/GUID-org-id/settings/peopleinsights> (insert your tenant identifier to create a complete URI). If you run a query against the URI, the API returns the following information:

- *isEnabledInOrganization*: This setting controls if people insights are available within an organization. By default, the setting is enabled (true). If disabled, no insights are created.

- *disabledForGroup*: This setting holds the GUID (object identifier) for an Azure AD group. The Microsoft Graph excludes the members of the group when it generates people insights.

The PowerShell code below shows how to query the Graph to retrieve the current settings for people insights using the *Invoke-RestMethod* cmdlet. The second part of the code fetches a set of accounts stored in a CSV file and updates the Azure AD group used to exclude accounts for people insights. The code used to connect to a registered Azure AD application and secure an access token is omitted.

```
[PS] C:\> $TenantId = (Get-MgOrganization).Id
$Uri = "https://graph.microsoft.com/beta/organization/" + $TenantId + "/settings/peopleinsights"
$Settings = Invoke-RestMethod -Uri $Uri -Method Get -ContentType "application/JSON" -Headers
$Headers -UseBasicParsing

If ($Settings.isEnabledInOrganization -ne $True) {
    Write-Host "Insights control setting not set for the tenant" ; break }
Else {
    $DisabledGraphInsightsGroup = $Settings.disabledForGroup }

[array]$CurrentMembers = Get-MgGroupMember -GroupId $DisabledGraphInsightsGroup

Write-Host "Adding users to the Disabled Graph Insights Group"
$Users = Import-CSV $InputCSV
ForEach ($User in $Users) {
    If ($User.ObjectId -notin $CurrentMembers.Id) {
        Write-Host "Adding" $User.Name
        New-MgGroupMember -GroupId $DisabledGraphInsightsGroup -DirectoryObjectId $User.ObjectId }
}
```

## The Microsoft Graph Explorer

There is nothing like using a tool to understand what it can do. To understand the capabilities of the Microsoft Graph, you can fire up the [Graph Explorer](#), a web application that allows users to execute Graph API requests (or queries) against real data. The Graph Explorer supports GET, POST, PUT, PATCH, and DELETE requests, and exposes the responses to those requests, including details like request bodies, responses, code, and headers. The application generates code snippets for API requests in C#, Java, JavaScript, and Go to copy and use elsewhere. In some cases, PowerShell examples using the cmdlets from the Microsoft Graph PowerShell SDK are available for Graph requests.

The idea behind the Graph Explorer is that it helps programmers understand the kind of calls they need to construct to interact with various endpoints and the format of the results returned through their interaction. It is particularly helpful in understanding the format of Graph API requests, including how filters are used, and the type of data returned by requests.

The Graph Explorer supports running requests against:

- Test data. You don't need to sign in to use these requests.
- Tenant data. You must sign in to use these requests. If you use an administrator account, you can see tenant-wide data. Otherwise, you're limited to your own data. If necessary, you can grant the Graph Explorer consent to use the permissions necessary to run the requests.
- The V1.0 and the beta endpoints.

The input field for the Graph query to run includes an IntelliSense-like auto-complete capability, so you can see the available calls as you type in the input box. For instance, assume that we select *https://graph.microsoft.com/beta/me/* as a starting point. Underneath the field, you see a list of calls organized (somewhat alphabetically) starting with *https://graph.microsoft.com/beta/me/activities* and then *https://graph.microsoft.com/beta/me/agreementacceptances* and so on.

To help users experiment with Graph requests, the Graph Explorer includes:



- **Getting started queries**, a small set of requests against personal data such as *my mail* or *my photo*.
- **Sample queries** organized into capabilities (like notifications) and product functionality (like Teams, Planner, OneDrive, and Outlook Calendar). some interesting requests to investigate include:

<https://graph.microsoft.com/V1.0/groups>: List all the groups in a tenant.

[https://graph.microsoft.com/v1.0/me/transitiveMemberOf/microsoft.graph.group?\\$count=true](https://graph.microsoft.com/v1.0/me/transitiveMemberOf/microsoft.graph.group?$count=true): List all the Microsoft 365 Groups that the logged-on user belongs to in a tenant.

<https://graph.microsoft.com/V1.0/me/joinedTeams>: List all the Teams that the logged-on user belongs to in a tenant.

- **Resources**, a list of the available Graph APIs. If you select a resource, the Graph Explorer loads the basic URI for the resource into the query. For example, if you select the *groupLifecyclePolicies* resource, Explorer loads the <https://graph.microsoft.com/v1.0/groupLifecyclePolicies> query. Some of the requests (like this one) can run as is. Others need additional information, like a tenant identifier.

After you enter a specific query or select one of the “Getting Started” examples, click **Run Query** to have the Graph return the requested data. For instance, Figure 23-4 shows the results of a call to fetch data from <https://graph.microsoft.com/v1.0/organization>, a query that returns details of the tenant (organization). You can see details of some of the SKUs (products) available in the tenant.

The screenshot shows the Microsoft Graph Explorer interface. The top navigation bar includes the Microsoft logo, 'Microsoft Graph', and various utility links like 'Explore', 'Graph Explorer', 'Docs', 'API', 'Learn', 'Developer Program', and 'Support'. The main area is titled 'Graph Explorer' and shows a tenant 'Redmond & Associates'. The query bar contains the URL 'https://graph.microsoft.com/v1.0/organization' and a 'Run query' button. Below the query bar, the 'Request body' is empty, and the 'Response preview' shows a successful GET request with a status of 'OK - 200 - 640ms'. The response is a JSON object with the following structure:

```

{
  "technicalNotificationMails": [
    "Tony.Redmond@"
  ],
  "tenantType": "AAD",
  "directorySizeQuota": {
    "used": 1427,
    "total": 300000
  },
  "privacyProfile": null,
  "assignedPlans": [
    {
      "assignedDateTime": "2022-02-28T19:48:47Z",
      "capabilityStatus": "Enabled",
      "service": "exchange",
      "servicePlanId": "b74d57b2-58e9-484a-9731-aecbba954f0"
    },
    {
      "assignedDateTime": "2022-02-28T19:48:47Z",
      "capabilityStatus": "Enabled",
      "service": "exchange",
      "servicePlanId": "c815c93d-0759-4bb8-b857-bc921a71be83"
    },
    {
      "assignedDateTime": "2022-01-24T14:55:49Z",
      "capabilityStatus": "Enabled",
      "service": "exchange",
      "servicePlanId": "c815c93d-0759-4bb8-b857-bc921a71be83"
    }
  ]
}

```

Figure 23-4: Using the Graph Explorer to test Graph API calls

The Graph Explorer retains details of requests you run for 30 days. You can go back and rerun a query at any time or download the set of requests and responses to a [.HAR file](#), which can be opened by any text editor.

When you list sets of objects like Groups or Teams, you see an identifier for each object. You need the identifier (a GUID) to navigate to the next level. For instance, this URL includes a group identifier <https://graph.microsoft.com/V1.0/groups/72ee570e-3dd8-41d2-bc84-7c9eb8024dd4>. If you include the URL in a query executed by the Explorer, you get back the properties of the group in JSON format. For example:

```

{
  "@odata.context": "https://graph.microsoft.com/V1.0/$metadata#groups/$entity",
  "id": "72ee570e-3dd8-41d2-bc84-7c9eb8024dd4",
  "deletedDateTime": null,

```

```

"classification": "External Access",
"createdDateTime": "2016-09-29T18:37:00Z",
"description": "Exchange Grumpy Old Men (and Women too)",
"displayName": "Exchange's Grumpy Old Men",
"groupTypes": [
  "Unified"
],
"mail": "ExchangeMVPs@office365itpros.org",
"mailEnabled": true,
"mailNickname": "exchangegoms",
"membershipRule": null,
"membershipRuleProcessingState": null,
"onPremisesLastSyncDateTime": null,
"onPremisesProvisioningErrors": [],
"onPremisesSecurityIdentifier": null,
"onPremisesSyncEnabled": null,
"preferredLanguage": null,
"proxyAddresses": [
  "SMTP:ExchangeMVPs@office365itpros.org",
  "smtp:exchangegoms@office365itpros.onmicrosoft.com"
],
"renewedDateTime": "2016-09-29T18:37:00Z",
"resourceBehaviorOptions": [],
"resourceProvisioningOptions": [],
"securityEnabled": false,
"theme": null,
"visibility": "Private"
}

```

The set of properties returned for a group by the Graph is different from the set returned by the PowerShell *Get-UnifiedGroup* cmdlet. This is because Exchange Online incorporates some extra information about the group mailbox when *Get-UnifiedGroup* returns data for a group (one of the reasons why the cmdlet is slow).

You can also copy the Graph code used as examples in Microsoft's documentation and run the API requests in the Graph Explorer to see how the Graph responds. For instance, the URI below returns sign-in data for Azure AD accounts.

*https://graph.microsoft.com/beta/users?\$select=displayName,userPrincipalName,mail,id,CreatedDateTime,signInActivity,UserType&\$top=999*

Full documentation about the various Microsoft Graph APIs is [available online](#).

## Graph Explorer Permissions

The Graph Explorer works with delegated permissions. In other words, you can access your own data (or test data), and if your account holds administrative roles, API requests run by the Graph Explorer can access data available to those roles too.

Requests executed through the Graph Explorer might require additional permissions to the set already assigned before they can run successfully. For instance, by default, the Explorer does not have permission to read items in a user's mailbox. Before it can proceed, the user must modify its assigned permissions and provide consent to the Graph Explorer app to use the *Mail.ReadBasic*, *Mail.Read*, and *Mail.ReadWrite* permissions.

To see the set of permissions needed to allow a query to run, click the Modify permissions tab. If you have an administrator account, you can consent on behalf of the organization to allow the Graph Explorer to use the required permission, such as *Mail.Read*. The signed-in user can also consent to permissions to allow the Graph Explorer to access their data. Understanding app permissions and assigning only the permissions necessary for an app to access the data it needs to process are fundamental to working with Graph API-based apps.

## Using the Graph Explorer for Real Work

If you sign into the Graph Explorer with an administrator account, you can manipulate tenant data. Microsoft does not intend the Graph Explorer to be an administrative tool, but situations exist where the Graph Explorer can do real work and avoid the necessity to write some code. Two examples are:

- [Customizing the Office 365 profile card](#) to expose additional information.
- [Updating privacy controls for document insights](#) derived from signals collected in the Graph.

In both cases, the only supported method to make changes in a tenant is via the Graph. Microsoft does not provide access in an administrative GUI like the admin center nor does a PowerShell cmdlet exist that can update these settings. However, the process to use the Graph Explorer to apply the updates is straightforward:

- Know your tenant identifier. The tenant identifier is available in the overview section of the Azure AD admin center. The suggested approach is to use the *Get-MgOrganization* cmdlet from the Microsoft Graph PowerShell SDK. For example:

```
[PS] C:\> $TenantId = (Get-MgOrganization).Id
```

- Know the URI for the endpoint to access. For example, to update the profile card, you interact with <https://graph.microsoft.com/beta/organization/{tenantid}/settings/profileCardProperties> (insert the tenant identifier in the placeholder to create the real URI). If you've been using a beta endpoint, it's a good idea to check periodically to see if Microsoft has upgraded an API to a fully-supported V1.0 version. If they do, you should change your code to use the supported version, always being aware that the change could affect how a query runs or the data it returns.
- Know the type of request you need to run. A GET command retrieves information, and a PUT command creates something, while you use PATCH commands to update existing data.
- Run the query.
- Check that an OK response (often 200, but other values are used) is returned. Check that the right data now exists.

Given the size and distributed nature of Office 365, it can take up to 24 hours before a change made to Graph settings is effective within applications.

## Using the Microsoft Graph PowerShell SDK

The [Microsoft Graph PowerShell SDK](#) is a compound module (it spans over 30 sub-modules) supporting modern authentication to many Graph APIs from PowerShell, using both delegate and application ([certificate only](#)) permissions. The SDK supports both the V1.0 (generally available) and beta Graph APIs. SDK cmdlets support PowerShell 5.1 and 7 on Windows, macOS, and Linux.

Microsoft's direction is to focus on Graph APIs for access to Microsoft 365 and other data (like Azure AD). As the process unfolds, Microsoft will deliver PowerShell access to Microsoft 365 data via the SDK rather than increasing the number of modules available for individual workloads. Some major workloads like Exchange Online and Teams are likely to continue with their separate modules while others, like Planner, will use the SDK. This might be because of the number and type of cmdlets in the module or the lack of support for certain functionality in the Graph APIs. For instance, the Exchange Online management module supports hundreds of cmdlets, some of which (like *Get-ExoMailboxFolderStatistics*) have no equivalent Graph API. The same is true of Teams, where many of the policy management cmdlets don't have a matching Graph API.

Microsoft's adoption of new platforms and deprecation of old software will force organizations to move away from older modules to the SDK. For example, the [license management cmdlets in the Azure AD module will cease working](#) after August 26, 2022, when Microsoft 365 moves to a new license management platform.

Because of the deprecation of the Azure AD Graph API, Microsoft plans to retire and remove support for the Microsoft Services Online (MSOL) and Azure AD modules. To help developers transition from these modules to the Microsoft Graph PowerShell SDK, Microsoft publishes a [cmdlet map](#), a way to find a suitable replacement for an older Azure AD or MSOL cmdlet in the SDK. Although useful to have a cmdlet map, it's important to understand that knowing which SDK cmdlet Microsoft recommends for an Azure AD cmdlet does not mean that the cmdlet parameters are the same or the cmdlet returns the same data in the same format. Careful testing is necessary to validate that upgraded scripts work as expected.

## SDK Modules

The Microsoft Graph PowerShell SDK spans over 30 sub-modules and several thousand cmdlets. Each module contains the cmdlets needed to interact with a different Graph API focusing on a specific area such as reporting or user accounts. Among the sub-modules are:

- Calendar: Microsoft.Graph.Calendar
- Compliance: Microsoft.Graph.Compliance.
- Device Management: Microsoft.Graph.DeviceManagement.
- Sign-ins: Microsoft.Graph.Identity.Signins.
- Mail: Microsoft.Graph.Mail.
- People: Microsoft.Graph.People.
- Planner: Microsoft.Graph.Planner.
- Reports: Microsoft.Graph.Reports.
- Security: Microsoft.Graph.Security.
- Teams: Microsoft.Graph.Teams.
- Users: Microsoft.Graph.Users.

When you download and install the SDK from the [PowerShell Gallery](#), the installation includes all the sub-modules. To install the SDK, run the command:

```
[PS] C:\> Install-Module Microsoft.Graph -Scope AllUsers -Force
```

## SDK Cmdlets

To find the cmdlets available for specific tasks, you can use the *Get-Command* cmdlet with some inspired guesswork. For example, to find the cmdlets available for groups, sign into PowerShell on a workstation where the SDK is installed and run the command:

```
[PS] C:\> Get-Command -Module Microsoft.Graph* *MgGroup*
```

Microsoft uses a process called *AutoRest* to generate SDK cmdlets from Graph APIs automatically. The process also generates the cmdlet documentation, which accounts for why the documentation is sometimes difficult to understand. The cmdlets follow the normal PowerShell naming convention of *verb-(optional prefix)noun* and map the http request made to the Graph APIs to the verb:

- GET https requests result in Get- PowerShell cmdlet, like *Get-MgUser*.
- POST and PUT https requests result in *New*, *Add*, or *Update* PowerShell cmdlets, like *New-MgUser*.
- PATCH https requests map to *Update*- or *Set*- PowerShell cmdlets, like *Update-MgUser*.
- DELETE https requests map to *Remove* PowerShell cmdlets, like *Remove-MgUser*.

Microsoft publishes new versions of the SDK modules monthly. Currently, the SDK doesn't include cmdlets for every Graph API call, but as explained earlier, you can make Graph API calls from PowerShell without needing to use the SDK. When upgrading scripts, the choice is to use either the SDK cmdlets or Graph API calls. Although usually a matter of personal choice, most developers seem to prefer using direct calls to Graph APIs in scripts. This situation will change over time as the SDK matures, Microsoft removes some of the unfinished

edges, and the SDK supports more scenarios. Because the SDK is under active development, you should always download the latest version from the PowerShell gallery before evaluating it as the basis for a new project.

## Connecting with Connect-MgGraph

The first step is to connect to the Graph using the *Connect-MgGraph* cmdlet. The SDK supports modern authentication, so you can connect using any supported method, including multi-factor authentication, using a FIDO2 key, or with a certificate. When you run *Connect-MgGraph*, be sure to specify the identifier of the tenant to which you want to connect. If you don't specify a tenant, *Connect-MgGraph* will choose the last tenant you signed into during the current session (which might not be the one you want to connect to). A session lasts until you run *Disconnect-MgGraph* (see below) and can be reinitiated multiple times by running *Connect-MgGraph*. If you work across multiple tenants, it's important to disconnect your session after you finish working with a tenant as otherwise you might end up inadvertently running commands in one tenant when you expect that the session is connected to another.

Behind the scenes, the Graph SDK keeps an encrypted token cache and will refresh the token as needed to allow you to continue working with Graph commands.

```
[PS] C:\> Connect-MgGraph -TenantId "828e1143-88e3-492b-bf82-24c4a47ada63"
```

After connecting, the session signs into your account. The permission scope for the connection comes from the well-known service principal (enterprise app) registered in Azure AD for the SDK. If you've never signed in with the Graph SDK before, it creates a service principal called *Microsoft Graph PowerShell* with an AppId of *14d82eec-204b-4c2f-b7e8-296a70dab67e* and requests a limited set of permissions. If you're an administrator, you can grant consent for these permissions on behalf of the organization. The service principal is the object that runs Graph API requests, not the account used to sign into PowerShell, so any administrative roles held by the account have no influence over application permissions usable during the session; they only affect delegate permissions.

Unlike other PowerShell modules that inherit the permissions held by the signed-in account, the SDK operates on a least-privilege basis. In other words, if you want to do something with an SDK cmdlet, you must check if the service principal used by the SDK has the permission to perform that action. If not, an administrator must grant consent to allow the action to proceed. This principle holds for interactive sessions with the SDK, in which case the Microsoft Graph PowerShell service principal receives the permission, or registered Azure AD apps used with certificate-based authentication using something like an Azure Automation runbook. In the latter case, the service principal for the app receives the permission.

For example, permission to read directory information is necessary to use the *Get-MgUser* cmdlet to retrieve a set of Azure AD accounts. With an interactive session, a user can request the set of permissions they need (for all the SDK cmdlets they plan to use) using the *Scope* parameter when connecting to the Graph as follows:

```
[PS] C:\> $RequiredScopes = @"Directory.AccessAsUser.All", "Directory.ReadWrite.All"  
Connect-MgGraph -Scopes $RequiredScopes
```

Connecting with a scope containing permissions not already held by the Microsoft Graph PowerShell service principal causes Azure AD to prompt for consent for those permissions. If granted, Azure AD adds the extra permissions to the set already held by the service principal. If an administrator does not grant consent, the permissions available in the session are those already consented plus those that the user has through any administrative roles assigned to their account. Users can grant consent to permissions necessary to access their own data.

Over time, the permissions held by the service principal grow from the initial set granted at the creation of the service principal plus any other permissions granted subsequently. In other words, the service principal

collects aggregated permissions over time. For this reason, it's not recommended to use the Graph SDK cmdlets interactively because if you do, over time a distinct possibility exists that the service principal will become very over-permissioned and therefore becomes a security risk. The same risk of permission creep exists for registered apps but is less evident than for the service principal used by the Microsoft Graph PowerShell SDK, if only because the way the SDK interacts with many different types of data.

Apart from going through and removing individual permissions, the only resolution for an over-permissioned service principal is its removal and recreation, at which time an administrator can grant consent for limited permissions to the new service principal. Here's how to remove the service principal using Graph SDK cmdlets (naturally):

```
[PS] C:\> $Sp = Get-MgServicePrincipal -top 999 | ? { $_.AppId -like "14d82eec-204b-4c2f-b7e8-296a70dab67e" }
Remove-MgServicePrincipal -ServicePrincipalId $sp.Id
```

In some situations, you'll need to use the beta profile to have access to certain commands and data. This is equivalent to specifying the beta endpoint when issuing Graph API requests. To select the beta profile, run the *Select-MgProfile* cmdlet:

```
[PS] C:\> Select-MgProfile beta
```

To check that you're connected to the right tenant with the right profile and permissions, we can extract information about the tenant with the *Get-MgOrganization* cmdlet, the current connection with the *Get-MgContext* cmdlet, and the profile used with the *Get-MgProfile* cmdlet and display some useful information:

```
[PS] C:\> Select-MgProfile beta
$Details = Get-MgContext
$Scopes = $Details | Select -ExpandProperty Scopes
$Scopes = $Scopes -Join ", "
$ProfileName = (Get-MgProfile).Name
$OrgName = (Get-MgOrganization).DisplayName
CLS
Write-Host "Microsoft Graph Connection Information"
Write-Host "-----"
Write-Host " "
Write-Host ("Connected to Tenant {0} ({1}) as account {2}" -f $Details.TenantId, $OrgName,
$Details.Account)
Write-Host "+-----+
-----+"
Write-Host ("Profile set as {0}. The following permission scope is defined: {1}" -f $ProfileName,
$Scopes)
Write-Host ""

Microsoft Graph Connection Information
-----

Connected to Tenant a662313f-14fc-43a2-9a7a-d2e27f4f3475 (Office 365 for IT Pros) as account
Global.Administrator@office365itpros.com
+-----+
Profile set as beta with the following scopes defined: Directory.AccessAsUser.All,
Directory.ReadWrite.All, openid, profile, User.Read, email, Group.Read.All, Group.ReadWrite.All
```

The permissions listed above include those inherited from the service principal and any others requested by the user for the session.

When you're finished interacting with the Graph, remember to close off the session by running *Disconnect-MgGraph* to sign the session out from the Graph. Disconnecting the session removes the encrypted token cache and prevents a session from being reinitialized.

```
[PS] C:\> Disconnect-MgGraph
```

The description above covers an interactive session. This is a good way to get to know the Graph SDK cmdlets and debug scripts in preparation for operational use. Because of the issues with consent and the service principal, Microsoft recommends that operational scripts have separate apps registered in Azure and use certificate-based authentication ([app-only access](#)). You can certainly write and test the code using the interactive client, but once the code is complete, it's time to run it using a separate app with a scoped permission set. This approach means that you can restrict the permissions assigned to apps to only the set needed by the processing done by the script. The downside of using separate apps with scoped permissions is that over time you might accumulate many registered apps in Azure AD which require management.

## Paging

Earlier, we discussed the topic of pagination for Graph API requests. Pagination allows developers to retrieve pages of data until they fetch all items matching a Graph API query. Unlike Graph requests, SDK cmdlets handle pagination automatically. Cmdlets that can handle large quantities of objects include an `All` parameter. The `All` parameter instructs the cmdlet to fetch pages of data asynchronously until no more pages are available. This approach delivers better performance than waiting for all pages to be collected in memory. For example:

```
[PS] C:\> # Fetch all Azure AD user accounts and all groups
[array]$AllUsers = Get-MgUser -All
[array]$AllGroups = Get-MgGroup -All
```

You can limit the amount of data returned by specifying the `Top` parameter. For instance, to fetch the first 10 users, run:

```
[PS] C:\> Get-MgUser -Top 10
```

If you use the `Invoke-MgGraphRequest` cmdlet to run a Graph query, your code must deal with pagination. For example, this command runs a query to return the set of users in a tenant and includes a parameter to request the count of the objects found:

```
[PS] C:\> [array]$UserCount = Invoke-MgGraphRequest -Uri
"https://graph.microsoft.com/beta/users?`$count=true" -Headers @{ConsistencyLevel='eventual'}
```

If we examine the number returned for the count, we see a value:

```
$UserCount
Name Value
----
@odata.context https://graph.microsoft.com/beta/$metadata#users
@odata.count 517
@odata.nextLink
https://graph.microsoft.com/beta/users?$count=true&$skiptoken=m~AQAn0zc3MGE2YmZjZTE5...
Value {System.Collections.Hashtable, System.Collections.Hashtable,
System.Collections.Hash...
```

The user account data is stored in the `$UserCount.Value` hash table. However, if you examine the count of objects, you'll see it is 100. This is the default number of objects returned by the Users Graph API (the number of returned items differs from API to API). To get more data, use the `@odata.nextlink` link returned in `$UserCount` to fetch successive pages until the nextlink value is null. For instance:

```
[PS] C:\> [Array]$MoreUsers = Invoke-MgGraphRequest -Uri $UserCount."@odata.nextlink" -Method Get
```

## Permissions

As noted earlier, the permissions available in an SDK session depend on those granted to the service principal used in the connection. Given the number of Graph permissions available, it can be hard to know what

permissions are needed unless you're told. The *Find-MgGraphPermission* cmdlet helps. For example, let's assume that you want to find permissions related to users. This command does the trick and shows all the permissions requiring administrator consent (output edited for space):

```
[PS] C:\> Find-MgGraphPermission User | ? {$_.Consent -eq "Admin"} | Format-Table Consent, Name, Description
```

Consent	Name	Description
Admin	CrossTenantUserProfileSharing.Read	Allows the application to list and query user
Admin	CrossTenantUserProfileSharing.Read.All	Allows the application to list and query any
Admin	CrossTenantUserProfileSharing.ReadWrite	Allows the application to list and query user
Admin	CrossTenantUserProfileSharing.ReadWrite.All	Allows the application to list and query any
Admin	User.Export.All	Allows the app to export data (e.g. customer
Admin	User.Invite.All	Allows the app to invite guest users to the
Admin	User.ManageIdentities.All	Allows the app to read, update and delete
Admin	User.Read.All	Allows the app to read the full set of profile
Admin	User.ReadWrite.All	Allows the app to read and write the full set

To find more details about the permission, use the *Find-MgGraphPermission* cmdlet to examine its properties:

```
[PS] C:\> Find-MgGraphPermission "User.ReadWrite.All" | ? {$_.PermissionType -eq "Application"} | fl
```

```
Id           : 656f6061-f9fe-4807-9708-6a2e0934df76
PermissionType : Application
Consent       : Admin
Name         : IdentityRiskyUser.ReadWrite.All
Description   : Allows the app to read and update identity risky user information for your
                organization without a
                signed-in user. Update operations include dismissing risky users.

Id           : 741f803b-c850-494e-b5df-cde7c675a1ca
PermissionType : Application
Consent       : Admin
Name         : User.ReadWrite.All
Description   : Allows the app to read and update user profiles without a signed in user.
```

The *Find-MgGraphCommand* cmdlet reports what SDK cmdlets are available to process objects and the permissions they use. The cmdlet takes a URI parameter to know what objects you're interested in. This corresponds to the Graph endpoint for an object, like */groups/* for Groups and */users/* for user accounts. For example, this command shows that two cmdlets are available. One gets a set of users; the other creates a new account:

```
[PS] C:\> Find-MgGraphCommand -Uri /users/ -ApiVersion V1.0
```

```
APIVersion: v1.0
```

Command	Module	Method	URI	OutputType	Permissions
Get-MgUser	Users	GET	/users	IMicrosoftGraphUser1	{DeviceManagementApps.Read.All,
New-MgUser	Users	POST	/users	IMicrosoftGraphUser1	{DeviceManagementApps.ReadWrite.All,

If we add a qualifier to indicate that we want to see commands which process individual user accounts, we see three cmdlets:

```
[PS] C:\> Find-MgGraphCommand -Uri "/users/{id}" -ApiVersion V1.0
```

```
APIVersion: v1.0
```

Command	Module	Method	URI	OutputType	Permissions
Get-MgUser	Users	GET	/users/{user-id}	IMicrosoftGraphUser1	{DeviceManagementApps.Read.All,
Remove-MgUser	Users	DELETE	/users/{user-id}		{DeviceManagementApps.ReadWrite.All
Update-MgUser	Users	PATCH	/users/{user-id}		{DeviceManagementApps.ReadWrite.All



The output truncates the permissions, so to find out what permissions a cmdlet supports, we run a command like (output truncated for space):

```
[PS] C:\> Find-MgGraphCommand -Command Get-MgUser -ApiVersion V1.0 | Select -ExpandProperty Permissions
```

Name	IsAdmin	Description
Directory.AccessAsUser.All	True	Access the directory as
Directory.Read.All	True	Read directory data
Directory.ReadWrite.All	True	Read and write directory
User.Read	False	Sign you in and read your
User.Read.All	True	Read all users' full
User.ReadBasic.All	False	Read all users' basic
User.ReadWrite	False	Read and update your
User.ReadWrite.All	True	Read and write all users' full

You now have a reasonable idea of what permission you need for any specific action.

In some instances, applications restrict access to data so that the only data available is that which the user is entitled to access through the application. For instance, a user must be a member of a team to be able to access the tags available to that team. Therefore, if they use the *Get-MgTeamTag* cmdlet to retrieve tags for a team, the cmdlet happily returns tags for any team the signed-in user is a member of but won't for any team they don't belong to. In Graph terms, the cmdlet uses delegate permission to return information the user can see whereas application permission is necessary to return data for all teams.

The solution is to create a registered Azure AD application, assign the necessary permission to the application, and use certificate-based authentication to connect the app to the SDK endpoint. The SDK recognizes that the app has consent to use specific application permissions and allows access to all data. See [this article for more information](#).

## Invoke-MgGraphRequest

The *Invoke-MgGraphRequest* cmdlet makes Graph API calls where a suitable SDK cmdlet is unavailable or doesn't return the same information as an API call does. One advantage of this cmdlet is that you can use the example code (Uri) shown in the Graph API documentation. For example, this call retrieves the groups an Azure AD user is a member of (using the [List memberOf call](#)):

```
[PS] C:\> $Uri = "https://graph.microsoft.com/beta/users"
$GroupType = "microsoft.graph.group"
$User = Get-MgUser -UserId Chris.Bishop@office365itpros.com
[Array]$Groups = (Invoke-MgGraphRequest -Uri "$Uri/$(($user.Id)/memberOf/$GroupType)").Value
```

The result stored in the *\$Groups* array is the first 100 groups the user belongs to. If more groups are available, the *@odata.nextLink* value contains the Uri to the next page. The command retrieves the information about groups from the Value property returned by the query.

## The Use of \$Null

*\$Null* is an [automatic PowerShell variable](#) that can be used in comparisons and to set values. It's often used to reset other variables or to clear values set on object properties. SDK cmdlets are not particularly good at handling *\$Null* and many reject its use. For example, despite being legitimate PowerShell code, this command to clear the mobile phone for a user won't work:

```
[PS] C:\> Update-MgUser -UserId Terry.Hegarty@office365itpros.com -MobilePhone $Null
```

Instead, to clear the property, it's necessary to write a space into it:

```
[PS] C:\> Update-MgUser -UserId Terry.Hegarty@office365itpros.com -MobilePhone " "
```

If you want to use a Null value instead of a space, you can do so by running *Invoke-MgGraphRequest*. For example:

```
[PS] C:\> $Uri = "https://graph.microsoft.com/v1.0/users/Terry.Hegarty@office365itpros.com"
Invoke-MgGraphRequest -Uri $Uri -Method Patch -Body @{mobilePhone' = $Null}
```

Problems using \$Null also exist when filtering or searching for objects that have null values (or spaces) in properties. For example, these commands don't work.

```
[PS] C:\> Get-MgUser -Filter "Country eq $Null"
Get-MgUser -Search "Country:$Null" -ConsistencyLevel Eventual
```

Needless to say, the way that the SDK handles PowerShell \$Null values is not a good situation to encounter, especially when converting scripts to replace old Azure AD cmdlets that handle \$Null properly. It's certainly something to consider during the conversion process.

## Azure AD Account Management

To perform Azure AD account management operations, connect to the SDK endpoint with *User.ReadWrite.All* and *Directory.ReadWrite.All* permissions. First, let's create a new account. To do this, we must create a password object in a hash table. The table contains the password for the account and a setting to force the user to change their password when they first sign into Microsoft 365:

```
[PS] C:\> $NewPassword = @{}
$NewPassword["Password"] = "!NewPassword2022!"
$NewPassword["ForceChangePasswordNextSignIn"] = $True
```

To create the new account, run the *New-MgUser* cmdlet.

```
[PS] C:\> New-MgUser -UserPrincipalName "June.Brown@Office365ITPros.com" -DisplayName "June Brown
(Sales Operations)" -PasswordProfile $NewPassword -AccountEnabled -MailNickName June.Brown -City NYC
-CompanyName "Office 365 for IT Pros" -Country "United States" -Department "Sales Operations" -
JobTitle "GM Operations" -BusinessPhones "+1 676 830 1101" -MobilePhone "+1 617 4466615" -State "New
York" -StreetAddress "1, Avenue of the Americas" -Surname "Brown" -GivenName "June" -UsageLocation
"US" -OfficeLocation "NYC"
```

The usage location is a [two-character ISO-3166 country code](#) to show where the account consumes services, and it's important to set the value correctly so that the license assignment works properly. After creating a new account, you'll need to assign it some licenses. We'll get to that topic soon.

To make changes to the account, run the *Update-MgUser* cmdlet:

```
Update-MgUser -UserId June.Brown@Office365itpros.com -JobTitle "Senior Sales Manager" -State CA
```

The Microsoft Graph currently treats email proxy addresses as read-only properties for users and groups. You cannot use the *Update-MgUser* cmdlet to add or remove a proxy address from an Azure AD account. However, if you update an account's user principal name, *Update-MgUser* automatically updates the account's primary SMTP address to match the user principal name. This is likely due to adherence to Microsoft's best practice that the two properties should match whenever possible. If you need to update the set of proxy addresses for an account or change its primary SMTP address, use the *Set-Mailbox* cmdlet. These commands add a new proxy address and set it as the primary SMTP address:

```
[PS] C:\> Set-Mailbox -Identity $UserId -EmailAddresses @{Add="Johnnie.West@Office365itpros.com"}
Set-Mailbox -Identity $UserId -WindowsEmailAddress Johnnie.West@Office365itpros.com
```

To update a user's manager, we must find the identifier for the manager's Azure AD account and pass it to the *Set-MgUserManagerByRef* cmdlet. Here's how:

```
[PS] C:\> $ManagerId = (Get-MgUser -UserId Chris.Bishop@office365itpros.com).Id
Set-MgUserManagerByRef -UserId June.Brown@office365itpros.com `
```

```
-AdditionalProperties @{
  "@odata.id" = "https://graph.microsoft.com/v1.0/users/$ManagerId" }
```

Because the manager property is resolved by reference, we check its value by running *Get-MgUser* and then fetching the information retrieved by expanding the Manager property:

```
[PS] C:\> $ManagerData = Get-Mguser -UserId $UserId -ExpandProperty Manager
$ManagerData.Manager.AdditionalProperties['displayName']
```

Chris Bishop

Deleting an account is done using the *Remove-MgUser* cmdlet. The cmdlet doesn't prompt for confirmation:

```
[PS] C:\> Remove-MgUser -UserId June.Brown@Office365itpros.com
```

Deleted accounts go into the Azure AD deleted items container and remain there for 30 days before Azure AD deletes them permanently. See this article for information about how to [list the set of soft-deleted user accounts](#). After finding the identifier of the account to restore, run the *Invoke-MgGraphRequest* cmdlet to perform the action, like this:

```
[PS] C:\> $Uri = "https://graph.microsoft.com/beta/directory/deleteditems/388d29d7-4c72-476d-be96-53060043122e/restore"
Invoke-MgGraphRequest -Method Post -Uri $Uri
```

To delete an account permanently and make it irrecoverable, use a Delete command:

```
[PS] C:\> $Uri = "https://graph.microsoft.com/beta/directory/deleteditems/388d29d7-4c72-476d-be96-53060043122e"
Invoke-MgGraphRequest -Method Delete -Uri $Uri
```

To retrieve sets of users, run the *Get-MgUser* cmdlet with parameters to tell Azure AD what to fetch. The parameters are:

- **All:** Fetch all matching objects.

```
[PS] C:\> [array]$Users = Get-MgUser -All
```

- **Top:** Fetch between 1 and 999 matching objects.

```
[PS] C:\> [array]$Top500 = Get-MgUser -Top 500
```

- **Filter:** Apply a server-side filter against account properties. For example, here's how to get the set of guest accounts in a tenant:

```
[PS] C:\> [array]$Guests = Get-MgUser -Filter "usertype eq 'Guest'" -All
```

Conversely, to get the set of tenant accounts:

```
[PS] C:\> [array]$Users = Get-MgUser -Filter "usertype eq 'Member'" -All
```

To find the accounts in a specific country:

```
[PS] C:\> Get-MgUser -filter "Country eq 'Ireland'"
```

Here's another example of how to use the *startsWith* filter. In this case, the filter compares two different properties to find objects:

```
[PS] C:\> Get-MgUser -Filter "startsWith(displayname, 'Ken') or startsWith(UserPrincipalName, 'Kim')"
```

This command returns the set of member accounts created between two dates. The trailing Z is important as without its inclusion, the Graph will not consider the dates as valid:

```
[PS] C:\> Get-MgUser -Filter "createdDateTime ge 2022-01-01T00:00:00Z and createdDateTime le 2022-04-01T00:00:00Z and usertype eq 'Guest'"
```

- **Search:** Search a property for a value and return all matching accounts. This example applies a search to the *DisplayName* property to return all accounts with the string "Paul" found in that property. You must enclose the name of the property and the value to search for in single quotes. It's also necessary to include the *ConsistencyLevel* parameter because searching is an advanced query.

```
[PS] C:\> Get-MgUser -Search 'DisplayName:Paul' -ConsistencyLevel Eventual
```

You can combine Search and a Filter. For example:

```
[PS] C:\> Get-MgUser -Search "Country:Ireland" -ConsistencyLevel Eventual -Filter "startsWith(displayname, 'James')"
```

**Graph X-Ray:** The Graph X-Ray extension for Edge and Chrome gives developers a peek into the Graph API commands issued by the Azure AD admin center. It can help you to understand the format of a command and offers the chance to download commands in a PowerShell script. See [this article](#) for more information.

## Assigning Azure AD Roles to User Accounts

The *Get-MgDirectoryRole* returns the set of Azure AD roles known in the tenant. To assign a role to a user account, retrieve the identifiers for the role and the user to whom to assign the role. For example:

```
[PS] C:\> $ExoAdminRoleId = Get-MgDirectoryRole | ? {$_.displayName -eq "Exchange administrator"} |
Select -ExpandProperty Id
$UserId = (Get-MgUser -UserId Brian.Weakliam@Office365itpros.com).Id
```

Then assign the role using the *New-MgDirectoryRoleMemberByRef* cmdlet, passing the identifier of the role to assign and the user:

```
[PS] C:\> New-MgDirectoryRoleMemberByRef -DirectoryRoleId $ExoAdminRoleId -BodyParameter
@{"odata.id" = "https://graph.microsoft.com/v1.0/directoryObjects/($UserId)"}
```

To check that the user now has the role assignment, use the *Get-MgDirectoryRoleMember* cmdlet to fetch the set of accounts assigned the role, and check the names:

```
[PS] C:\> RoleMembers = Get-MgDirectoryRoleMember -DirectoryRoleId $ExoAdminRoleId
ForEach ($RoleMember in $RoleMembers) {
    Write-Host $RoleMember.AdditionalProperties["displayName"]}
```

Currently, the Microsoft Graph PowerShell SDK doesn't include a cmdlet to remove a role assignment. You can remove role assignments through the Azure AD admin center, or by [running a Graph query](#).

## Azure AD Group Management

The types of groups included in Azure AD are:

- Microsoft 365 groups/Teams (including groups used by Yammer).
- Security groups.
- Dynamic Azure AD groups (including those used by Microsoft 365 groups/Teams).
- Distribution lists.
- Mail-enabled security groups.

This section discusses how to use SDK cmdlets to create, update, remove, and restore Azure AD Groups.

### Creating New Azure AD Groups

The *New-MgGroup* cmdlet creates Microsoft 365 groups, dynamic groups, and security groups. It can't create distribution lists or mail-enabled security groups (these are essentially Exchange Online rather than Azure AD

objects). Although *New-MgGroup* can create groups of different types, it is often better to use the dedicated cmdlet for a particular type of group to ensure that Microsoft 365 performs all the necessary provisioning, like *New-UnifiedGroup* or *New-Team*. Here's an example of how to create a new Microsoft 365 Group with *New-MgGroup*:

```
[PS] C:\> $Group = New-MgGroup -Description "Banking Team" -DisplayName "Banking Managers and Others" -MailEnabled:$True -SecurityEnabled:$True -MailNickname "BankingGroup" -GroupTypes "Unified"
```

It's a good idea to capture the result of the cmdlet in a variable. If the command is successful, you'll then have a variable containing properties of the new group including its identifier. As we'll see, you'll need the identifier to interact with the group using other SDK cmdlets.

Although this command successfully creates a new Microsoft 365 group, we don't recommend that you use *New-MgGroup* in this way. It's better to use the *New-UnifiedGroup* cmdlet because it gives more control over group properties such as its privacy setting, sensitivity label, and visibility. It doesn't make sense to create a new group and then run the *Set-UnifiedGroup* cmdlet to update its properties. The same logic applies if you want to create a new team as the easiest method is to use the *New-Team* cmdlet.

## Dynamic Microsoft 365 Groups

The *New-UnifiedGroup* cmdlet doesn't support the creation of dynamic Microsoft 365 groups. Instead, you use the *New-MgGroup* cmdlet and mark the group as being a unified (Microsoft 365 group) with dynamic membership. For example:

```
[PS] C:\> $Group = New-MgGroup -DisplayName "Salespeople in France" -Description "Dynamic group containing all salespeople based in France" -MailEnabled:$True -SecurityEnabled:$False -MailNickname SalesPeopleFrance -GroupTypes "DynamicMembership", "Unified" -MembershipRule "(User.Department -eq ""Sales"" and User.Country -eq ""France"")" -MembershipRuleProcessingState "On"
```

Looking at the *New-MgGroup* command, we see:

- Exchange Online uses the value passed in the *MailNickname* parameter to create the SMTP address of the group (in this case, SalesPeopleFrance@Office365ITPros.com) and the group's alias. If the alias is not unique, Groups adds a four-digit suffix to make it unique.
- The *MailEnabled* and *SecurityEnabled* switches are both defined. *MailEnabled* must be True, but you can set *SecurityEnabled* to be True or False.
- Specifying *Unified* in the *GroupTypes* parameter tells Azure AD to create a Microsoft 365 group. You must include both *DynamicMembership* and *Unified* in *GroupTypes* to create a dynamic Microsoft 365 group.
- *MembershipRuleProcessingState* determines whether Azure AD evaluates group membership on an ongoing basis to find the most up-to-date member set. If you are in the middle of a migration, you can pause processing (set the value to Paused) to stop the evaluation of group membership against Azure AD. In this case, the most recent member set is used.

If you're used to building Exchange dynamic distribution lists, you might find that [a limited set of properties](#) are available to build the membership rule for a dynamic Microsoft 365 group. This is because the properties available are those defined in the Azure AD schema rather than the scheme for the Exchange Directory. Even so, because the request used to retrieve the membership of a dynamic group can become quite complex, it is usually easier to use the GUI in the Azure AD admin center to create dynamic Microsoft 365 groups.

To team-enable the dynamic group, run the *New-Team* cmdlet and specify the identifier of the newly-created group:

```
[PS] C:\> New-Team -GroupId $Group.Id
```

Alternatively, a group owner can team-enable a dynamic Microsoft 365 group through the Teams desktop client.

## Searching for Groups

The *Get-MgGroup* cmdlet fetches details of Azure AD groups. To retrieve a single group, use its display name as a filter:

```
[PS] C:\> Get-MgGroup -Filter "DisplayName eq 'Leadership Team'"
```

If you add the *All* parameter, you'll get all the groups in the tenant.

```
[PS] C:\> [array]$Groups = Get-MgGroup -All
```

The command returns groups of all types. To filter out the various types of groups, we can check different properties to identify each type of group. Table 23-5 lists some useful properties to check.

<b>Property</b>	<b>Used by</b>
MailEnabled = True	Distribution lists Microsoft 365 groups Mail-enabled security groups
SecurityEnabled = True	Security groups Mail-enabled security groups
GroupTypes = Unified	Microsoft 365 groups
GroupTypes = DynamicMembership	Dynamic Azure AD groups
GroupTypes = Unified, DynamicMembership	Dynamic Microsoft 365 groups
ResourceProvisioningOptions = Team	Team-enabled Microsoft 365 groups

Table 23-5: Filter possibilities for different types of Azure AD group

The simplest filters are those which find groups based on a property. For example, to find all security-enabled groups:

```
[PS] C:\> Get-MgGroup -Filter "securityEnabled eq true" -All
```

Find all mail-enabled groups:

```
[PS] C:\> Get-MgGroup -Filter "mailEnabled eq true" -All
```

The *GroupTypes* and *ResourceProvisioningOptions* properties require complex filters with [Lambda operators](#). For example, to find the set of Microsoft 365 groups in the tenant:

```
[PS] C:\> [array]$M365Groups = Get-MgGroup -Filter "groupTypes/any(c:c eq 'unified')" -All
```

To find the set of dynamic Microsoft 365 Groups:

```
[PS] C:\> Get-MgGroup -Filter "groupTypes/any(c:c eq 'dynamicmembership') and groupTypes/any(x:x eq 'unified')" -All
```

To find the set of Microsoft 365 groups enabled for Teams:

```
[PS] C:\> [array]$Teams = Get-MgGroup -Filter "resourceProvisioningOptions/Any(x:x eq 'Team')" -All
```

In addition, client-side filter can refine the results returned by the server. For instance, after fetching all security-enabled groups, we use a client-side filter to find the set with dynamic membership:

```
[PS] C:\> [array]$SecurityGroups = Get-MgGroup -Filter "securityEnabled eq true" -All |  
[array]$DynamicSecurityGroups = $SecurityGroups | ? {$_.GroupTypes -eq "DynamicMembership"}
```

The filters used by *Get-MgGroup* are server-side, meaning that the data is filtered when the server returns it to PowerShell. Because they're Graph-based and return fewer properties than cmdlets like *Get-UnifiedGroup*, these commands are very fast, which makes them worth considering if you have scripts that fetch subsets of groups for processing.

## Updating Group Properties and Membership

PowerShell modules like Exchange Online and Azure AD usually include *Set-* cmdlets to update the properties of objects. The SDK uses *Update-Mg\** cmdlets, so to update a group, you run the *Update-MgGroup* cmdlet. For example:

```
[PS] C:\> Update-MgGroup -GroupId $Group.Id -Description "Everyone involved with sales in France"
```

You cannot use *Update-MgGroup* to add or remove email proxy addresses for a group because the Microsoft Graph treats the *proxyAddresses* attribute as read-only. If you need to change a proxy address for a group, use an Exchange Online cmdlet like *Set-UnifiedGroup* instead.

The *New-MgGroupMember* cmdlet populates group membership (for non-dynamic groups). In this example, we fetch the group identifier, use *Get-MgUser* to find a set of suitable group members, and finally add them to the group:

```
[PS] C:\> $GroupId = (Get-MgGroup -Filter "displayName eq 'Sales Operations Team').Id
[array]$Users = Get-MgUser -Filter "department eq 'Sales'"
ForEach ($User in $Users) {
    New-MgGroupMember -GroupId $GroupId -DirectoryObjectId $User.Id }
```

Unfortunately, checking that the right members are present afterward is not as simple as it should be (and is with either the *Get-AzureADGroupMember* or *Get-UnifiedGroupLinks* cmdlets). Instead of *Get-MgGroupMember* returning a list of group members, meaning that to see information like the display names or user principal names for group members, you must pipe the output to the *Get-MgUser* cmdlet:

```
[PS] C:\> Get-MgGroupMember -GroupId $GroupId -All | ForEach {Get-MgUser -UserId $_.Id}
```

Adding a group owner is a little complicated because the owner is stored by reference to its object rather than as a simple property. The *New-MgGroupOwnerByRef* cmdlet, which adds the manager for a group, requires the identifier for the owner's account to be passed in a hash table:

```
[PS] C:\> New-MgGroupOwnerByRef -GroupId $GroupId -AdditionalProperties
@{"@odata.id"="https://graph.microsoft.com/v1.0/users/2a3b60f2-b36b-4758-8533-77180031f3d4"}
```

Remember that synchronization to workloads is necessary before owner or member deletions take full effect across Microsoft 365.

## Removing Group Members

To remove a group owner, you need to know the group identifier and the identifier for the owner's account. You can then run the *Remove-MgGroupOwnerByRef* cmdlet. For example:

```
[PS] C:\> $GroupId = (Get-MgGroup -Search 'DisplayName:Ultra Fans' -ConsistencyLevel Eventual).Id
$UserId = (Get-MgUser -UserId Sean.Landy@Office365itPros.com).Id
Remove-MgGroupOwnerByRef -DirectoryObjectId $UserId -GroupId $GroupId
```

The cmdlet won't allow you to remove the last owner of a group. If you want to remove the user as a group member, run the *Remove-MgGroupMemberByRef* cmdlet:

```
[PS] C:\> Remove-MgGroupMemberByRef -DirectoryObjectId $UserId -GroupId $GroupId
```

## Removing and Restoring Groups

The *Remove-MgGroup* cmdlet removes a Microsoft 365 group or security group. For example:

```
[PS] C:\> Remove-MgGroup -GroupId f6dd8a3e-d50c-4af2-a9cf-f4adf71ec82b
```

The cmdlet can't remove distribution lists or mail-enabled security groups.

## Finding Deleted Groups

Azure AD uses a two-phase deletion process for groups. First, a deleted group enters a soft-deleted state where it remains for 30 days. During this time, administrators can restore the group using options in the Microsoft 365 and Azure AD admin centers. After the 30-day retention period lapses, Azure AD removes the group permanently and it is then irrecoverable.

During the soft-deleted retention period, it's possible to restore the group using the *Restore-MgGroup* cmdlet after using the *Get-MgDirectoryDeletedItem* cmdlet to return find the correct group identifier to use. For example:

```
[PS] C:\> Restore-MgGroup -GroupId 314bab9d-af9b-43b8-b627-af38075d40b0
```

However, some issues exist with these cmdlets in the current version of the SDK. Instead, we can use a Graph API query to fetch the set of deleted groups. In this example, the query executed by the *Invoke-MgGraphRequest* cmdlet retrieves details for up to 100 deleted groups and displays them using the *Out-GridView* cmdlet.

```
[PS] C:\> Connect-MgGraph
Select-MgProfile Beta
$uri =
"https://graph.microsoft.com/beta/directory/deleteditems/microsoft.graph.group?`$select=id,displayName,groupTypes,deletedDateTime&`$orderBy=displayName%20asc&`$top=100"
[array]$Groups = (Invoke-MgGraphRequest -Uri $Uri).Value
If (!$Groups) { write-Host "No deleted groups available for recovery" ; break }
$Report = [System.Collections.Generic.List[Object]]::new(); $Now = Get-Date
ForEach ($Group in $Groups) {
    $PermanentRemovalDue = Get-Date($Group.deletedDateTime).AddDays(+30)
    $TimeTillRemoval = $PermanentRemovalDue - $Now
    $ReportLine = [PSCustomObject]@{
        Group           = $Group.DisplayName
        Id              = $Group.Id
        Deleted         = $Group.deletedDateTime
        PermanentDeleteOn = Get-Date($PermanentRemovalDue) -format g
        DaysRemaining  = $TimeTillRemoval.Days
    }
    $Report.Add($ReportLine)
}
$Report | Sort {$_.PermanentDeleteOn -as [datetime]} | Out-GridView
```

To restore a deleted group, run a POST query and pass the identifier for the group to restore. Once again, the *Invoke-MgGraphRequest* cmdlet runs the query:

```
[PS] C:\> $Uri = "https://graph.microsoft.com/beta/directory/deleteditems/8783e3dd-66fc-4841-861d-49976f0617c0/restore"
Invoke-MgGraphRequest -Method POST -Uri $Uri
```

The newly restored group should be available and ready for users within an hour.

## Using Advanced Azure AD Queries Against Groups

*Get-MgGroup* supports [advanced queries against Azure AD](#). As discussed earlier, an advanced query is one that the Microsoft Graph resolves against a separate property store. It's not always obvious when an advanced query is necessary. Microsoft could hide this need in code instead of forcing PowerShell coders to remember when they must add the *ConsistencyLevel* parameter to mark a query as advanced. Searching the display name of groups for a term is an example of an advanced query.

```
[PS] C:\> [array]$Groups = Get-MgGroup -Search "'displayname:Office 365'" -ConsistencyLevel Eventual
```

Another example is to use the *Filter* parameter to find groups which start with a value. For instance, we might want to find groups whose display name starts with Office:



```
[PS] C:\> [array]$Groups = Get-MgGroup -Filter "startsWith(DisplayName, 'Office')" -ConsistencyLevel Eventual
```

## Managing Account Licenses

To start our discussion about how to manage licenses through PowerShell, let's review some terms used by Microsoft:

- A **SKU** (traditionally, a "stock keeping unit") is something a customer can purchase, like Office 365 E3. SKUs can be bought on a subscription or trial basis. A SKU is often referred to as a **product** or **license**. To demonstrate the intermingling of terms, when you manage user accounts through the Microsoft 365 admin center, the GUI shows that you assign licenses to accounts.
- A **service plan** is a license for an individual feature included in a SKU. For example, Microsoft Bookings, Microsoft Planner, and Mobile Device Management for Office 365 are service plans included in the Office 365 E3 SKU. You cannot buy a service plan as these items are always part of products. In the Microsoft 365 admin center, the individual service plans listed for a user account are called **apps**.
- Organizations can buy **add-on licenses** to access an individual capability. For example, you can purchase Viva Topics add-on licenses to use this capability to build a knowledge network. Add-on licenses are assigned to user accounts to allow access to the capability.

Because product names and service plan names differ from language to language, Microsoft 365 uses GUIDs to identify SKUs and service plans. Microsoft documents the [product names and their SKU identifiers](#) online. This is a very useful page to bookmark because you'll often need to know what the identifier is for a product or service plan. The page also lists the services plans that are included in each product, including the service plan identifiers.

Now that we understand the structure of licenses, we can proceed to interact with those available in a tenant using cmdlets from the Graph SDK. To see license information in user accounts, you must connect to the Graph with the beta profile. Your session needs the *Directory.ReadWrite.All* permission to set licenses.

```
[PS] C:\> $RequiredScopes = @(Directory.ReadWrite.All)
Connect-MgGraph -Scopes $RequiredScopes
Select-MgProfile Beta
```

### Fetching License Information

The *Get-MgSubscribedSku* cmdlet reports the set of products known in a tenant. Before a SKU shows up in a tenant, the organization must subscribe for at least one license (free or paid-for). This command shows the list of SKU identifiers, the SKU names, and the consumed units (allocated licenses).

```
[PS] C:\> Get-MgSubscribedSku | Format-Table SkuId, SkuPartNumber, ConsumedUnits
```

SkuId	SkuPartNumber	ConsumedUnits
1f2f344a-700d-42c9-9427-5cea1d5d7ba6	STREAM	3
b05e124f-c7cc-45a0-a6aa-8cf78c946968	EMSPREMIUM	4
4016f256-b063-4864-816e-d818aad600c9	TOPIC_EXPERIENCES	9
6fd2c87f-b296-42f0-b197-1e91e994b900	ENTERPRISEPACK	25
f30db892-07e9-47e9-837c-80727f46fd3d	FLOW_FREE	9
a403ebcc-fae0-4ca2-8c8c-7a907fd6c235	POWER_BI_STANDARD	7
26d45bd9-adf1-46cd-a9e1-51e9a5524128	ENTERPRISEPREMIUM_NOPSTNCONF	4

The output for the examples in this section are edited. Typically, many more product and service plans are available within a tenant.

Among the list of SKUs, we see products like Office 365 E3 (*EnterprisePack*), Viva Topics (*Topic\_Experiences*), and Office 365 without PSTN (*EnterprisePremium\_NoPSTNConf*). To discover the set of service plans covered

in a product, filter for the product and select the *ServicePlans* property. The *AppliesTo* property indicates if the service plan applies:

- Tenant-wide (like Microsoft Search). These service plans are marked as **Company** can only be enabled or disabled for the complete tenant.
- Per-user basis (like Whiteboard). You can disable service plans marked as **User** to control the functionality available to individual user accounts.

```
[PS] C:\> Get-MgSubscribedSku | ? {$_.SkuPartnumber -eq 'ENTERPRISEPACK' } | Select -ExpandProperty ServicePlans | Format-Table AppliesTo, ServicePlanId, ServicePlanName
```

AppliesTo	ServicePlanId	ServicePlanName
User	b76fb638-6ba6-402a-b9f9-83d28acb3d86	VIVA_LEARNING_SEEDED
Company	db4d623d-b514-490b-b7ef-8885eee514de	Nucleus
Company	2b815d45-56e4-4e3a-b65c-66cb9175b560	ContentExplorer_Standard
User	94a54592-cd8b-425e-87c6-97868b000b91	WHITEBOARD_PLAN2

At a user account level, assigned licenses are in the *AssignedLicenses* property. In this output, we see that the user has three products, and that some service plans are disabled in two products.

```
[PS] C:\> Get-MgUser -UserId Andy.Ruth@Office365itpros.com | Select -ExpandProperty AssignedLicenses | Format-List
```

```
DisabledPlans      : {aebd3021-9f8f-4bf8-bbe3-0ed2f4f047a1, a23b959c-7ce8-4e57-9140-b90eb88a9e97}
SkuId              : 6fd2c87f-b296-42f0-b197-1e91e994b900
AdditionalProperties : {}

DisabledPlans      : {bea4c11e-220a-4e6d-8eb8-8ea15d019f90}
SkuId              : b05e124f-c7cc-45a0-a6aa-8cf78c946968
AdditionalProperties : {}

DisabledPlans      : {}
SkuId              : 4016f256-b063-4864-816e-d818aad600c9
AdditionalProperties : {}
```

The easiest way to interpret the GUIDs is to search [Microsoft's product names and identifiers page](#), where we find that:

- 6fd2c87f-b296-42f0-b197-1e91e994b900 is Office 365 E3, and the two disabled service plans are Kaizala (aebd3021-9f8f-4bf8-bbe3-0ed2f4f047a1) and Sway (a23b959c-7ce8-4e57-9140-b90eb88a9e97).
- b05e124f-c7cc-45a0-a6aa-8cf78c946968 is Enterprise Mobility and Security E5, and the disabled service plan is Rights management (bea4c11e-220a-4e6d-8eb8-8ea15d019f90).
- 4016f256-b063-4864-816e-d818aad600c9 is Viva Topics.

After a while, you might even get to recognize some of the more common product and service plan identifiers.

Service plans can exist in multiple products licenses assigned to a single account. Microsoft 365 only cares that a service plan is available when it checks if an account is licensed to use some functionality. To make it easy to check the license status of a service plan, accounts have an *AssignedPlans* property. We can exploit this to find accounts that can use a certain capability. For instance, let's assume that we want to check on accounts that can use Teams. The Teams1 service plan is included in many Microsoft 365 products, and its service plan identifier is 57ff2da0-773e-42df-b2af-ffb7a2317929. This code creates a quick report.

```
[PS] C:\> $TeamsServicePlanId = "57ff2da0-773e-42df-b2af-ffb7a2317929" #Teams
$PlanUsers = [System.Collections.Generic.List[Object]]::new()
[array]$Users = Get-MgUser -All -Filter "Usertype eq 'Member'"
ForEach ($User in $Users) {
    If ($TeamsServicePlanId -in $User.AssignedPlans.ServicePlanId) {
```

```
$Status = ($User.AssignedPlans | ? {$_.ServicePlanId -eq $TeamsServicePlanId} | Select -
ExpandProperty CapabilityStatus )
$ReportLine = [PSCustomObject] @{
    User      = $User.DisplayName
    UPN       = $User.UserPrincipalName
    Department = $User.Department
    Country   = $User.Country
    Enabled   = $Status }
$PlanUsers.Add($ReportLine) }
}
```

## Find Unlicensed Accounts

Many accounts do not have or need licenses, including:

- Accounts used by Exchange Online shared mailboxes that are smaller than 50 GB or have an archive mailbox.
- Accounts used by Exchange Online room, equipment, or resource mailboxes.
- Guest accounts created by SharePoint Online in Azure AD to share documents with external users.
- Guest accounts created by Azure B2B collaboration for use in applications like Microsoft 365 Groups and Teams.
- Administration accounts (if they don't need to use a mailbox).

To find the set of members (non-guest) accounts that don't have assigned licenses, use this command:

```
[PS] C:\> [array]$UnlicensedAccounts = Get-MgUser -Filter "assignedLicenses/`$count eq 0 and
userType eq 'Member'" -ConsistencyLevel eventual -CountVariable Records -All
$UnlicensedAccounts | Sort DisplayName | Format-Table DisplayName, Mail
```

The Records count variable specified in the command returns the count of records in \$Records.

## Assigning Licenses to User Accounts

The *Set-MgUserLicense* cmdlet assigns licenses to user accounts. You need to know:

- The UPN or object identifier of the target account.
- The SKU identifier of the product you wish to license for the account.

With this information, we can assign a license. In this example, we assign a Viva Topics license (the product identifier is 4016f256-b063-4864-816e-d818aad600c9).

```
[PS] C:\> Set-MgUserLicense -UserId "Terry.Hegarty@Office365itpros.com" -AddLicenses @{SkuId =
'4016f256-b063-4864-816e-d818aad600c9'} -RemoveLicenses @()
```

You must pass an empty array in the *RemoveLicenses* parameter. This applies even if you don't want the command to remove any licenses. If the command completes without an error, the assignment worked (or the account already holds the license), but you can check by reviewing the set of product identifiers stored in the account:

```
[PS] C:\> $Products = (Get-MgUser -UserId Terry.Hegarty@Office365itpros.com).AssignedLicenses.SkuId
$Products
a403ebcc-fae0-4ca2-8c8c-7a907fd6c235
4016f256-b063-4864-816e-d818aad600c9
6fd2c87f-b296-42f0-b197-1e91e994b900
```

Table 23-6 lists some common error conditions that occur during license assignments.

<b>Reason</b>	<b>Example error</b>
No licenses for the specified product are available.	Subscription with SKU 26d45bd9-adf1-46cd-a9e1-51e9a5524128 does not have any available licenses.

The license is unknown in the tenant (a subscription for the license has never existed).	License e82ae690-a2d5-4d76-8d30-7c6e01e6022e does not correspond to a valid company License.
An invalid identifier for a service plan in a compound license is specified.	Service plan 2078e8df-cff6-4290-98cb-5408261a760a for license 26d45bd9-adf1-46cd-a9e1-51e9a5524128 is not valid.
Attempt to remove a license from an account that it doesn't have.	User does not have a corresponding license.

Table 23-6: License assignment errors

To find accounts where error conditions exist for license assignments, run this request:

```
[PS] C:\> $Uri =
'https://graph.microsoft.com/v1.0/users?`$filter=provisionedPlans/any(p:p/provisioningStatus eq
'Error')&`$count=true'
$Data = Invoke-MgGraphRequest -Uri $Uri -Method Get -Headers @{ConsistencyLevel='eventual'}
ForEach ($D in $Data.Value) { Write-Host ("Error found for {0} {1}" -f $D.displayname,
$D.UserPrincipalName) }
```

You'll need to check each account to determine what caused the error condition to exist.

## More Complex Assignments

Viva Topics is a simple license to assign. Products like Office 365 E3 or E5 include multiple service plans, some of which are for functionality that organizations do not want people to use. For example, in academic settings, colleges might not want students to use Yammer for Education, while business tenants might not want people to use Microsoft Bookings, Sway, Kaizala, and so on. To assign a compound license like Office 365 E3 without disabling any service plans, all you need is the product identifier. We know that is 6fd2c87f-b296-42f0-b197-1e91e994b900 (see above). To disable selected service plans, we need to know their identifiers. In this example, we use:

- Microsoft Bookings: 199a5c09-e0ca-4e37-8f7c-b05d533e1ea2
- Yammer Enterprise: 7547a3fe-08ee-4ccb-b430-5077c5041653
- Microsoft Kaizala Pro Plan 3: aebd3021-9f8f-4bf8-bbe3-0ed2f4f047a1

Be careful when you disable a service plan. Because of the interconnected nature of Microsoft 365, you might find that disabling a service plan affects the ability of other features to work properly. For instance, if you disable Stream, Teams won't be able to save meeting recordings.

To start, create a hash table containing the license information. Entries exist for the product and the disabled service plans:

```
[PS] C:\> $LicenseOptions = @{SkuId = "6fd2c87f-b296-42f0-b197-1e91e994b900"; DisabledPlans =
@("199a5c09-e0ca-4e37-8f7c-b05d533e1ea2", "7547a3fe-08ee-4ccb-b430-5077c5041653", "aebd3021-9f8f-
4bf8-bbe3-0ed2f4f047a1")}
```

Then run *Set-MgUserLicense* to assign the license:

```
[PS] C:\> Set-MgUserLicense -UserID Terry.Hegarty@Office365itpros.com -AddLicenses
@($LicenseOptions) -RemoveLicenses @()
```

You can also pass the disabled service plans in an array. For example:

```
[PS] C:\> $LicenseOptions = @("199a5c09-e0ca-4e37-8f7c-b05d533e1ea2", "7547a3fe-08ee-4ccb-b430-
5077c5041653", "aebd3021-9f8f-4bf8-bbe3-0ed2f4f047a1")
$Office365E3Sku = "6fd2c87f-b296-42f0-b197-1e91e994b900"
Set-MgUserLicense -UserID Terry.Hegarty@Office365itpros.com -AddLicenses @{SkuId = $Office365E3Sku;
DisabledPlans = $LicenseOptions } -RemoveLicenses @()
```

To reenable a disabled service plan, remove the service plan from the list of disabled plans and run the command again. In this example, the service plan for Microsoft Bookings is not present in the list of disabled plans specified in the `$LicenseOptions` variable, so it is enabled when we run `Set-MgUserLicense`:

```
[PS] C:\> $LicenseOptions = @"7547a3fe-08ee-4ccb-b430-5077c5041653", "aebd3021-9f8f-4bf8-bbe3-0ed2f4f047a1"

Set-MgUserLicense -UserId Terry.Hegarty@Office365itpros.com -AddLicenses @{SkuId = $Office365E3Sku; DisabledPlans = $LicenseOptions } -RemoveLicenses @()
```

To check the status of disabled service plans in an account, run the `Get-MgUserLicenseDetail` cmdlet:

```
[PS] C:\> Get-MgUserLicenseDetail -UserId Terry.Hegarty@office365itpros.com | ? {$_.SkuId -eq "6fd2c87f-b296-42f0-b197-1e91e994b900"} | Select-Object -ExpandProperty ServicePlans | Sort ProvisioningStatus
```

AppliesTo	ProvisioningStatus	ServicePlanId	ServicePlanName
User	Disabled	aebd3021-9f8f-4bf8-bbe3-0ed2f4f047a1	KAIZALA_O365_P3
User	Disabled	7547a3fe-08ee-4ccb-b430-5077c5041653	YAMMER_ENTERPRISE
User	Disabled	199a5c09-e0ca-4e37-8f7c-b05d533e1ea2	MICROSOFTBOOKINGS
User	Success	57ff2da0-773e-42df-b2af-ffb7a2317929	TEAMS1

To assign multiple licenses to an account in a single operation, create a hash table containing details of the licenses to add and pass the hash table to `Set-UserMgLicense` in the `BodyParameter` parameter instead of using the `AddLicenses` parameter. In this example, we assign licenses for Office 365 E3 and Viva Topics and disable the service plans for Bookings, Yammer, and Kaizala from the Office 365 E3 license:

```
[PS] C:\> $LicenseParams = @{
    AddLicenses = @(
        @{ DisabledPlans = @"199a5c09-e0ca-4e37-8f7c-b05d533e1ea2", "7547a3fe-08ee-4ccb-b430-5077c5041653", "aebd3021-9f8f-4bf8-bbe3-0ed2f4f047a1"
          SkuId = "6fd2c87f-b296-42f0-b197-1e91e994b900" }
        @{ DisabledPlans = @()
          SkuId = "4016f256-b063-4864-816e-d818aad600c9"
        })
    RemoveLicenses = @() }

Set-MgUserLicense -UserId John.Wilson@office365itpros.com -BodyParameter $LicenseParams
```

It's also possible to assign licenses to an account by copying the assignment information from another account. This approach works well when you have configured an account with all the necessary licenses and service plans and need to assign the same licenses to multiple accounts. For example:

```
[PS] C:\> $SourceLicenses = (Get-MgUser -UserId John.Wilson@Office365itpros.com).AssignedLicenses
[array]$TargetUsers = Get-MgUser -Filter "Country eq 'United States' and Usertype eq 'Member'" -All
ForEach ($User in $TargetUsers) {
    Write-Host "Processing" $User.DisplayName $User.Id
    Set-MgUserLicense -UserId $User.Id -AddLicenses $SourceLicenses -RemoveLicenses @() }
```

## Removing a Service Plan from Multiple Accounts

Microsoft often adds a service plan to a product to support the deployment of a new feature. For example, when they made Microsoft Bookings available to Office 365 E3 and E5 licensees, Microsoft introduced a service plan for Microsoft Bookings. It might be the case that an organization is unprepared for their users to access a new feature after deployment and need some extra time to prepare training and support. In this scenario, you could remove the service plan from assigned licenses.

This is an example of removing the Kaizala service plan from accounts assigned Office 365 E3 licenses. The code:

- Sets up the license information to remove.
- Runs `Get-MgUser` to find the set of tenant users.

- Checks each account to see if it has Office 365 E3, and if so, check that the Kaizala service plan is enabled.
- Updates the account to disable the Kaizala service plan.
- Reports the actions taken.

```
[PS] C:\> $RemovePlan = "aebd3021-9f8f-4bf8-bbe3-0ed2f4f047a1" #Kaizala
$LicenseOptions = @{SkuId = "6fd2c87f-b296-42f0-b197-1e91e994b900" ; DisabledPlans = $RemovePlan}
$Report = [System.Collections.Generic.List[Object]]::new()
$TenantUsers = Get-MgUser -Filter "UserType eq 'Member'" -All

ForEach ($User in $TenantUsers) {
    Write-Host "Checking licenses for" $User.DisplayName
    #Check that Office 365 E3 is assigned
    If ("6fd2c87f-b296-42f0-b197-1e91e994b900" -in $User.AssignedLicenses.SkuId) {
        Foreach ($Plan in $User.AssignedPlans) {
            If ($Plan.ServicePlanId -eq $RemovePlan -and $Plan.CapabilityStatus -eq "Enabled") {
                Write-Host ("Removing service plan {0} from account {1}" -f $Plan.Service, $M.DisplayName) -
                ForegroundColor Red
                Set-MgUserLicense -UserId $User.Id -AddLicenses @($LicenseOptions) -RemoveLicenses @(
                $LicenseUpdateInfo = "Kaizala service plan removed from account " + $User.UserPrincipalName + "
on " + (Get-Date) + " from Office 365 E3"
                Write-Host ("Service plan {0} removed from SKU {1} for {2}" -f $Plan.Service, $SelectedSku,
$User.DisplayName)
                $ReportLine = [PSCustomObject][Ordered]@{
                    DisplayName = $User.DisplayName
                    UPN = $User.UserPrincipalName
                    Info = $LicenseUpdateInfo
                    SKU = "Office 365 E3"
                    "Service Plan" = $Plan.Service
                    "ServicePlanId" = $Plan.ServicePlanId }
                $Report.Add($ReportLine)
            } # End if
        } # End ForEach
    } #End if for Office 365 E3 check
} # End ForEach mailbox
```

## Removing Licenses

To remove licenses from a user account, specify the license identifiers in an array and pass the array to the *Set-MgUserLicense* cmdlet in the *RemoveLicenses* parameter. For example:

```
[PS] C:\> [array]$LicensesToRemove = "4016f256-b063-4864-816e-d818aad600c9", "6fd2c87f-b296-42f0-
b197-1e91e994b900"
Set-MgUserLicense -UserId "John.Wilson@Office365itpros.com" -AddLicenses @() -RemoveLicenses
$LicensesToRemove
```

*Set-MgUserLicense* cannot remove a license from an account if the [assignment is via group membership](#).

Another way to remove multiple licenses from an account is to pass the license information in a hash table and pass it in the *BodyParameter* parameter:

```
[PS] C:\> $LicenseParams = @{ AddLicenses = @()
    RemoveLicenses = @("4016f256-b063-4864-816e-d818aad600c9",
    "6fd2c87f-b296-42f0-b197-1e91e994b900") }
Set-MgUserLicense -UserId John.Wilson@office365itpros.com -BodyParameter $LicenseParams
```

## Reporting Assigned Licenses

We can easily find the set of products (SKUs) used in a tenant. With this list, we can find the licenses assigned to user accounts for each product. For example, this code loops through the set of SKUs to report the users with assigned licenses:

```
[PS] C:\> $Report = [System.Collections.Generic.List[Object]]::new()
[array]$Skus = Get-MgSubscribedSku
ForEach ($Sku in $Skus) {
```

```
Write-Host "Processing license holders for" $Sku.SkuPartNumber
$SkuId = $Sku.SkuId
[array]$SkuUsers = Get-MgUser -Filter "assignedLicenses/any(x:x/skuId eq $SkuId)" -All
ForEach ($User in $SkuUsers) {
    $ReportLine = [PSCustomObject] @{
        User      = $User.DisplayName
        UPN       = $User.UserPrincipalName
        Department = $User.Department
        Country   = $User.Country
        SKU       = $Sku.SkuId
        SKUName   = $Sku.SkuPartNumber}
    $Report.Add($ReportLine) }}
$Report | Sort User | Out-GridView
```

Note the use of the filter with the *Get-MgUser* cmdlet to find accounts assigned a specific SKU identifier.

The report above lists accounts and the individual licenses each account holds. With a small change to the code, you could report the data from a license perspective by finding the set of accounts for each license:

```
[PS] C:\> [array]$Skus = Get-MgSubscribedSku
$Report = [System.Collections.Generic.List[Object]]::new()
ForEach ($Sku in $Skus) {
    $SkuId = $Sku.SkuId
    [array]$SkuUsers = Get-MgUser -Filter "assignedLicenses/any(x:x/skuId eq $SkuId)" -All
    $ReportLine = [PSCustomObject][Ordered]@{
        Sku           = $Sku.SkuId
        Product       = $Sku.SkuPartNumber
        'Consumed Units' = $Sku.ConsumedUnits
        'Calculated Units' = $SkuUsers.Count
        'Assigned accounts' = $SkuUsers.UserPrincipalName -Join ", " }
    $Report.Add($ReportLine)
} # End ForEach
$Report | Sort Product | Out-GridView
```

[This article](#) describes a more comprehensive treatment of reporting license assignments to user account.

Another way of reporting assigned licenses is to check against the *AssignedLicenses* property in user accounts. This code looks for accounts assigned the Office 365 E3 license and reports how many it finds:

```
[PS] C:\> [array]$AllUsers = Get-MgUser -Filter "UserType eq 'Member'" -All | Select Id,
AssignedLicenses, UserPrincipalName, DisplayName
[array]$O365E3Users = $AllUsers | Where-Object {$_.AssignedLicenses.SkuId -contains "6fd2c87f-b296-
42f0-b197-1e91e994b900"}
Write-Host ("You have {0} tenant accounts and {1} have Office 365 E3 licenses." -f $AllUsers.Count,
$O365E3Users.Count)
```

## Interacting with Teams

To explore how the SDK cmdlets can interact with Teams, this example shows how to post messages to a team channel. First, we use the *Get-MgUser* cmdlet to list all Azure AD accounts with a display name starting with "Paul." After finding the account we want to use, we then call the *Get-MgUserJoinedTeam* cmdlet using the object identifier for that account to find a list of the teams the chosen account belongs to:

```
[PS] C:\> $RequiredScopes = @("User.Read.All","Group.ReadWrite.All","ChannelMessage.Send")
Connect-MgGraph -Scopes $RequiredScopes -TenantId "828e1143-88e3-492b-bf82-24c4a47ada63"
Get-MgUser -Filter "startsWith(displayname, 'Paul')" -All | Format-Table DisplayName, Id
Get-MgUserJoinedTeam -Userid aff4cd58-1bb8-4899-94de-795f656b4a18
```

After selecting a team and storing its identifier in a variable, we find the identifier for the team's General channel and use it to post a high-priority message to the channel. To perform this operation, the scope for the Graph connection must include the *ChannelMessage.Send* permission, which is why it's included in the scope listed above.

```
[PS] C:\> $TeamId = "82ae842d-61a6-4776-b60d-e131e2d5749c"
```

```
$ChannelId = Get-MgTeamChannel -TeamId $TeamId | ? {$_.DisplayName -eq "General"} | Select -ExpandProperty Id
$Message = (New-MgTeamChannelMessage -TeamId $TeamId -ChannelId $ChannelId -Body @"Content="PowerShell and the Graph SDK can do great things" } -Subject "Greetings to PowerShell fans" -Importance "High")
```

Getting a little more complicated, let's assume that we create a report or similar information in HTML format that we want to publish to the channel. This can be done by specifying a *ContentType* in the *Body* parameter:

```
[PS] C:\> $Message = (New-MgTeamChannelMessage -TeamId $TeamId -ChannelId $ChannelId -Body @"Content = $HTMLContent; ContentType = "html" } -Subject "Report about Teams activity" -Importance "High")
```

In both cases where we post a message, the *\$Message* variable captures the response from Teams. The variable includes an identifier for the message posted to the channel which can be used to post a response. For example:

```
[PS] C:\> $Response = New-MgTeamChannelMessageReply -TeamId $TeamId -ChannelId $ChannelId -ChatMessageId $Message.Id -Body @"Content = "Invaluable!" }
```

## Using Azure Automation with Microsoft 365

Azure Automation allows background jobs such as PowerShell scripts to run on virtual sandbox machines managed by Azure. It is a convenient mechanism to execute long-running PowerShell scripts such as those which must process large quantities of mailboxes, groups, or teams, or to retrieve information from the audit log to export to an external repository. The basic concepts are:

- Jobs run under the control of an Azure Automation account. The account is associated with an Azure subscription to pay for the resources consumed by any processing. The automation account owns resources such as runbooks, credentials, schedules, and modules. It also has a [RunAs account](#), which leverages several assets created for the automation account, including:
  - An Azure AD application with a service principal and a self-signed certificate. The service principal can hold API permissions such as Graph API permissions and be a member of Azure AD role groups.
  - An automation certification asset called *AzureRunAsCertificate* holding the private key used by the Azure AD application. The certificate is valid for one year and is renewable.
  - An automation connection called *AzureRunAsConnection* holding the application identifier, certificate thumbprint, Azure subscription identifier, and tenant identifier. Runbooks use this connection to authenticate to access Azure resources. The information is also usable to secure an access token to connect to the Graph.
- A runbook is where developers create the PowerShell code they wish to run as background jobs. Writing code in a runbook is like writing code in the PowerShell ISE. You can paste code into the runbook, but some work is probably necessary to update the code to make it ready to run under Azure Automation. For example, code in a sandbox is non-interactive, so any prompts or outputs won't work. Developers can execute a runbook in a test pane to validate that the code works before releasing it for production use.
- Credentials can be anything from a username/password combination (stored securely in the automation account, but not recommended for production use) to certificates and certificate thumbprints used for certificate-based authentication. Many PowerShell modules support certificate-based authentication, so runbooks should use this method when possible.
- PowerShell modules adapted for Azure Automation are available for loading into an Automation account to make them available to PowerShell scripts executing in a runbook. Most Microsoft 365 modules (Exchange Online, SharePoint Online, SharePoint PnP, Azure AD, Teams, the Microsoft Graph PowerShell SDK, etc.) are available.



- When a runbook is ready for production use, it can link to a schedule in the automation account to execute on regular basis, such as daily or weekly. The schedule can have a predetermined end or be ongoing.

For more information on using Azure Information with Microsoft 365 workloads, see these articles:

- [Using Azure Automation with Exchange Online PowerShell.](#)
- [Using Azure Automation with the Microsoft Graph PowerShell SDK](#) (send email).
- [Using Azure Automation and PowerShell to create documents in SharePoint Online.](#)

Like the service principal used by interactive sessions with the Microsoft Graph PowerShell SDK, tenant administrators and developers should take care to ensure that the API permissions assigned to automation accounts do not accrue over time and result in overly-permissioned accounts.

## Office Connectors

Connectors allow you to link a feed from a cloud data source to Groups, Teams, and Yammer. The idea behind a connector is that it brings information from the source application to the attention of group members, who consume the information through connector cards. Each card contains a snapshot of information about an item available in the network source. The information is intended to allow the reader to know what's going on, but if they want the full story, they need to go to the source. Usually, the card contains a hyperlink to get to the full content. Connectors are linked to:

- **An Outlook group:** Cards are posted as conversations in the group inbox. Group members can read and respond to cards in the conversation or via email (if they subscribe to the group).
- **A channel in a team:** Cards are posted as new topics in the channel. Team members can read and respond to cards in the channel.
- **Yammer community:** Cards are posted as new topics in the feed. Community members can read and respond to cards like normal posts. This functionality is only supported for Yammer communities connected to Groups.

Connectors support an extensive set of network sources, including those featured on project activities (Trello, Asana, and Wunderlist), customer relationships (Salesforce, Dynamics 365, and Zendesk), news sources (Bing News and RSS feeds), and developer tools (GitHub and Visual Studio). In addition, the *Incoming Webhook* connector is available as a generic link to allow developers to fetch data from other services to a group or team. Programmers can use the webhook to create a link to a group for a company-specific system or some other network data source for which a connector does not exist.

## Actionable Cards

Some connectors, like those from GitHub or Trello, use [actionable or adaptive cards](#), which allow the reader to take quick actions using buttons in the card without having to navigate to the source application. Many inbound messages have a link to bring the user to a website to conduct a transaction such as responding to an inbound customer request or deciding how to process an item. The idea behind actionable cards is that you avoid the need for the reader to switch context by enabling them to execute tasks within the message. Clients like Outlook and Teams can interpret code held in the message to make interactive UI elements available to the reader. Cards can have buttons, dropdown lists, date pickers, and text input boxes.

For more information, see the [Adaptive Cards website](#).

## Adding a New Connector to a Group

To create a new connector for a group, select the group from the list of groups shown in OWA resources, and then open the [...] menu. Select Settings from the menu to display the Group Settings pane and then click the

Connectors link. OWA opens the Office connectors page, where you can see the list of available connectors and can choose to add a new connector or configure one that's already connected to the group.

Connectors often need some form of authentication for the network source. For instance, if you connect Yammer to a group, you must use a valid Microsoft 365 account. Other connectors, like the RSS feed, don't need to be authenticated. For instance, to set up an RSS feed for blog posts relating to this book, all you need to enter is the URL for the blog feed (<https://office365itpros.com/feed/>). After saving the connector, you can close the Office connectors page and return to OWA. Most connectors post a welcome or initial message to the group to inform members that the connector.

You can configure a group with multiple connectors of the same type and use multiple connectors for different sources to gather information from multiple places into a single group. There are many interesting ways that you might consider using connectors with Groups. For instance, you might set up connectors to link to a corporate RSS feed so that group members automatically see updates about important corporate initiatives without having to constantly check the company blog. In addition, the group then serves as an index for blog posts and because the items are in the group mailbox, the messages are discoverable and easily searched. Remember that the messages posted by connectors are only snippets of the information that exists in the network source. Even so, the snippet might be helpful in an eDiscovery investigation.

A key point to understand is that the messages created by connectors simply inform group members about new information. If they find an interesting topic, the group members can discuss that information in group conversations, perhaps after checking out the full content from the original source.

## Using the Incoming Webhook Connector

Connecting something like an RSS feed to a group is a good use for a connector. Being able to use a data source as a source to post cards to a group (or to a channel in a team) is even more interesting. Microsoft makes this possible with the Incoming Webhook connector, which receives incoming HTTP requests holding cards formatted in JSON (the payload) and routes the information to a destination group or team. The *webhook address* generated when the Incoming Webhook connector is linked to a group or team is used by programs to find the connector for future communications. When the group receives inbound information, it uses the information to create a new card. If received by a team, the team creates a new conversation in the target channel. All we need to do to connect a data source to a group or team is:

1. Find the data source to use.
2. Decide on the destination – an Outlook group or a channel in a team.
3. Create an Incoming Webhook connector for the destination.
4. Extract data from the source and format the data as a JSON payload.
5. Post the data to the webhook address.

Any data source that you can access to extract relevant data and format for submission to the connector can be used to send items to Groups or Teams via the incoming webhook connector. You can use most programming languages to extract, format, and send the data. Sending information to a group has the advantage that the cards can be emailed to group members so that they know when action is needed. Items sent to Teams need people to open the host channel to review the content.

The example presented here posts tickets to a group created for support staff when user mailbox quotas exceed a set threshold. For example, the mailbox size is more than 95% of the assigned quota. The information needed to check mailbox sizes and quotas is easily achieved with the *Get-ExoMailbox* and *Get-ExoMailboxStatistics* cmdlets (see the script to report mailbox quotas described earlier).

## Creating an Incoming Webhook Connector

To create the connector, select the group that you want to use and access it from OWA. Select **Manage Connectors** from Group settings (gear icon), and then **Incoming Webhook** from the list of available connectors. Click **Add** to start creating the connector and add (Figure 23-5):

- The name of the connection.
- An image file if you do not want to use the default icon for the connector. OWA seems to ignore the custom image while Outlook displays it.

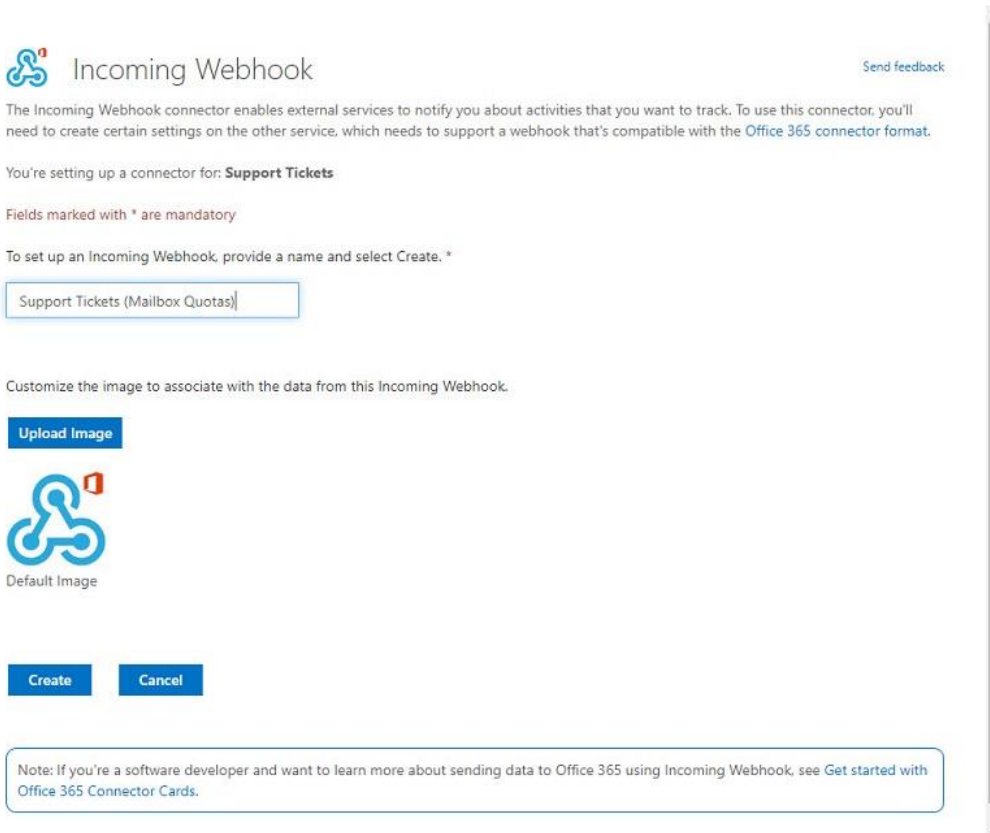


Figure 23-5: Creating a new Incoming Webhook connector

Click **Create** when ready to go ahead. The connector is then created along with its webhook URL, which you can use with the source application to route data through the connector. The easiest way to get the URL is to copy it to the clipboard using the option available in the form. If you forget to save the webhook URL, you can always retrieve it by selecting the connector and editing its properties. Click **Done** to complete the creation process.

### Webhook URL Format

The format of a webhook looks like this:

```
https://office365itpros.webhook.office.com/webhookb2/7aa49aa6-7840-443d-806c-08ebe8f59966@c662313f-14fc-43a2-9a7a-d2e27f4f3478/IncomingWebhook/8592f62b50cf41b9b93ba0c0a00a0b88/eff4cd58-1bb8-4899-94de-795f656b4a18
```

The parts of the webhook are:

- **Tenant name** (taken from Azure AD): In this case, Office365itpros.
- **Webhook root**: webhook.office.com/webhookb2.
- **Group identifier**: The Azure AD identifier for the Microsoft 365 group. In the example, this is 7aa49aa6-7840-443d-806c-08ebe8f59966.

- **Tenant identifier:** The Azure AD identifier for the tenant. In the example, this is c662313f-14fc-43a2-9a7a-d2e27f4f3478. The group and tenant identifiers are separated by an at (@) sign.
- **Provider name:** This is always *"IncomingWebhook."*
- **Alternate identifier:** Another GUID to make the URL unique because a group or channel can support multiple webhooks. In this case, it's 8592f62b50cf41b9b93ba0c0a00a0b88.
- **Group owner:** The Azure AD identifier (GUID) for the group owner. Or rather, the account of the owner who adds the webhook to the group.

To update the URL for a connector, select it and use the Update URL option. Then copy the new URL and update the source application (or PowerShell script) used to generate the inbound data.

## Setting up a Webhook Connector for Teams

Much the same approach is followed to create the webhook connector for a team channel. Go to the channel that you want to use to hold the messages and select Connectors from the menu. Then input the same kind of information given to create a connector for a group. The webhook uses the same format for Teams, with the only difference being that the webhook points to a channel in a team instead of the inbox in a group mailbox.

One thing to remember when posting to Teams is that the maximum message size is 25 KB. If you try to post a larger message, you'll see a *"Microsoft Teams endpoint returned HTTP error 413"* error.

## Creating payloads for an Incoming Webhook connector

We want to send information about service incidents to the group. To collect the information to insert into the payload, we use PowerShell to analyze mailbox quotas and report any over the threshold. Creating the payload means that we format the input data to make it ready for submission to the webhook. In doing so, it is important to remember the following points:

- It can take some time and practice to learn how to use PowerShell to create the JSON content. It is best to start with a basic card and then gradually build up the full content for the cards you want to create.
- Posting to the connector depends on making sure to format the payload correctly. You might be able to convert your payload to JSON successfully using the *ConvertTo-JSON* cmdlet only to find that the *Invoke-RestMethod* cmdlet rejects the payload when processing it for submission to the connector. All of which leads to hours of fun debugging payloads.

This code takes details about a mailbox that's over the threshold and creates a suitable payload in a PowerShell variable. The different elements that we want to report are formed into name-value pairs.

```
[PS] C:\> # $M is the current mailbox being processed
If ($QuotaPercentUsed -gt $Threshold) {
    Write-Host $M.DisplayName "current mailbox use is above threshold at" $QuotaPercentUsed -
Foregroundcolor Red
    $ErrorText = "Mailbox quota over threshold"
    # Log support incident
    $NotificationBody = ConvertTo-Json -Depth 4 @{
        Text = "Support Ticket"
        Title = "New Support Ticket Logged at " + (Get-Date -Format U)
        Summary = "The mailbox size of " + $M.DisplayName + " exceeds the quota warning threshold"
        sections = @(
            @{}
            facts = @(
                @{}
                name = "Mailbox name"
                value = $M.DisplayName
            },
            @{}
            name = "Current mailbox size (GB):"
            value = $QuotaUsedGB
        ),
    },
}
```

```
@{
  name = "Current mailbox quota (GB):"
  value = $MbxQuotaGB
},
@{
  name = "Percent of mailbox quota used:"
  value = $QuotaPercentUsed
},
@{
  name = "Number of mailbox items"
  value = $MbxStats.ItemCount } ) } ) }
# Post support incident via webhook
```

After populating the payload with information about the mailbox, the *Invoke-RestMethod* cmdlet sends the payload to the connector identified by the webhook address stored in the *\$Webhook* variable. The webhook address is generated by Office 365 when it creates the connector. The same approach is taken to post to webhooks pointing to a group or a team channel.

```
[PS] C:\> $PostStatus = (Invoke-RestMethod -ContentType "Application/Json" -Method Post -Body
$NotificationBody -Uri $WebHook) Submitting the Payload
```

If everything goes well, the cmdlet responds with 1 (one) to show that the post was successful. And of course, a new card should appear in the target group or team channel (Figure 23-6).



Figure 23-67: Details of a support ticket posted to a team channel

The example described here is incomplete. It has no error checking nor does not check whether a support ticket already exists for a mailbox. Some additional code will easily address these deficiencies and create an [operational-quality script](#). However, it serves to illustrate the principle of using the Incoming Webhook connector to deliver data from a source to a group or team.

This is a simple example of the information that can be included in a card. Other elements can be included to enhance the card further. For example, the code described in [this blog post](#) creates message cards for new updates fetched from the RSS feed for the Microsoft 365 roadmap. The cards include a More Info button with a hyperlink to the item on the roadmap site. For a full explanation of how to format the JSON payload for an Office connector, see the [Reference for the Adaptive card format](#).

## Using a Template

Another method used to create messages for posting to a connector is to create a template with fields that are replaced with values when you want to post a message. For example, this template has four fields that we replace for each service incident we post.

```
[PS] C:\> Notification = @"
{
  "@type": "MessageCard",
  "@context": "https://schema.org/extensions",
  "summary": "Exchange Online Notification",
```

```

    "themeColor": "0072C6",
    "title": "Mailbox Quota Threshold Exceeded",
    "sections": [
      {
        "facts": [
          {
            "name": "Mailbox:",
            "value": "DISPLAYNAME"
          },
          {
            "name": "Current Quota (GB):",
            "value": "MBXQUOTAGB"
          },
          {
            "name": "Quota used (GB):",
            "value": "QUOTAUSEDGB"
          },
          {
            "name": "Percent Quota used (%):",
            "value": "QUOTAPERCENTUSED"
          }
        ],
        "text": "Exchange Online Mailbox Support Ticket. Please increase the quota for this mailbox or remove some items to get its size under the threshold"
      }
    ]
  }
"@

```

To replace the fields, we extract information for a mailbox into variables and use the PowerShell `replace` method to insert values into the template. This example processes a set of mailbox objects in an array, updates the template with information for each mailbox, and posts the message.

```

[PS] C:\> ForEach ($M in $Mbx) {
    $Body =
    $Notification.Replace("DISPLAYNAME","$M.Mailbox").Replace("MBXQUOTAGB","$M.MbxQuotaGB").Replace("QUOTAUSEDGB","$M.QuotaUsedGB").Replace("QUOTAPERCENTUSED","$M.QuotaPercentUsed")
    $Command = (Invoke-RestMethod -uri $WebHook -Method Post -Body $Body -ContentType 'application/json') }

```

If you're interested in exploring the possibilities that exist in posting messages to channels via the inbound webhook connector, consider the use of the `PSTeams` module ([available in GitHub](#)), which is designed to assist in the process.

## Disabling Office Connectors

The wide range of available connectors makes it easy to see how groups and teams can serve as one-stop collections for inbound information gathered from different cloud sources. However, if you do not want to use connectors, you can disable them for a specific group or the entire tenant. To disable connectors for an individual group, run the `Set-UnifiedGroup` cmdlet to set the `ConnectorsEnabled` property to `$False`. Disabling connectors for a group also disables the ability to add connectors to the team associated with the group, if a team exists for the group.

```

[PS] C:\> Set-UnifiedGroup -Identity "Service Communications" -ConnectorsEnabled:$False

```

If you want to disable connectors for a complete tenant, run the `Set-OrganizationConfig` cmdlet to set the `ConnectorsEnabled` property to `$False`. This command blocks the ability to create connectors for all Office components.

```

[PS] C:\> Set-OrganizationConfig -ConnectorsEnabled:$False

```

Disabling connectors on a tenant-wide basis naturally takes precedence. You cannot block connectors for a tenant and then selectively enable connectors for some groups. Unfortunately, apart from opening each

group to see whether it has any configured connectors, there is no quick way to scan all the groups within a tenant to find which groups have connectors and the sources for the connectors.

## Managing Self-Service Purchases

As described in Chapter 5, Microsoft 365 includes a feature that allows users to buy licenses for certain products. Individual users receive the bills for these licenses, but the data generated with the licenses remains part of the organization and is subject to all the organization's retention, security, and eDiscovery policies. Despite this, many administrators don't want users buying licenses outside of the normal IT controls. You can prevent them from doing so, but it requires PowerShell.

Cmdlets in the *MSCommerCe* PowerShell module control the ability of users to make self-service purchases in a tenant. To make changes, download the latest version of the module from [the PowerShell gallery](#), and then log into PowerShell using a global administrator or billing administrator account. You can then install the module and connect to the *MSCommerCe* endpoint:

```
[PS] C:\> Install-Module -Name MSCommerCe
Import-Module MSCommerCe
Connect-MSCommerCe
```

Run the *Get-MSCommerCeProductPolicies* cmdlet to see what products are available for self-purchase:

```
[PS] C:\> Get-MSCommerCeProductPolicies -PolicyId AllowSelfServicePurchase
```

ProductName	ProductId	PolicyId	PolicyValue
Project Plan 3	CFQ7TTC0KXNC	AllowSelfServicePurchase	Disabled
Visio Plan 1	CFQ7TTC0KXN9	AllowSelfServicePurchase	Disabled
Project Plan 1	CFQ7TTC0KXND	AllowSelfServicePurchase	Disabled
Power Apps	CFQ7TTC0KPOP	AllowSelfServicePurchase	Disabled
Power BI Pro	CFQ7TTC0L3PB	AllowSelfServicePurchase	Disabled
Power Automate	CFQ7TTC0KPON	AllowSelfServicePurchase	Disabled
Visio Plan 2	CFQ7TTC0KXN8	AllowSelfServicePurchase	Disabled

You can disable self-service selectively for any product by running the *Update-MSCommerCeProductPolicy* cmdlet, or, as in this example, disable self-service purchases for all products:

```
[PS] C:\> Get-MSCommerCeProductPolicies -PolicyId AllowSelfServicePurchase | ? {$_.PolicyValue -eq "Enabled" } | ForEach {Update-MSCommerCeProductPolicy -PolicyId AllowSelfServicePurchase -ProductId $_.ProductId -Enabled $False }
```

# Chapter 24: What don't you know you don't know about Office 365 and Azure AD?

*Matthew Vinton, Jeff Shahan, Bryan Patton, and Sofya Serna Perez*

For over 20 years, Quest Software has added value to the Microsoft platform. We offer a wide range of solutions that complement the functionality that Microsoft provides, from streamlining the process of modernizing to newer platforms like Office 365 to making it far easier to manage and secure an on-premises, cloud or hybrid environment. Here, we focus on just four of the challenges you need to be prepared for as you move to the Microsoft cloud:

- **Azure Active Directory recovery:** If you have a cloud-only IT infrastructure, you might think that the native Azure AD and Office 365 tools are sufficient to ensure fast recovery of any accounts, groups or attributes that might be accidentally or deliberately deleted. If you have a hybrid IT infrastructure, you might think that native tools combined with your on-premises backup and recovery solution are sufficient. But the reality is far more complex. In the first section of this chapter, we'll explain what native tools and on-prem solutions do and do not cover, and how to get the enterprise-level recovery capabilities you need.
- **Tenant-to-tenant migration:** With the rapid adoption of Office 365 around the world and the explosion in merger and acquisition (M&A) activity in recent years, more and more IT pros face the challenge of merging two or more tenants, often on extremely tight timelines. There are no native tools to help, and the process is even more complicated if any of the entities involved in the M&A has a hybrid AD environment. In the second part of this chapter, we'll review the challenges involved in tenant-to-tenant migrations and explain how Quest solutions can help ensure a secure and successful project.
- **Hybrid auditing:** With Microsoft 365, organizations get multiple advancements in auditing, including capable searching and alerting over a unified audit stream of normalized data. However, many IT pros still face a steep challenge: effectively auditing the on-premises Active Directory being synched to their tenant and accurately correlating that on-prem data with the Microsoft 365 audit stream to get a unified view of the IT ecosystem. In the third part of this chapter, we'll review why overcoming this challenge is difficult with native tools, why SIEM solutions also fall short, and how Quest solutions deliver the hybrid auditing you need for strong security.
- **Group management:** While Microsoft has greatly improved the capabilities of Microsoft 365 Groups, most IT professionals also oversee hundreds or even thousands of groups originating from their on-premises Active Directory. Unfortunately, there are no native tools to help you manage your on-premises AD groups and cloud groups together in a consistent manner. Instead, you're stuck addressing them separately with different tools. This fractured approach makes it difficult to avoid group sprawl and authorization creep, leaving your organization vulnerable to security risk. In the final section of this chapter, we'll discuss the challenges of managing groups across a hybrid environment and explain how Quest can simplify group management, regardless of where the group originates.



# Azure Active Directory Recovery

As discussed earlier in this book, Microsoft offers several strategies for backing up and recovering data in Office 365 that might be changed or lost, from database availability groups with lagged copies for Exchange Online to true backups for SharePoint Online. However, you share responsibility for backup and recovery with Microsoft, as detailed in your service agreement. While Microsoft provides tools that help you recover some items, you should consider investing in software like Quest On Demand Recovery to get the comprehensive recovery capabilities you need to ensure productivity and business continuity.

## Why Azure Active Directory Needs to Be Part of Your Recovery Strategy

Azure Active Directory is the authentication and authorization system for Microsoft 365. If a user accesses any Microsoft 365 services, they have an Azure Active Directory account. The core processes of authenticating a user and authorizing them to access the correct resources are quite similar in Azure AD and on-premises AD.

However, there are important differences. In particular, Azure AD does not have equivalents to Group Policy, discretionary access control lists (DACLS) or hierarchical objects such as organizational units. Moreover, Azure AD has objects and properties that do not exist in on-premises AD, including:

- **Roles:** Roles enable users to perform specific tasks such as adding or changing users, resetting user passwords, managing user licenses, and managing domain names. Examples of roles include Application Developer, Cloud Application Administrator and Directory Reader.
- **Licenses:** A user cannot use a licensed service until their account has been assigned a license. If that license attribute is lost, the user cannot access those services until it is restored.
- **Multi-Factor Authentication (MFA) settings:** MFA settings control whether a user must provide an additional form of verification, such as a code from their cell phone or a fingerprint scan, during sign-in. MFA is deployed using Conditional Access policies.
- **Conditional Access policies:** These policies provide additional control over access by requiring MFA under certain conditions or enforcing other restrictions. For example, a policy might block or grant access from specific locations or require the use of organization-managed devices for certain applications.
- **Dynamic group definitions:** These complex, attribute-based rules automatically adjust a user's group membership. For example, if a user transfers from the Sales department to Marketing (and their Azure AD user object is updated accordingly), a dynamic group definition could automatically remove them from groups associated with Sales and add them to groups associated with Marketing.
- **Applications and service principals:** Applications and service principals define how non-Office 365 applications interact with Azure AD. They are used for external applications that need access to Azure AD data, and they are also used to provide single sign-on (SSO) for federated applications. Service principals rely on applications to act as a defining template for their configuration. Every service principal has an application somewhere, either in the same tenant as the service principal or in another tenant.

## What Can and Cannot Be Recovered from the Recycle Bin

Any of these Azure AD objects and properties could get deleted and need to be restored. To help, Microsoft offers the Azure AD Recycle Bin. However, the Recycle Bin was never intended to be an enterprise-level recovery solution, and it's critical to understand what it can and cannot do.

First, not all objects go through the Recycle Bin when they are deleted. Certain types of objects are “soft deleted,” which means they do get put into the Recycle Bin, but other objects are “hard deleted” — they are not put into the Recycle Bin and therefore cannot be recovered from it.

Azure AD objects that are soft deleted include:

- User and guest accounts.
- Microsoft 365 groups (including associated data such as properties, members, e-mail addresses, Exchange Online shared inbox and calendar, SharePoint Online team site and files, OneNote notebook, Planner, Teams, and Yammer group and group content).
- Azure AD applications.

Azure AD objects that are immediately hard deleted include:

- Security groups.
- Distribution groups.
- Service principals.
- Conditional access policies.
- Devices.

Second, soft-deleted objects remain in the Recycle Bin for only 30 days. After that, they are permanently deleted.

Third, many Azure AD objects have complex configurations or specific interactions with other systems. Those details are not captured by the Recycle Bin and cannot be restored from it.

Finally, the Recycle Bin is for deleted objects only. If an object has been changed rather than deleted, the Recycle Bin cannot help you restore the object to its previous state.

## How to Recover Soft-Deleted Objects from the Recycle Bin

You can recover soft-deleted users, Microsoft 365 groups and Azure AD applications from the Recycle Bin using PowerShell. You can also recover users and Microsoft 365 groups — but not Azure AD applications — using the Microsoft 365 admin center.

### How to Recover a Soft-Deleted User using PowerShell

The simplest way to recover a user using PowerShell is to use the older Microsoft Online cmdlets:

```
[PS] C:\> Connect-MsolService
Get-MsolUser -ReturnDeletedUsers
Restore-MsolUser -UserPrincipalName <UPN of deleted user>
```

### How to Recover a Soft-Deleted Microsoft 365 Group using PowerShell

To recover a soft deleted group, we will use the newer Graph-based PowerShell cmdlets:

```
[PS] C:\> Connect-AzureAD
Get-AzureADMSDeletedGroup
Restore-AzureADMSDeletedDirectoryObject -Id <ObjectID of deleted group>
```

### How to Recover a Soft-Deleted Azure AD Application using PowerShell

Just like with the Microsoft 365 group, we will again use the Graph-based PowerShell cmdlets. This procedure is like the one for recovering a group, but there is a specific cmdlet for executing application recovery.

```
[PS] C:\> Connect-AzureAD
Get-AzureADDeletedApplication
Restore-AzureADDeletedApplication -ObjectId <ObjectID of deleted application>
```

## How to Recover a Soft-Deleted User or Group using the GUI

To recover a soft-deleted user from the Microsoft 365 admin center, go to **Users \ Deleted users**, select the user and then click **Restore user**. To restore a soft-deleted group, go to **Groups \ Deleted Groups**, select the group and then click **Restore group**.

## Recovery in Hybrid Environments

Many organizations use a hybrid model where they synchronize their on-premises Active Directory to the cloud using Azure AD Connect. They often believe that the recovery strategy they have in place for their on-premises Active Directory will also cover Azure AD due to the synchronization, thereby making up for the limitations of the Azure AD Recycle Bin. The truth is more complex than that.

When a user or a group is deleted from the on-premises Active Directory (or excluded from Azure AD Connect synchronization), what happens to the associated hybrid cloud object will differ based on the type of object:

- **Hybrid user objects** will be soft-deleted from Azure AD. If the on-premises user object is recovered or brought back into synchronization scope, the hybrid user will be restored from the soft-deleted state. If the user object has aged out of the soft-delete status and become permanently deleted, Azure AD Connect will create a new hybrid cloud user object based on the on-premises user object.
- **Security groups** will immediately be hard deleted from Azure AD. If the on-premises group is recovered or brought back into synchronization scope, Azure AD Connect will create a new hybrid cloud security group based on the on-premises group.

When Azure AD Connect creates the new user object or security group, that object will lack the cloud-only attributes of the deleted object, such as roles, conditional access policies and licenses. This can have serious real-world consequences.

## Potential Disruptions to Users and Continuity of Service

Let's consider just two real-life examples of how relying on just the Azure AD Recycle Bin in concert with your on-prem solution can cause serious problems for your users — and your service continuity.

### Scenario 1: Security Group Deprovisioned

Suppose you have several Active Directory OUs with security groups that are being synched to Azure AD to grant certain users access to specific SharePoint Online sites. You also have a rule that whenever a security group exceeds its attestation period, it is moved into a deprovisioning OU to await final decommissioning.

Now suppose a security group that enables a key team to access critical data misses its attestation deadline and is therefore moved. As soon as Azure AD Connect detects that the group is no longer in scope, it will remove the security group in the cloud. As a result, users who were members of that group will lose their access to the sites they need and start logging tickets.

To correct the issue, you'll move the security group back into the proper OU, and Azure AD Connect will create a new security group in Azure AD. That security group will have the same name and the same members — but it will be assigned a new ID, so the process will not restore the users' access to the SharePoint Online data.

To restore access using native tools, you'll need to consult your documentation about which SharePoint Online sites the group needs to have access to (or scramble to figure it out), and then re-apply the new group to those sites. This process takes time and effort, adding to IT workload and hurting user productivity.

### Scenario 2: Conditional Access Policy Modified

Suppose you are using Azure AD Conditional Access policies to require an extra form of authentication if a user from outside the corporate network attempts to access a sensitive application and to block access entirely from

certain IP addresses. Then you discover that the authentication controls you want are no longer in place for that application. Where does the issue lie? Is the application no longer associated with the proper Conditional Access policy? Has the policy changed? Has the security group that controls who the Conditional Access policy applies to changed?

The Recycle Bin offers no benefit here, since it is for deleted objects only; improper modifications to objects are not stored in the Recycle Bin and therefore cannot be recovered from it. To ensure you can restore any object that might be improperly changed, you'll need complete, current documentation of the configuration of all your Azure AD objects — which is virtually impossible to create and maintain using manual methods.

## Going Beyond the Recycle Bin: True Enterprise-level Recovery

As we have seen, recovering data in Azure AD using native tools works well if the scenario fits two fairly rigid guidelines:

- You want to recover an Azure AD user, Microsoft 365 group or Azure AD application that was deleted (not modified).
- No more than 30 days have passed since the object was deleted.

For other scenarios, you need a backup and recovery solution like [Quest® On Demand Recovery](#).

### Quest On Demand Recovery

On Demand Recovery records the configuration of most objects in Azure AD and recreates them when you request their recovery. This enables the recovery of objects that cannot be restored from the Recycle Bin, including the following:

- Users that were deleted more than 30 days ago.
- Azure AD applications that were deleted more than 30 days ago.
- The configuration of Microsoft 365 groups deleted more than 30 days ago.
- Security groups.
- Distribution lists.
- Devices.
- Conditional access policies.
- Service principals.

Moreover, On Demand Recovery backs up and restores the individual properties of each of these objects, so you can easily and precisely revert many configuration changes as well as outright deletions. On Demand Recovery backs up Azure AD three times per day, enabling the recovery of many Azure AD objects to a specific point in time.

For hybrid environments, On Demand Recovery integrates with [Quest Recovery Manager for Active Directory](#), which protects your on-premises AD. If a portion of a given recovery needs to be completed on premises and a portion in the cloud, On Demand Recovery can orchestrate the entire recovery with a single click.

There are a few things that On Demand Recovery is not able to back up and therefore cannot recover. In particular, there is no way to access a password for a user or certificates for service principals; you will need to reset the password or upload the certificate again.

### Avoiding Disruptions to Users and Service Continuity

Let's consider how On Demand Recovery handles the problematic recovery scenarios discussed earlier.

#### **Scenario 1: Security Group Deprovisioned**

When the users in the security group that was deprovisioned lose their access to the SharePoint Online site they need and start logging tickets, you can get them back to work quickly and easily with On Demand Recovery.

Log into On Demand Recovery and unpack a backup made before the OU was moved. On Demand Recovery clearly shows the changes that can be recovered (Figure 24-1). To restore the group that was deleted from Azure AD, simply choose the group and click **Restore**.

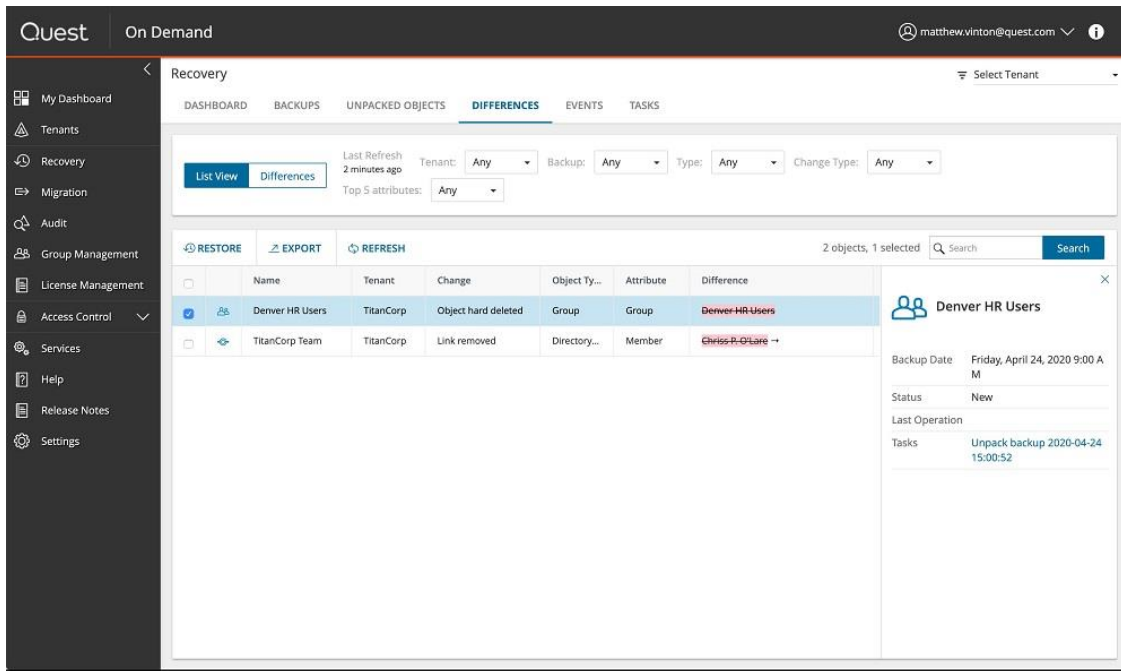


Figure 24-1: On-Demand Recovery enables the recovery of a wide range of deleted objects in a single click

On Demand Recovery performs the following actions to make the group whole again:

1. Since the Azure AD group is based on an on-premises AD group that needs to be reverted to an earlier state, On Demand Recovery dispatches a command to the on-premises instance of Recovery Manager for Active Directory, telling it to restore the AD group to a state prior to the recovery point, where it is in its normal OU.
2. On Demand Recovery then initiates a delta sync of Azure AD Connect so that the group is re-created in the cloud. As noted earlier, even though the group looks the same as it did before, it has a new object ID, so it won't be recognized by any of the services that need to use it to secure information or provide access to application.
3. On Demand Recovery also backs up many locations where security groups are used, such as service principals, SharePoint sites and conditional access policies. Therefore, the solution configures the SharePoint site to use the new group to secure access to the data, restoring access to the end users.

Remember, all this entire process happens automatically; all you need to do is click **Restore**.

### **Scenario 2: Conditional Access Policy Modified**

On Demand Recovery can also help you out when you discover that the expected authentication controls are no longer in place for a specific application. Log into On Demand Recovery and unpack a backup made prior to the controls being missing. The Differences tab (Figure 24-2) shows the differences between the live Azure AD and the backup. You can clearly see that the Conditional Access policy was changed. To revert the change, simply choose it and click **Restore**.

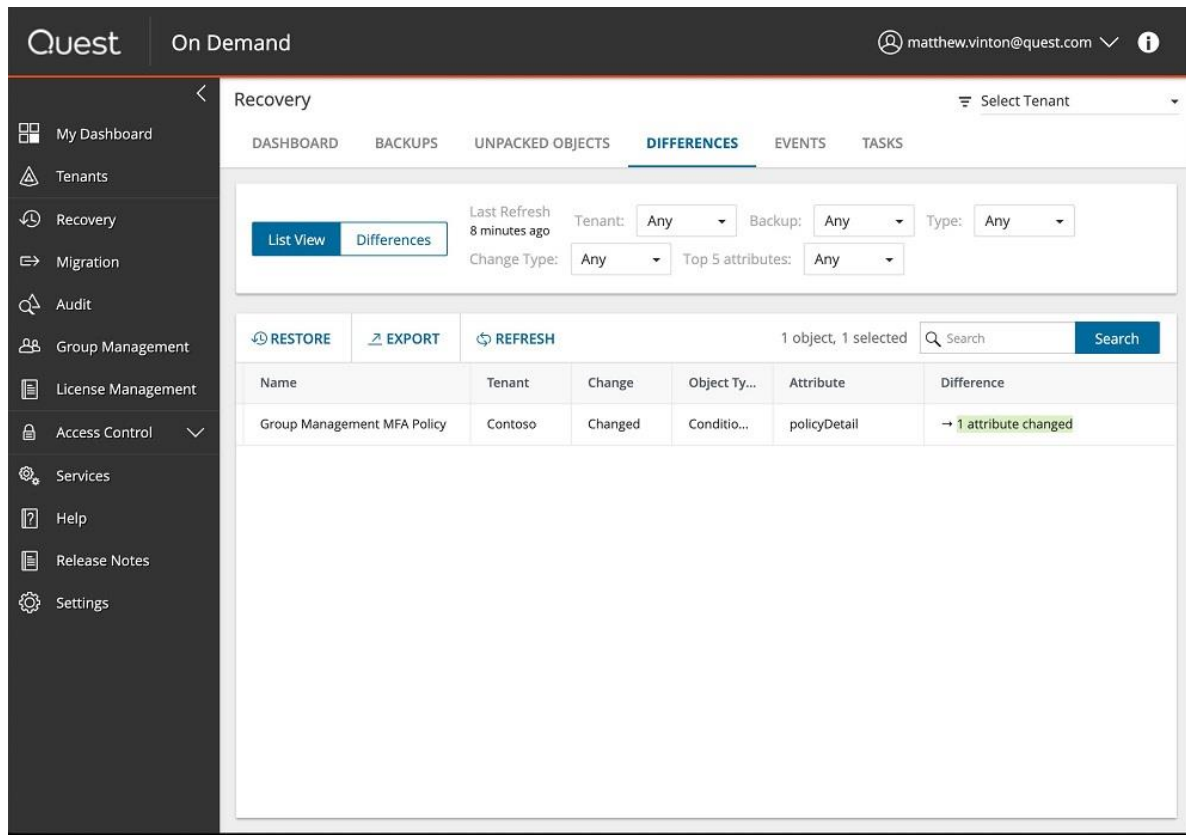


Figure 24-2: On-Demand Recovery enables reversion of unwanted changes to an object's attributes

## Summary – Azure Active Directory Recovery

Azure Active Directory provides authentication and authorization for all the critical Office 365 resources your organization depends upon every day, so it needs to be fully accounted for in your backup and recovery strategy. Native tools were never intended to be a comprehensive recovery solution, and even the best on-prem solution cannot protect the cloud-only objects and attributes it never sees. Be sure you have true enterprise-level recovery capabilities across your entire IT environment.

## Tenant-to-Tenant Migration

Now let us shift focus to another key challenge of Office 365 and Azure AD deployments: migrations. For more than 20 years, Quest has been a leader in IT migration. As Microsoft turned its attention to the cloud, Quest continued to fill the gaps left behind, enabling customers to migrate to the Office 365 platform.

As adoption of the platform increased, a new need came to the forefront: migrating users and content from one tenant to another. Indeed, the recent dramatic increase in mergers and acquisitions has made tenant-to-tenant migrations a top concern for organizations around the world. Since there are no native tools to facilitate this process, Quest again stepped up to help. Here, we detail the technical challenges involved in a tenant-to-tenant migration and explain how Quest solutions ensure a quick and automated project.

### Overview

In any migration, the business requirements present two competing forces: technical requirements and political requirements. These forces drive your decisions and will influence your migration process. The political forces will drive your technical challenges, and the technical challenges will push back on the political requirements. This section will focus on the technical challenges and point to where they can push back on the political

requirements. For example, if a larger entity buys a smaller entity, the political requirement will be to maintain the branding for the UPN and SMTP Reply-To addressing. Unless the entity is very small and can be migrated in a big-bang migration with an outage over a weekend, you are going to need to push back in order to execute the migration in a phased process that minimizes risk and impact to the end users.

## Technical Challenges

### Timelines

Many migration projects have a target end date handed down by the business, often driven by the need to present a unified branding to the external world as soon as possible. But these imposed deadlines often fail to consider factors such as the volume of data and workloads involved and the need to execute the migration without hurting user productivity. For example, a big-bang migration over a weekend is feasible only if the entity being acquired is very small and some amount of downtime is acceptable. Most of the time, migrations need to be executed in a phased process that minimizes risk and impact to the users and service continuity.

Tight migration timelines are especially an issue for SharePoint Online data and Exchange Online archive mailboxes:

- **SharePoint Online (and OneDrive for Business) Data:** Migration tools cannot directly migrate data into SharePoint Online without being throttled to a snail's pace to maintain service quality. Specifically, when SharePoint data is moved using the Microsoft migration API, the data is stored temporarily in Azure storage and then moved to SharePoint Online. This process is transparent to both users and the migration tool. While a migration tool could move data from the source into the migration pipeline, the process for publishing the data into SharePoint Online is a shared process and runs at its own pace. Execution time depends on the type of content and metadata being migrated. Large files with minimal metadata will migrate faster than small files with heavy metadata.
- **Exchange Online Archive Mailboxes:** Exchange Online allows for an unlimited size of the archive mailbox for some licensing plans. Initially, the archive mailbox is provisioned for 100GB. Today there is no automated way to force the archive mailboxes to increase the available storage. It can take up to 30 days for the additional storage space to be provisioned. There is currently no way to migrate an archive mailbox larger 100GB.

### Distributed Migration Teams

In larger organizations, the migration team will be segregated in ways that can present technical and political challenges. For example, the on-prem Active Directory to on-prem Active Directory migration workstream will be handled by the identity management team; workstation migration and OneDrive for Business will be handled by the desktop team; Exchange Online will be handled by the messaging team; and SharePoint Online will be handled by the SharePoint team. This distribution of work will require project teams to communicate and work together effectively, since one team's tasks will need to be completed to allow another team to start their work.

### Azure AD Connect (or Other On-Prem to Tenant Sync Engine)

Most tenant-to-tenant migrations today involve at least one hybrid environment that uses Azure AD Connect or some other synchronization engine like Microsoft Identity Manager. This presents a technical challenge. Since nearly all an object's attributes are mastered in the on-prem Active Directory, the object should be created on premises and synchronized to the tenant prior to migrating the workload. It could be created by your migration tool and matched to the on-prem object after the workload migration, but there are fewer moving parts to deal with if you complete the directory object migration before starting the workload migration.

## Federation

Since all objects authenticated via federation are authenticated back to the authoritative repository, they must be created in those repositories prior to their creation in Azure AD. This will present challenges for distributed project teams working independently. It will also impact forest trusts and UPN suffix routing during a DNS domain move between AD domains/forests and tenants.

## DNS Domains

One common political requirement in M&As is maintain the current login and email addresses, since they reflect the current branding of the organization being migrated. It is common for the primary SMTP address to be the same as the login ID. However, there is only one technical option for meeting this political requirement: a big-bang migration. This is because a DNS domain can be registered in only a single tenant at time.

If a big-bang migration is not an option, you are going to have to reset the political requirements based on what is technically possible, which is to maintain the public-facing SMTP address during the migration while using a temporary login ID, and at the end of the migration, move the DNS domain and change the login ID and SMTP addresses.

Quest On Demand Migration offers a Domain Coexistence add-on that can maintain brand or rebrand SMTP addressing during a phased migration.

## DNS Domain Move

The DNS domain must be completely removed from one tenant before it can be registered in another tenant. This presents a window during which inbound SMTP messages could bounce. Microsoft states that the removal of DNS domain from a tenant can take anywhere from five minutes to a day to execute.

To minimize the time this process will take for the domain to be removed, there are two options:

1. Delete all the objects from the tenant before removing the DNS domain from the tenant.
2. Remove all references to the domain in the attributes of objects before removing the domain from the tenant.

If Azure AD Connect or another sync engine is in use, these changes must be made on premises and allowed to replicate to the cloud.

There are several ways to eliminate the worry about bounced messages during this process. If the MX record points host that is unavailable, each message will be queued on the sending SMTP server until it expires.

## Domain Move Process

This is a rough outline of how to move a DNS domain; things like an active on-prem AD-to-AD migration or any type of directory synchronization running between the source tenant's AD and target tenant's AD are not considered. This process will remove (DELETE) all objects in the scope of the domain to be migrated from the source tenant.

1. Change the TTL for the current MX records to 1 or 2 minutes so the changes will quickly propagate when the process is complete.
2. Wait for the old TLL time plus 10 mins to ensure that TTL is applied to all cached sending servers.
3. Change the MX record to point to a host that does not exist.
4. Wait for the current TLL time plus 10 mins to ensure that all sending servers are queueing the messages.
5. Suspend the source Azure AD Connect from synchronizing.
6. Extract all source attribute values for user and group objects:
  - a. ProxyAddresses
  - b. Mail
  - c. UserPrincipalName



- d. msRTCSIP-PrimaryUserAddress
7. Reconfigure the source Azure AD Connect to exclude all objects in scope of the migration from the sync scope.
8. Back up the source tenant using a solution such as Quest On Demand Recovery, so that you can roll back from the soft and hard deletes in the next step and ensure that the proper entitlements are restored.
9. Execute a full replication cycle to remove all objects from the tenant.
10. Remove the DNS domain from the source tenant.
11. If a trust needs to be maintained, remove the UPN suffix from the forest.
12. Stop any directory sync running between the source AD and the target AD, since it could create issues when you make changes to the target objects as part of this procedure.
13. Suspend the target Azure AD Connect from synchronizing.
14. Add the source UPN suffix as a UPN suffix in the target forest.
15. Using the attributes extracted from the source objects, write the following to the target objects:
  - a. ProxyAddresses
  - b. Reply-to address
  - c. UserPrincipalName
  - d. msRTCSIP-PrimaryUserAddress
16. Set the *TargetAddress* attribute to **@%TargetTenant%.mail.onmicrosoft.com**.  
Note: The format of the target address attribute is "SMTP:LocalPart@domain".
17. Add the DNS domain to the target tenant.
18. Enable Azure AD Connect.
19. After the sync process is completed, validate that each ProxyAddress has been properly populated.
20. Change the MX and TTL to point back to Microsoft's shared service.

## Licensing

You cannot move licenses from one tenant to another, so before you begin your migration, you need secure the number and types licenses needed to cover the users involved, along with licenses for any additional services and workloads they will require. Failure to do this has stopped several migration efforts dead in their tracks.

## What Cannot Be Migrated

There are Office 365 workloads and settings that cannot be migrated by any tool, including the following:

### Licensing Entitlements

While a migration tool might be able to migrate a workload and enable the license type associated with it (or enable specific entitlements inside the license type), there is no way to enumerate the source license entitlement for an object and apply it for the target users. This may be because most tenants do not have the same type of licensing available.

### Application Entitlements

Azure AD application entitlements are not migrated. For example, a migration cannot preserve who has an entitlement for Box.com or SalesForce.com. However, Quest On Demand Migration can enable users to retain access to applications from their target accounts until the service is established in the target tenant.

### Certain Microsoft Teams Content

Microsoft Teams is an open platform, so organizations can make use of additional services and entitlements. Unfortunately, not all these content types can be migrated. For example, you can add a Microsoft Stream channel or video as a tab in Microsoft Teams, but you cannot migrate Stream video content because there is no

API or other method to access the data for migration. As the Teams platform matures, migration of more of its content will likely be fully supported.

## Tenant-Joined Workstations

There is no way to move a tenant-joined workstation to another tenant. Most organizations still maintain an on-prem Active Directory, so is not usually an issue for migrations. However, we are starting to see more customers asking about this type of support.

## Settings Outside Exchange Online Mailboxes

Most data in Exchange Online mailboxes is migrated. However, settings outside of the mailboxes are not migrated, including:

- **Litigation holds** — Litigation holds on Exchange Online mailboxes should be enabled after the content is migrated and before users access the mailbox.
- **Retention policies** — These are system settings that need to be configured by system administrators.

## What Can Be Migrated

[Quest On Demand Migration](#) is a SaaS migration tool that enables you to easily and securely consolidate Office 365 tenants — while enabling users to continue to communicate and collaborate seamlessly throughout the migration. On Demand Migration supports the following crucial workloads:

### Accounts and Groups

On Demand Migration supports tenant-to-tenant migration including hybrid organizations. User and group objects from the source tenant can be created directly in the target tenant or merged into existing objects in the target tenant.

### Exchange Online Mailboxes

On Demand Migration can migrate both primary mailbox content and archive mailbox content. You can perform address rewriting to maintain a single domain brand for all or selected users. Moreover, it enables you to avoid user frustration and a flood of helpdesk tickets by migrating the following:

- Recoverable items (the hidden data for users on litigation hold).
- Outlook Web Access (OWA) rules.
- Mailbox delegation.
- Folder permissions.

Note that Exchange Online requires that a user be mailbox-enabled in order to be assigned access to an Exchange resource.

### Resources and Applications

On Demand Migration can enable users to maintain access to source resources and applications. This is done by creating a guest account and granting that guest account the same level of access to the resource. Supported resources include:

- SharePoint Online.
- Resource roles.
- Applications assigned.

### Outlook Profiles

During the mailbox migration, each user's Outlook profile must be reconfigured to point to the new tenant and mailbox. On Demand Migration makes this easy: You simply deploy a small utility called the Client Update Agent

to the users' machines, and it connects to your migration project to know when the profile needs to be updated and what settings to use. One good option is to send users a notification with a link to a self-extracting archive that executes the agent.

Note: Autodiscover can present a challenge based on the domain/forest membership of the client host. We have found that setting Outlook to exclude the service connection point lookup (`ExcludeScpLookup=1`) can help address this issue.

## OneDrive for Business

On Demand Migration can migrate users' OneDrive data to the new tenant, giving them continued access to their data. The process preserves document versions, as well as the metadata and permissions associated with each document. You can filter the data based on folder, type, date, or size to exclude unwanted content and increase migration speed.

## SharePoint Online

On Demand Migration migrates documents and other valuable data stored in classic and modern SharePoint sites, document libraries and lists. The process ensures continued access and compliance by preserving site and document permissions and metadata. For even more advanced SharePoint migration capabilities, Quest offers [Metalogix Essentials for Office 365](#). For instance, Essentials for Office 365 supports the migration of SharePoint workflows, SharePoint Designer workflows, Nintex workflows and Nintex forms.

## Public Folders

On Demand Migration makes it easy to preserve shared data stored in public folders, along with all associated permissions and delegation. You can filter data by folder, type and date to speed the migration.

## Microsoft Teams

Microsoft Teams has become the hub of communication and collaboration in Office 365, so it's vital to migrate it effectively. On Demand Migration can:

- Discover all Teams in the source tenant, including team members and groups.
- Provision Teams and channels in the target tenant.
- Preserve user and group access and permissions.
- Migrate files stored in Teams channels.
- Migrate all conversations.
- Rename or merge Teams during the migration.

## Considerations and Tips

Quest has been in the migration business for years, and this extensive exposure to real-world scenarios has enabled us to gain some insights and observations that can be helpful when planning your tenant-to-tenant migration.

### Check Microsoft throttling limits and regulations.

In order to ensure service availability and strong performance for users, Office 365 is designed to throttle discretionary workloads like migrations. When use of the Office 365 platform skyrocketed early in the Covid-19 pandemic, Microsoft tightened its throttling limits for all data migrations into Office 365 tenants during business hours. Evening and weekend hours for the region of your tenant will not be impacted.

Therefore, we recommend scheduling migration jobs to run outside of the business hours of your tenants. Staying on top of the latest throttling policies will help you avoid migration delays, as well as measure the benchmark speeds properly so you can forecast the timeline for the project more accurately.

### Back up your source tenant before shutting it down.

In an ideal world, completing your migration and conducting user acceptance testing to validate the results of the project would be enough to safely shut down the source tenant. In the real world, there are many scenarios where a safety net is needed: You migrated a large volume of data, you had limited resources to perform thorough post-migration verification, the timelines were tight. In addition, there is always the chance of human error. Any of these could result in some data being lost or corrupted during the migration, but the issue not being immediately apparent.

Therefore, it's wise to have a solid backup of the source tenant. Having an easy way to restore the data you need can save you a lot of time, as well as all the headaches that you'd face due to business disruptions.

### Take advantage of the opportunity to clean up.

Avoid taking stale and unused content into the new environment. Analyze your source environment to identify the content that can be archived during the migration or simply left behind.

Also look for ways to clean up permissions. For instance, if permissions were poorly maintained in SharePoint Online, they can be stripped off during the migration. If the migration is being driven by a merger, acquisition or divestiture, be sure to analyze and plan the desired permissions structure and address it during the migration process.

In many cases, organizational changes will also require changes to the taxonomy or information architecture of the tenant, so a copy-paste type of migration will not suffice. A robust migration solution will help you reorganize teams, departments and groups in the process of the migration.

Migrations are also great in addressing data sprawl. The Office 365 platform is all about empowering users. However, a lack of governance can lead to many types of data sprawl, including a proliferation of groups, teams, SharePoint Online sites. The problem with sprawl is not just the space that the data takes up and the effort involved in managing it; sprawl can put the security of the whole company at risk and make collaboration difficult. Doing a thorough inventory and identifying excessive objects prior to the migration and merging the entities can help minimize data sprawl across Office 365.

### Over-communicate with your users about the migration.

Business users are often very busy and tend to skim through communications sent by IT teams. Make sure your messages about the upcoming migration are very clear and include detailed directions for users. This will help reduce the volume of help desk tickets and calls after the migration. Communication is especially important in scenarios like One Drive for Business migrations, since users may feel their privacy has been violated when there is new data in their One Drive and they do not know where it came from.

### Consider the level of fidelity you need.

Determine how you want to migrate your content: Do you need all metadata and versions to be migrated, or would just the latest version of each document with basic metadata be sufficient? If retaining all metadata and versioning is required, be sure to choose a robust migration solution that supports full-fidelity migrations.

## Summary – Tenant to Tenant Migration

With M&A activity, organizational restructuring and other business changes on the rise, IT pros need to be prepared to deliver accurate, complete, and efficient tenant-to-tenant migrations. Balancing the technical challenges and political requirements is tough enough without struggling with complex and error-prone manual processes. A comprehensive solution that automates the migration process, time after time, can be a very wise investment.

# Practical Auditing in a Hybrid World

Chapter 20 of this book provides a detailed explanation of the native Microsoft 365 auditing capabilities. In particular, Microsoft 365 customers can leverage a unified audit stream with normalized data, and the Microsoft 365 Security and Compliance Center provides capable searching and alerting of that audit stream.

However, as we've seen, many organizations use a hybrid model in which their on-premises Active Directory is synched to the cloud using Azure AD Connect. Since on-prem Active Directory enjoys none of the auditing benefits of the Microsoft 365 platform, maintaining security in these hybrid environments is a steep challenge. Let's briefly review why effective hybrid auditing is crucial to security and then explore the options for getting the insight you need.

## How Active Directory Can Impact Azure Active Directory

A hybrid model, where identity information in Azure Active Directory is controlled by on-premises Active Directory, can open Azure Active Directory to security vulnerabilities. For example, here are three examples of attacks against Azure AD that use on-premises Active Directory.

**Example 1: An on-prem service account that is made hybrid can be used to compromise Azure AD.**

Applications in Azure Active Directory normally use special accounts called service principals to obtain privilege in the directory. Sometimes, however, the required access cannot be provided by a service principal (for example, because there is no API for accomplishing what the application needs that can be used with a service principal).

In those situations, organizations often have the application use a regular user account — either an Azure AD user account or a hybrid AD service account (an on-premises AD service account that's synchronized to Azure AD). Using an Azure AD user account is less than optimal: Because user accounts used by applications typically cannot use multifactor authentication, they are more vulnerable than properly protected administrative accounts. But using a hybrid account is even worse. While it might seem convenient, leveraging a hybrid service account opens Azure AD up to all the vulnerabilities that plague on-premises AD, including Kerberoasting, Golden Ticket generation and password spraying.

To mitigate this risk, it is important to have visibility into both Azure AD and Active Directory. You can use simple searches to alert you whenever a service account is made hybrid so you can prevent its use in the cloud. If there are already hybrid service accounts in use, it is critical to keep a close eye on the activity of those accounts in both directories.

**Example 2: A hybrid account that is made the owner of an Azure AD service principal could misuse that service principal.**

Azure AD service principals can be quite powerful, so it's essential to monitor them. A service principal can be assigned an owner, who can manage many aspects of it. There is no restriction on the owner who may be assigned to a service principal — so a hybrid account could be made the owner of an Azure AD service principal.

In that case, the hybrid account would gain a great deal of power in Azure AD. In particular, an owner may assign an application-specific password (sometimes known as a secret) or a certificate so that the service principal can log in. While a service principal couldn't be used to log into an administrative portal, it could be used to manipulate the directory through Azure AD PowerShell cmdlets.

To manage this risk, it is important to have clear visibility into both Azure Active Directory and Active Directory. Identify any service principal that has a hybrid user as its owner, and either remove that owner or configure

greater scrutiny of the activity of the service principal. In addition, alerts on any change to the ownership of powerful service principals should be created.

### Example 3: An on-prem password spraying attack can enable attackers to gain a foothold in Azure AD.

In a password spraying attack, hackers test multiple likely passwords across a long list of accounts, attempting to log in to any account to establish a foothold and begin reconnaissance. To avoid tripping account lockout policies, attackers usually try only a few passwords on each account at irregular and infrequent intervals. As a result, password spraying can be difficult to detect.

Azure AD has several features that help mitigate the risk of these attacks. Multifactor authentication can make an account password only part of what is needed to log on. In Azure AD Premium P1, weak passwords can be detected and blocked by Azure AD Password Protection. And in Azure AD Premium P2, password spray attacks can be spotted through anomaly detection.

But for most organizations, Azure AD in hybrid mode bypasses most of these mitigations. To make work more convenient for users, they employ Azure AD Connect, ADFS, Azure AD Pass-through Authentication or other tools that allow on-premises credentials to be used for access in the cloud. This architecture makes on-premises Active Directory a tempting target for password spray attacks. Unless Azure AD Password Protection is deployed on-premises, users will be able to set weak passwords. Moreover, any endpoint in Active Directory that is exposed to an attacker could be used to carry out password spraying attacks away from the detection capabilities of Azure AD.

Managing this risk involves not only auditing on-premises authentication activity but having the ability to analyze and correlate it on a wide-scale, statistical level in order to detect any increase in failed authentications across the entire domain, which is a tell-tale sign of a password spraying attack.

## The Trouble with Active Directory Auditing

To properly secure a hybrid Microsoft 365 tenant, therefore, administrators need visibility into Active Directory as well as Azure AD. The challenges with Active Directory auditing will be familiar to many administrators with a history in Active Directory administration, but here is a brief summary:

- **Active Directory audit logs are very decentralized.** To get an accurate picture of Active Directory activity, administrators must analyze the Security event log on each domain controller where auditing is enabled.
- **Active Directory audit logs have a poor signal to noise ratio.** Multiple events are sometimes generated for a single audited action. Events can often contain irrelevant or obfuscated information, such as GUIDs rather than recognizable object names.
- **Log data can be short-lived.** Directory auditing information is written into the Security event log, which is highly active and regularly overwritten.
- **Active Directory audit logs are incomplete.** Critical aspects of Active Directory, such as Group Policy, are either partially audited or not audited at all.

## Hybrid Directory Auditing Adds Another Wrinkle

Organizations with a hybrid Azure Active Directory face an additional challenge: correlation.

Because Active Directory can affect Azure Active Directory (and in limited circumstances, Azure AD can affect on-premises Active Directory), understanding the behavior of the complete system requires administrators to not only review and act on security signals coming from each platform, but to correlate those signals.

Typically, this task involves manually correlating activity from two different auditing systems. In a worst-case scenario, it requires looking directly into Active Directory domain controller logs in order to find audit data corresponding to the activity found in the Microsoft 365 audit stream.

## ADFS Introduces Even More Dependencies on Active Directory

Microsoft has been recommending that organizations shift away from Active Directory Federation Services (ADFS) for some time, as it leverages Active Directory to authenticate Azure Active Directory accounts, which has significant ramifications for both reliability and security.

But, for many organizations, migrating away from ADFS may take some time. While ADFS remains in place, the organization will have additional requirements for hybrid Azure Active Directory auditing, since both Active Directory changes and authentications need to be audited. Furthermore, relying on ADFS reduces the capability of Azure AD risk detection, which helps administrators identify and act on unusual or risky sign-in activity.

## Overcoming the Challenges of Hybrid Auditing

To tackle the challenges of auditing a hybrid Azure Active Directory environment, most organizations require more than the native capabilities. Many of them leverage a security information and event management (SIEM) solution or a dedicated auditing system like Quest Change Auditor for Active Directory.

### SIEM Solutions

SIEM tools, such as Splunk and Microsoft Azure Sentinel, are powerful systems designed to aggregate security signals from many different sources. They are often used to correlate on-premises Active Directory activity with activity in the Microsoft 365 platform, but they can also accept data from network devices or low-level endpoint auditing tools like Sysmon.

However, SIEM systems can be very expensive, as well as complex to configure and operate. In larger organizations, they are often managed exclusively by security teams — which means that AD administrators have little or no access to the data.

Additionally, SIEM systems aggregate and analyze security events, but they don't produce them. This makes them susceptible to the same noise and blind spots inherent in native Active Directory auditing.

### Hybrid Active Directory Auditing Systems

Another common approach to hybrid auditing is to use a purpose-built solution such as On Demand Audit from Quest. These systems produce independent audit information from on-premises Active Directory, ingest audit events from the Microsoft 365 unified audit log, and then normalize and combine both of those sources into a single audit stream.

For instance, On Demand Audit agents capture activity from Active Directory domain controllers that it is not possible to "see" natively, such as changes to Group Policy settings and DCSync style attacks. These events are collected in an on-premises server and then forwarded to the cloud-based On Demand Audit platform, where they are combined with information from one or more Microsoft 365 tenants. As a result, IT pros can include information from both on-premises and Microsoft 365 signals in a single search or alert, such as the following:

- A search that encompasses a user's activity both on-premises and in the cloud
- An alert that is triggered whenever an account with "service" somewhere in its name or description is made hybrid
- A search that reveals who modified a hybrid security group that controls SharePoint data and when the change was made

Some hybrid Active Directory auditing systems even include pre-defined indicators of critical or risky activity. For instance, On Demand Audit considers all the following events to be indicative of “critical activity”:

- A change to AD schema configuration
- An unusual increase in sign-in failures for Azure AD or on-premises Active Directory
- A change to the membership of a critical AD group
- A change to a critical Azure AD role

## Summary — Hybrid Auditing

To secure a hybrid tenant, IT administrators need to be able to effectively audit both Azure Active Directory and the attached on-premises Active Directory, and correlate and analyze the data from both environments. Options include native tools, SIEM products and third-party auditing solutions.

# Group Management

## Microsoft 365 Groups

Office 365 adds exciting new capabilities to and poses new challenges for group management. Just like Active Directory, Office 365 has both security and distribution groups. These groups can exist as cloud-only objects managed either through the Microsoft 365 admin center, the Azure AD Portal, or the Exchange Online admin center (and their corresponding PowerShell and API interfaces). But they can also be hybrid. Hybrid security and distribution groups can be used for all the same things as cloud-only groups can, but they are managed by traditional on-premises tools and APIs (Active Directory and Exchange, respectively).

In 2014, Office 365 introduced a new group type called Office 365 Groups (now Microsoft 365 Groups). These groups are cloud-only. A Microsoft 365 group cannot be hybrid. It can provide authorization to resources like a security group can, it can function as a distribution list like a distribution group can, and it can also connect resources such as a SharePoint team site and plan. In addition, Microsoft 365 Groups underpin Teams.

## The Pain of Groups

While groups are a foundational and very useful aspect of directory services, managing groups can be problematic. The fundamental problem with groups is well understood: Over time, the accuracy of both the membership of a group and the purpose of the group tend to drop. This can cause everything from annoyances, like having the Exchange Global Address List (GAL) populated with long list of defunct entries, to serious problems, such as the very real security risk of having users collecting group memberships for every new role they take on but never being removed from old ones.

## Successfully Managing Groups

The fundamental ingredients to making groups less painful are also well understood. There are three techniques that play a part in successful group management:

### Reporting on Group Usage

There are a wide variety of tools that provide some level of reporting on group utilization across an environment. These tools can find where a group is being used, such as to secure a folder in a file system or to secure access to a SharePoint list.

Being able to report on groups is important. Unfortunately, group reporting has limitations. No tool is able to report on everywhere a group could potentially be used. And even more fundamentally, group reporting is



mostly useful in the context of a group clean-up project as performed by IT. When the project is over, groups resume their drift toward inaccuracy.

## Dynamic Groups

The membership of dynamic groups is determined through rules. For example, a group might be defined to contain any user whose Department attribute is "HR" and whose Office attribute is "Denver". When the user's attributes change, their group membership changes automatically.

Dynamic groups are very useful for making group management less painful. Their main weakness is that they are appropriate only for certain types of groups; many groups simply cannot be defined in any other way than ad-hoc.

## Self-Service Group Management

The core concept of self-service group management is that management of groups is pushed outside of information technology and into the general organizational process.

End users can make a request to join a group, and owners of those groups can decide whether to allow the user to join. End users can request the creation of groups and can manage the membership of groups that they have created. A scheduled process regularly prompts group owners to certify the accuracy of the membership of each of their groups, as well as the need for the continued existence of the group itself. All these activities are done without having to have IT staff directly involved.

Taking a self-service approach to group management has its shortcomings. If group ownership is not well defined, then you can end up with orphaned and abandoned groups. Moreover, it is possible that group owners will not carry out their responsibilities properly, even when prompted to — the "rubber stamping" problem.

Ultimately, though, a combination of these three techniques is the best way to minimize group sprawl and atrophy.

## Native Microsoft 365 Group Management

The addition of Office 365 to an organization's technical stack often makes group management harder. Usually, their on-premises Active Directory groups still exist and still have to be managed. Some of those groups will likely become hybrid groups that act as the "source of truth" for the cloud. And, as explained earlier, Office 365 adds a new cloud-only group type (the Microsoft 365 group) that not only proliferates quickly but can directly contain the organization's data.

To assist with these challenges, Microsoft has introduced several new capabilities. Most of these capabilities are tied to specific Azure AD license tiers, which will be denoted below (accurate as of the date of writing).

If your organization is licensed for Azure AD Premium P1, Microsoft enables dynamic membership rules for security groups and Microsoft 365 groups. This lets the group membership be determined by combinations of user or device properties and kept up to date automatically. This option works just for cloud-only security and Microsoft 365 groups; distribution groups and hybrid groups need not apply.

Azure AD Premium P1 also adds group expiration policies, naming policies and self-service group membership for Microsoft 365 and cloud-only security groups. Just as with dynamic groups, this works only for cloud-based groups, and only for security and Microsoft 365 groups.

For organizations with Azure AD Premium P2, Microsoft adds one more additional capability: group attestation, which Microsoft calls "access review." Group owners are reminded on a regular basis to review the membership of the groups they manage for accuracy.

## Where Native Group Management Breaks Down

The native group management capabilities covers all three of the key techniques for controlling group pain: reporting, dynamic rule-based groups, and self-service and attestation. So, what's the hitch? The core problem is clear: The native capabilities cover only a portion of the group types that are going to be in nearly every organization's production environment.

Distribution groups, hybrid groups and on-prem-only groups need to be covered by a different solution.

## Quest On Demand Group Management

On Demand Group Management from Quest is designed to fill these functionality gaps. This software-as-a-service (SaaS) solution adds straightforward self-service group management and attestation for the full complement of groups in the Microsoft technology stack.

It delivers self-service and attestation for cloud-only security groups and Microsoft 365 groups, similar to the capabilities in Azure AD Premium P1 and P2. And on top of that, it provides support for cloud-only distribution groups, hybrid groups of all types, and even on-prem-only security and distribution groups. Here are some of the key features of this innovative solution:

### Group Lifecycle Management

On Demand Group Management is designed to manage the lifecycle of groups from inception to removal. Each step of group management can involve independent approval processes.

When a group is created On Demand Group Management prompts the requester to classify the group. This classification determines several aspects of the group:

- What its naming pattern needs to be
- What kind of group it is (on-premises, cloud-only, security, Microsoft 365, etc.)
- What its sensitivity label is
- Whether the group is available for self-service request
- Whether and how often it needs to be attested to

If the group is included in the list of groups that are available for users to request, then end users can request to join the group, request to manage the group, or request to leave the group.

And if the group category is configured to have an attestation or review period, users that are responsible for the group will be periodically asked two key questions: Is this group still necessary? Is the membership of this group still correct? When the group is deemed no longer necessary by its owners, it will be removed from whichever directory it exists in, ending its lifecycle.

### Support for All Kinds of Groups

On Demand Group Management treats all Microsoft groups in the same way. It interfaces with the management APIs for Azure AD, Exchange Online, Active Directory and on-premises Exchange, depending on the type of group.

Hybrid and on-prem-only group management is particularly sophisticated. In order to provide lifecycle management for these groups, ODGM utilizes an agent that is installed on an on-premises member server that interfaces with Active Directory and Exchange APIs.

### Simple Implementation

There are many reasons why an organization might not have solved its group management problem, but complexity is often a big one.

Therefore, On Demand Group Management is designed to be simple to implement. It is a SaaS solution, so there is no need to size and deploy a server platform to support it. And it also has sensible default configuration settings, which make it easier to start using the solution effectively.

## Summary – Group Management

Your groups are essential to your business — from strong security to effective collaboration and user productivity, so you need to have control over all them throughout their full lifecycle. In particular, be sure you have reporting, dynamic rule-based groups, and self-service and attestation for all your various groups, including the types not covered by native group management capabilities: distribution groups, hybrid groups and on-prem-only groups.

## About Quest

Master the challenges of managing Office 365 with Quest, your go-to vendor for moving, managing and securing Azure Active Directory, Exchange Online, OneDrive for Business, SharePoint Online and Teams. Quest delivers the most comprehensive set of Office 365 and hybrid management solutions, including solutions from [recently acquired Quadrotech](#), [Binary Tree](#) and [Metalogix](#). With Quest solutions, you can:

- Move all your workloads to Office 365 with little to no disruption to end users.
- Reduce the time, clicks and scripts needed to manage your Office 365 or hybrid environment.
- Extend your existing security and compliance framework to your ever-evolving Office 365 environment.

To learn more about Quest solutions for Office 365, visit [quest.com/solutions/office-365](https://quest.com/solutions/office-365).

# Appendix

This appendix holds interesting information about Microsoft 365 that fits better here rather than interrupting the flow of text in a chapter.

## Annualized Run Rate for the Microsoft Cloud

Table: A-1 lists the annualized run rate for the Microsoft Cloud business segment as reported in Microsoft quarterly results. The big leap in FY19 is accounted for [Microsoft's September 2018 announcement](#) that they would include results for LinkedIn in the commercial cloud segment on an ongoing basis (LinkedIn accounted for [\\$6.75 billion revenue](#) in Microsoft's FY19 results).

<b>Quarter results</b>	<b>Reported annualized revenue run rate (ARR) for the Microsoft Cloud</b>
FY15 Q3 (April 2015)	\$6.3 billion
FY15 Q4 (July 2015)	\$8.0 billion
FY16 Q4 (July 2016)	\$12.1 billion
FY17 Q1 (October 2016)	Over \$13 billion
FY17 Q2 (January 2017)	Over \$14 billion
FY17 Q3 (April 2017)	\$15.2 billion
FY17 Q4 (July 2017)	\$18.9 billion
FY18 Q1 (October 2017)	\$20.4 billion
FY18 Q2 (January 2018)	\$21.2 billion (based on \$5.3B revenue reported for the quarter)
FY18 Q3 (April 2018)	\$24 billion (\$6B revenue)
FY18 Q4 (July 2018)	\$27.6 billion (\$6.9B revenue)
FY19 Q1 (October 2018)	\$39 billion (\$9.77B revenue)
FY19 Q2 (January 2019)	\$40.4 billion (\$10.1B revenue)
FY19 Q3 (April 2019)	\$40.96 billion (\$10.24B revenue)
FY19 Q4 (July 2019)	\$44 billion (\$11B revenue)
FY20 Q1 (October 2019)	\$46.4 billion (\$11.6B revenue)
FY20 Q2 (January 2020)	\$50 billion (\$12.5B revenue)
FY20 Q3 (April 2020)	\$53.2 billion (\$13.3B revenue)
FY20 Q4 (July 2020)	\$57.2 billion (\$14.4B revenue)
FY21 Q1 (October 2020)	\$60.8 billion (\$15.2B revenue)
FY21 Q2 (January 2021)	\$66.8 billion (\$16.7B revenue)
FY21 Q3 (April 2021)	\$70.8 billion (\$17.7B revenue)
FY21 Q4 (July 2021)	\$78 billion (\$19.5B revenue)
FY22 Q1 (October 2021)	\$82.8 billion (\$20.7B revenue)
FY22 Q2 (January 2022)	\$88.4 billion (\$22.1B revenue)
FY22 Q3 (April 2022)	\$93.6 billion (\$23.4 B revenue)

Table: A-1: Annualized revenue run rate for the Microsoft Cloud

## Growth in Office 365 User Numbers

Table A-2 details the growth in Office 365 user numbers since November 2015. For several years Microsoft reported the number of monthly active users every six months during their April and October earnings. In

April 2020, Microsoft switched to reporting the number of paid seats instead, citing 258 million then. You could assume that a paid seat is an active seat, but that isn't always the case. In the October 2021 earnings call, Microsoft said that the number of paid commercial Office 365 seats grew 17% year-over-year (the same percentage cited in January and July 2021). If usage tracks paid seats, we can therefore estimate the number of active seats as shown below.

<b>Date</b>	<b>Microsoft number for monthly active Office 365 users</b>	<b>Microsoft number for paid seats</b>
November 2015	60 million	
April 2016	70 million	
October 2016	85 million	
April 2017	Over 100 million	
October 2017	120 million	
April 2018	135 million	
October 2018	155 million	
April 2019	180 million	
October 2019	200 million	
April 2020	230 million (estimated)	258 million
October 2020	245 million (estimated)	278 million (estimated)
April 2021	264.5 million (estimated)	296.7 million
July 2021	280 million (estimated)	315 million (estimated)
October 2021	290 million (estimated)	325 million (estimated)
April 2022	321 million (estimated)	345 million

Table A-2: Growth in Office 365 user numbers over time

*Note: Estimated numbers shown since April 2020 are based on long-term growth for Office 365 seats as reported by Microsoft when they share this data in quarterly results. The estimates use the growth rate per month since the last reported number.*

## Quarterly Performance Against SLA

Table A-3 details the performance against SLA since figures first became available in 2013. The highlighted figure is the most recent published quarterly result.

<b>Q1 2013</b>	<b>Q2 2013</b>	<b>Q3 2013</b>	<b>Q4 2013</b>	<b>Q1 2014</b>	<b>Q2 2014</b>	<b>Q3 2014</b>	<b>Q4 2014</b>
99.94%	99.97%	99.96%	99.98%	99.99%	99.95%	99.98%	99.99%
<b>Q1 2015</b>	<b>Q2 2015</b>	<b>Q3 2015</b>	<b>Q4 2015</b>	<b>Q1 2016</b>	<b>Q2 2016</b>	<b>Q3 2016</b>	<b>Q4 2016</b>
99.99%	99.95%	99.98%	99.98%	99.98%	99.98%	99.99%	99.99%
<b>Q1 2017</b>	<b>Q2 2017</b>	<b>Q3 2017</b>	<b>Q4 2017</b>	<b>Q1 2018</b>	<b>Q2 2018</b>	<b>Q3 2018</b>	<b>Q4 2018</b>
99.99%	99.97%	99.985%	99.988%	99.993%	99.98%	99.97%	99.98%
<b>Q1 2019</b>	<b>Q2 2019</b>	<b>Q3 2019</b>	<b>Q4 2019</b>	<b>Q1 2020</b>	<b>Q2 2020</b>	<b>Q3 2020</b>	<b>Q4 2020</b>
99.97%	99.97%	99.98%	99.98%	99.98%	99.99%	99.97%	99.97%
<b>Q1 2021</b>	<b>Q2 2021</b>	<b>Q3 2021</b>	<b>Q4 2021</b>	<b>Q1 2022</b>			
99.97%	99.98%	99.985%	99.976%	99.98%			

Table A-3: Office 365 SLA performance since 2013